

Acoustica 7.5

User Guide

Acon Digital AS

Acoustica 7.5 User Guide

© 2023 Acon Digital AS

All rights reserved. No parts of this work may be reproduced in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems - without the written permission of the publisher.

Products that are referred to in this document may be either trademarks and/or registered trademarks of the respective owners. The publisher and the author make no claim to these trademarks.

While every precaution has been taken in the preparation of this document, the publisher and the author assume no responsibility for errors or omissions, or for damages resulting from the use of information contained in this document or from the use of programs and source code that may accompany it. In no event shall the publisher and the author be liable for any loss of profit or any other commercial damage caused or alleged to have been caused directly or indirectly by this document.

Table of Contents

| | |
|---|-----------|
| Part I Introduction | 5 |
| 1 What's new in Acoustica 7.5 | 6 |
| 2 Purchase and Authorization | 6 |
| 3 Requirements | 8 |
| Part II Working with Digital Audio | 8 |
| 1 Sampling | 9 |
| 2 Quantization | 9 |
| 3 The Decibel Unit (dB) | 10 |
| 4 Visualization of Audio | 10 |
| Part III Basic Audio Editing | 11 |
| 1 The Acoustica Workspace | 11 |
| 2 Loading Audio Files | 12 |
| 3 Inserting Audio Files | 13 |
| 4 Saving Audio Files | 14 |
| 5 Playing a Recording | 14 |
| 6 Selecting Regions | 15 |
| 7 Selecting Channels | 16 |
| 8 Zooming and Scrolling | 18 |
| 9 Drag and Drop Editing | 20 |
| 10 Editing using the Clipboard | 21 |
| 11 Avoiding Clicks when Editing | 21 |
| 12 Drawing Waveforms | 22 |
| 13 Audio Scrubbing | 22 |
| 14 Labels and Regions | 24 |
| 15 Adding Loops for MIDI Samplers | 25 |
| 16 Caption Editing | 27 |
| 17 Automatic Editing Tools | 29 |
| 18 Saving and Loading Workspace Files | 30 |
| 19 Edit History Pane | 30 |
| Part IV Recording Audio | 32 |
| 1 Re-recording | 33 |
| 2 Timer Record | 34 |
| Part V Audio Analysis | 35 |

| | |
|---|---------------|
| 1 Time and Frequency Domains | 36 |
| 2 Analysis Tools | 37 |
| Spectrogram | 37 |
| Wavelet | 40 |
| Spectrum | 42 |
| Spectral Histogram | 44 |
| Statistics | 47 |
| 3 Realtime Analyzers | 48 |
| Level Meter | 48 |
| Loudness Meter | 50 |
| Spectrum Analyzer | 52 |
| Phase Correlation Meter | 52 |
| Time Display | 53 |
| Part VI Spectral Editing | 53 |
| 1 Selections in Time and Frequency | 54 |
| 2 The Spectrogram View Settings | 56 |
| 3 The Retouch Tool | 58 |
| Part VII Audio Processing | 60 |
| 1 Tools | 61 |
| Dynamics | 62 |
| Multiband Dynamics | 69 |
| Limit | 73 |
| Dither | 75 |
| Convert Sample Format | 77 |
| Insert Silence | 77 |
| Signal Generator | 77 |
| Remix | 79 |
| Equalize | 80 |
| Equalize Preferences | 88 |
| Equalize Phase Modes | 90 |
| Equalize Light | 92 |
| Phono Filter | 96 |
| Remove DC Offset | 97 |
| Rotate Phase | 97 |
| Invert Phase | 98 |
| 2 Volume | 98 |
| Mute | 99 |
| Normalize | 99 |
| Volume | 100 |
| Channel Mixer | 100 |
| Volume Curve | 101 |
| Fade | 102 |
| Quick Fades | 103 |
| 3 Effects | 103 |
| Verberate | 103 |
| Reverb (Standard Edition) | 107 |
| Convolve | 109 |
| Echo | 111 |

| | |
|-------------------------------------|------------|
| Chorus | 113 |
| Modulate | 116 |
| Time Stretch | 118 |
| Transpose | 119 |
| Harmonizer | 120 |
| Reverse | 121 |
| 4 Enhancement | 123 |
| Interpolate | 123 |
| DeClip | 123 |
| DeClick | 124 |
| DeHum | 126 |
| About De-noising | 128 |
| Analyze Noise | 129 |
| Automatic DeNoise | 129 |
| DeNoise | 129 |
| DeNoise Light | 133 |
| Extract:Dialogue | 135 |
| DeWind:Dialogue | 136 |
| DeRustle:Dialogue | 138 |
| DeBuzz:Dialogue | 140 |
| DePlosive:Dialogue | 142 |
| DeClick:Dialogue | 144 |
| DeEss:Dialogue | 146 |
| DeBird | 148 |
| Vitalize | 149 |
| 5 Using Audio Plug-Ins | 151 |
| The Plug-in Manager | 151 |
| Accessing the Plug-Ins | 153 |
| 6 The Processing Chain | 153 |

Part VIII Multitrack Sessions 155

| | |
|---|------------|
| 1 Creating A Multitrack Session | 156 |
| 2 Working with Tracks | 157 |
| 3 Adding Clips | 160 |
| 4 Looping and Stretching Clips | 161 |
| 5 Moving and Grouping Clips | 161 |
| 6 Crossfading Clips | 161 |
| 7 Clip Fades | 162 |
| 8 Editing Clips with Clip Editor | 162 |
| 9 Saving, Loading and Exporting | 163 |
| 10 Automation Curves | 163 |
| 11 Using Sends and the Auxiliary Bus | 165 |
| 12 Importing Stems from a Mix | 165 |

Part IX Working With Audio CDs 166

| | |
|--|------------|
| 1 Creating Audio CDs | 166 |
| 2 Importing Audio Tracks from CDs | 168 |

| | |
|---|------------|
| Part X Batch Processing | 169 |
| 1 The Source Page | 169 |
| 2 The Destination Page | 170 |
| Part XI Using The Cleaning Wizard | 171 |
| 1 Import | 171 |
| Record Audio | 172 |
| Import Files | 173 |
| 2 Repair | 174 |
| Track Splitting | 175 |
| Restoration | 175 |
| Further Editing and Processing | 176 |
| 3 Export | 176 |
| Burn a CD | 177 |
| Export to Audio Files | 178 |
| Part XII Customizing the Workspace | 178 |
| 1 Using Factory Layouts | 179 |
| 2 Docking Window Panes | 179 |
| 3 Loading and Saving Custom Layouts | 180 |
| Part XIII Preferences | 180 |
| 1 General | 180 |
| 2 Keyboard shortcuts | 181 |
| 3 Audio Device Settings | 182 |
| Part XIV Included Plug-ins | 184 |
| 1 The Acoustica ARA Plug-in | 185 |
| 2 The Pro Tools Transfer Tool | 188 |
| Index | 189 |

1 Introduction

Acoustica is a comprehensive audio editor that covers a whole range of use cases - recording, editing, restoration work, mixing and mastering. The intuitive user interface was designed with speed, accuracy and ease-of-use in mind. Acoustica runs on Windows as well as on Mac, and the Mac version runs natively on both Apple Silicon and Intel processors.

You can edit clips with ease in the clip editor or mix multiple tracks in multitrack sessions. Acoustica is available in a *Standard* and a *Premium Edition* and the latter also features a powerful spectral editing mode that allows precise restoration work since processing can be limited to isolated regions in time and frequency. You can create and edit CD Projects that let you assemble and burn Redbook compatible audio CDs within the application. The integrated batch processing tool automates processing and format conversion of multiple files.

A large collection of integrated high quality processing tools covers the most demanding needs in mixing, mastering and restoration work. Cutting edge processing tools based on deep learning are included, such as the Remix tool for stem separation, or the Premium Edition tools Extract:Dialogue. Both editions include high quality time stretching and transposing tools.

The Premium Edition is also shipped with a large collection of plug-ins that you can use in 3rd party host applications that support VST, VST3:

- **Restoration Suite 2** — DeClick 2, DeClip 2, DeHum 2 and DeNoise 2
- **Mastering Suite** — Dynamics, Multiband Dynamics, Limit, Equalize 2 and Dither
- Extract:Dialogue
- DeClick:Dialogue
- DePlosive:Dialogue
- DeWind:Dialogue
- DeRustle:Dialogue
- DeBuzz:Dialogue
- DeEss:Dialogue
- DeBird
- Verberate 2
- Phono Filter
- Vitalize
- Convolve
- ARA plug-in that offers the full functionality of Acoustica's clip editor in third party applications with ARA2 support
- Transfer — Pro Tools to Acoustica audio transfer plug-in

Acoustica 7.5 introduces speech recognition for automatic captioning and transcription, a new *Edit History Pane* with audit trail export (Premium Edition only), freehand waveform drawing tool, and much more.

A large range of real-time metering tools are offered, such as a spectrum analyzer, peak and true peak metering, phase correlation metering and EBU R 128 and ITU-R BS.1770 compliant metering. The dockable window panes let you set up the workspace according to own preferences. Acoustica Premium Edition 7 handles surround formats up to 7.1 both in the single track editor and in multitrack sessions.


1.1 What's new in Acoustica 7.5

Version 7.5 includes major new features as well as a large number of smaller usability improvements and bug fixes. Here are some of the most interesting new features:

- Automatic captioning and transcription based on automatic speech recognition
- *Edit History Pane* with audit trail export (Premium Edition only)
- New *DeEss:Dialogue* plug-in
- Updated *Remix* tool that can be inserted in the Processing Chain or the Batch Editor
- Updated *Extract:Dialogue* to version 1.5 (Premium Edition only)
- Draw waveform tool
- Search bar in keyboard shortcut and possibility to save and load keyboard shortcuts
- Support for *Broadcast Wave Metadata*, including time offsets
- Tab to transients in clip editor

1.2 Purchase and Authorization

You can use Acoustica in trial mode for up to 30 days free of charge. During the trial period, you can choose whether you want to evaluate the standard or premium edition. You will be greeted with the welcome window below when you open Acoustica while in trial mode:



The trial mode dialog appears when starting Acoustica prior to authorization.

To continue in trial mode, please click either "Try Acoustica Premium Edition 7" or "Try Acoustica Standard Edition 7". You can purchase a license using either of the *Purchase* buttons and follow the checkout procedure in your web browser. You will receive a license key per email after completing checkout. Please enter your name, company name (if applicable) and license key in the bottom part of the window to authorize Acoustica. The *Authorize* button gets activated as soon as you have entered a valid license key.

Upgrading from Standard to Premium Edition

If you have an authorized version of Acoustica Standard Edition and have purchased the upgrade to Acoustica Premium Edition, you will need to deauthorize the Standard Edition license first. Please open Acoustica Standard Edition, choose *Deauthorize...* from the *Help* menu and confirm that you wish to deauthorize. Now you can restart Acoustica and the trial mode dialogue where you can enter your new Acoustica Premium Edition license key appears again.

Authorizing the Plug-in Pack from Acoustica Premium Edition

Acoustica Premium Edition comes with a large collection of plug-ins for use in third party applications with VST, VST3, AAX or Audio Unit (Mac only) support. The plug-in run in demo when installed and the demo mode is fully functional with exception of short passages with muted audio output at irregular intervals. The most convenient way to authorize the plug-ins is to authorize Acoustica Premium Edition, which will automatically authorize all the included plug-ins.

1.3 Requirements

Before you install Acoustica Premium Edition, please make sure your computer fulfills the following requirements:

PC Version (Windows)

- Windows 10 / 11 - 64 Bit
- Intel Core i5 or AMD multi-core processor
- 1366 x 768 display resolution (1920 x 1080 or higher recommended)
- 4 GB RAM
- 10 GB free HD space

Macintosh Version (macOS)

- OS X 10.13 or later
- 4 GB RAM
- 10 GB free HD space

2 Working with Digital Audio

Before audio can be edited on computers it must be digitized. The output from most audio equipment such as tape recorders, microphones or record players is analog. Analog means that the audio signal is represented by an alternating electrical voltage. The voltage is analog to the air pressure changes in the air during the performance, hence the term analog signals. An audio interface connected to your computer is needed to convert the constantly changing electrical voltage to a stream of numbers at fixed rate intervals. This process is done in two steps called sampling and quantization.


2.1 Sampling

The conversion from a continuously changing measure to a series of measured values at discrete time instances is called *sampling*. The rate at which the sampling is done, is along with the quantization depth the most important quality factor of digital recording equipment. You will not be able to record the highest audible frequencies if it is set too low. A CD quality recording is measured at a rate of 44 100 samples per second. We say that the sampling rate (or sampling frequency) is 44 100 Hertz (or short Hz).

In fact, all frequencies above half the sampling frequency, which is known as the Nyquist frequency, are substituted by frequencies below the Nyquist frequency unless the audio input is filtered. This effect is called *aliasing*. To avoid aliasing a sampling system contains of a low pass filter which ideally filters out all frequencies above the Nyquist frequency and leaves all frequencies below unaffected. In the case of CD audio, the highest frequency that can theoretically be recorded is 22 050 Hz.

2.2 Quantization

After measuring an analog input signal at fixed time intervals we have a stream of samples. The samples exist in terms of a voltage measured at a certain point in time. The voltage can usually be one of an infinite number of possible voltages within the legal voltage range. Computers cannot accurately describe every single one of the infinite number of possibilities, so it is necessary to divide the voltage range of interest into fixed sized regions. All voltages within one region are given a certain number during the quantization process. If we have a large number of regions which implies a larger number of discrete voltage levels, we can describe a voltage more accurately than with fewer voltage levels. The audio CD is quantized with 65536 voltage levels, which is the maximum number of levels possible to achieve with a binary number with 16 bits. Thus we say that the Audio CD has 16 bit resolution. Modern recording studios are frequently using 24 bit resolution or even higher during the mastering process.




The digital representation of a sine wave.

Acoustica works with 32 bit floating point resolution internally. That allows enough precision even through several processing steps and the audio signal won't be clipped or otherwise distorted before playback or saving the audio file with a lower resolution.

2.3 The Decibel Unit (dB)

When the volume of the recorded sound is changed, the degree of change is usually expressed in terms of decibels or short dB. This is a common unit in connection with audio. In Acoustica, decibel is used to express the extent of change relative to the original level.


Special for the decibel unit is that it is based on a logarithmic scale. Zero dB represents no change, whereas an increase of six dB represents a doubling of the signal amplitude. Reducing by six dB results in half the signal amplitude.



The decibel scale is chosen to suit the sensitivity curve of the human ear which have the same logarithmic property.

2.4 Visualization of Audio

The normal wave plot shown when making a recording in Acoustica is a time domain representation of the signal. When recording, Acoustica has taken samples of the signal at certain intervals, quantized them, and stored them as series of digitized values. The wave plot is the result of drawing these samples on the screen with the time evolving along the horizontal axis.



The waveform visualization of an audio signal in Acoustica.

The waveform visualization is very convenient in audio editing because it provides a good overview of the recording while allowing you to select time regions.

3 Basic Audio Editing

This chapter describes the basic audio editing capabilities of Acoustica such as loading and saving files as well as editing using the clipboard or drag and drop.

3.1 The Acoustica Workspace

About the Acoustica workspace

The workspace of Acoustica 7 is highly customizable. Most of the windows can be adjusted in size and you can also change the position of the individual windows. In the screen shot below, you will see the default workspace.



The Acoustica 7 workspace

The elements indicated with the red digits are explained below:

1. Toolbar

The toolbar section with short cut icons for commonly used commands.

2. Waveform Overview.

The portion of the waveform that is shown between the two yellow markers is shown below in the clip editor window. You can use the yellow markers to specify how much of the waveform should be shown in the clip editor window.

3. Clip Editor

The clip editor containing an audio file. The audio is visualized with a curve corresponding to the recorded audio.

4. **Level meter**

The level meters show the current output level during audio playback.

5. **Channel mode buttons**

You can enable/disable the channel selection by clicking the channel icon.

Clicking the arrow up and down icon will expand a channel for a more detailed view.

6. **Region List / Label List / Media File Browser**

Three different panes are docked together. The *Media File Browser* lets you browse and open audio files while *Region List* and *Label List* show lists of regions and labels in the recording. You can choose which one to show by clicking the corresponding tab.

7. **Processing Chain**

The Processing chain editor allows you to create a chain of processing tools and plugins. The chains can be saved including the processor settings for later use.

Furthermore, each element can easily be bypassed and the order of the elements changed using drag and drop. You can click the title of an element to open its graphical editor window.

8. **Loudness Meter**

The Loudness Meter provides you with three different loudness readings (Momentary, Short term and Integrated) and Loudness Range according to the EBU R 128 and ITU-R BS.1770 recommendations.

9. **Spectrum Analyzer**

This shows the frequency spectrum of the output audio.

10. **Correlation Meter**

This meter displays the phase relationship between the left and right channels of a stereo signal. If the meter goes below zero, this could be an indication that the mono compatibility is compromised.

3.2 **Loading Audio Files**

How to load an audio file in Acoustica 7

There are multiple ways to load an audio file into Acoustica 7

Option 1 (Main menu)

- Choose *Open...* from the *File* menu (or use Ctrl/⌘ + O)
- Browse to a folder where one or more audio files are located
- Select an audio file which you would like to load and press the *Open* button in the bottom right corner
- The audio file will now be loaded into Acoustica and will be shown in the audio editor window

Option 2 (Media File Browser)

- Go to the Media File Browser tab. If disabled, enable the Media File Browser tab by selecting *Show or hide file browser* from the View menu.
- In the Media File Browser, click on the arrows in the pull down menu and select the right hard drive.
- Browse to a folder where one or more audio files are located
- Double click on the audio file which you would like to load
- The audio file will now be loaded into Acoustica and will be shown in the clip editor

Option 3 (Drag & Drop)

- Select an audio file on your desktop or in your file explorer
- Drag this audio file to the Acoustica main window and drop it
- The audio file will now be loaded into Acoustica

3.3 Inserting Audio Files

Instead of opening an audio file as a separate clip in the clip editor, you can insert audio from a file into the active clip in two different ways:

Option 1 (Main menu)

- Move the cursor to the position in clip where you want to insert the audio from the file
- Choose *Insert Audio File...* from the *File* menu
- Browse to a folder where one or more audio files are located
- Select an audio file which you would like to insert and press the *Open* button in the bottom right corner

Option 2 (Media File Browser)

- Go to the Media File Browser tab. If disabled, enable the Media File Browser tab by selecting *Show or hide file browser* from the View menu.
- In the Media File Browser, click on the arrows in the pull down menu and select the right hard drive.
- Browse to a folder where one or more audio files are located

- Click on the audio file which you would like to insert, keep the mouse button pressed while you move the mouse to the position in the clip where you want to insert the audio and release the mouse button

3.4 Saving Audio Files

Quick Save

Acoustica allows you to quick save an audio file, which is a quick and easy way if you don't need to change the name or location of the file, or any of the audio settings. Choose *Save* from the *File* menu, press **Ctrl/⌘ + S** or click the save icon in the main toolbar

Save As a New File


If you wish to save the current audio with a different name, in a different folder or with different audio settings, this is of course also possible:

- Choose *Save as...* from the *File* menu or press **F12**
- Browse to a folder where you wish to save the audio file (optional)
- Enter a name for your audio file
- Choose the file format of your audio file from the **Save as type** drop-down list
- You can set the bit rate of the audio file by selecting the **Options** button
- Select the **Save** button to save the audio file

3.5 Playing a Recording

How to play a recording

There are several ways to play a recording. You can choose to play the whole recording, a selection of the recording, or toggle looped playback mode.



The transport toolbar with the hidden options visible

Play the Complete Recording

Click on the down arrow on the right side of the play button in the navigation toolbar to reveal the hidden options and select *Play all* (or press Shift + Spacebar). Alternatively, you can choose *Play all* under *Audio* on to the main menu.

Play Selection

Click on the down arrow on the right side of the play button in the navigation toolbar to reveal the hidden options and select *Play selection* (or press Spacebar). Alternatively, you can choose *Play* under *Audio* on to the main menu.

Play Selection without Filtering (Spectral Mode)

If you are in the spectral editing mode (see [Spectral Editing](#)^[53]), Acoustica plays the filtered audio per default. You can play the same time range without filtering. Click on the down arrow on the right side of the play button in the navigation toolbar to reveal the hidden options and select *Play selection without filtering* (or press Shift+Spacebar). Alternatively, you can choose *Play selection without filtering* under *Audio* on to the main menu.

The Play Looped Toggle

You can choose whether or not the selection should be played looped. To toggle the loop mode, click loop icon in the toolbar. Alternatively, you can choose *Play selection in loop* under *Audio* on to the main menu (or use Ctrl/⌘ + spacebar).

Note: To stop the playback, use the stop button in the navigation toolbar (or use spacebar). Alternatively, you can choose Stop under Audio on to the main menu.

3.6 Selecting Regions

How to select regions

Acoustica does all processing on the selected region and the selected channels only. The selected region is highlighted.

Select Region

- Click the beginning of the region you wish to select and keep the mouse button down.
- Move the mouse cursor to the end of the region you wish to select while keeping the mouse button down.
- Release the mouse button.
- The newly selected region should now be highlighted.

Modify selected Region

As shown in the screen shot below, you can modify the selected region by selecting the beginning or the end of the region and drag it to the desired location.



This demonstrates the ability to modify the selected region


3.7 Selecting Channels

How to Select Channels

There are multiple ways to select a specific channel in an audio file. How this works depends on the properties of the audio. It goes without saying that there is no specific channel select when you work on a single channel (mono) audio file. However, as soon as it has two or more channels you can select or deselect channels. Only selected channels will be affected by processing and be audible during playback.

Selecting the Left or Right Channel in Stereo Recording

In the screenshot below, you will notice we have used the color red, blue and green to indicate three different areas in the waveform overview.



Channel selection areas in stereo files

Selecting the Left Channel of a Stereo Recording

To edit or select a region in the left audio channel of a stereo file, you'll have to select the upper half of the first channel, which is indicated as a **red** area in the screenshot above. Alternatively, you can disable the right channel by using the enable/disable right channel button.


Selecting the Right Channel of a Stereo Recording

To edit or select a region in the right audio channel of a stereo file, you'll have to select the lower half of the second channel, which is indicated as a **green** area in the screenshot above. Alternatively, you can disable the left channel by using the enable/disable left channel button.

Selecting both Channels of a Stereo Recording

To select or edit a region in both channels simultaneously, you'll have to select between the upper half of the first channel and the lower half of the second channel, which is indicated as a **blue** area in the screenshot above. Make sure that the enable/disable button for both channels is enabled.

Select Channels in Multichannel Recordings



5.0 surround audio file where only the center channel is selected.

Selecting Channels in a Multichannel Audio File

To edit or select a region in one of the available audio channels of a multichannel audio file, press and hold Ctrl while selecting your region in the desired channel. You can move the mouse cursor to other channels to create a contiguous selection of channels.

Selecting Channels using the Channel Buttons

You can also toggle the selection status of a channel using the channel selection buttons to the left of the waveform. Click one of the channel buttons to toggle the selection of that specific channel.

3.8 Zooming and Scrolling

How to zoom and scroll in Acoustica 7

You can zoom in to get a more detailed view in an editing window. You can zoom either horizontally to view a smaller part of the recording or vertically to show a smaller amplitude range. In those cases where only a part of the recording is shown, a scroll bar is visible below the visualization of the wave form. You can use the scroll bar to view other parts of the recording.

Horizontal zoom

There are three ways to zoom horizontally:

Option 1 (mouse)

Use the mouse scroll wheel upwards to zoom in, or downwards to zoom out.

Option 2 (keyboard)

Press the arrow up key to zoom in, or the arrow down key to zoom out.

Option 3 (zoom in/out buttons)

Use the zoom-in or zoom-out buttons which are located underneath the waveform overview in the right corner (see screenshot below).

Vertical zoom

There are three ways to zoom vertically:

Option 1 (mouse)

Press and hold Ctrl and use the mouse scroll wheel upwards to zoom in, or downwards to zoom out.


Option 2 (keyboard)

Press and hold Ctrl and use the arrow up key to zoom in, or the arrow down key to zoom out.

Option 3 (zoom in/out buttons)

Use the zoom-in or zoom-out buttons which are located on the right side of the waveform overview, in the bottom right corner (see screenshot below).

Note: For the keyboard arrows to work, the waveform overview windows has to be selected



This screenshot shows the vertical and horizontally zoom-in and zoom-out buttons

3.9 Drag and Drop Editing

How to use drag and drop editing in Acoustica

You can move or copy regions using so-called drag and drop editing.

To move a region to a another location in the same editing window or to another editing window

- Select the region you wish to move.
- Click somewhere within the highlighted region and keep the mouse button down.

- Press the Shift key while moving the mouse cursor to the new start position and release the mouse button.

To insert a copy of a region using drag and drop

- Select the region you wish to copy.
- Click somewhere within the highlighted region and keep the mouse button down.
- Press the Ctrl key while moving the mouse cursor to the insert position and release the mouse button.

3.10 Editing using the Clipboard

How to edit using the clipboard

The clipboard concept provides a common way of editing documents for all Windows applications.

Standard Copy and Paste

You can edit your recordings by copying the selected region to the Acoustica clipboard using the *Copy* command from the *Edit* menu (or press Ctrl/⌘ + C) and paste the region into another location using the *Paste* command (or press Ctrl/⌘ + V). The *Paste Insert* command is equivalent to the normal paste command common in most Window applications.

Acoustica offers two additional ways of pasting:

- *Overwrite paste* (or press Ctrl/⌘ + Alt + V), which substitutes a selected region with the content of the clipboard.
- *Mix paste* (or press Ctrl/⌘ + Shift + V), which mixes the selected region with the content of the clipboard.

The command *Cut* command copies the selected region to the clipboard before deleting it from the source. For the purpose of deleting parts of the recording, use one of the following two commands:

- *Delete* (or press Delete), which deletes the selected region
- *Crop* (or press Ctrl/⌘ + Shift + C), which deletes everything but the selected region.

3.11 Avoiding Clicks when Editing

Clicks due to discontinuities can be problem when editing audio. You can avoid this in Acoustica by either making sure that you make cuts where the audio signal crosses zero, or by smoothing the signal at the splice points.

Expanding Selection to Nearest Zero Crossings

You can expand selections to the nearest zero crossings in Acoustica with the menu command *Edit > Selection > Expand to Zero Crossings*. Alternatively, you can let Acoustica

do this automatically after every selection change by activating *Automatically Expand to Zero Crossings* from the *Edit* menu.

Automatic Splice De-click

Another way to avoid clicks during editing is to activate the automatic splice-declick function in Acoustica using the menu command *Edit > Automatic Splice De-click*. When activated, Acoustica will use digital signal processing to eliminate clicks at every splice point after any editing operation.

3.12 Drawing Waveforms

You can manually draw waveforms to remove clicks or other unwanted elements. We generally recommend to use the [Interpolate](#)¹²³ tool from the *Enhancement* menu to manually remove clicks, but freehand drawing could be a last resort if the automatic interpolation doesn't work as expected. To enable the draw waveform mode, please click the pencil button in the toolbar as indicated below:



Click the Pencil icon to toggle the waveform drawing mode.

Waveform drawing is only possible when you have zoomed in on a very short time period so that the waveform is clearly visible. The mouse cursor changes from a prohibition sign to a pencil when you have zoomed in sufficiently.


How to Draw Waveforms Manually

- Select the Draw Waveform mode by clicking the Pencil icon in the toolbar
- Zoom in on the anomaly that you want to correct until the mouse cursor switches to a pencil
- Click the waveform and keep the mouse button pressed while drawing the waveform shape
- Release the mouse button when done.

3.13 Audio Scrubbing

Audio Scrubbing in Acoustica

Sometimes it is hard to find a specific part of a recording using the visualization only. Acoustica 7 allows you to listen to a specific part of the audio recording using the position marker, which doubles as a scrubbing tool.



The position marker which can be used as a scrubbing tool

How to use audio scrubbing

- Click and hold the position marker, as indicated in the screenshot above.
- Move your mouse to the left or right to activate the scrubbing capabilities of the position marker.

3.14 Labels and Regions

About Labels and Regions

You can simplify the editing process by inserting anchors such as labels and regions to your recording. Labels can be perfect for adding a notification on a specific point in time of the recording, like "Interview ends" or "Backing vocals start" for example. While labels have a specific point in time, Regions have a beginning and an end. This makes them perfect for describing which part of the recording contains the verse or chorus for example.



This screenshot shows an example of three regions and two labels

How to add a Label

- Move the cursor position to where you want the to insert a label
- Press the "L" key on your keyboard
- A label appears at the cursor position, indicated by a purple anchor

Note: You can also press the "L" key on your keyboard during playback, this will add the label at the location of your playback position.

How to add a Region

- Select the part of the recording where you want a region marker
- Press the "R" key on your keyboard
- A region appears at the selection, indicated by a transparent blue overlay between the beginning and the end of the region

Changing the anchor properties

It is possible to change the properties of the anchors, for example if you want to give a custom description, or to change the position of a label or region.

How to change the description of an anchor


- Click on the current anchor description inside the waveform overview window, this will open up a window where you can change the description
- Alternatively, you can go to the Region List or Label List window and double click on an anchor description to change the description

How to change the position of an anchor

- Click on an anchor (label or region) and drag it to the desired position

3.15 Adding Loops for MIDI Samplers

Audio files in the WAVE and AIFF formats can contain specialized loop information for MIDI Samplers and software samplers. Acoustica supports this standard and you can add loops and define the base note and fine tuning of the recorded sample. To add a sample loop, choose *Add Sampler Loop...* from the *Edit* menu or press "O". The *Add Sampler Loop* dialog box appears:



The add sampler loop helper window

You can specify either an infinite loop that is repeated until the MIDI note is released or a specific loop count. Acoustica also allows you to crossfade the ending and the beginning of the loop to avoid clicks at the loop edge. You can set the length of the cross fade regions relative to the loop length in percent and choose between constant amplitude and constant power cross fades. Constant amplitude is recommended for tonal sounds whereas constant power usually works better with noisy sounds. The cross fading is performed when you click the *Apply* buttons and you can undo the last cross fade by clicking the arrow back button.

The information in Global MIDI Sample Settings is not stored for each loop, but are global settings for the complete recording. You can edit these settings when adding loop, because loop and note with fine tuning information is usually required when creating

loops for MIDI samplers. However, you can also change these parameters in the *Tempo and Key* tab of the recording properties (*File > Edit Properties...*).

3.16 Caption Editing

Acoustica lets you add captions that you can export as subtitles for video production software or as text transcripts. Captions are edited in the *Caption Editor* that you show or hide by choosing *Show or Hide Caption Editor* from the *View* menu. Alternatively, you can choose a window pane layout optimized for caption editing by choosing *Caption Editing* from the *Factory Layouts* in the *View* menu.




The Caption Editor lets you edit captions that you can later export as subtitle files for video production or as transcribed text.

In addition to the caption text, captions can be associated with different *Actors* and each actor can be assigned a unique color.

Setting Up a Captioning Project

Before starting your captioning project, we recommend to add meta-data such as titles and the language. If you want to use automatic speech recognition engine, please make sure you select the correct language first.



Before you start your captioning project, you should select the language and add meta-data.

Adding Captions Manually


Adding captions is similar to how region markers are added:

- Select the part of the recording where you want a caption
- Press the "C" key on your keyboard
- The caption text field now receives keyboard focus and you can enter the caption text
- If you need to change the Actor, you can enter a new actor name or click the drop-down list to choose from one of the previously used actor names. You can quickly access the first 9 actors by pressing Ctrl and the digit inside the color indication for the actor
- Press the *Enter* key or click "Apply" to add the caption

Adding Captions Using Automatic Speech Recognition

Adding captions is similar to how region markers are added:

- Make sure you have selected the correct language in the caption editing properties.
- Select the part of the recording where you want a caption
- Click the magic wand button in the caption editor toolbar



You can click the Magic Wand Icon in the toolbar to use automatic speech recognition.

- The suggestion appears in the caption text field for your review
- If you need to change the Actor, you can enter a new actor name or click the drop-down list to choose from one of the previously used actor names. You can quickly access the first 9 actors by pressing Ctrl and the digit inside the color indication for the actor
- Press the *Enter* key or click "Apply" to add the caption

Modifying Captions


To modify the caption text or the selected actor of an existing caption:

- Click to select the caption you want to modify in the captions list
- Change the caption text or choose another actor
- Press the *Enter* key or click the *Apply* button to update the caption

If you want to modify the time region of the caption, you will need to unlock the caption by clicking the lock icon in the caption list first. When unlocked, you can click and drag the start or end point of the caption directly in the waveform display. We recommend to lock the time region again by clicking the same lock icon when done in order to avoid accidental changes.

Deleting Captions

You can delete captions by selecting the caption entry in the caption list and either click the *Delete* button or press the delete key. If you want to clear all caption, you can click the *New* icon in the caption editor's toolbar:



You can click the New Icon in the toolbar to clear all captions.

Exporting and Importing Captions


Captions are automatically saved and loaded as .captions files along with your audio files. You can also import captions from other caption files using the *Import* button in the caption editor toolbar. Similarly, you can export captions in a variety of output formats using the *Export* button. The following export formats are supported:

- Acoustica Caption Files (.captions) – the proprietary caption file format that Acoustica uses
- SubRip File (.srt) – a common caption file format which is accepted by most video editing software for video sub-titling
- SubRip File with Color Codes (.srt) – same as normal SubRip files, but with HTML-style color codes for colored captions
- Transcription File (.rtf) – a rich text format file containing the transcription generator based on the captions
- EDIUS Marker File (.xml) – For Marker export to the EDIUS Video Editing application
- Text file (.txt) – Text file with the transcription encoded as UTF8

3.17 Automatic Editing Tools

Acoustica offers convenient auto-editing tools that automates tedious editing tasks like removing unwanted silence. You can find the automatic editing tools under *Edit > Automatic Editing* where you can find the following tools:

- **Crop Automatically** detects and removes silence at the beginning and end of the recording
- **Remove Silence Automatically** detects silent regions in the selection and deletes them



The Automatic Silence Removal tool has an adjustable silence threshold and you can set the minimum duration and add padding to the remaining pauses.

- **Automatic Track Splitting** adds region markers for each track detected. These can easily be exported to separate files using the command *File > Save Regions to Files...*
- **Automatic DeNoise** automatically estimates the noise level of the selection and opens *DeNoise* with the estimated background noise profile.

3.18 Saving and Loading Workspace Files

Acoustica lets you save and load workspace files. The workspace files contain the following:

- All open documents (clips, multitrack sessions and CD projects)
- The Processing Chain
- The current window pane layout

To save a workspace file, please choose *File > Save Workspace...* and *File > Load Workspace...* to reopen a previously saved workspace file.

3.19 Edit History Pane

The *Edit History* pane keeps track of all the changes you have made to a clip in the clip editor.

| Processing Chain ✕ | | Edit History ✕ | |
|----------------------|--------------|----------------|--|
| Command | Begin | End | |
| Loaded "Cordoba.wav" | | | |
| Delete | 00:00:00:000 | 00:00:00:442 | |
| Apply Dynamics | 00:00:00:000 | 00:00:25:707 | |
| Apply Equalize 2 | 00:00:00:000 | 00:00:25:707 | |
| Apply Verberate 2 | 00:00:00:000 | 00:00:25:707 | |
| Apply Limit | 00:00:00:000 | 00:00:25:707 | |

The Edit History window pane with the list of all the changes made in the Clip Editor.

You can easily revert to a previous editing state by double clicking an entry in the history list. After reverting to a previous editing state, the yellow line will indicate the current state and possible redo states will appear in a gray color as in the screenshot below:

| Processing Chain ✕ | | Edit History ✕ | |
|----------------------|--------------|----------------|--|
| Command | Begin | End | |
| Loaded "Cordoba.wav" | | | |
| Delete | 00:00:00:000 | 00:00:00:442 | |
| Apply Dynamics | 00:00:00:000 | 00:00:25:707 | |
| Apply Equalize 2 | 00:00:00:000 | 00:00:25:707 | |
| Apply Verberate 2 | 00:00:00:000 | 00:00:25:707 | |
| Apply Limit | 00:00:00:000 | 00:00:25:707 | |

The current state is indicated with a yellow line and possible redo states are indicated with a dimmed color.


You can right click elements in the history list to reveal further options:

- **Switch to State** — Loads the clicked editing state (same as double clicking the element).
- **Recall Selection** — Recalls the selection state from the clicked editing state
- **Show Processor Settings** — If the editing step involved a processor, this will open the processor with the same settings
- **Export Audit Trail** (Premium Edition only) — Lets you export an audit trail in the HTML format that contains a detailed description of all the editing steps

4 Recording Audio

Please follow the steps below to record audio from audio equipment such as record players, tape decks or microphones through your sound card:


1. Make sure the audio source is properly connected to the audio input of your audio interface or computer.
2. Create a new an empty recording window by selecting *File > New...* or position the cursor where you want to insert recorded audio in an existing recording.
3. Select *Audio > Record...* or press Ctrl/Cmd+R.
4. If you are recording to an empty editing window Acoustica needs to know what sample format you wish to use (see [Working with Digital Audio](#)⁸ for more information). The following dialog box appears:



The sample format dialog box in Acoustica

Please choose the desired recording format and click the *OK* button.

5. The Recording dialog box now opens:




The Recording dialog box

Now make sure that the correct Input source (1) is selected. You can monitor the input level using the level meters (2) or by ears if the *Monitor* option is enabled (3). The level meter should never be in the red area in order to avoid clipping.

6. Click the record button (4) to start the recording.
7. When you are done recording, click the *Insert* (5) or *Overwrite* (6) button to either insert the recorded audio at the cursor position or overwrite existing content with the recorded audio.

4.1 Re-recording

Acoustica makes it easy to re-record a time region if the first take wasn't successful. Select the time range you want to re-record in the clip editor and open the recording dialog as described in [Recording Audio](#)^[32]. You can use the pre-roll to let Acoustica playback the last couple of seconds prior to the selected time range:




The pre-roll lets you instruct Acoustica to playback a couple of seconds of audio prior to the part you want to re-record.

With the pre-roll enabled, Acoustica plays the selected number of seconds from the audio prior to the audio you want to replace. The actual recording starts immediately after the pre-roll time has elapsed. This makes it much simpler for e.g. voice over artists to synchronize their performance. You should click the *Overwrite* button when the recording is done. If the recorded audio is longer than the selection to replace, Acoustica will ask you if you would like to crop the recorded audio to fit the selected time range.

4.2 Timer Record

The timer record feature allows you to start and stop recording after a certain period of time or depending on the presence of an input signal. To start timer record, click the button labelled *Timer* in the recording dialog. The following dialog box appears:



The Timer Record settings.

You can choose to start the recording immediately (after clicking the *Start timer record* button), after a specified period of time or when an input signal is present. The threshold value for the input signal detection can be defined using the *Silence threshold* slider at the bottom of the dialog.

The recording can also be stopped automatically, either after a certain period of silence or after a certain period of time.

5 Audio Analysis

This chapter covers Acoustica's audio analysis features. There are realtime analyzers that gives visual feedback on the audio that are played as well as tools that analyze audio selections in the clip editor.

5.1 Time and Frequency Domains


Most of the analysis tools in Acoustica are related to the time and frequency content of the recorded audio. The chapter [Working with Digital Audio](#)^[8] describes how sound is captured and converted to a series of samples. We call this the time domain, since each sample is part of a time series that represents the audio signal. Our perception of audio, however, is more like a combination of tones with different pitches and timbres along with noise. Spectrum analysis allows us to present audio in a different way and the signal is represented as sum of frequency components which is called the frequency domain.

All natural sounds can be described as an infinite sum of sine waves. The frequency of a sine wave is related to what we sense as pitch. Our ears are not able to hear frequencies above 20 kHz (a sine wave with 20 thousand completed wavelengths within one second), so the mentioned infinite sum turns into a finite sum which is possible to handle on a computer. The signal in the frequency domain is represented through the weight of each sine wave needed to recreate the signal, rather than the sampled values from the time series. These weights can be visualized in Acoustica using the [Spectrum](#)^[42] analysis tool or the real-time [Spectrum Analyzer](#)^[52]. The mathematical tool that converts a time series to the frequency domain is called the *Fourier transform* (the computer optimized version thereof is called *Fast Fourier Transform* or *FFT*).

Combining Time and Frequency

Now we have a tool for examining the frequency content (the spectrum) of our recording and we have the normal wave plot visualization for examining how our recording evolves over time. Is there a possibility to combine these features, in order to study how the frequency content evolves over time? Actually, Acoustica features two ways of displaying so called time-frequency plots. The spectrogram and the wavelet transform (based on the Morlet class of wavelets for the advanced reader).

The spectrogram is created by creating spectra of slices of audio at regular intervals and creating a two dimensional intensity map with the time on the horizontal and frequency on the vertical axis as depicted below:



Spectrogram analysis of percussion instruments.


The length of each slice of audio determines the frequency and time resolution. Time slices with a longer duration result in a better frequency resolution, but the time resolution gets poorer.

Analysis Windows

Some care is required when analyzing time slices as described above due to the fact that the Fourier transform is cyclic by nature and considers each time slice as an infinitely repeating sequence. To avoid discontinuities between the start and end of the time slices we apply an *analysis window*. There are several different analysis windows available in the literature. Acoustica uses a *Dolph Chebyshev* window that is well suited because the noise caused by the edge discontinuities can be reduced to a user specified level. The downside is that the frequency resolution decreases as the attenuation increases.

Wavelet Analysis

The wavelet analysis is similar to the spectrogram, but the duration of each time slice depends on the frequency so that higher frequencies get a better time resolution (at the cost of a poorer frequency resolution).




Wavelet analysis of percussion instruments. Notice the better time resolution at higher frequencies.

5.2 Analysis Tools

Acoustica features several different tools for audio analysis that you can find under *Analysis* in the main menu.


5.2.1 Spectrogram

To show a spectrogram as described in [Time and Frequency Domains](#)^[36], please choose the section in the clip editor that you wish to analyze. Then choose *Spectrogram...* from the *Analysis* menu. The following window appears after the analysis has been completed:



The spectrogram analysis window

You can change different aspects of the analysis by clicking the *Settings...* button:



The spectrogram settings

The different settings are described below:

- **Block size (samples)**

The duration in samples of the time slices that are used create the spectrogram (see [Time and Frequency Domains](#)^[36] for more information).

- **Vertical resolution (pixels)**

The number of pixels along the vertical axis.

- **Horizontal resolution (pixels)**

The number of pixels along the horizontal axis which corresponds to the number of time slices used to create the spectrogram.

- **Window attenuation (dB)**

The attenuation in the analysis window specified in Decibels. Please see [Time and Frequency Domains](#)^[36] for details.


- **Frequency scale**

You can choose between different scales for the frequency axis: *Linear*, *logarithmic* or *Mel scale*. The latter is optimized according to our perception of sound.

- **Minimum frequency (Hz)**
The lowest frequency to show in the spectrogram.
- **Maximum frequency (Hz)**
The highest frequency to show in the spectrogram.
- **Minimum amplitude (dB)**
The minimum amplitude to show in the spectrogram.
- **Maximum amplitude (dB)**
The maximum amplitude to show in the spectrogram.
- **Color palette**
You can choose among different color palettes according to your own liking.

5.2.2 Wavelet

To show a wavelet transform as described in [Time and Frequency Domains](#)^[36], please choose the section in the clip editor that you wish to analyze. Then choose *Wavelet...* from the *Analysis* menu. The following window appears after the analysis has been completed:



The wavelet analysis window

You can change different aspects of the analysis by clicking the *Settings...* button:

The screenshot shows the 'Wavelet Settings' dialog box. It contains several input fields and a dropdown menu. At the bottom, there are three buttons: 'OK', 'Cancel', and 'Apply'.

| Setting | Value |
|---------------------------------|----------------|
| Vertical resolution (pixels): | 400 |
| Horizontal resolution (pixels): | 800 |
| Coherence factor: | 20 |
| Minimum frequency (Hz): | 20 |
| Maximum frequency (Hz): | 22050 |
| Minimum amplitude (dB): | -96 |
| Maximum amplitude (dB): | 0 |
| Color palette: | Black to green |


The wavelet settings

The different settings are described below:

- **Vertical resolution (pixels)**
The number of pixels along the vertical axis.
- **Horizontal resolution (pixels)**
The number of pixels along the horizontal axis.
- **Coherence factor:**
The coherence factor adjusts the time versus frequency resolution. Higher coherence factors will improve frequency resolution at the cost of time resolution and increased processing time.
- **Minimum frequency (Hz)**
The lowest frequency to show in the wavelet analysis.
- **Maximum frequency (Hz)**
The highest frequency to show in the wavelet analysis.
- **Minimum amplitude (dB)**
The minimum amplitude to show in the wavelet analysis.
- **Maximum amplitude (dB)**
The maximum amplitude to show in the wavelet analysis.
- **Color palette**
You can choose among different color palettes according to your own liking.

5.2.3 Spectrum

To show a spectrum analysis, please choose the section in the clip editor that you wish to analyze or remove the selection and move the cursor to the time position you would like to analyze. Then choose *Spectrum...* from the *Analysis* menu. The following window appears after the analysis has been completed:



The spectrum analysis window

You can change different aspects of the analysis by clicking the *Settings...* button:

The figure shows a screenshot of a software window titled "Spectrum Settings". It contains a list of parameters for the spectrum analysis, each with a corresponding input field. The parameters are: Block size (samples) set to 4096, Resolution (bins) set to 1600, Window attenuation (dB) set to 120, Frequency scale set to Logarithmic, Minimum frequency (Hz) set to 20, Maximum frequency (Hz) set to 24000, Minimum amplitude (dB) set to -96, and Maximum amplitude (dB) set to 0. At the bottom, there are two buttons: "OK" and "Cancel".

The spectrum settings

The different settings are described below:

- **Block size (samples)**

The duration of the time slices that are used create the spectrum (see [Time and Frequency Domains](#)^[36] for more information) in samples. The spectra from the analyzed time sliced are averaged to create the shown spectrum.

- **Resolution (bins)**

The number of curve points along the horizontal axis.

- **Window attenuation (dB)**

The attenuation in the analysis window specified in Decibels. Please see [Time and Frequency Domains](#)^[36] for details.

- **Frequency scale**

You can choose between different scales for the frequency axis: Linear, logarithmic or *Mel scale*. The latter is optimized according to our perception of sound.

- **Minimum frequency (Hz)**

The lowest frequency to show in the spectrum.

- **Maximum frequency (Hz)**

The highest frequency to show in the spectrum.

- **Minimum amplitude (dB)**


The minimum amplitude to show in the spectrum.

- **Maximum amplitude (dB)**

The maximum amplitude to show in the spectrum.


5.2.4 Spectral Histogram

The spectral histogram shows the distribution of the amplitude levels per frequency component over time. Amplitude levels that occurred more frequently will have a brighter intensity. To show a spectral histogram, please choose the section in the clip editor that you wish to analyze. Then choose *Spectral Histogram...* from the *Analysis* menu. The following window appears after the analysis has been completed:



The spectral histogram analysis window

You can change different aspects of the analysis by clicking the *Settings...* button:



The spectral histogram settings

The different settings are described below:

- **Vertical resolution (pixels)**
The number of pixels along the vertical axis.
- **Horizontal resolution (pixels)**
The number of pixels along the horizontal axis.
- **Intensity (%)**
The brightness of the resulting image in percent.
- **Block size (samples)**
The duration in samples of the time slices that are used create the spectral histogram (see [Time and Frequency Domains](#)^[36] for more information).
- **Window attenuation (dB)**
The attenuation in the analysis window specified in Decibels. Please see [Time and Frequency Domains](#)^[36] for details.

- **Frequency scale**

You can choose between different scales for the frequency axis: *Linear*, *logarithmic* or *Mel scale*. The latter is optimized according to our perception of sound.

- **Minimum frequency (Hz)**

The lowest frequency to show in the spectral histogram.

- **Maximum frequency (Hz)**

The highest frequency to show in the spectral histogram.

- **Minimum amplitude (dB)**

The minimum amplitude to show in the spectral histogram.

- **Maximum amplitude (dB)**

The maximum amplitude to show in the spectral histogram.

- **Color palette**

You can choose among different color palettes according to your own liking.

5.2.5 Statistics

Acoustica can show different statistical parameters extracted from the current selection, such as peak values, RMS levels and loudness levels. To run the statistics analysis, please choose the section in the clip editor that you wish to analyze and choose *Show Statistics...* from the *Analysis* menu. The following window appears after the analysis has been completed:

Svanes (feat King Milo) - Up to You.wav - Statistics ✕

| Description | Left | Right |
|---------------------------------|--------------|--------------|
| Global RMS (AES) | -10.51 dB | -10.64 dB |
| Global RMS (Mathematical) | -13.52 dB | -13.65 dB |
| Maximum sample peak level | -0.70 dB | -0.70 dB |
| Maximum sample peak position | 00:02:26:125 | 00:01:22:499 |
| Minimum sample peak level | -0.70 dB | -0.70 dB |
| Minimum sample peak position | 00:03:23:059 | 00:01:03:922 |
| Maximum true peak level | -0.27 dB | 0.21 dB |
| Maximum true peak position | 00:01:28:990 | 00:03:03:118 |
| Minimum true peak level | -0.36 dB | -0.16 dB |
| Minimum true peak position | 00:03:12:675 | 00:01:50:120 |
| Maximum short-term RMS level | -6.97 dB | -6.69 dB |
| Maximum short-term RMS position | 00:02:28:725 | 00:01:59:775 |
| Average short-term RMS level | -14.48 dB | -14.59 dB |
| DC offset (average) | 0.033% | 0.033% |
| Number of zero crossings | 475258 | 489230 |
| Zero crossing rate | 1971.37 Hz | 2029.33 Hz |

| | |
|--|--------|
| Loudness (EBU R-128): | -11.81 |
| Loudness range (EBU R-128): | 8.02 |
| True peak (EBU R-128): | 0.21 |
| Maximum momentary loudness (EBU R-128): | -7.49 |
| Maximum short-term loudness (EBU R-128): | -9.19 |

OK

The Statistics command shows a range of useful statistical parameters calculated based on the current selection.

5.3 Realtime Analyzers


The integrated realtime analyzers allow you to analyze the output audio signal in real time during playback. You can hide or show the analyzers by choosing *View > Analyzers* and selecting one of the analyzers from the sub menu.

5.3.1 Level Meter

About Level Meter


The level meter lets you analyze the output level in terms of true peak, sample peak, peak hold and RMS values. The sample peak value is the maximum sample within a short analysis interval and is the value defining the height of the level meter bars. The true peak value takes into account how digital to analog convertors (DACs) would reconstruct an audio waveform, where it is possible for digital clipping to occur, even when the sample peak is never beyond 0 dB. The peak hold value is the maximum sample level over a longer period of time. It is indicated as a white line above or at the top of the level meter

bar. RMS stands for root-mean-square and is calculated by the root of the sum of the squared sample values during the analysis interval. The RMS level is calibrated according to the AES17-1998 standard which is 3 dB higher than the mathematical RMS level.



The level meter analyzer showing the peak hold value (1), peak value (2) and RMS value (3). The maximum true peak encountered since playback was started is displayed at the bottom (4)

You can configure the level meter to use different scales or change ballistics by clicking the left mouse button somewhere in the level meter. The following dialog box appears:



The settings dialog for the level meters.

Acoustica supports the K-System metering standard proposed by the audio engineer Bob Katz. The K-System is an attempt to standardize leveling practices throughout the audio industry. Three standards are available, K-20, K-14, and K-12 which are intended for different listening environments. You can choose to use one of the K-System meters or use the digital full scale meter as in earlier versions of Acoustica.

5.3.2 Loudness Meter

About Loudness Metering

Peak and RMS level metering don't necessarily match what we perceive as loudness very well. Short term RMS levels are related to loudness, but the sensitivity of the human auditory system is frequency dependent and that needs to be taken into account when measuring perceived loudness. The EBU R128 (based on the ITU-R BS.1770) recommendation defines a more suitable way of measuring loudness and deals with important issues such as how to react to loudness changes over time and how to measure loudness in multichannel audio. These two recommendations are becoming increasingly popular and are important to ensure a consistent listening experience when switching between audio tracks, television channels, radio programs and similar. Most music streaming providers and broadcasting organizations now specify loudness requirements that content providers need to adhere to.

Note: The Normalize processor in Acoustica (see **Volume > Normalize...**) supports EBU R128 / ITU-R BS.1770 as well and lets you set a target integrated loudness level in LUFS.

The Loudness Meter in Acoustica

Acoustica has loudness metering built in that follows the EBU R128 and ITU-R BS.1770 recommendations. You can enable and disable the meter using **View > Analyzers > Show or Hide Loudness Meter**.



The loudness metering in Acoustica showing the loudness curves from the last 30 seconds of output audio along with the current momentary loudness, short-term loudness, integrated loudness and the loudness range as defined by the EBU R128 recommendation.

There are four different loudness measures defined by the EBU R128 recommendation and all of these are displayed in the loudness meter in Acoustica:

- **Momentary (LUFS)**

The momentary loudness is calculated based on audio from the past 400 milliseconds and is visualized with a yellow curve in the loudness history graph. The maximum momentary loudness since start of playback is displayed in parenthesis.

- **Short-term (LUFS)**

The short-term loudness is calculated based on audio from the past 3 seconds and is visualized with a blue curve in the loudness history graph. The maximum short-term loudness since start of playback is displayed in parenthesis.

- **Integrated (LUFS)**

The integrated loudness is a measurement of the loudness of a complete audio track or program. Absolute and relative gates are used to avoid silent or very soft periods from affecting the loudness measurement. The integrated loudness is visualized with a white or red horizontal line in the loudness history graph. The indicator is white when the integrated loudness is below the target level and red if it is above. You can set the target level in the analyzer settings (see description below).

- **Loudness Range (LU)**

The loudness range is a measure for the loudness variance over time. Large changes in loudness levels over time will result in an increased loudness range.

Acoustica also indicates the illegal loudness range with a red background color in the loudness history. You can change the target level by clicking anywhere on the background of the loudness meter and the following window appears:




The Loudness Meter Settings lets you select the target level in LUFS

5.3.3 Spectrum Analyzer

About Spectrum Analyzer

The Spectrum Analyzer shows the frequency content of short analysis time frames, using FFT, which stands for Fast Fourier Transform. FFT is an efficient way of calculating the frequency domain of a signal (see [Time and Frequency Domains](#)^[36]).




The Spectrum Analyzer shows the frequency content of the output audio signal.

5.3.4 Phase Correlation Meter

About Phase Correlation Meter


The phase correlation meter shows the phase relationship between the left and the right audio channel in a stereo recording and is an important tool when mastering stereo recordings. If both channels contain exactly the same signal, the phase correlation meter will show a vertical line. If one channel is exactly the opposite of the other channel, the phase correlation meter shows a horizontal line. Normal stereo recordings will show a cloud of dots spread out vertically and horizontally (see the picture below). In a properly mastered recording, the cloud of dots should not be wider than it is tall and the correlation value indicated in the gauge at the bottom should be positive.



The phase correlation meter shows the relationship between the left and the right channel in a stereo recording.

5.3.5 Time Display

The big time display shows the current playback position in a dockable window pane



Note: if you don't see the Time Display, you can go to View > Analyzers > Show or hide Time Display to enable it.

6 Spectral Editing

Premium Edition Only


The spectral editing mode in Acoustica Premium Edition allows you to make selections in the time-frequency domain (see [Time and Frequency Domains](#)^[36] for more information). Processing is applied to the selection in time-frequency so that you can easily edit out isolated problems visible in the spectrogram. This is also a very powerful tool for sound design. A specialized [Retouch](#)^[58] tool is available in the spectral editing mode that you can use to attenuate noise while using information from surrounding "clean" areas in the spectrogram.

You can activate the spectral editing mode by clicking the spectrogram button on the toolbar:



The waveform (left) and spectral editing (right) mode buttons in the main toolbar.

Alternatively, you can choose *View Spectrogram* from the *View* menu or press Shift+S. The waveform view is changed to a spectrogram after activating the spectral editing mode:



The spectral editing mode in Acoustica Premium Edition.

6.1 Selections in Time and Frequency

Premium Edition Only

There are several selection tools available in the spectral editing mode. You can switch between these using the corresponding buttons in the main toolbar:



The selection tools available in the spectral editing mode (left to right): Area selection, freehand selection, brush selection, magic wand and the zoom tool.

- **Area Selection**

The area selection tool allows you to select rectangular shapes in the spectrogram. It can also be limited to time or frequency ranges by clicking the down arrow and choosing the desired mode from the drop-down menu.

- **Freehand Selection**

You can draw freehand selection shapes with the freehand selection tool. Click and hold down mouse button while you draw the desired shape. Release the mouse button when you are done. You can add additional shapes to the selection by holding down the shift key while repeating the procedure.

- **Brush Selection**

The brush tool allows you to "paint" a selection like you would use a paint brush. Click and hold down the mouse while you paint over the region you want to select. Release the mouse button when you are done. You can add additional shapes to the selection by


holding down the shift key while repeating the procedure. You can change the size of the brush by clicking the down arrow and choosing a brush size from the drop-down menu.

- **Magic Wand**

The magic wand automatically selects an area with similar intensity when you click somewhere in the spectrogram. You can adjust the tolerance in dB by clicking the down arrow and choosing the desired tolerance from the drop-down menu.

Selecting Harmonics

If you want to process tonal events in the spectrogram, you will frequently need to process harmonic overtones as well. You can add harmonics to the selection automatically by holding down the Alt key and moving the mouse wheel. Alternatively, you can hold down the Alt key and use the arrow up and down keys on the keyboard to add or remove harmonics.




Harmonics automatically added to the selection using Alt+key up or key down.

Sometimes, it will be easier to select the shape of an harmonic component than of the fundamental frequency shape. In those cases, you can select the harmonic shape and shift the selection down one octave using the command *Selection | Shift Octave Down* from the *Edit* menu or by pressing Alt+/. You can shift the selection up one octave using the *Selection | Shift Octave Up* or by pressing Alt+*.

Resizing and Moving Spectral Selections

You can move a spectral selection by moving the mouse cursor over the selected time-frequency geometry so that a hand cursor appears:



The hand cursor indicates that you can move the selection in time and frequency.


When you see the hand shaped cursor, you can click and keep the mouse button down while moving the geometry to a new location in time and frequency.

You can also resize spectral selections by clicking the edge of the bounding box or the corners. Keep the mouse button pressed while moving the cursor to define the new size. The corners allow you to resize along both the time and frequency axes, whereas the edges lock the resizing to one dimension.

6.2 The Spectrogram View Settings


Premium Edition Only

You can change the view settings of the spectrogram by clicking the button with the wrench icon:



The view settings button.

The *View Settings* window appears:



The view settings in the spectral editing mode.

The different settings are described below:

- **Frequency scale**

You can choose between different scales for the frequency axis: *Linear*, *logarithmic* or *Mel scale*. The latter is optimized according to our perception of sound.

- **Block size (samples)**

The duration in samples of the time slices that are used create the spectrogram (see [Time and Frequency Domains](#)^[36] for more information).

- **Temporal resolution enhancement:**

Normally, there is a trade-off between time and frequency resolutions in spectrograms. Acoustica overcomes this shortcoming with the temporal resolution enhancement and you can set the time resolution enhancement factor.

- **Color palette**

You can choose among different color palettes according to your own liking.

- **Minimum magnitude level (dB)**

The minimum magnitude to show in the spectrogram.

- **Maximum magnitude level (dB)**

The maximum magnitude to show in the spectrogram.

- **Window attenuation (dB)**

The attenuation in the analysis window specified in Decibels. Please see [Time and Frequency Domains](#)^[36] for details.

6.3 The Retouch Tool


Premium Edition Only

The spectral editing mode is a very powerful tool in post processing and audio restoration since many problems can be better isolated in the time frequency domain than using a time range alone. One of the most interesting applications is to attenuate unwanted noise based on information from the surrounding areas in time and frequency. The *retouch* tool is a specialized tool for this purpose. Retouch reduces the magnitudes in the selection based on a reference that you can position freely. Any noise present in the selection that is not present in the reference signal will thus be attenuated. You can process both horizontally (keeping tonal information) or vertically (keeping time events). Horizontal processing is the most common usage since it preserves tonal information.

The first step of the process is to select the smallest region possible that covers the unwanted noise completely. You can choose whichever selection tool is more suited (see [Selections in Time and Frequency](#)^[54]). When done, please activate the retouch tool using the corresponding button in the main toolbar:




Alternatively, you can choose *Edit > Active Tool > Retouch Tool* to enable or disable the retouch tool. After activation, you can see a process button and a green *reference signal* indicator as depicted below:



The retouch tool in Acoustica allows you to attenuate unwanted sounds that are isolated in time and frequency. In this example there is a disturbing beep from a camera auto-focus.


You can press *Ctrl* and move the mouse to position the reference signal indicator. The reference should be a clean signal that is as similar to the wanted signal as possible, either before or after the noise you want to eliminate. You can also process along the vertical axis by positioning the reference above or below the unwanted noise. When you have positioned the reference signal indicator, click the arrow button to process.



The disturbing beep is completely removed after processing with the retouch tool.

Processing Options

You can click the arrow down button on the retouch tool button in the toolbar to change the retouch tool's processing settings:



You can change the Retouch processing mode and strength.


The *Attenuate* mode will attenuate all signal components that are lower than the reference whereas *Replace* will use the signal component levels from the reference even if they are higher. You can reduce the strength if you want more gentle processing.

7 Audio Processing

Most processing tools in Acoustica have some properties in common such as preset management, A/B comparisons and bypassing. You can find a common set of buttons and controls in the header section of the included processors and these are described below.

Preset Manager

Acoustica is shipped with a set of factory presets that serve as a starting point for further adjustments. You can browse through preset categories and presets as well as create and manage your own presets using the preset management section:



The preset management section available in all the integrated processors.

You can browse through the presets using the arrow buttons. Alternatively, you can click the current preset name and a drop-down menu appears. You can also save your own presets by choosing "Save user preset file..." from the menu. A file chooser dialog box appears where you can enter the name of the preset you wish to save. You can create sub folders and place your preset files inside, and these will appear as categories in the user presets.

Undo and Redo

You can undo (or redo) changes to the parameter settings by clicking the circular arrow back or forward buttons:



Undo and redo buttons

A / B Comparisons

It is frequently useful to be able to quickly compare different parameter settings. You can do this using the A / B comparison buttons:



The A / B comparison buttons allows you to quickly compare different settings

You can keep two independent sets of parameter settings, the A and B settings, and switch between them using the corresponding buttons. The arrow button copies the settings from A to B or the other way around depending on which parameter set that is currently active.

The Processor Menu

The last button in the header section displays the processor menu:



You can click the processor menu button for the processor specific menu

The processor menu allows you to access the processor help topic directly along with other processor specific features.

Using Sliders to Adjust Parameters

Horizontal, vertical and rotary sliders (or knobs) are frequently used in the built-in processors in Acoustica. Clicking with the mouse on a knob and moving up or down will allow you to modify the setting. You can also use the scroll wheel of your mouse. For more precise control, you can hold down the Ctrl key on your keyboard at the same time as you use your mouse to modify the setting.

Hearing the Effect of a Processor in Real Time

You can listen to the effect of the current processor or plug-in settings and make changes while immediately hearing the results. This is controlled using the transport buttons in the lower right corner of the processor window:



The "preview" transport controls (play, stop and bypass) in the processor window

You can start or stop the playback with the two first buttons. The third button is a bypass toggle that allows you to disable the processing so that you can easily compare the processed and unprocessed versions.

7.1 Tools


The *Tools* menu in Acoustica contains the most common audio processing tools such as dynamic processing, equalization, sample format conversions and more.

7.1.1 Dynamics

About Dynamic Processing


A dynamic processor is used to alter the dynamic properties of the recording. To understand how a dynamic processor works, imagine a sound engineer trying to maintain as steady a volume level as possible while doing a recording. When the input level increases he pulls down the volume fader, and he pushes it up when the input level decreases. A dynamic processor does the same thing automatically according to its settings, only with a much faster reaction time.

Modern dynamic processors allow you to set a ratio between the input levels and the output levels. This ratio is visualized as a curve where the horizontal axis represents the input level and the vertical axis represents the output level. A straight line as shown below represents a 1:1 ratio.




With such a setting, no change is made to the level as it is processed. Changes are made to the dynamics by altering the ratio and the threshold. In the example below, all signal levels above -20 dB are attenuated with a 100:1 ratio, so that the output level only raises 1dB when the input raises 100dB above the threshold. This setting would be comparable to a piece of hardware known as a limiter. You can see from the graph below that once

the input level reaches -20 dB, the output level is barely going any higher than -20 dB, even when the input level increases.



If the dynamic processor changes the level too fast, low frequency signal components might become distorted. How quickly the dynamic processor adapts to changes in the input level is called response time. The response time is divided into the amount of time when the input level rises (the attack time) and when it falls (the release time).

When applying compressor ratios that lead to extreme changes in the dynamics, audible artifacts might become noticeable (often referred to as "pumping and breathing"). Smoother compression curves will generally reduce the artifacts of a dynamic processor. Soft kneeing automatically softens the curve to reduce such artifacts. A high level of soft kneeing was used in the image below.



User Interface



Parameter Settings

- **Threshold level (-60 dB to 0 dB)**

When the input level exceeds the set threshold, the compressor starts to apply compression. For example, if the threshold is set at -12 dB, it will only respond to incoming audio that exceeds -12 dB. However, as soon as you activate the Soft knee, it's possible that compression takes place even before the input level exceeds the set threshold, due to the nature of the soft knee.

- **Ratio (0.01:1 to 100:1)**

By setting up a ratio, you decide how much gain reduction should take place for every dB the input level exceeds above the set threshold. For example, if the threshold is set on 2.00:1, there will only be a 1 dB output increase for every 2 dB the input level exceeds above the set threshold. You could also say that the compressor will do 1 dB of gain reduction for every 2 dB the input level exceeds the set threshold.

It's also possible to set a ratio that is below 1.00:1. This basically turns the compressor into a gate or expander. Audio levels below the threshold will now be reduced.

- **Range (0 to 60 dB)**

The range control sets the maximum gain reduction of the dynamic processor.

- **Soft knee (0 dB to 24 dB)**

A soft knee allows you to create a smoother compression curve, which can help to prevent or reduce audible artifacts. With the soft knee set to 0 dB, there is no soft knee applied to the compression curve. When the soft knee is set to 8 dB for example, the compression curve will be smoothly interpolated for input levels in the range -8 dB to +8 dB relative to the threshold level.

- **Attack time (0.01 ms to 200 ms)**

You can use the attack time knob to adjust how quickly the compressor will react to incoming audio that is above the threshold. The faster the attack time, the faster the compressor will react. If compression is applied on a snare for example which has a very loud attack, setting up a very fast attack time might help to reduce the loud attack of the snare, as the compressor can almost instantly apply gain reduction. However, in many scenarios a slightly slower attack time might provide you with a more natural sounding compression behavior, for example during mastering.

- **Release time (ms)**

You can use the release time knob to adjust how quickly the compressor will stop processing audio that is no longer above the threshold. A fast release time can greatly enhance the sustain/decay of certain softer elements of the audio, but it can also introduce a "pumping" effect. Slow release times (>200ms) often result in a more smoother/natural compression behavior.

- **Auto-Release button**

You can enable auto-release mode by clicking the button labeled "A" on the left side of the numerical entry field for the release time. In auto-release mode, the release time specifies the maximum release time and the dynamic processor shortens the release time for sharp transients.

- **Hold time (0 ms to 500 ms)**

As soon as audio is below the set threshold, the release time kicks in. However, the hold time allows the compressor to postpone the activation of the release time for up to 500ms, which might provide you with a smoother compression behavior in certain audio material. Although you are free to experiment with the hold function, we would like to suggest to use a hold time that is significantly shorter than the release time.

- **Make-up (0 dB to 32 dB)**

You can use the make-up knob to compensate for the level reduction caused by the compressor. It's possible to set the make-up in auto mode, by pressing the A button in front of the knob. When the A button is yellow, the auto make-up is activated and the make-up knob is disabled.

- **Latency (0 ms to 15 ms)**

It is possible to reduce harmonic distortion of low frequency content even with short attack and release times if the dynamic processor is allowed to examine the signal ahead of time. The downside is a slightly increased latency. You can adjust the maximum allowed latency in milliseconds.

- **Channel Linking (0% to 100%)**

With the channel linking set to 100%, the amount of gain reduction will be the same for the left and right channels, even if there is a difference in level between these channels. The more you dial the knob away from 100%, the more the channels will be independently treated, with complete independence between the channels once you have a 0% channel linking. Channel linking values below 100% can result in shifts in the stereo image. During mastering, the best results will be most likely achieved with the channel linking set between 70% and 100%.

- **Mix (0% to 100%)**

The mix parameter controls the unprocessed and processed signal mix. Normally, this will be set to 100%, lower settings allow for convenient parallel compression.

- **Oversampling (off, 2x, 4x)**

The internal sample rate of the dynamic processor can be multiplied by 2 or 4, depending on the sample rate of the project. For example, if you have a project with a sample rate of 48 kHz, setting the oversampling to 4x will result in an internal sample rate of 192 kHz.

Side Chain Filtering

The signal that is used to calculate the gain signal in dynamic processing is referred to as the *side chain* signal. You can filter the side chain signal in *Dynamics* so that the dynamic processing is more sensitive in certain frequency regions. This works similar to an equalizer. Click the button labeled *Side chain filter* in the bottom left corner of the Dynamics user interface to open (or close) the side chain filtering settings:



The side chain filter in Dynamics allows you to filter the side chain signal similar to the way an equalizer works.

You can click one of the filter buttons to toggle the filter and the following filter types are available:

-  : Low cut, removes frequency content below the band frequency
-  : Low shelf, boosts or attenuates frequency content below the band frequency
-  : Peak filter, boosts or attenuates frequency content around the band frequency
-  : High shelf, boosts or attenuates frequency content above the band frequency
-  : High cut, removes frequency content above the band frequency

The *Side chain filter* button turns orange as soon as filtering is active.


7.1.2 Multiband Dynamics

Premium Edition Only

About Multiband Dynamics


Multiband Dynamics features four independent dynamic processors in parallel, each working on a different frequency band of the audible spectrum. This allows you to apply dynamic processing with different settings for each of the four bands, making it a perfect tool for addressing certain problem frequencies in your mix, like low frequency rumble or sibilance for example.

There are four bands in *Multiband Dynamics*, which means that there are three cross-over frequencies. These are between the four frequency bands called *Low Band*, *Low Mid Band*, *High Mid Band* and *High Band*. Each band can be set between 20 Hz and 20 kHz, but a cross-over point can never cross another cross-over point. In the example below you see the three cross-overs set at 100 Hz, 400 Hz and 4 kHz.



In the same example you also see that each cross-over uses a different slope setting. The first cross-over uses a slope of 6 dB/octave, it's the slowest slope that you can set for each band. The second cross-over uses a pretty common 24 dB/octave slope, while the last cross-over uses the steepest slope that you can set for each band, which is 96 dB/octave. To set your preferred slope, select the desired cross-over marker to make the slope setting visible. You can drag either one of the slope markers to modify the slope to your liking. A steeper slope provides you with more precision at the cost of potential pre-ringing artifacts, which can distort transient rich material. Therefore, we encourage you to start with a medium slope and modify to taste.

User Interface



Global Parameter Settings

- **Phase mode**

You can choose between two different phase modes in *Multiband Dynamics*:

Linear phase: The advantage of the linear phase mode is that it completely avoids phase issues when the bands are recombined after processing. However, linear phase processing introduces pre-ringing effects that distorts transients and also increases latency.

Mixed phase: The unique mixed phase mode combines the best of minimum and linear phase processing. You can adjust the latency and hence the maximum pre-ringing time so that it is masked by the temporal masking effect in the human auditory system. At the same time, phase issues that plague multiband processing with minimum phase filters are reduced to a minimum.

You can read more about the phase modes in the [Phase Modes chapter](#)⁹⁰.

- **Latency (ms)**

If the mixed phase mode is activated, you can choose the latency and maximum pre-ringing time freely in the range 3 ms to 20 ms. Low latencies will result in a phase distribution more like minimum phase while higher latencies result in a more linear

phase distribution. We recommend a settings around 5 ms which is sufficiently short to mask pre-ringing.

- **Channel Linking (%)**

With the channel linking set to 100%, the amount of gain reduction will be the same for the left and right channels, even if there is a difference in level between these channels. The more you dial the knob away from 100%, the more the channels will be independently treated, with complete independence between the channels once you have a 0% channel linking. Channel linking values below 100% can result in shifts in the stereo image.

- **Input level (-16 dB to 16 dB)**

You can use the input level to adjust the level prior to all processing.

- **Output level (-16 dB to 16 dB)**

You can use the output level to adjust the level after all processing.

Band Parameter Settings

The following parameters are available for each of the four bands. Press and hold the *Shift* key while using the sliders to modify a parameter for all four bands.

- **Band solo**

You can solo the output from a band by clicking the solo button ()

- **Band on / off toggle**

Bands can be enabled or disabled with the power on / off button

- **Input gain (-16 dB to 16 dB)**

You can use the input gain to increase or decrease the signal level at the band's audio input.

- **Threshold level (-60 dB to 0 dB)**

When the input level exceeds the set threshold, the compressor starts to apply compression. For example, if the threshold is set at -12 dB, it will only respond to incoming audio that exceeds -12 dB. However, as soon as you activate the Soft knee, it's possible that compression takes place even before the input level exceeds the set threshold, due to the nature of the soft knee.

- **Ratio (0.01:1 to 100:1)**

By setting up a ratio, you decide how much gain reduction should take place for every dB the input level exceeds above the set threshold. For example, if the threshold is set on 2.00:1, there will only be a 1 dB output increase for every 2 dB the input level

exceeds the set threshold. You could also say that the compressor will do 1 dB of gain reduction for every 2 dB the input level exceeds the set threshold.

It's also possible to set a ratio that is below 1.00:1. This basically turns the compressor into a gate or expander. Audio levels below the threshold will now be reduced.

- **Soft knee (0 dB to 24 dB)**

A soft knee allows you to create a smoother compression curve, which can help to prevent or reduce audible artifacts. With the soft knee set to 0 dB, there is no soft knee applied to the compression curve. When the soft knee is set at 8 dB for example, the compression curve will be smoothly interpolated for input levels in the range -8 dB to +8 dB relative to the threshold level.

- **Attack time (0.01 ms to 200 ms)**

You can use the attack time knob to adjust how quickly the compressor will react to incoming audio that is above the threshold. The faster the attack time, the faster the compressor will react. If compression is applied on a snare for example which has a very loud attack, setting up a very fast attack time might help to reduce the loud attack of the snare, as the compressor can almost instantly apply gain reduction. However, in many scenarios a slightly slower attack time might provide you with a more natural sounding compression behavior, for example during mastering.

- **Release time (ms)**

You can use the release time knob to adjust how quickly the compressor will stop processing audio that is no longer above the threshold. A fast release time can greatly enhance the sustain/decay of certain softer elements of the audio, but it can also introduce a "pumping" effect. Slow release times (>200ms) often result in a more smoother/natural compression behavior.

- **Auto-Release button**

You can enable auto-release mode by clicking the button labeled "A" on the left side of the numerical entry field for the release time. In auto-release mode, the release time specifies the maximum release time and the dynamic processor shortens the release time for sharp transients.

- **Output Level (-16 dB to 16 dB)**

You can use the level handle (colored circle) in the filter response display to compensate for the level increase or decrease through the dynamic processing.

7.1.3 Limit

The Purpose of a Limiter

Limit is a two-stage dynamic processor, designed to transparently increase the perceived loudness of audio, while at the same time making sure no audio clipping occurs. First, there is a compressor stage (Pre-Compressor) that helps to keep the number of exceeding peaks within an acceptable range, which helps to prevent or reduce audible distortion. The second stage is the actual peak limiting stage (Peak Suppression) which makes sure that no audio will pass through above the threshold, also known as *brickwall* limiting. In this stage, peaks are reduced with a very quickly responding algorithm in the most transparent way as possible.

User Interface



Parameter Settings

- **Input (0 dB to +32 dB)**

The internal threshold in Limit is always 0 dB. In order to increase the perceived loudness, you need to increase the input level of the audio. As soon as the input audio will exceed the internal threshold, gain reduction will be applied in order to keep the audio from clipping. Excessive amounts of gain reduction can lead to audible distortion

and/or a pumping effect. For mastering duties, we recommend to keep the amount of gain reduction to a minimum.

- **Output (0 dB to -32 dB)**

The output of Limit is set to 0 dB by default, but can be used to compensate for the increased level when the input is boosted.

- **Attack time (0.01 ms to 500 ms)**

You can use the attack time knob to adjust how quickly the Pre-Compressor stage will react to incoming audio that is above the internal threshold. The faster the attack time, the faster the compressor will react. A fast attack time will be more transparent, at the cost of a slightly lower increase in perceived loudness. A slow attack time will increase the perceived loudness, at the risk of introducing audible distortion or a pumping effect. For mastering duties, we recommend to start with an attack time of 40 ms and adjust it to taste.

- **Release time (1 ms to 5000 ms)**

You can use the release time knob to adjust how quickly the compressor will stop processing audio that is no longer above the internal threshold. A fast release time will greatly increase the perceived loudness at the risk of audible distortion. A slow release time will be more transparent, at the cost of a lower increase of perceived loudness. For mastering duties, we recommend to start with a release time of 150 ms and adjust it to taste.

- **Auto-Release button**

You can enable auto-release mode by clicking the button labeled "A" on the left side of the numerical entry field for the release time. In auto-release mode, the release time specifies the maximum release time and the pre-compressor shortens the release time for sharp transients.

- **Pre-Compressor Channel Linking (0% to 100%)**

With the channel linking set to 100%, the amount of gain reduction will be the same for the left and right channels, even if there is a difference in level between these channels. The more you dial the knob away from 100%, the more the channels will be independently treated, with complete independence between the channels once you have a 0% channel linking. Channel linking values below 100% can result in shifts in the stereo image. During mastering, the best results will be most likely achieved with the channel linking set between 70% and 100%.

- **Look ahead time (0 ms to 15 ms)**

Limit uses a sophisticated look ahead algorithm to ensure the highest possible transparency. The higher this value, the more transparent the processing will be. Please

keep in mind that look ahead time will also introduce latency, which will normally be automatically compensated for by your host application.

- **Peak Suppression Channel Linking (0% to 100%)**

With the channel linking set to 100%, the amount of limiting will be the same for the left and right channels, even if there is a difference in level between these channels. The more you dial the knob away from 100%, the more the channels will be independently treated, with complete independence between the channels once you have a 0% channel linking. Although channel linking values below 100% can result in shifts in the stereo image, it is not as evident as during the pre-compressor stage. During mastering, the best results will be most likely achieved with the channel linking set between 50% and 100%.

- **Oversampling (off, 2x, 4x)**

The internal sample rate of Limit can be multiplied by 2 or 4, depending on the sample rate of the project. For example, if you have a project with a sample rate of 48 kHz, setting the oversampling to 4x will result in an internal sample rate of 192 kHz.

7.1.4 Dither

Introduction to Dither

Whenever you reduce the resolution (bit depth) of an audio signal you will introduce truncation errors, which can be a very unpleasant artifact if audible. The quantization noise is correlated to the audio signal which we perceive as more disturbing than uncorrelated noise. Dither introduces low level random noise to de-correlate the noise. Additionally, it is possible to alter the frequency distribution of the noise signal. The human ear has a frequency dependent sensitivity and it makes sense to "move" the noise to frequency regions where the ear is less sensitive. This process is called noise shaping and the *Dither* plug-ins offers detailed control over the noise shaping process.

User Interface



Parameter Settings

- **Target resolution**

Here you have to set the target resolution. If you want to change the resolution from 24 bit to 16 bit for example, you need to select 16.

- **Enable noise shaping**

Click to enable or disable the noise shaping filter

- **Dither level (%)**

Dither amplitude in percent of 1 LSB (least significant bit). This should be at 100% for complete de-correlation of quantization noise.

- **Filter length (ms)**

This allows you to set the length of the noise shaping filter in milliseconds. Longer filters allow more accurate shaping, but comes at CPU cost.


- **Max. noise shaping (dB)**

This sets the maximum gain allowed in the noise shaping filter. The higher the setting, the less audible the added noise will be, by moving the noise up in the frequency spectrum. However, too much energy concentration in the high frequency bands should be avoided as this can cause issues with certain digital to analog converters.

7.1.5 Convert Sample Format

The *Convert Sample Format* tool allows you to change the sample format of an audio recording. If you are not familiar with the term sample rate, please read [Working with Digital Audio](#)^[8] before proceeding.

To convert the sample format of a recording, select *Tools > Convert Sample Format...* A dialog box appears where you can define the sample rate and the number of channels in the new sample format. You can also choose if Acoustica should perform sample rate conversion or simply change the sample rate information. The latter will cause the playback speed and pitch to change. Click the *OK* button when you are done to start the conversion process.



The Convert Sample Format dialog box.

7.1.6 Insert Silence


The *Insert Silence* command lets you insert silence with a specified duration at the cursor position. The following window appears where you can enter the duration in seconds:



The Insert Silence tool lets you insert silence at the cursor position.

7.1.7 Signal Generator

The *Signal Generator* tool generates common test signals like frequency sweeps, sine waves and Gaussian noise with different frequency distributions (referred to as "colors").



The Signal Generator can create a range of different test signals.

The generated signal is inserted into the clip at the cursor position and you can set the duration of the signal using the *Duration* parameter.

Parameter Settings

- **Signal type**

The signal type can be *Sine* (sine wave with constant frequency), *Exponential sweep* (Sine wave with exponential transition of frequency from *Start frequency* to *End frequency*), *Linear sweep* (Sine wave with linear transition of frequency from *Start frequency* to *End frequency*), *White noise* (noise with Gaussian value distribution and a flat frequency distribution), *Pink noise* (noise with Gaussian value distribution and a $1/f$ type frequency distribution) and *Brown noise* (noise with Gaussian value distribution and a $1/f^2$ type frequency distribution).

- **Duration (seconds)**

The duration of the test signal to generate specified in seconds.

- **Amplitude (dB)**

The RMS (mathematical) amplitude of the signal to generate. Colored noise signals are defined at 1 kHz.

- **Start Frequency (Hz)**

The start frequency of linear and exponential sweeps or the frequency if *Signal type* is *Sine*.

- **End Frequency (Hz)**

The end frequency of linear and exponential sweeps.

7.1.8 Remix

Remix tool uses artificial intelligence (AI) to split a finished mix into five instrument groups, or so called stems - *Vocals*, *Piano*, *Bass*, *Drums* and *Other*. The *Other* stem contains everything that's not in the prior four stems. Despite the internal complexity, *Remix* is very easy to use and the user interface shows a mixer layout that should be familiar to most users.



The Remix separates a mix into stems and lets you adjust the levels of the stems in real time.

Mixer Settings

- **Stem Gain (dB)**

The mixer contains sliders that set the output level for each of the available stems in Decibel (dB). The output levels for each stem is visualized in the level meter to the right of every stem gain slider. You can enter dB values for the stem gains numerically by clicking the numerical entry below the slider.

- **Solo (S) and Mute (M)**

Stems can be auditioned separately (S) or muted completely (M) by clicking the respective buttons.

- **Sensitivity (%)**

You can control the sensitivity of the stem signal detection independently for each stem. The sensitivity refers to the sensitivity of the stem separation, and higher values lead to more signal being detected. Thus, positive values will increase naturalness at the cost of additional spill and negative values will reduce spill at the cost of possible artifacts. We recommend to start with 0% at which the neural network works as trained without user influence, and then adjust according to preference if necessary.

7.1.9 Equalize

Premium Edition Only

Introduction

Equalize 2 is a versatile, user friendly and great sounding equalizer with several unique features. Unlike other equalizers, you can freely adjust not only center frequency, gain and bandwidth, but also the filter slope for each band. The filter slope can be set anywhere from 3 to ultra sharp 120 dB per octave. Needless to say, you can operate Equalize as a zero latency processor when using the minimum phase mode or choose to preserve the phase relationships in the linear phase mode. Equalize goes one step further, though, and introduces the unique mixed phase mode that allows you to set the latency freely in the range 5 to 120 milliseconds while preserving the phase relationships as far as possible. That gives a unique control over potential pre-ringing artifacts which is a common problem with linear phase filtering.

Latency values below 20 milliseconds ensure that any pre-ringing is masked by the temporal masking of the human hearing while preserving the time-alignment across the audible frequency range. Great care has been taken to provide a user interface that is straight forward to use. Band parameters can be adjusted using handles directly in the graphical representation of the frequency response, including bandwidths and filter slopes. A flexible real-time analyzer lets you monitor every aspect of the processing. You can easily switch between full, mid, side, left or right channel processing for each band and Equalize automatically routes the audio signal internally to ensure the best results and lowest possible latency.

User Interface

The graphical user interface of *Acon Digital Equalize 2* is designed to hide the complexity of the plug-in and provide an efficient workflow as well as intuitive control over the plug-in parameters.



Adding and Removing Bands

You can add up to 12 individual bands in Equalize. Each equalizer band represents a filter type, which can be either low cut, low shelf, bell, notch, high shelf or high cut. To add an additional band, click the + button below the curve display or double click in the curve display. A bullet shaped handle appears in the curve that you can move around using the mouse. The currently selected band also shows additional handles that you can use to manipulate the bandwidth (if applicable) and filter slope. You can remove the band by clicking the x button or by double clicking the handle.

Soloing and Bypassing Bands

You can solo or temporarily bypass an equalizer band in order to monitor its effect on the audio signal. You can enable the solo mode by pressing the Ctrl key while moving an equalizer band handle. The Shift key enables the bypass band mode. Alternatively, you can use the buttons to enable or disable the solo and bypass modes:

-  : Solo mode
-  : Bypass mode

Band Parameters

The parameters that are related to a specific band are placed within the band group and the header indicates which band is currently active. The parameters apply to the currently active band only. To change the currently active band you can either click its bullet handle in the frequency response visualization or use the arrow buttons to browse through the active bands.

Frequency (Hz)

The center (for peak and notch filters) or threshold frequency of the currently active equalizer band in Hertz.

Band gain (dB)

The gain of the currently active equalizer band in decibels. This parameter is not available for high or low cuts, or the notch filter.

Bandwidth (oct.)

The bandwidth of the currently active equalizer band in octaves. This parameter is only available for peak and notch filters.

Slope (dB/oct.)

The filter slope of the currently active equalizer band in decibels per octave. This controls the steepness of the filter.

Gain to Bandwidth Link

The perceived bandwidth of an equalizer band is dependent on the gain setting. You can link the bandwidth to the gain setting, so that the bandwidth is automatically adjusted when you change the gain to preserve the perceived bandwidth.

 : Click this button to activate or deactivate the gain to bandwidth link

Filter Type Buttons

You can choose between eight different filter types:

 : Low cut, removes frequency content below the band frequency

 : Low shelf, boosts or attenuates frequency content below the band frequency

 : Peak filter, boosts or attenuates frequency content around the band frequency

 : Tilt filter, tilts the frequency content around the band frequency so that one side get boosted while the other side gets attenuated

 : Bandpass filter, removes all frequency content except for the band surrounding the band frequency

 : Notch filter, removes frequency content around the band frequency

 : High shelf, boosts or attenuates frequency content above the band frequency

: High cut, removes frequency content above the band frequency

Channel Mode Buttons

You can choose any of the following channel modes for each equalizer band individually:

L : Apply to left channel only

M : Apply to mid channel only. The mid channel is the sum of the input channels multiplied by a scaling factor.

 : Apply to both channels

S : Apply to side channel only. The side channel is the difference of the input channels multiplied by a scaling factor.

R : Apply to right channel only

Global Parameters

Master gain (dB)

The master gain controls the overall gain of the equalizer in decibels. Equalize can automatically adjust the overall gain to compensate for any gain changes in the equalizing process. This is based on an average distribution typically found in music and cannot be completely accurate. To activate the automatic gain compensation, click the button labelled A to the left of the master gain knob.

Phase mode

You can choose between three different phase modes in Equalize. Since Equalize is a 'clean' equalizer in the sense that it doesn't apply any saturation or similar effects to simulate analog artifacts, the resulting filter applied by Equalize is fully described by the amplitude response (which is the visualization of the filter curve response that is shown in the upper part of the plug-in window) and the phase response.

Minimum phase: This is the way analog equalizers process audio and has the great advantage of zero latency and no pre-ringing. Sharp filters can cause time alignment problems and crosstalk between tracks can cause unpredictable results when tracks are processed separately.

Linear phase: The advantage of the linear phase mode is that you can process track independently even if there is crosstalk between them. Linear phase processing introduces high latencies and can cause pre-ringing effects that aren't desirable.

Mixed phase: The unique mixed phase mode combines the best of the two phase distributions. You can adjust the latency and hence the maximum

pre-ringing time so that it is masked by the temporal masking effect in the human auditory system.

You can read more about the phase modes in the [Phase Modes chapter](#)^[90].

Latency (ms)

If the mixed phase mode is activated, you can choose the latency and maximum pre-ringing time freely in the range 5 ms to 120 ms. Low latencies will result in a phase distribution more like minimum phase while higher latencies result in a more linear phase distribution. We recommend to keep this value around 20 ms to ensure that the pre-ringing is masked by the temporal masking effect of the human auditory system while still preserving the time alignment of the signal component in the audible frequency range.

Note: The latency is doubled if you combine mid or side bands with left or right channel bands and use the mixed phase mode.

Analyzers and the Frequency Response Visualization

Equalize visualizes the frequency response of the current equalizer settings and contains two separate spectrum analyzers that you can set up to monitor the effects of the processing.

Analyzer 1 & 2

You can choose to analyze different signal sources using the drop-down lists under the Analyzer 1 and Analyzer 2 headers. Either input or output signals can be analyzed and you can also choose if the left, right, mid or side signal should be analyzed.

Range Buttons

You can change the range of the frequency response visualization or the spectrum analyzers by clicking the buttons at the lower end of the level axes. A drop down list appears with the alternatives.

Adjusting the Analyzer Frequency Range

The frequency range of the equalizer curve can be altered conveniently thus making it possible to zoom in on specific frequency ranges for more precise control. Move the mouse cursor to either the left or right edge of the frequency axis below the frequency curve editor. The mouse cursor turns into a left-to-right arrow. Now click and keep the left mouse button pressed while adjusting either the start or end frequency of the visible frequency range. You can also move the visible frequency range up or down moving the mouse to the center of the frequency axis. The mouse cursor turns into a hand. Click and keep the left mouse button pressed while moving the visible frequency range up or down in frequency.

Piano Keyboard Display

You optionally let Equalize display a piano keyboard in the bottom of the equalizer curve editor and you can find this setting in the [Preferences](#)^[88].



Acon Digital Equalize 2 with the piano keyboard display activated.

The piano keyboard makes it easy to map the band frequencies to the frequency of a specific piano key. Click a key on the piano keyboard to assign the frequency of the currently active band to the frequency of the clicked key.

Adding and Removing Bands

You can add up to 12 individual bands in Equalize. Each equalizer band represents a filter type, which can be either low cut, low shelf, bell, notch, high shelf or high cut. To add an additional band, click the + button below the curve display or double click in the curve display. A bullet shaped handle appears in the curve that you can move around using the mouse. The currently selected band also shows additional handles that you can use to manipulate the bandwidth (if applicable) and filter slope. You can remove the band by clicking the x button or by double clicking the handle.

Soloing and Bypassing Bands

You can solo or temporarily bypass an equalizer band in order to monitor its effect on the audio signal. You can enable the solo mode by pressing the Ctrl key while moving an equalizer band handle. The Shift key enables the bypass band mode. Alternatively, you can use the buttons to enable or disable the solo and bypass modes:

-  : Solo mode
-  : Bypass mode

Band Parameter Settings

The parameters that are related to a specific band are placed within the band group and the header indicates which band is currently active. The parameters apply to the currently active band only. To change the currently active band you can either click its bullet handle

in the frequency response visualization or use the arrow buttons to browse through the active bands.

- **Frequency (Hz)**

The center (for peak and notch filters) or threshold frequency of the currently active equalizer band in Hertz.

- **Band gain (dB)**

The gain of the currently active equalizer band in decibels. This parameter is not available for high or low cuts, or the notch filter.

- **Bandwidth (oct.)**

The bandwidth of the currently active equalizer band in octaves. This parameter is only available for peak and notch filters.

- **Resonance (dB)**

Controls the resonance in decibels at the threshold frequency for cut and shelving filters. Resonance occurs in analog filters and results in a boost around the threshold frequency.

- **Slope (dB/oct.)**

The filter slope of the currently active equalizer band in decibels per octave. This controls the steepness of the filter.

- **Gain to Bandwidth Link**

The perceived bandwidth of an equalizer band is dependent on the gain setting. You can link the bandwidth to the gain setting, so that the bandwidth is automatically adjusted when you change the gain to preserve the perceived bandwidth.



: Click this button to activate or deactivate the gain to bandwidth link

- **Filter Type Buttons**

You can choose between six different filter types:



: Low cut, removes frequency content below the band frequency



: Low shelf, boosts or attenuates frequency content below the band frequency



: Peak filter, boosts or attenuates frequency content around the band frequency



: Notch filter, removes frequency content around the band frequency



: High shelf, boosts or attenuates frequency content above the band frequency



: High cut, removes frequency content above the band frequency

- **Channel Mode Buttons**

You can choose any of the following channel modes for each equalizer band individually:

L : Apply to left channel only

M : Apply to mid channel only. The mid channel is the sum of the input channels multiplied by a scaling factor.

Ⓓ : Apply to both channels

S : Apply to side channel only. The side channel is the difference of the input channels multiplied by a scaling factor.

R : Apply to right channel only

Global Parameters

- **Master gain (dB)**

The master gain controls the overall gain of the equalizer in decibels. Equalize can automatically adjust the overall gain to compensate for any gain changes in the equalizing process. This is based on an average distribution typically found in music and cannot be completely accurate. To activate the automatic gain compensation, click the button labelled A to the left of the master gain knob.

- **Phase mode**

You can choose between three different phase modes in Equalize. Since Equalize is a 'clean' equalizer in the sense that it doesn't apply any saturation or similar effects to simulate analog artifacts, the resulting filter applied by Equalize is fully described by the amplitude response (which is the visualization of the filter curve response that is shown in the upper part of the plug-in window) and the phase response.

Minimum phase:

This is the way analog equalizers process audio and has the great advantage of zero latency and no pre-ringing. Sharp filters can cause time alignment problems and crosstalk between tracks can cause unpredictable results when tracks are processed separately.

Linear phase:

The advantage of the linear phase mode is that you can process track independently even if there is crosstalk between them. Linear phase processing introduces high latencies and can cause pre-ringing effects that aren't desirable.

Mixed phase:

The unique mixed phase mode combines the best of the two phase distributions. You can adjust the latency and hence the maximum pre-ringing time so that it is masked by the temporal masking effect in the human auditory system.

- **Latency (ms)**

If the mixed phase mode is activated, you can choose the latency and maximum pre-ringing time freely in the range 5 ms to 120 ms. Low latencies will result in a phase distribution more like minimum phase while higher latencies result in a more linear phase distribution. We recommend to keep this value around 20 ms to ensure that the pre-ringing is masked by the temporal masking effect of the human auditory system while still preserving the time alignment of the signal component in the audible frequency range.

Note: The latency is doubled if you combine mid or side bands with left or right channel bands and use the mixed phase mode.

Analyzers and the Frequency Response Visualization

Equalize visualizes the frequency response of the current equalizer settings and contains two separate spectrum analyzers that you can set up to monitor the effects of the processing.

- **Analyzer 1 & 2**

You can choose to analyze different signal sources using the drop-down lists under the Analyzer 1 and Analyzer 2 headers. Either input or output signals can be analyzed and you can also choose if the left, right, mid or side signal should be analyzed.

- **Drop (dB/s)**


The drop parameter sets how quickly the analyzer adapts to lower signal levels and is specified in decibels per second.

- **Range Buttons**

You can change the range of the frequency response visualization or the spectrum analyzers by clicking the buttons at the lower end of the level axes. A drop down list appears with the alternatives.


7.1.9.1 Equalize Preferences

Equalize 2 introduces a new *Preferences* page where you can adjust settings like language, tooltips and analyzer resolution. Click the menu button in the upper right corner and choose **Show Preferences...** to show the *Preferences* page:



Select the *Show Preferences* command as indicated above to show the *Preferences* page.

The following *Preferences* page appears:



The new *Preferences* page shows *language, tooltips, analyzer and piano keyboard settings*.

Settings available from the Preferences Page

| | |
|----------------|--|
| UI language: | Allows you to switch the user interface language. You will need to reopen the plug-in before the change has an effect. |
| Show tooltips: | Allows you to enable or disable the tooltips that appear when you rest the mouse cursor over a user interface control for a while. |

| | |
|--------------------------------|---|
| Analyzer block size (samples): | Affects the frequency versus time resolution in the analyzer. Larger block sizes will result in a better frequency resolution, but will at the same time reduce the time resolution. |
| Analyzer tilt (db/oct.): | You can tilt the frequency spectrum with an adjustable attenuation per octave. White noise will appear like a flat horizontal line when the tilt is set to 0 dB/octave. Pink noise can be set to appear like a flat horizontal line with the tilt set to 3 dB/octave. |
| Analyzer drop rate (dB/s): | Defines how quickly the analyzer adapts to lower signal levels and is specified in decibels per second. |
| Show piano keyboard: | Shows or hides the new piano keyboard display. |
| Tuning frequency (Hz): | The tuning frequency for the piano keyboard display. You can enter the center frequency of the note A4 in Hertz (440 is the most common tuning). |


7.1.9.2 Equalize Phase Modes

The mixed phase mode in *Acon Digital Equalize 2* is a new and unique feature and the difference between the available phase modes are described in detail in this chapter. The effect of a clean equalizer that doesn't emulate analog artifacts like saturation or noise can be fully described by the amplitude and the phase response. The amplitude response is what Equalize visualizes in the curve editor in the upper part of the plug-in window. The phase response is not visualized, however, because its impact on the processed signal wouldn't be visualized intuitively in a spectrum analyzer. However, the importance of the phase response is revealed when we measure the so called impulse response of the equalizer. The impulse response is measured by feeding the equalizer with a Dirac impulse. In the digital world with sampled time series, a Dirac impulse is nothing more than a zero signal consisting of one single full scale sample. The output is what we describe as the impulse response.


The illustrations below visualize the impulse responses obtained from a steep band pass filter (100 - 200 Hz) using the three phase modes that are available in Equalize along with the input signal at the top:




Input test signal consisting of a single pulse at 200 ms (Dirac impulse).



Minimum phase mode. Notice that the peak amplitude is shifted backwards in time.



Linear phase mode. The extensive non-zero signal prior to the peak is noticeable and disturbing. Another disadvantage is the long latency introduced by the filtering.



Mixed phase mode. The pre-ringing is limited to fit within the temporal masking of the human auditory system and the latency introduced is short. The amplitude peak is perfectly time aligned.

Traditionally, only the minimum and the linear phase modes have been available. The minimum phase mode is the only mode that ensures zero latency and is the way analog equalizers work. The downside is that the time alignment between different frequency components is lost.

The linear phase mode was introduced with the advent of digital processing, but the pre-ringing effect can distort transient rich material severely. Linear phase filters show advantages when processing tracks separately and there is crosstalk between the tracks. In such cases, linear phase processing has predictable results whereas minimum phase processing can alter the frequency response in unintentional ways.

The new mixed phase mode is a unique feature in Equalize and allows you to set the latency and hence the maximum pre-ringing time. The human auditory system masks signals before and after an impulsive noise, something which is called temporal masking. The masking is stronger after the impulsive noise than before, which is why extensive pre-ringing is undesirable. By choosing latencies around 20 ms in Equalize, the pre-ringing is masked by our hearing. At the same time, the maximum peak is perfectly time aligned.

7.1.10 Equalize Light

Standard Edition Only

Introduction

Equalize Light is a versatile, user friendly and great sounding equalizer with several unique features. Unlike other equalizers, you can freely adjust not only center frequency, gain and bandwidth, but also the filter slope for each band. The filter slope can be set anywhere from 3 to ultra sharp 120 dB per octave. Equalize Light is a zero latency plug-in with a minimum phase response.

Great care has been taken to provide a user interface that is straight forward to use. Band parameters can be adjusted using handles directly in the graphical representation of the frequency response, including bandwidths and filter slopes. A flexible real-time analyzer lets you monitor every aspect of the processing. You can easily switch between full, mid, side, left or right channel processing for each band and Equalize automatically routes the audio signal internally to ensure the best results and lowest possible latency.

User Interface

The graphical user interface of Equalize Light is designed to hide the complexity of the plug-in and provide an efficient workflow as well as intuitive control over the plug-in parameters.



Adding and Removing Bands

You can add up to 12 individual bands in *Equalize Light*. Each equalizer band represents a filter type, which can be either low cut, low shelf, bell, notch, high shelf or high cut. To add an additional band, click the + button below the curve display or double click in the curve display. A bullet shaped handle appears in the curve that you can move around using the mouse. The currently selected band also shows additional handles that you can use to manipulate the bandwidth (if applicable) and filter slope. You can remove the band by clicking the x button or by double clicking the handle.

Soloing and Bypassing Bands

You can solo or temporarily bypass an equalizer band in order to monitor its effect on the audio signal. You can enable the solo mode by pressing the Ctrl key while moving an equalizer band handle. The Shift key enables the bypass band mode. Alternatively, you can use the buttons to enable or disable the solo and bypass modes:

-  : Solo mode
-  : Bypass mode

Band Parameter Settings

The parameters that are related to a specific band are placed within the band group and the header indicates which band is currently active. The parameters apply to the currently active band only. To change the currently active band you can either click its bullet handle in the frequency response visualization or use the arrow buttons to browse through the active bands.

- **Frequency (Hz)**

The center (for peak and notch filters) or threshold frequency of the currently active equalizer band in Hertz.

- **Band gain (dB)**

The gain of the currently active equalizer band in decibels. This parameter is not available for high or low cuts, or the notch filter.

- **Bandwidth (oct.)**

The bandwidth of the currently active equalizer band in octaves. This parameter is only available for peak and notch filters.

- **Resonance (dB)**

Controls the resonance in decibels at the threshold frequency for cut and shelving filters. Resonance occurs in analog filters and results in a boost around the threshold frequency.

- **Slope (dB/oct.)**

The filter slope of the currently active equalizer band in decibels per octave. This controls the steepness of the filter.

- **Gain to Bandwidth Link**

The perceived bandwidth of an equalizer band is dependent on the gain setting. You can link the bandwidth to the gain setting, so that the bandwidth is automatically adjusted when you change the gain to preserve the perceived bandwidth.



: Click this button to activate or deactivate the gain to bandwidth link

- **Filter Type Buttons**

You can choose between six different filter types:



: Low cut, removes frequency content below the band frequency



: Low shelf, boosts or attenuates frequency content below the band frequency



: Peak filter, boosts or attenuates frequency content around the band frequency

-  : Notch filter, removes frequency content around the band frequency
-  : High shelf, boosts or attenuates frequency content above the band frequency
-  : High cut, removes frequency content above the band frequency

- **Channel Mode Buttons**

You can choose any of the following channel modes for each equalizer band individually:

- L : Apply to left channel only
- M : Apply to mid channel only. The mid channel is the sum of the input channels multiplied by a scaling factor.
-  : Apply to both channels
- S : Apply to side channel only. The side channel is the difference of the input channels multiplied by a scaling factor.
- R : Apply to right channel only

Global Parameters

- **Master gain (dB)**

The master gain controls the overall gain of the equalizer in decibels. Equalize can automatically adjust the overall gain to compensate for any gain changes in the equalizing process. This is based on an average distribution typically found in music and cannot be completely accurate. To activate the automatic gain compensation, click the button labelled A to the left of the master gain knob.

Analyzers and the Frequency Response Visualization

Equalize visualizes the frequency response of the current equalizer settings and contains two separate spectrum analyzers that you can set up to monitor the effects of the processing.

- **Analyzer 1 & 2**

You can choose to analyze different signal sources using the drop-down lists under the Analyzer 1 and Analyzer 2 headers. Either input or output signals can be analyzed and you can also choose if the left, right, mid or side signal should be analyzed.

- **Drop (dB/s)**

The drop parameter sets how quickly the analyzer adapts to lower signal levels and is specified in decibels per second.

- **Range Buttons**

You can change the range of the frequency response visualization or the spectrum analyzers by clicking the buttons at the lower end of the level axes. A drop down list appears with the alternatives.

7.1.11 Phono Filter

About Phono Filter

The phono filter emulates the effect of a phono preamplifier (deemphasis filter) or the opposite process applied when creating a master record (emphasis filter). It can also apply PCM type emphasis and deemphasis as found on DAT and early CD recordings. The analog amplitude frequency response is calculated accurately using analytic calculations and applied using a minimum phase filter to match the analog counterpart as closely as possible.

User Interface



Parameter Settings

- **Filter mode**

Choose deemphasis mode if you have recorded an LP record without a phono preamplifier. Choose Emphasise if you want to prepare an audio file for an LP master.

- **Output level (dB)**

Use the output level slider to compensate for the increase or decrease in audio level.

Advanced Settings

You can adjust the time constants as specified in analog emphasis and deemphasis circuitry. These constants (T1, T2 and T3) are often specified in literature about restoration of vintage recordings (before the RIAA standard). You can enter these time constants yourself under *Time constants*.

7.1.12 Remove DC Offset

About Remove DC Offset

Wrongly calibrated recording equipment may result in a signal that is not centered around zero as it should be. The *Remove DC Offset* tool calculates the DC offset of the selected region and subtracts the offset if necessary.

7.1.13 Rotate Phase

About Phase Rotation

Speech recordings often have asymmetric waveforms like in the example below.



Recordings of the human voice frequently exhibits asymmetric waveforms.

Rotating the phase equally over the full frequency bandwidth is usually inaudible and can restore the symmetry of the waveform. This has several advantages as you can increase

the volume without introducing clipping. Asymmetric waveforms also give limiters and compressors a tougher job than necessary. The *Rotate Phase* can rotate the phase with fixed rotation or adaptively rotate the phase for maximum symmetry.

User Interface



Parameter Settings

- **Adaptive**

The *Adaptive* mode button lets you toggle between adaptive and manual phase rotation. When the adaptive mode is enabled, the *Rotate Phase* tools will continuously analyze the audio and adjust the phase rotation for maximum symmetry.

- **Link channels**

The *Link channels* button lets you toggle between individual phase rotation per channel or phase rotations that are common to all channels.

- **Channel rotation knobs (deg.)**

When the adaptive mode is disabled, you can adjust the phase rotation manually in degrees using the channel rotation controls.

7.1.14 Invert Phase

The *Invert Phase* tools flips the polarity of the audio signal. This is equivalent to using the *Rotate Phase* tool with a fixed 180 degrees rotation.

7.2 Volume

The *Volume* menu in Acoustica contains a collection of audio processing tools related to volume manipulation.

7.2.1 Mute

The *Mute* command from the *Volume* menu mutes the current selection. You can also press Ctrl/Cmd+M.

7.2.2 Normalize

The *Normalize* tool in the *Volume* menu can be used to ensure a constant signal level in all your audio recordings. There are 5 different level measurement methods available to choose from.



The five different level measurement methods in Normalize

Measurement Methods:

- **ITU-R BS.1770 (EBU R-128)**

The ITU-R BS.1770 and the EBU R-128 recommendations aim at normalizing the perceived loudness and perform better than previous methods such as RMS. The ITU recommendation defines how to measure loudness using filters to match the frequency dependent sensitivity of the ears and how to apply absolute and relative gates. The EBU recommendation refers to the ITU standard and specifies specific values for the absolute gate (-70 LUFS) and the relative gate (-10 LU). The loudness is measured in LUFS (Loudness Unit Full Scale) or LKFS (in the older versions of the ITU standard).

- **True peak**

In the digital world, level meters and normalize tools often use sample peak detection. The problem with sample peak detection is that an audio signal that has no digital clipping might start to clip during the conversion back to analogue. These "in-between" peaks are often referred to as inter-sample peaks. The true peak level measurement method takes inter-sample peaks into account.

- **Sample peak**


The sample peak is the loudest sample value encountered.

- **Global RMS (AES)**

The Root-Mean-Square (RMS) level calibrated according to the AES17-1998 standard so that an input sine wave has the same peak level as the RMS level. This results in an AES-RMS level that is 3 dB higher than the mathematical RMS level.

- **Global RMS (Mathematical)**


The direct mathematical evaluation of Root-Mean-Square (RMS).



The Normalize tool offers several different level measurements methods.

7.2.3 Volume

The most basic volume manipulation tool in Acoustica is *Volume*. The only parameters are the volume change in Decibel and the *Invert signal* toggle:



The Volume level slider has a range from -96 dB to +32 dB

7.2.4 Channel Mixer


The channel mixer is a tool that works only on stereo and multichannel recordings. Each output channel is represented by a tab in a tab control and you mix several input channels to each output channel. The input channel levels are adjusted using the sliders. The input signal from each source channel can also be inverted using the *Invert* check boxes.



The Channel Mixer settings

7.2.5 Volume Curve

The *Volume Curve* tool lets you define volume curves by adding a set of curve points. You can add a curve point by moving the mouse cursor over the curve. The mouse cursor turns into a pointing hand. Now you can click the mouse button to create a new curve point and move it around with the using the mouse. Double click a curve point to remove it again.



The volume curve tool with the blue volume curve. Notice the focused curve point (white outline), the two arrow shaped handles at each side and the smooth curve before the focused curve point.

You can choose how to interpolate the curve between two curve points. Click one of the curve points so that it gets focused (indicated with a white outline). Now, two arrow shaped handles appear to the left and to the right of the curve point. You can drag these to get a smoother transition on either side.

7.2.6 Fade

The *Fade* tool lets you create fade-ins or fade-outs quickly. You can select a fade curve from one of mathematical functions from the *Fade function* drop-down list:

- Exponential
- Linear
- Sinusoidal
- Logarithmic
- Power

Select the direction of the fade using the *Fade in* or *Fade out* radio buttons. The range slider is only relevant to logarithmic and exponential fade curves and defines the minimum level in dB. The resulting fade curve is visualized in the graph at the bottom:



The Fade tool with a typical exponential fade-out.

7.2.7 Quick Fades

The commands *Quick Fade In* (Ctrl/Cmd + Alt + I) and *Quick Fade Out* (Ctrl/Cmd + Alt + O) adds a fade the current selection using and exponential curve with -40 dB range.

7.3 Effects

The *Effects* menu in in Acoustica contains a collection of common effect processors such as reverb, echo, chorus and more.

7.3.1 Verberate

Premium Edition Only

About Reverb

Reverberation occurs when sound is produced in an enclosed acoustical environment. Even outdoors, there is likely to be some level of reverberation, however subtle. The sound propagates through the air before it arrives at the listener, but the sound is also reflected when it hits walls or other objects. Due to the propagation time, these reflections arrive at the listener later than the sound from the direct path. After a certain build-up

time, there are usually so many reflections that no distinct echoes are distinguishable, but rather a smoothly decaying sound.

The first few reflections, usually called early reflections, are important cues for our perception of an acoustical environment. For that reason, most digital reverberation units differentiate between early reflections and the dense late reverberation.

Verberate, an algorithmic reverb from Acon Digital which is integrated in Acoustica Premium Edition 7, simulates both the early reflections and the dense reverberation in a way that comes extremely close to what can be measured in a real acoustical environment. Several conventional building blocks in traditional digital reverb design have been replaced by novel methods that allow Verberate to produce convincing early reflections and the dense reverb tails that characterize impulse response measurements of real acoustical surroundings.

User Interface



Parameter Settings

- **Dry level (dB)**

The amount of dry (unprocessed) signal to send to the output specified in decibel. You can use the toggle button to exclude the dry signal completely. You can lock the relation between the dry and reverb level by clicking the lock button between the dry and reverb level sliders.

- **Reverb level (dB)**

The amount of reverberation signal to send to the output specified in decibel. You can use the toggle button to exclude the late reverberation completely.

- **ER level (dB)**

The amount of early reflections to send to the output specified in decibel relative to the reverb level. The early reflections are important for our perception of distance to the sound source. Increasing the early reflection level gives the impression of getting closer to the source. You can use the toggle button to exclude the early reflections completely.

- **Early reflections mode**

You can choose the early reflection program from a set of predefined programs, such as rooms, chambers, halls and plates.

- **Reverb time (s)**

The reverb time specifies the duration of the reverberation and is specified by the number of seconds before the reverb tail level drops below -60 dB or 1/1000 of its initial amplitude.

- **Room size (%)**

The room size determines the size of the simulated acoustical environment. In general, the room size should match the reverberation time. Using a very long reverberation time combined with small room sizes may result in resonances causing a slightly "metallic" quality. In some cases, as with a couple of the plate presets, the resonances are desired to achieve a specific sonic quality.

- **Pre-delay (ms)**

The pre-delay slider allows you to adjust the time in milliseconds before the reverberation signal arrives.

- **Modulation rate (Hz)**

Modulation can be applied to the late reverberation to simulate fluctuations in the arrival times of the sound waves and you can control how rapid the fluctuations are, by using the modulation rate.

- **Modulation depth (%)**


The modulation depth controls the depth of the modulation ranging from no modulation (0%) to full modulation (100%).

- **Stereo spread (%)**

You can use the stereo spread parameter to control the stereo width of the reverberation signal. If the stereo spread is set to 0%, the reverberation signal will be mono and at 100% the full stereo width is achieved.

- **Decay Editor**


The decay editor allows you to set up reverberation times that are dependent of the frequency. Air absorbs high frequencies to a larger degree than lower frequencies, so the low frequencies will in general have a longer reverberation time in natural spaces. Verberate has two filters controlling the decay, a shelf filter and a low-pass filter. You can control the cut-off frequency of the low-pass filter by moving the **green** handle. The current cut-off frequency in Hertz is displayed as cursor information while you do the adjustment. Additionally, a shelf filter allows you to increase or reduce the reverberation time below a given cross-over frequency. You can control the gain (vertical movements) and the frequency (horizontal movements) using the red handle. As with the low-pass filter, cursor information is displayed with the cross-over frequency in Hertz as well as the reverberation time multiplication factor (relative to the current reverberation time) in percent. At all times, the current decay profile is displayed in terms of the relative reverberation time as a function of frequency.



- **Output Equalizer**

You can use the output equalizer to apply filtering to the reverberation signal. Both the early reflections and the dense reverberation are filtered using the output equalizer. The equalizer consists of high and low pass filters as well as high and low shelving filters. As in the decay editor, you can use handles to change the filter settings and the frequency in Hertz as well as the gain in dB are displayed as cursor information while

making adjustments. The filters in the output equalizer have variable filter slopes. You can modify the filter slope by clicking the filter section you want to edit. A small arrow appears in the same color as the filter section. You can move this arrow using the mouse to modify the slope as depicted below:



7.3.2 Reverb (Standard Edition)


Standard Edition Only

About Reverb

Reverberation occurs when sound is produced in an enclosed acoustical environment. Even outdoors, there is likely to be some level of reverberation, however subtle. The sound propagates through the air before it arrives at the listener, but the sound is also reflected when it hits walls or other objects. Due to the propagation time, these reflections arrive at the listener later than the sound from the direct path. After a certain build-up time, there are usually so many reflections that no distinct echoes are distinguishable, but rather a smoothly decaying sound.

The first few reflections, usually called early reflections, are important cues for our perception of an acoustical environment. For that reason, most digital reverberation units differentiate between early reflections and the dense late reverberation.

User Interface



Parameter Settings

- **Dry level (dB)**

The amount of dry (unprocessed) signal to send to the output specified in decibel. You can use the toggle button to exclude the dry signal completely. You can lock the relation between the dry and reverb level by clicking the lock button between the dry and reverb level sliders.

- **Reverb level (dB)**

The amount of reverberation signal to send to the output specified in decibel. You can use the toggle button to exclude the late reverberation completely.

- **ER level (dB)**

The amount of early reflections to send to the output specified in decibel relative to the reverb level. The early reflections are important for our perception of distance to the sound source. Increasing the early reflection level gives the impression of getting closer to the source. You can use the toggle button to exclude the early reflections completely.

- **Early reflections mode**

You can choose the early reflection program from a set of predefined programs, such as rooms, chambers, halls and plates.

- **Reverb time (s)**

The reverb time specifies the duration of the reverberation and is specified by the number of seconds before the reverb tail level drops below -60 dB or 1/1000 of its initial amplitude.

- **Room size (%)**


The room size determines the size of the simulated acoustical environment. In general, the room size should match the reverberation time. Using a very long reverberation time combined with small room sizes may result in resonances causing a slightly "metallic" quality. In some cases, as with a couple of the plate presets, the resonances are desired to achieve a specific sonic quality.

- **Pre-delay (ms)**

The pre-delay slider allows you to adjust the time in milliseconds before the reverberation signal arrives.

- **Output Equalizer**

You can use the output equalizer to apply filtering to the reverberation signal. Both the early reflections and the dense reverberation are filtered using the output equalizer. The equalizer consists of high and low pass filters as well as high and low shelving filters. As in the decay editor, you can use handles to change the filter settings and the frequency in Hertz as well as the gain in dB are displayed as cursor information while making adjustments. The filters in the output equalizer have variable filter slopes. You can modify the filter slope by clicking the filter section you want to edit. A small arrow appears in the same color as the filter section. You can move this arrow using the mouse to modify the slope as depicted below:



7.3.3 Convolve

Introduction to Convolve

Convolve is a convolution reverb which can be used to apply recorded impulse responses from real acoustic spaces. This makes it fundamentally different from algorithmic reverbs, like *Verberate*. Convolution reverb is perfect for adding lifelike ambiance to a recording, like the ambiance from a room, a hall or a parking garage for example.


However, compared to an algorithmic reverb which often offers control over the decay time and room size for example, *Convolve* doesn't offer you this control, it totally relies on the impulse response you load.

There are a lot of amazing free impulse responses available on the net. The following two websites are recommended:

<http://www.openairlib.net/>

<http://www.echothief.com/>

User Interface



Parameter Settings

- **Dry level (-48 dB to 12 dB)**

With this slider, you can control the amount of dry (unprocessed) signal to send to the output, specified in decibel.

- **Reverb level (-48 dB to 12 dB)**

With this slider, you can control the amount of processed signal to send to the output, specified in decibel.


- **Pre-delay (0 ms to 1000 ms)**

The pre-delay knob allows you to adjust the time in milliseconds before the reverberation signal arrives.

- **Topology**

This drop down menu offers three different topologies. *Mono input* will sum the input and processes the sum using the left and right channel of the IR file to create a stereo output. *Parallel channels* will process the left input channel with the left channel of the IR file and the right channel with the right channel of the IR file independently. *Matrix processing* requires one IR file for each input channel where each IR file defines how to process the input channel to the set of output channels. This is the only option that is "true stereo", which means that you will be able to hear the panning of the input signal in the reverb signal in a natural way.

When you select Matrix processing (true stereo) under Topology, you will need to load two IR files, one for the left channel and one for the right channel. This will look like this:



- **IR file:**

Here you can load an impulse response file. This file can be of any audio file format supported by Acoustica.

7.3.4 Echo

About Echo

The echo effect is a multi-tap delay effect. Multi-tap means that you can add several delays (up to eight) with arbitrary delay times and gains. Two different timing modes are offered, the BPM (Beats per minute) mode or the milliseconds mode. In the BPM mode, the time delay of each tap is specified in beats.

User Interface



Parameter Settings

- **Dry Level**
The amount of unprocessed signal in the output mix.
- **Echo Level**
The amount of processed signal in the output mix.
- **Echogram**
This depicts the gain versus time for the configured taps. You can add new delay taps by clicking the left mouse button where you want the new delay tap to appear. To move a delay tap, click an existing point and keep the mouse button down while moving the mouse pointer to the new location. You can remove an existing delay tap by double clicking the tap you want to remove.
- **Range**
This button allows you to modify the amount of time shown in the echogram. With the time mode in seconds, this can be set between 0.1 second and 5 seconds. With the time mode in beats, this can be set between 1/16 beat and 8 beats.

- **Delay (s) or Delay (beats)**

This lets you specify the amount of time audio is delayed (in seconds or beats) before it is sent to the output. In the echogram, it affects the horizontal position of the tap.
- **Gain (dB)**

The individual tap is present by this amount in the output. In the echogram, it affects the vertical position of the tap.
- **Feedback (dB)**

The feedback percentage specifies the amount of attenuation since the last delay interval.
- **High cut frequency (Hz)**

This allows you to change the cut-off frequency of the low pass filter in the feedback loop. The low pass filter can be enabled or disabled by clicking the check box underneath the knob.
- **Ping pong echo**

This check box provides for a bouncing stereo delay.
- **Time Mode**

Select either beats per minute for the BPM mode or seconds.
- **Beats per minute**

Here you can specify the tempo by entering the number of beats per minute.
- **Snap mode**

To make it easier to adjust the delay of each tap without losing the alignment to the tempo, you can optionally turn on the snap mode. This option will only be enabled when the time mode is in beats.

7.3.5 Chorus

Introduction to Multiply

Acoustica 7 comes with Acon Digital Multiply, which is a versatile chorus effect with a unique twist. Each simulated voice is processed with a phase randomizing filter so that unpleasant comb filter effects are omitted. The effect can be used to simulate the effect of several performers playing the same tones simultaneously, to widen the spatial image or to create special effects for sound design. Multiply can simulate up to 6 additional voices and both the pitch and the loudness of the voices can be modulated. There is also an integrated equalizer consisting of a low cut, low shelf, high shelf and high cut filters that

can be applied to the effect signal. An integrated pre-delay section makes it possible to create modulated and diffuse echo effects.

The graphical user interface of Multiply is designed to give quick and intuitive access to all the parameter settings defining the sonic quality of the effect.

User Interface



Parameter Settings

- Dry level (dB)**
 The amount of dry (unprocessed) signal to send to the output specified in decibel. You can use the toggle button to exclude the dry signal completely. You can lock the relation between the dry and effect levels by clicking the lock button between the dry and effect level sliders.
- Effect level (dB)**
 The amount of processed signal to send to the output specified in decibel. You can use the toggle button to exclude the effect signal completely.
- Frequency modulation rate (Hz)**

The frequency modulation rate controls how rapid the tone fluctuations in the simulated voices should be. It is specified in Hertz.

- **Frequency modulation depth (%)**

The frequency modulation depth controls the amount of tone fluctuations in the simulated voices. It is specified in percent, ranging from no modulation (0%) to full modulation (100%).

- **Amplitude modulation rate (Hz)**

The amplitude modulation rate controls how rapid the loudness variations in the simulated voices should be. It is specified in Hertz.

- **Amplitude modulation depth (%)**

The amplitude modulation depth controls the amount of loudness variations in the simulated voices. It is specified in percent, ranging from no modulation (0%) to full modulation (100%).

- **Voice count (#)**

The amount of simulated voices.

- **Stereo spread (%)**


You can use the stereo spread parameter to control the stereo width of the effect signal. If the stereo spread is set to 0%, the effect signal will be mono when the input signal is mono and at 100% the full stereo width is achieved.

- **Pre-delay (ms)**

The pre-delay slider allows you to adjust the time in milliseconds before the effect signal arrives.

- **Equalizer**

You can use the output equalizer to apply filtering to the effect signal. The equalizer consists of high and low pass filters as well as high and low shelving filters. You can use handles to change the filter settings and the frequency in Hertz as well as the gain in dB are displayed as cursor information while making adjustments. The filters in the equalizer have variable filter slopes. You can modify the filter slope by clicking the filter section you want to edit. A small arrow appears in the same color as the filter section. You can move this arrow using the mouse to modify the slope as depicted below:



7.3.6 Modulate

About Modulate

Modulate is a modulation tool which offers a flanger and a phaser effect.

The flanging effect occurs when two sources playing exactly the same recording with a short time delay are mixed together. The result is that some frequencies are canceled, while others are amplified. Changing the time delay between the two sources will result in other frequencies being canceled or amplified (the name comes from years ago when studio engineers using open-reel tape would manually turn or hold one reel or "flange" to create this effect, distorting the tape speed past the playback heads).

The phaser effect uses several notch filters with a time-varying center frequency to create a sweeping or electronic "whooshing" effect.

User Interface



Parameter Settings

- Dry Level (dB)**
 Use the dry level slider to adjust the amount of unprocessed signal in the output mix.
- Effect Level (dB)**
 Use the effect level slider to adjust the amount of processed signal in the output mix.
- Effect type**
 You can select the modulator mode by selecting *Flanger* or *Phaser* from the list.
- Modulation shape**
 The delay (in the flanger mode) or the notch center frequency (in phaser mode) can be modulated by one of four different modulation sources. To change the modulation source, select sinusoidal, triangular, square, or random from the list.
- Modulation depth (ms)**
 Use the modulation depth knob to set the depth or amount of the modulation.
- Modulation frequency (Hz)**
 The modulation frequency knob allows you to set the frequency or the speed of the modulation function. A higher modulation frequency will result in faster changes in the tone quality.
- Delay (ms)**


You can add an additional fixed delay to the effect by adjusting the delay slider.

7.3.7 Time Stretch

About Time Stretch

Time Stretch allows you to change the length of an audio recording without changing the perceived pitch.

User Interface



Parameter Settings

- **New duration (25% to 400%)**

Use this slider to change the duration of the audio. Anything below 100% will decrease the duration (speed up), while anything above 100% will increase the duration (speed down).

- **Original duration**

In the left field you can fill in the original duration. By filling in the desired duration in the right field, *Time Stretch* will automatically be adjusted to accommodate the new duration.

- **Original tempo (BPM)**


In the left field you can fill in the original tempo in BPM. By filling in the desired tempo in the right field, *Time Stretch* will automatically be adjusted to accommodate the new tempo.

7.3.8 Transpose

Introduction to Transpose

The transpose tool allows you to change the pitch of a recording, with or without changing the tempo. In many cases, large pitch changes lead to unnatural sounding results. The effect is especially pronounced when changing the pitch of the human voice, where higher pitched voices sound more like Disney's chipmunks than a human voice. The transpose tool has an optional *Maintain timbre* option that reduces this artifact. When *Maintain timbre* is checked, Acoustica creates a smooth spectral envelope estimation of the signal and whitens the signal before pitch shifting. The original smooth spectral envelope is applied after transposing and the original timbre is preserved

User Interface



Parameter Settings

- **Interval**
The musical interval to transpose. Use the radio buttons to set the transpose direction to up or down.
- **Fine tune (-100 cents to 100 cents)**
With this slider you can fine tune the pitch shift factor in cents which are 1/100 of a semitone.
- **Maintain duration**
Check this to keep the duration of the audio intact.
- **Maintain timbre**
Check this to keep the timbre of the audio, this is especially useful for preventing the "chipmunk" effect when transposing a vocal.

- **Timbre smoothing (Hz)**

This knob adjusts the smoothing when *Maintain timbre* is enabled and larger values will lead to more smoothing.

7.3.9 Harmonizer

About Harmonize

Harmonize mixes several pitch shifted voices to create interesting harmonies. You can mix up to four pitch shifted voices. The often experienced "chipmunk" effect which occurs when transposing the human voice or musical instruments can be reduced using the *Maintain timbre* option. When *Maintain timbre* is checked, Acoustica creates a smooth spectral envelope estimation of the signal and whitens the signal before pitch shifting. The original smooth spectral envelope is applied after transposing and the original timbre is preserved.

User Interface




Parameter Settings (identical for each voice)

- **Voice #**
Check this to activate the current voice
- **Interval**
The musical interval to transpose. Use the radio buttons to set the transpose direction to up or down.
- **Fine tune (-100 cents to 100 cents)**
With this slider you can fine tune the pitch shift factor in cents which are 1/100 of a semitone.
- **Volume (-48 db to 12 dB)**
Volume of the current voice in dB.
- **Pan (%)**
Left / right panning of the current voice in percent.
- **Maintain timbre**
Check this to keep the timbre of the audio, this is especially useful for preventing the "chipmunk" effect when transposing a vocal.
- **Timbre smoothing**
This knob adjusts the smoothing when *Maintain timbre* is enabled and larger values will lead to more smoothing.


7.3.10 Reverse

About Reverse

The reverse tool in Acoustica 7 allows you to reverse an entire audio recording, or just a selection of a recording.



Original audio



Reversed audio

7.4 Enhancement

The *Enhancement* menu in Acoustica contains a collection audio restoration and enhancement tools such as noise reduction, de-clipping and more.

7.4.1 Interpolate

You can remove clicks, discontinuities or other short impulsive noise using the *Interpolate* tool which substitutes the selected time range with an estimation of the signal based on the signal before and after the the selection. You should select the shortest possible time range which includes noise before you choose *Interpolate* from the *Enhancement* menu or press Ctrl/⌘+I.

Note: Interpolation is only possible on time ranges that are shorter than 200 ms.

7.4.2 DeClip

About DeClip

DeClip 2 restores audio recordings distorted by analog or digital clipping. Clipping occurs during recording when the recording level is too high and the highest peaks cannot be correctly recorded. *DeClip 2* substitutes such distorted peaks with an estimation of the signal curve in such a way that the frequency content obtained from the sampled values in the reliable range (i.e. the parts of the waveform not affected by clipping) is preserved as far as possible.



The DeClip 2 user interface.

DeClip 2 shows a histogram of the signal level distribution in order to visualize restoration process and simplify the adjustment of the threshold values. *DeClip 2* will substitute all recorded signals above the upper and below the lower threshold value with

an estimate of the signal. The threshold values can be adjusted using their corresponding slider controls or in the histogram by dragging the colored handles.

Settings

- **Positive threshold (dBFS)**

All samples values above the positive threshold value specified in dBFS are substituted by an estimation of the signal.

- **Negative threshold (dBFS)**

All samples values below the negative threshold value specified in dBFS are substituted by an estimation of the signal.

- **Link threshold values**

Usually, the clipping introduced during recording will be symmetrical, which means that the upper and lower thresholds will have the same absolute value. By activating the upper and lower threshold link, the adjustment of the declipper is simplified in the case of symmetrical clipping.

- **Detect**

By clicking this button, *DeClip 2* will automatically detect the threshold values. Please note that this button requires audio input to work properly.

- **Input gain (dB)**

The input gain is specified in dB and useful for adjusting the signal level before the de-clipping process and adding enough headroom for the restoration process.

- **Output gain (dB)**

The output gain in dB allows you to make up for peak level changes caused by the de-clipping process.

- **Quality (%)**

The quality slider controls the quality at the cost of CPU usage of the algorithm. Increased quality factors lead to higher CPU usage.

7.4.3 DeClick

About DeClick

DeClick 2 removes impulsive noise such as clicks and crackle. These distortions are very frequently encountered on LP and 78 RPM records, but also occur in digital recordings to drop-outs or distorted data packets. DeClick contains two different algorithms to deal with clicks and crackle. The actual declicker algorithm eliminates large clicks and pops in the recording, while the decrackler algorithms eliminates the frequent, but short clicks

that the human ear perceives as crackle. DeClick removes clicks by substituting the recorded signal in the short period of time during the click with a signal estimated using the undistorted audio surrounding each click.

DeClick also features a third algorithm, which is designed to reduce plosives and thumps. These distortions are typically found in recordings from scratched vinyl records or in vocal recordings, where words that start with a P or B can cause a strong blast of air to reach the microphone diaphragm, resulting in a thump sound.



The DeClick 2 user interface.

The DeClick user interface contains a reduction meter that gives visual feedback of the restoration process. It shows a history of the reduction activity during the past ten seconds. The red line shows the number of clicks removed per second, where the green line shows the number of crackles removed per second.

Settings

- Crackle sensitivity (%)**
 Sets the sensitivity of the decrackler algorithm. Higher crackle sensitivity values result in more crackle reduction. You can use the toggle button to activate or deactivate the decrackler algorithm.
- Click sensitivity (%)**
 Sets the detection sensitivity of the declicker algorithm. Higher click sensitivity values result in more click reduction. You can use the toggle button to activate or deactivate the declicker algorithm.
- Plosive & thump sensitivity (%)**

Sets the sensitivity of the dethump algorithm. Higher plosive & thump sensitivity values result in more thump reduction. You can use the toggle button to activate or deactivate the dethump algorithm.

- **Click expand (ms)**

The amount of smoothing applied to the click detection signal. This will expand the average duration of the detected clicks and should be kept as small as possible, while still covering the clicks. This can be set between 0.1ms and 5.0ms.

- **Click time skew (%)**

The shape of the clicks, where 0% represents a symmetrical clip. Positive values should be used when the clicks cause more post ringing.

- **Plosive & thump cut-off (Hz)**

The upper frequency limit for the dethump algorithm. This can be set between 20 Hz and 350 Hz.

- **M/S mode**

The Mid / Side (M/S) mode is only available when working with source material in stereo. When enabled, the incoming audio is converted to a M/S signal prior to processing and decoded back to normal left / right stereo at the output.

- **Solo noise**

You can use this to listen to the removed signal, which can be considered an audible representation of the reduction meter.

7.4.4 DeHum

About DeHum

DeHum 2 is designed to remove hum and buzz typically introduced by poorly grounded electrical equipment, but also other tonal noise sources like electrical motor noise. The dehum algorithm also supports adaptive hum reduction so that the algorithm adapts to fluctuations in the fundamental frequency of the hum signal. Real life hum noise is likely to consist of a fundamental frequency and a set of harmonic frequencies. These are multiples of the fundamental frequency. DeHum allows you to set the number of harmonics to remove and also has the option to address only odd harmonics, since hum noise with only odd harmonics are frequently encountered.

There are two different operating modes and you can set the operating mode using the *Notch filter* toggle button. When *Notch filter* mode is disabled, DeHum subtracts a hum signal reconstructed using a sinusoidal re-synthesis technique in order to minimize distortions of the wanted signal. With *Notch filter* mode enabled, the hum reduction is performed using conventional notch filters.



The *DeHum 2* user interface with green and red lines indicating the hum harmonics. The lines turn green when hum is detected.

Settings

- **Frequency (Hz)**

The *Frequency* knob sets the fundamental frequency of the hum or buzz noise. If the hum originates from the power distribution net, the fundamental frequency should be set to either 60 Hz (American standard) or 50 Hz (European standard), depending on frequency of the AC power distribution in the country the recording was made. You can let *DeHum 2* fine-tune this fundamental frequency automatically using the *Scan* button on the lower left hand side of the *Frequency* control knob.

- **Sensitivity (%)**

The sensitivity parameter is only available when the *Notch filter* mode is deactivated. Higher sensitivity values cause the hum reduction algorithm to classify more frequency components as hum.

- **Adaptivity (Hz / s)**

The adaptivity knob controls the maximum fluctuation of the fundamental frequency in number of Hertz per second that is allowed in the detection of the fundamental frequency. This value should be as low as possible while still detecting the fluctuations of the hum signal being removed.

- **Harmonics**

The number of harmonics to detect and remove. This should be as low as possible while still removing all the harmonics present in the hum noise.

- **Odd harmonics**

Click this toggle button to reduce only odd harmonics. Hum noise consisting of a fundamental frequency with only odd harmonics are frequently encountered in real life situations and are typically the result of sine wave signal with a symmetrical non-linear distortion.

- **Notch filter mode**

The *Notch filter* button toggles between the notch filter and the sinusoidal re-synthesis modes. The advantages of *Notch filter* mode is that it doesn't introduce any latency and it consumes less CPU. However, the sinusoidal re-synthesis mode introduces considerably less distortion to the wanted signal (you can enable the *Solo hum* button to hear the difference).

- **M/S mode**

The Mid / Side (M/S) mode is only available when working with source material in stereo and it avoids drifting in the stereo image after processing. When enabled, the incoming audio is converted to a M/S signal prior to processing and decoded back to normal left / right stereo at the output.

- **Solo hum**

Click this button if you wish to monitor the signal removed by the hum reduction algorithm.

7.4.5 About De-noising

The integrated *DeNoise* tool targets stationary noise such as broadband noise, hiss, wind noise, buzz and camera noise. Great efforts have been put into preserving as much of the original signal as possible during the noise reduction process. As a result, *DeNoise* can reduce or in many cases completely remove the noise in a very transparent manner without loss of transients, attacks or "air" in the recording.

In the Premium Edition, *DeNoise* operates either in an adaptive mode or by learning from a selection containing noise only. The *Standard Edition* does not offer the adaptive mode. Regardless of the operating mode, the noise reduction algorithm requires an estimate of the expected frequency distribution of the noise called noise profile. In the adaptive mode, the noise profile is evaluated constantly using advanced statistical methods.

There are several ways to estimate the noise profile and access DeNoise in Acoustica which are described in the following chapters.

7.4.6 Analyze Noise

If your recording contains pauses with pure noise without any other signal, the pause can be used to create a noise profile automatically. Select *Enhancement > Analyze Noise...* after selecting a time range containing noise only. Acoustica automatically opens the *DeNoise* tool after the analysis phase with the result of the analysis automatically loaded as the noise profile. Now you can select the part you want to de-noise. Please see [DeNoise](#)^[129] for more information about how to use the *DeNoise* tool.

7.4.7 Automatic DeNoise

The easiest way of removing stationary noise is to use the *automatic noise reduction*. Select *Enhancement > Automatic DeNoise...* Acoustica then performs a statistical analysis of recording in order to estimate the noise profile. This process might take a little while, depending on the length of your recording. After the analysis phase, the *DeNoise* tool is opened with the estimated noise profile. Please see [DeNoise](#)^[129] for more information about how to use the *DeNoise* tool.

7.4.8 DeNoise

Premium Edition Only

DeNoise 2 reduces noise such as hiss, wind noise, buzz and camera noise. It operates either in an *adaptive mode* or in *noise profile mode*, by learning from a selection containing noise only. Regardless of the operating mode, an accurate estimation of the frequency content of the noise is crucial for a good result.

Adaptive De-Noising

In the adaptive mode, the noise frequency spectrum is constantly estimated using statistical methods. You can choose to address only broadband noise (noise type set to *broadband*) or both tonal and broadband noise (noise type set to *combined*). Broadband noise doesn't have a perceivable pitch and tape hiss is one example. Addressing only broadband noise reduces the probability of affecting the wanted signal, but noise that has tonal elements such as buzz will not be removed unless the noise type is set to *combined*.

You can choose how quickly *DeNoise 2* responds to changes in the noise floor by adjusting the *Adaptation time* slider. However, the risk of affecting the wanted signal increases with

lower adaptation times. For dialogue, one second is usually sufficient, but we recommend longer adaptation times for music.

De-Noising Based on Noise Measurements

If there's a section of the noise without the wanted signal available, you can guide the noise estimation algorithm by creating a *Noise Profile*. Traditionally, these *noise profiles* (sometimes also called *noise finger prints*) are stationary, which means that their frequency distribution doesn't vary over time. However, Version 2 of *DeNoise* introduces the new *Dynamic profiles* that measure how noise fluctuates over time. A common example is wind noise. *DeNoise 2* is able to gather statistics that is used to estimate a noise floor that changes rapidly over time.



The *DeNoise 2* plug-in window. The graph shows an estimation of the noise (yellow curve) as well as frequency spectrum (filled gray curve) of the input signal.

General Settings

- **Reduction (dB)**

The Reduction knob controls the noise estimation level and allows you to remove more (positive values) or less (negative values) noise than the estimation algorithm detected.

- **Soft knee (%)**

The soft knee parameter reduces the steepness of the transition between noise and the wanted signal. Higher soft knee values result in a more natural transition.

- **Maximum attenuation (dB)**

Maximum attenuation allows you to adjust a maximum attenuation factor for each frequency band, and thus control the noise floor after processing.

- **Reaction time (ms)**


The reaction time controls the maximum time for the algorithm to respond to level changes in each frequency band. Longer reaction times effectively reduces tonal and burst-like processing artifacts.

- **M/S Mode**

The Mid / Side (M/S) mode is only available when working with source material in stereo and it avoids drifting in the stereo image after processing. When enabled, the incoming audio is converted to a M/S signal prior to processing and decoded back to normal left / right stereo at the output.

- **Emphasis**

The emphasis filter allows you to apply frequency weighting to the noise profile estimate. This is very useful if you wish to make manual corrections to the estimated noise profile. The frequency weighting curve consists of a low shelf filter, two peak filters and a high shelf filter, similar to a parametric equalizer. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed.. You can also change the filter slope of the shelving filters or the bandwidth of the peak filters. Click the handle of the filter you wish to modify. Arrows surrounding the handle will appear. Move these to change the bandwidth for peak filters or the filter slope for the shelving filters.



The Emphasis filter curve boosts or reduces the noise reduction in frequency bands similar to an equalizer.

- **Solo Noise**

Enable this option if you wish to monitor the signal removed by the noise reduction algorithm.

Noise Profile Section

The *Noise Profile* section contains user interface elements for de-noising based on a noise measurement. The first step is to measure the noise to remove using the learn button:




DeNoise 2 in Noise Profile mode with the Learn button activated

Now play a section that contains the noise you want to remove in your host application. No noise reduction is performed when the learning mode is activated. You can now click the button with the *power on* icon to enable the profile based de-noiser. If your noise varies over time, you can click the *Dynamic profile* button to enable the new dynamic noise profile mode in *DeNoise 2*. In the dynamic mode, *DeNoise 2* will constantly adapt the noise floor to fit the input signal.

Adaptive Section

The *Adaptive* section contains user interface elements for adaptive de-noising:



DeNoise 2 in Adaptive mode

The Noise type can be set to one of two types:

- *Broadband*, which is noise without a perceivable pitch such as tape hiss. When this noise type is selected, *DeNoise 2* will not try to reduce tonal noise
- *Combined*, which is noise that can have tonal components such as hum and buzz. The risk of affecting the wanted signal is higher when choosing this noise type

The *Adaption time* controls how quickly *DeNoise 2* will respond to changes in the noise floor and is specified in seconds. A shorter adaptation time will reduce the time before

DeNoise 2 is able to reduce noise effectively after an increase in noise level, but at the cost of a higher risk of affecting the wanted signal.

7.4.9 DeNoise Light

Standard Edition Only

The *DeNoise Light* tool targets stationary noise such as broadband noise, hiss, wind noise, buzz and camera noise (please see [About De-noising](#)^[128]).

User Interface



The *DeNoise* plug-in window in *Acoustica Standard Edition*. The graph shows the current noise profile as well as frequency spectrum of the input signal.

General Settings

- **Reduction (dB)**

Reduction factor scales the estimated noise profile and allows you to remove more (positive values) or less (negative values) noise than the analysis algorithm detected.

- **Soft knee (%)**

The soft knee parameter reduces the steepness of the transition between noise and the wanted signal. Higher soft knee values result in a more natural transition.

- **Maximum attenuation (dB)**

Maximum attenuation allows you to adjust a maximum attenuation factor for each frequency band component. By leaving a certain noise floor, you can mask artifacts from the noise reduction algorithm.

- **Listen to removed signal**


Enable this option if you wish to monitor the signal removed by the noise reduction algorithm.

- **Learn button**

Enables or disables the noise profile learning mode. When you use the learning mode, you should select a region containing only noise in your audio editor or digital audio workstation (DAW). Play the section containing noise only through Acon Digital DeNoise. No noise reduction is performed when the learning mode is activated.

- **Use emphasis filter**

The emphasis filter allows you to apply frequency weighting to the noise profile estimate. This is very useful if you wish to make manual corrections to the estimated noise profile. The frequency weighting curve consists of a low shelf filter, two peak filters and a high shelf filter, similar to a parametric equalizer. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed.. You can also change the filter slope of the shelving filters or the bandwidth of the peak filters. Click handle of for the filter you wish to modify. Arrows appear surrounding the handle. Move these to change the bandwidth for peak filters or the filter slope for the shelving filters.



Acon Digital DeNoise running with the emphasis filter enabled.

7.4.10 Extract:Dialogue

Premium Edition Only

Extract:Dialogue separates dialogue from common types of background noise such as wind, rustle, traffic, hum, clicks and pops. The algorithm is based on deep learning and has been trained on thousands of high quality voice recordings on one hand and an equally extensive set of recordings of unwanted sound. The extensive training enables the artificial intelligence to automatically distinguish dialogue from noise without user interaction. This makes *Extract:Dialogue* extremely easy to use. Despite the internal complexity, *Extract:Dialogue* is extremely easy to use. The default setting will immediately and fully automatically remove background noise. In many cases, however, you might not want to remove the noise entirely. You can also adjust the sensitivity of the dialogue extraction. The sensitivity is adjustable for the full frequency band or in frequency bands. A spectrum analyzer analyzes the incoming audio signal (the filled gray curve) and indicates the removed signal (the yellow curve):



Extract:Dialogue detects dialogue and removes background noise. The spectrum analyzer shows the frequency spectrum of the input audio (filled gray curve) and the frequency spectrum of the removed signal (yellow curve).


Parameter Settings

- **Sensitivity (%)**

You can adjust the relative sensitivity of the noise detection with the *Sensitivity* slider. Positive values will increase the noise sensitivity and remove more noise. Negative values will reduce the noise sensitivity and leave more noise.

- **Sensitivity Curve**

The sensitivity filter allows you to control the noise sensitivity in adjustable frequency bands. There are three filters available that you can toggle on or off using the buttons in the upper right corner of the spectrum analyzer: A low shelf filter (☞), a peak filter (⊖) and a high shelf filter (☜). These work the same way as in parametric equalizers. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed. You can also change the filter slopes or the bandwidth of the peak filters. Click the handle of the filter you wish to modify. Arrows surrounding the handle will appear that you can use to change the bandwidth or the filter slope of the selected filter.



The sensitivity curve lets you change the noise sensitivity in separate frequency bands.

- **Maximum attenuation (dB)**

In many cases, it is desirable to leave some noise and you can use the *Maximum attenuation* slider to control this. The *Maximum attenuation* is set in dB specifies the maximum attenuation if the frequency components, thus control the noise floor after processing.

- **Solo Noise**


Enable this option if you wish to monitor the noise signal that was removed

7.4.11 DeWind:Dialogue

Premium Edition Only

DeWind:Dialogue reduces wind noise in dialogue recordings. The algorithm is based on deep learning and has been trained on thousands of high quality voice recordings and wind noise recordings. The default setting will immediately and fully automatically remove wind noise. In many cases, however, you might not want to reduce the noise

entirely. You can also adjust the sensitivity of the noise detection. The sensitivity is adjustable for the full frequency band or in frequency bands. A spectrum analyzer analyzes the incoming audio signal (the filled gray curve) and indicates the removed signal (the yellow curve):



DeWind:Dialogue automatically reduces wind noise from dialogue recordings. The spectrum analyzer shows the frequency spectrum of the input audio (filled gray curve) and the frequency spectrum of the removed signal (yellow curve).


Parameter Settings

- **Sensitivity (%)**

You can adjust the relative sensitivity of the noise detection with the *Sensitivity* slider. Positive values will increase the noise sensitivity and remove more noise. Negative values will reduce the noise sensitivity and leave more noise.

- **Sensitivity Curve**

The sensitivity filter allows you to control the noise sensitivity in adjustable frequency bands. There are three filters available that you can toggle on or off using the buttons in the upper right corner of the spectrum analyzer: A low shelf filter (↔), a peak filter (⊖) and a high shelf filter (↔). These work the same way as in parametric equalizers. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed. You can also change the filter slopes or the bandwidth of the peak filters. Click the handle of the filter you wish to modify. Arrows surrounding the handle will appear that you can use to change the bandwidth or the filter slope of the selected filter.



The sensitivity curve lets you change the noise sensitivity in separate frequency bands.

- **Maximum attenuation (dB)**

In many cases, it is desirable to leave some noise and you can use the *Maximum attenuation* slider to control this. The *Maximum attenuation* is set in dB specifies the maximum attenuation if the frequency components, thus control the noise floor after processing.

- **Solo Noise**

Enable this option if you wish to monitor the noise signal that was removed

7.4.12 DeRustle:Dialogue

Premium Edition Only

DeRustle:Dialogue reduces the rustle noise frequently encountered in lavalier microphone recordings. The algorithm is based on deep learning and has been trained on thousands of high quality voice recordings and recordings of rustle and microphone bump noises. The default setting will immediately and fully automatically reduce rustle noise. In many cases, however, you might not want to remove the noise entirely. You can also adjust the sensitivity of the noise detection. The sensitivity is adjustable for the full frequency band or in frequency bands. A spectrum analyzer analyzes the incoming audio signal (the filled gray curve) and indicates the removed signal (the yellow curve):



DeRustle:Dialogue automatically reduces rustle noise from lavalier microphone recordings. The spectrum analyzer shows the frequency spectrum of the input audio (filled gray curve) and the frequency spectrum of the removed signal (yellow curve).


Parameter Settings

- **Sensitivity (%)**

You can adjust the relative sensitivity of the noise detection with the *Sensitivity* slider. Positive values will increase the noise sensitivity and remove more noise. Negative values will reduce the noise sensitivity and leave more noise.

- **Sensitivity Curve**

The sensitivity filter allows you to control the noise sensitivity in adjustable frequency bands. There are three filters available that you can toggle on or off using the buttons in the upper right corner of the spectrum analyzer: A low shelf filter (☞), a peak filter (⊖) and a high shelf filter (☜). These work the same way as in parametric equalizers. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed. You can also change the filter slopes or the bandwidth of the peak filters. Click the handle of the filter you wish to modify. Arrows surrounding the handle will appear that you can use to change the bandwidth or the filter slope of the selected filter.



The sensitivity curve lets you change the noise sensitivity in separate frequency bands.

- **Maximum attenuation (dB)**

In many cases, it is desirable to leave some noise and you can use the *Maximum attenuation* slider to control this. The *Maximum attenuation* is set in dB specifies the maximum attenuation if the frequency components, thus control the noise floor after processing.


- **Solo Noise**

Enable this option if you wish to monitor the noise signal that was removed

7.4.13 DeBuzz:Dialogue

Premium Edition Only

DeBuzz:Dialogue reduces the typical buzz and hum noise from dialogue recordings. The algorithm is based on deep learning and has been trained on thousands of high quality voice recordings and recordings of neon light buzz, AC hum and similar noises. The default setting will immediately and fully automatically reduce buzz noise. In many cases, however, you might not want to remove the noise entirely. You can also adjust the sensitivity of the noise detection. The sensitivity is adjustable for the full frequency band or in frequency bands. A spectrum analyzer analyzes the incoming audio signal (the filled gray curve) and indicates the removed signal (the yellow curve):



DeBuzz:Dialogue automatically reduces buzz noise from neon lights, AC power or similar. The spectrum analyzer shows the frequency spectrum of the input audio (filled gray curve) and the frequency spectrum of the removed signal (yellow curve).


Parameter Settings

- **Sensitivity (%)**

You can adjust the relative sensitivity of the noise detection with the *Sensitivity* slider. Positive values will increase the noise sensitivity and remove more noise. Negative values will reduce the noise sensitivity and leave more noise.

- **Sensitivity Curve**

The sensitivity filter allows you to control the noise sensitivity in adjustable frequency bands. There are three filters available that you can toggle on or off using the buttons in the upper right corner of the spectrum analyzer: A low shelf filter (👉), a peak filter (👁️) and a high shelf filter (👈). These work the same way as in parametric equalizers. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed. You can also change the filter slopes or the bandwidth of the peak filters. Click the handle of the filter you wish to modify. Arrows surrounding the handle will appear that you can use to change the bandwidth or the filter slope of the selected filter.



The sensitivity curve lets you change the noise sensitivity in separate frequency bands.

- **Maximum attenuation (dB)**

In many cases, it is desirable to leave some noise and you can use the *Maximum attenuation* slider to control this. The *Maximum attenuation* is set in dB specifies the maximum attenuation if the frequency components, thus control the noise floor after processing.


- **Solo Noise**

Enable this option if you wish to monitor the noise signal that was removed

7.4.14 DePlosive:Dialogue

Premium Edition Only

DePlosive:Dialogue attenuates excessive plosives, pops and low frequency noise that typically occur when recording dialogue without a pop filter or wind shield. The algorithm is based on deep learning and has been trained on thousands of high quality voice recordings. The default setting will immediately and fully automatically reduce plosives, however, you might not want to remove the noise entirely. You can also adjust the sensitivity of the noise detection. The sensitivity is adjustable for the full frequency band or in frequency bands. A spectrum analyzer analyzes the incoming audio signal (the filled gray curve) and indicates the removed signal (the yellow curve):



DePlosive:Dialogue automatically reduces excessive plosive sounds, pops and low frequency noise. The spectrum analyzer shows the frequency spectrum of the input audio (filled gray curve) and the frequency spectrum of the removed signal (yellow curve).


Parameter Settings

- **Sensitivity (%)**

You can adjust the relative sensitivity of the noise detection with the *Sensitivity* slider. Positive values will increase the noise sensitivity and remove more noise. Negative values will reduce the noise sensitivity and leave more noise.

- **Sensitivity Curve**

The sensitivity filter allows you to control the noise sensitivity in adjustable frequency bands. There are three filters available that you can toggle on or off using the buttons in the upper right corner of the spectrum analyzer: A low shelf filter (☞), a peak filter (⊖) and a high shelf filter (☜). These work the same way as in parametric equalizers. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed. You can also change the filter slopes or the bandwidth of the peak filters. Click the handle of the filter you wish to modify. Arrows surrounding the handle will appear that you can use to change the bandwidth or the filter slope of the selected filter.



The sensitivity curve lets you change the noise sensitivity in separate frequency bands.

- **Maximum attenuation (dB)**

In many cases, it is desirable to leave some noise and you can use the *Maximum attenuation* slider to control this. The *Maximum attenuation* is set in dB specifies the maximum attenuation if the frequency components, thus control the noise floor after processing.

- **Solo Noise**


Enable this option if you wish to monitor the noise signal that was removed

7.4.15 DeClick:Dialogue

Premium Edition Only

Mouth clicks and lip smacks can be a major annoyance in voice recordings.

DeClick:Dialogue is better suited for these problems than the standard [DeClick](#) ¹²⁴ tool which is optimized for vinyl and 78 RPM recordings.



The *DeClick:Dialogue* user interface.

The *DeClick:Dialogue* user interface shows the frequency range of the processing and you can adjust it using optional low and high cut filters with adjustable filter slopes.

Parameter Settings

- **Sensitivity (%)**

You can adjust the sensitivity of the click detection with the *Sensitivity* slider between 0% (no clicks detected) and 100% (maximum detection).

- **Click expand (ms)**


This setting lets *DeClick:Dialogue* expand the duration of the interpolation time range after each detected click. It should be kept as short as possible, while still covering the clicks.

- **Solo Noise**

Enable this option if you wish to monitor the click signal that was removed

- **Frequency Range**

The frequency range curve allows you to control the frequency range of the processing. You can enable low (↖) and high (↗) cut filters with adjustable filter slopes. These filters work the same way as in parametric equalizers. You can modify the filter characteristics by clicking the handles (colored squares) in the curve and move them around and the current cut-off frequency settings of the frequency band is displayed. Arrows surrounding the handle will appear that you can use to change the filter slope of the selected filter.



The frequency range control lets you adjust the frequency range of the processing. The yellow handle controls the detection frequency described below.

- **Detection Frequency**

To avoid removing wanted transients that occur in conjunction with sounds like 'p' and 't', *DeClick:Dialogue* only removes clicks with most energy in high frequency bands. The detection frequency controls the threshold for this mechanism, and only clicks with a spectral center of mass (also called spectral centroid) above the detection frequency are processed. You can control the detection frequency with the yellow handle in the frequency range curve.

7.4.16 DeEss:Dialogue

DeEss:Dialogue is an audio processing tool designed to reduce or eliminate excessive sibilance in vocal recordings, which refers to harsh "s" and "sh" sounds. It specifically targets and controls the high-frequency content associated with sibilant sounds. By employing frequency-selective compression, a *DeEss:Dialogue* allows for the reduction of sibilance without affecting the overall quality of the vocal performance. This tool is particularly useful in audio production, ensuring that vocals remain clear, smooth, and free from distracting or piercing sibilant noises.



The DeEss:Dialogue user interface.

The *DeEss:Dialogue* user interface shows a spectrum analyzer of the incoming signal (gray) along with the changes in the spectrum indicated with a red color. You can control the frequency range of the processing along with the threshold level directly in the spectrum analyzer.

Parameter Settings

- **Threshold (dB)**

Frequency components above the threshold level will be reduced according to the Reduction settings.

- **Reduction (dB)**

Sets the amount of reduction for signal components that are above the threshold level.

- **Tame (dB)**

The Tame control evens out resonances in the frequency spectrum.

- **Channel Linking (%)**

With the channel linking set to 100%, the amount of gain reduction will be the same for the left and right channels, even if there is a difference in level between these channels. The more you dial the knob away from 100%, the more the channels will be

independently treated, with complete independence between the channels once you have a 0% channel linking. Channel linking values below 100% can result in shifts in the stereo image.

- **Audition**

Click the Audition button to listen to the audio changes resulting from the de-essing.

7.4.17 DeBird

Premium Edition Only

DeBird reduces unwanted bird song from recordings. The algorithm is based on deep learning and has been trained to separate bird noise from other sounds using an extensive training set. The default setting will immediately and fully automatically reduce bird noise. In many cases, however, you might not want to remove the noise entirely. You can also adjust the sensitivity of the bird noise detection. The sensitivity is adjustable for the full frequency band or in frequency bands. A spectrum analyzer analyzes the incoming audio signal (the filled gray curve) and indicates the removed signal (the yellow curve):



DeBird detects and reduces bird noise. The spectrum analyzer shows the frequency spectrum of the input audio (filled gray curve) and the frequency spectrum of the removed signal (yellow curve).


Parameter Settings

- **Sensitivity (%)**

You can adjust the relative sensitivity of the bird noise detection with the *Sensitivity* slider. Positive values will increase the detection sensitivity and remove more bird noise. Negative values will reduce the noise sensitivity and leave more noise.

- **Sensitivity Curve**

The sensitivity filter allows you to control the noise sensitivity in adjustable frequency bands. There are three filters available that you can toggle on or off using the buttons in the upper right corner of the spectrum analyzer: A low shelf filter (☞), a peak filter (⊖) and a high shelf filter (☜). These work the same way as in parametric equalizers. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around and the current frequency and gain settings of the frequency band is displayed. You can also change the filter slopes or the bandwidth of the peak filters. Click the handle of the filter you wish to modify. Arrows surrounding the handle will appear that you can use to change the bandwidth or the filter slope of the selected filter.



The sensitivity curve lets you change the noise sensitivity in separate frequency bands.

- **Maximum attenuation (dB)**

In many cases, it is desirable to leave some noise and you can use the *Maximum attenuation* slider to control this. The *Maximum attenuation* is set in dB specifies the maximum attenuation if the frequency components, thus control the noise floor after processing.

- **Solo Noise**

Enable this option if you wish to monitor the noise signal that was removed


7.4.18 Vitalize

About Vitalize

Vitalize is tool which allows you to add extra presence and/or sparkle in a different way than with an equalizer. When you use an equalizer to boost a certain frequency range, everything in that range is amplified. In many situations this might be perfectly fine, but

there is a chance you also boost unwanted elements of the audio, like background noise or vocal sibilance for example. Vitalize doesn't boost certain frequencies, but adds harmonics instead. These harmonics are not based on distortion, but are artificially created by a sophisticated algorithm. The result is added presence and sparkle in the most smooth and natural sounding way as possible.

User Interface



Parameter Settings

- **Dry level (-48 dB to 12 dB)**

You can use the dry level slider to increase or decrease the level of the unprocessed audio.

- **Vitalize level (-48 dB to 12 dB)**

You can use the vitalize level slider to increase or decrease the processed audio. For a natural sounding result, we recommend to have the vitalize level always at a lower level than the Dry level. If you have the dry level at 0 dB, a good starting point for the vitalize level would be around -6 dB from where you can adjust it to taste.

- **Output Equalizer**

The built-in frequency graph acts as an output equalizer and consists of a low cut, low shelf, high shelf and a high cut filter. You can modify the filter characteristics by clicking the handles (colored bullets) in the curve and move them around. The current frequency and gain settings of the frequency band is displayed. You can also change the


filter slopes. Click handle of the filter you wish to modify. Arrows appear surrounding the handle. Move these to change the filter slope.

7.5 Using Audio Plug-Ins

Acoustica supports the plug-in formats VST, VST3 and Audio Units. The latter is only available on Mac. You can scan for and manage plug-ins using the Plug-in Manager.

7.5.1 The Plug-in Manager

The *Plug-in Manager* lets you scan for new plug-ins and manage the existing ones. Choose *Plug-in Manager...* from the *Plug-ins* menu to access the plug-in manager.



The plug-in manager in Acoustica with a list of detected plug-ins.

You can scan for new plug-ins using the *Scan for changes* button. Only newly added plug-ins will be validated and no longer existing plug-ins will be removed. If you click the *Rescan all* button, Acoustica will validate all plug-ins including those who have already been validated successfully.


Note: You can remove one or more plug-ins from the list by clicking it's entry in the list and press the delete key.

Menu Sorting Options

You can also choose how Acoustica should sort plug-ins in the plug-in menu. Click the drop-down list for different sorting options.

Plug-in Folders

The VST3 and AU plug-in standards define the directory where plug-ins are situated. That's not the case for VST prior to version 3, though, and you can add or remove VST search folders by clicking the *Folders...* button. The following window appears:




List of VST search folders.

You can click the button with the plus icon to add a folder and the button with the minus icon to remove an existing entry. Furthermore, you can rearrange the search order using the up and down arrow buttons. Click *OK* when you are done.

7.5.2 Accessing the Plug-Ins

After you have scanned for plug-ins you can access the plug-ins through the *Plug-Ins* menu if the clip editor is active. They are automatically sorted according to the brand names. The plug-ins also appear in the add processor drop-down menu of the chain editor (see [The Processing Chain](#)^[153]).


Plug-ins behave just like internal processing tools in Acoustica. One slight difference is the possibility to save preset files in the common preset exchange format defined by the different plug-in standards. To save or load presets for exchange with other host applications, open the plug-in (using the main menu or in the Processing Chain) and click the plug-in button at the bottom of the plug-in window:



The plug-in format specific commands. Here you can save and load presets for exchange with other host applications or access preset banks defined by the plug-in.

7.6 The Processing Chain

The *Processing Chain* allows you to create a chain of processors or plug-ins. The processing chains can be saved including the processor settings for later use. Furthermore, each element can easily be bypassed or soloed and the order of the elements changed using drag and drop. The *Processing Chain* window also contains an output level slider and a sample rate conversion (SRC) indicator. The SRC indicator gets highlighted with the text *SRC on* if Acoustica needs to perform sample rate conversion in real time during playback or recording. If SRC is active, you can also see the sample rate of the audio device below the SRC indicator.



The Processing Chain editor in Acoustica.

Audio playback from the clip editor in Acoustica is routed through the Processing Chain in real-time if the Processing Chain is enabled. You can enable or disable the Processing Chain using the power on/off button as depicted below:



The power on/off button toggles processing in the processing chain. The right button toggles the mono compatibility test mode.

The mono compatibility mode makes it easy to test your audio for mono compatibility prior to rendering.

Adding Processors to the Chain

To add a new processor to the chain, click the plus icon in the list. A pop-up menu appears where you can choose among internal processors or plug-ins.

Removing Effects from the Chain

To remove a processor, click the red X icon in the processor item in the list.

Editing the Effect Settings of an Element in the Chain

Double click the name of an entry in the list to open the processor window of an element in the chain.

Bypassing an Element

You can bypass a single element in the chain by clicking the *power on* icon () in the processor list entry.

Solo an Element

You can solo a single element in the chain by clicking the *solo* icon () in the processor list entry. When one or more elements are soloed, all other elements in the chain will be bypassed.

Saving and Loading Processing Chains

You can store a complete processing chain including all parameter settings for later use. To store the processing chain, click the save button. A standard file save dialog box appears where you can enter the file name. To open an effect, click the load chain button and select the file in the file browser.




The load and save chain buttons.

Applying the Processing Chain

You can apply the complete processing chain to a selection in an open clip editor window by clicking the *Apply* button in the toolbar at the top of the *Processing Chain* window.

Render Selection to a New Audio File

You can render the current selection through the processing chain and save the results to a new audio file automatically. This is achieved by clicking the drop down arrow on the *Apply* button and the following drop down menu appears:



The Apply button has a drop-down menu with additional options.


Now, please choose *Render to New File...*, specify a file name and location, and click OK.

8 Multitrack Sessions

Acoustica can edit multitrack audio in addition to the clip based editing described earlier. An arbitrary number of audio tracks can be created which are mixed in real-time during playback. Each track contains can contain multiple audio clips.

8.1 Creating A Multitrack Session

To create a new multitrack session, choose *New > Multitrack Session...* from the *File* menu. A dialog box appears where you can enter the name of the project, the folder location, audio formats and other session specific settings:




The screenshot shows the 'New Session' dialog box with the following settings:

- Session Name & Location:**
 - Session name: The Rise
 - Location: C:\Users\Peter\Documents\Acoustica Sessions
- Audio Format:**
 - Sample rate (Hz): 48000
 - Resolution (bits): 16
 - Channel format: 6 (5.1 surround)
- Tempo and Meter:**
 - Meter numerator: 4
 - Meter denominator: 4
 - Tempo (BPM): 120

Buttons: OK, Cancel

The New Session dialog in Acoustica where you can specify name, location and format of the session.

Acoustica will create a sub-folder with the name specified in *Session name* in the folder defined in *Location*. Clips that are imported, recorded or modified in the multitrack editor will be placed in this folder along with the session file itself. Choose *OK* when you have entered a session name and are ready to create the multitrack session window. The new multitrack session window now shows up in Acoustica containing a single audio track:



The newly created multitrack session as shown in Acoustica

8.2 Working with Tracks

Adding Tracks


To add a new track, choose *Insert new Track* from the *Session* menu. Alternatively, you can press **Ctrl/⌘+T**. A new track will appear before the selected track or as the last track if no track is selected.

Removing Tracks

To remove a track, click the track you want to remove in the multitrack editor to select it. Then choose *Delete Selected Track* from the *Session* menu. Alternatively, you can press **Ctrl/⌘+Alt+T**.


Track Settings

You can change the volume and see the current playback level directly from the track pane on the left side in the multitrack editor:



The track pane with volume, pan and send level controls as well as level metering and record, solo, mute and processing chain buttons.

You can rename the track by clicking the name and entering a new name. Each track also has a color assigned and you can click the colored rectangle to the left of the track title to change the color. A color picker shows up where you can choose a new track color:




The track color picker.

Tracks can be muted or soloed (listening to a single track) during playback using the corresponding buttons:



The track strip buttons from left: Record, Solo, Bypass and Processing Chain


The *Processing Chain* button opens the a processing chain control (see [The Processing Chain](#)¹⁵³) with processors that are applied to the track (track inserts):



The track insert editor where you can add and remove track processors.

Working in Surround

You can create multitrack sessions for 5.1 or 7.1 surround in *Acoustica Premium Edition*. The pan controls are substituted with surround pan controls if you choose a surround format as rendering format when you create your multitrack session (see [Creating A Multitrack Session](#)¹⁵⁸):



Multitrack sessions targeting surround formats show a surround pan control instead of the rotary pan slider.

Click the surround pan control to adjust the surround pan. A large surround pan controller appears:



The surround pan control allows positioning in a 2D space. Single channels can also be enabled or disabled by clicking the channel icons.

You can grab the circular pan handle with the mouse and position it freely in the 2D space. It is also possible to enable or disable single channel output by clicking the loudspeaker icons or the wave icon representing the LFE (Low-Frequency Effects) channel. The LFE channel is disabled by default.

8.3 Adding Clips

Adding an Existing Audio File

To add an existing audio file as a clip in the multitrack session, you can simply double click the location where you want to add the clip. A file open dialog appears where you can choose the file you want to insert. Alternatively, you can drag and drop audio files to the multitrack sessions using the file browser of the operating system.

Recording Clips

To record a new clip, click the record button located on the left side of the track pane:



The track pane buttons. Start recording by clicking the record button (the leftmost button). Click the down arrow to choose recording channel format (mono or stereo).

A new clip will now be recorded and at the same time you will hear the audio from other tracks in the session. The clip will use the same resolution and sample rate as defined as rendering format of the session. You can choose to record either mono or stereo clips by clicking the down arrow at the right hand side of the recording button.

8.4 Looping and Stretching Clips

You can change the duration of clips either by looping or stretching the content. To loop or set the duration shorter than the content of the clip, move the mouse cursor to the upper part of the right clip border so that the normally arrow shaped mouse cursor turns into a loop cursor. Now hold the mouse button down while you move the cursor to the left or to the right to extend or shorten the clip. When you release the mouse button, the clip length is fixed. If you extend the clip you can see that the repeating sequences are always colored alternately light and dark.

As mentioned, you can also stretch the clip to match the tempo of the project. This involves processing the clip content with the time stretching algorithm in Acoustica. The procedure is the same as for looping, but now move the mouse cursor to the button part of the right clip border so that the mouse cursor turns into a double sided arrow pointing left and right before holding down the mouse button.

8.5 Moving and Grouping Clips


You can move a clip in time or to another track by pressing and holding down the left mouse button. Now move the clip to its new location in time or another track. If you want to move several clips, you can use the hold down the Ctrl button on your keyboard while clicking the clips you want to move. The clips are now highlighted and you can use the same procedure as for single clips to move the group of clips.

Grouping Clips

If you have a set of clips that should always be moved together, you can *group* clips. Highlight the clips you want group using the procedure described above and choose *Group Clips* from the *Clips* menu. To ungroup, simply click one of the clips in the group to highlight all the clips and choose *Ungroup Clips* from the *Clips* menu.

8.6 Crossfading Clips

You can easily crossfade clips in a multitrack session. Press and hold down the left mouse button on a clip and move it so that it overlaps the clip it should be crossfaded with. The crossfaded region is shaded and the fade-in (red) and fade-out (blue) curves are visually indicated:




Crossfaded clips in the multitrack editor. Click the transition region to modify the fade curves.

To choose different fade curves, click the shaded transition region and a pop-up menu appears where you can choose between a set of predefined crossfade characteristics.

8.7 Clip Fades

You can easily add fade-ins or fade-outs to clips in multitrack sessions by moving the fade handles. When you import or record a clip, the fade-in and fade-out handles will be positioned at the top left and top right corners of the clip. Click a handle and keep the mouse button down while moving the mouse to change the clip duration:



The fade handles let you add fade-ins or fade-outs very quickly.

The circle shaped handles let you control the curvature of the fades from exponential through linear to logarithmic.

8.8 Editing Clips with Clip Editor

You can easily edit clips in a multitrack session using the clip editor in Acoustica:

- Double click the clip you want to edit or right click the clip and choose *Edit Clip(s)...*
- The clip editor opens with the selected clip and you can make your edits
- Click the *Render* button in the toolbar (first button) to render the changes back to the session.

If you want to discard the changes, you can simply close the clip editor. Acoustica will ask you if you want to render the clip, and you can click the *No* button to discard the edits.

8.9 Saving, Loading and Exporting

You can save and load multitrack sessions similar to saving and loading of ordinary audio files in Acoustica.

Loading Multitrack Session Files

To open an existing multitrack session file,

1. Select from the *File* menu the command *Open...*
2. Choose the folder in which your file is located.
3. Click the multitrack session file you wish to open and click the *OK* button.

Saving Multitrack Session Files

Select from the *File* menu the command *Save* or press *Ctrl/⌘+S* to save the content of a multitrack session with its original name. If you wish to save the content of a multitrack session with a different name, in a different folder or with different settings:

1. Select from the *File* menu the command *Save as...*
2. Choose the folder in which you wish to save the file.
3. Enter a name for your multitrack session file and click the *Save* button.

Exporting Audio Files from Multitrack Sessions


You can save the rendered audio from a multitrack session as an audio file:

1. Select from the *File* menu the command *Export to Audio File...*
2. Choose the folder in which you wish to save the file.
3. Choose the file format of your audio file from the *Save as type* drop-down list.
4. Enter a name for your audio file and click the *Save* button.

Most export filters offer different settings such as encoding bit rate or number format. To change the settings, click the *Options* button.


8.10 Automation Curves

It is possible to automate volume, pan and send levels for each track by adding automation curves. Choose *Automation Curves...* from the *Session* menu and the *Automation Curves* window appears:



Automation Curves lets you have volumes, panning and send levels that change along the timeline.

Automation curves are visible and editable directly on each audio track in the multitrack editor:



Track with activated volume automation curve.

You can add a curve point by moving the mouse cursor over the curve. The mouse cursor turns into a pointing hand. Now you can click the mouse button to create a new curve point and move it around with the using the mouse. Double click a curve point to remove it again. You can also choose how to interpolate the curve between two curve points. Click one of the curve points so that it gets focused (indicated with a white outline). Now, two arrow shaped handles appear to the left and to the right of the curve point. You can drag these to get a smoother transition on either side.


Moving Clips with Automation Curves

If you move clips after you have created an automation curve, you can choose if you want the automation points to follow the clip. The default action is that automation points

follow the clip movement. If you want the automation points to remain at their original locations when moving clips in time, hold down the Alt key while moving the clip.

8.11 Using Sends and the Auxiliary Bus

Multitrack sessions in Acoustica are automatically set up with an Auxiliary Bus (Aux). This makes it possible to route audio from each track for effects processing. The most common usage is to add reverbation on auxiliary bus. This reduces CPU usage since only one instance of the reverb processor is required and at the same time, it's possible to set the amount of reverb for each track using the *Send* knobs in each track pane:



Each track pane has a *Send* knob (rightmost) that controls the output level to the auxiliary bus.

The Auxiliary bus has controls for volume, pan, effect inserts just like the master bus:




The Auxiliary (Aux) bus is indicated with a red rectangle.

Just like effects inserted in a track, you can add effects using the *Processing Chain* button which opens the a processing chain control (see [The Processing Chain](#)¹⁵³).

8.12 Importing Stems from a Mix

Acoustica can automatically separate an audio file with a complete mix, split it into stems using artificial intelligence (AI) and add a separate track for each stem to a multitrack session. You can do this by choosing *File > Import Stems from File...* in a multitrack session. You will now be asked which stem format to use. Click OK when ready and Acoustica lays the stems out in separate tracks as shown below:



A multitrack session with automatically created stems.

9 Working With Audio CDs

You can create Audio CDs (see Creating Audio CDs) that you can play in normal CD players directly within Acoustica if you have a supported CD burner. Furthermore, Acoustica allows you to import audio tracks from existing CDs digitally and without quality loss for further editing or archiving on the computer (see Importing Audio Tracks from CDs).

9.1 Creating Audio CDs

Acoustica allows you to create audio CDs containing your edited recordings. The first step towards your own CD is to create a CD-Project by selecting *File > New > CD Project...* or by clicking the *New* button in the main toolbar and select *CD Project...* from the drop-down list. The CD Project contains a list of the audio tracks to be written on the CD. At the bottom you can also specify a media catalog number (MCN) and see the current fill status in percent:



An empty CD Project window.

You can add audio files to the CD project by choosing *Import Audio File...* from the *CD* menu or use the corresponding button in the toolbar:




The CD related buttons on the toolbar (from left): Add track, remove track, burn CD and erase CD

You can remove tracks from the CD project using the *Remove track* toolbar button or through the *Remove Selected Track(s)* command from the *CD* menu.

Burning an Audio CD

When you have compiled the audio CD you can burn it to an empty CD-R or CD-RW disc. Choose *Burn Audio CD...* from the *CD* menu or use the *Burn CD* toolbar button and the following Window appears:




The *Burn CD* dialog where you can choose CD recording device and recording speed.

Choose a CD recording device and the desired recording speed, and click the *Burn* button.

Erasing a CD-RW Disc


You can erase the content of CD-RW discs directly within Acoustica. Choose *Erase CD-RW* from the *CD* menu or use the corresponding icon in the toolbar.



The Erase CD-RW dialog where you can choose CD recording device and whether to perform a full erase (slow) or quick erase.

9.2 Importing Audio Tracks from CDs

It is possible to digitally import audio data from audio CDs with most CD-ROM readers. This feature is part of the operating system on Mac where you can use choose *Open...* from the *File* menu and navigate to a mounted audio CD and choose the tracks you wish to import. On Windows please choose *Import Tracks from Audio CD...* from the *File* menu and the track import window appears:



The Import CD Audio tool under Window lets you import one or more audio tracks from redbook compatible audio CDs.


You can choose one or more tracks (press Ctrl to select several tracks) and import them by clicking the *Extract* button. You can also listen to tracks using the play and pause buttons in the lower right corner of the window.

10 Batch Processing

The *Batch Processor* in Acoustica allows you to process and convert the file format of a large number of files without user interaction. Even complete directory structures can be processed at once. Select *Batch Processor...* from the *File* menu to open the batch processor.

10.1 The Source Page

The batch processor opens up with the *Source* page active as depicted below:



The empty Source page of the Batch Processor.

The first step is to add the source files for processing. The most straight forward way to add files is to click the *Import files* button. A file open dialog opens where you can browse and choose one or more files to open. You can repeat this procedure if you wish to add files from several folders.

Note: You can listen to the imported files by clicking the play and pause buttons on the right hand side above the source file tree.

Importing Folder Structures

You can also choose to import an folder structure by clicking the *Import folder* button. The OS specific folder picker appears. Acoustica imports scans through the folder and all sub-folders and adds all known audio files to the source tree.

Creating Target Folders

It is possible to create target folders manually and place files within them. Click the *Create folder* button to create a sub-folder. Imported items are inserted into the selected target sub-folder.

Proceeding to the Target Files Page


When you are done adding the source files it is time to move on to the *Destination* page where you can set up the target file format and set up a processing chain. Click the *Next* button in the top right corner to proceed:



The wizard style navigation buttons in the Batch Processor.

10.2 The Destination Page

The *Batch Processor* lets you process the files with internal processing tools or with external plug-ins and includes a processing chain that you can modify as described in [The Processing Chain](#)¹⁵³:



The Destination page of the Batch Processor with a processing chain set up.

The *Target Format* section lets you select an output file format by choosing the desired file format from the *File format* drop-down list. If the file format offers custom parameters settings, like target resolution or encoding bitrate for mp3 files, you can edit these settings

by clicking the *Options...* button next to the *File format* drop-down list. Furthermore, you can define the sample rate and channel format of the target files or choose to use the original sample rate and channel format of the files.

The *Destination* section lets you select an output folder. You can either enter a folder path in the edit box or click the *Browse* button to select a folder from the folder structure of your computer.

Click the *Process* button in the top right corner of the *Batch Processor* window to start processing. Now is a great time for a cup of fair trade coffee...


11 Using The Cleaning Wizard

The *Cleaning Wizard* simplifies the transfer of analog recordings to computer or CD for users new to digital audio recording. It guides you through all the steps from recording, track splitting, restoration and CD burning. Choose *Cleaning Wizard...* from the *File* menu to open the *Cleaning Wizard*.

The *Cleaning Wizard* leads you step by step through the process of transferring your analog audio to CD. You can, however, choose to go back to an earlier step or skip one or more steps by clicking at the *Import*, *Repair* or *Export* tabs at the top of the wizard.

11.1 Import


The first screen you see after opening the *Cleaning Wizard* lets you choose whether to record from an analog source or open an existing file:



The Start Page of the Cleaning Wizard

11.1.1 Record Audio

Cleaning Wizard will proceed to the recording page if you choose to record from an analog source:



The recording page in the Cleaning Wizard.


The *Input level* meter shows the current input level. If you have connected your audio equipment and started playback, the meter should show a constantly changing input level. If the level is low and not changing, there is probably something wrong with the connection or the wrong input line is selected. You can usually choose between several different input sources, like microphone or line in.

Recording Step by Step

1. Make sure your audio equipment is properly connected to your computer
2. Check that the *Input source* is correctly set. You should see the name of your external audio interface or the internal sound card on your computer if you use that.
3. Check that the input level is in the correct range. You can easily check the input level using *Input level* meter. The meter should never go up to 0 dB, otherwise digital clipping will be introduced. Check with the loudest part of the record or cassette tape you are recording that the input level meter doesn't go higher than about -12 dB to allow some headroom. Most audio interfaces will have an input level knob where you can adjust the level, otherwise adjust the output level of your source.
4. Click the record button [•] to start the recording.
5. Press play on your tape deck or record player.
6. Press the *Next* button in the upper right part of the *Cleaning Wizard* when you have recorded the whole record or cassette tape.

11.1.2 Import Files

If you choose to import an existing audio file, a file browser window appears where you can select the audio file(s) you want to open.



The File Import page in the Cleaning Wizard.


To import one or more audio files, please do the following:

1. Choose the folder in which your file is located from the drop-down list above the file list. You can go up one level or browse the folders on your computer using the buttons to the right.

2. Click the audio file you wish to open. If you want to open several files, press Ctrl while selecting multiple files. Click the button *Next* button in the upper right corner when done.

11.2 Repair

The Restoration Page allows you to adjust the settings of the audio restoration tools and split the recording into several tracks.



The Restoration Page contains a waveform view of the recording and list of the tracks, as well as audio restoration and processing options.

The Repair Screen Elements

1. The waveform view shows you a graphical representation of the recording. Tracks splits are indicated with a green line and the name of the track.
2. The toolbar with transport and command buttons.
3. The track list shows you the tracks defined. The *Cleaning Wizard* automatically suggests tracks, however, you can easily add, move or remove track markers.
4. The audio restoration tools, DeClick, DeCrackle, DeClip, DeNoise and DeHum. You can adjust the amount of processing using the sliders and activate or deactivate a tool using the toggle buttons to the left.
5. You can add further effects and processing tools, like equalizing or reverb in the *Processing Chain*. See [The Processing Chain](#)^[153] for more information.

11.2.1 Track Splitting

The *Cleaning Wizard* automatically searches for pauses and suggests track split positions when recording or importing audio files. However, if the recording is very noisy or tracks are blended seamlessly into each other, the tracks suggested by the Cleaning Wizard might not be identical to the original tracks on the source record or cassette.

Moving the Position of an Existing Track Split

1. Move the mouse cursor to an existing track split indicator in the waveform. The mouse cursor turns into a left-right arrow.
2. Keep the mouse button pressed while moving the mouse cursor to the new position.
3. Release the mouse button.

Adding a Track Split

1. Move the mouse cursor the beginning of the track you want to add in the waveform view.
2. Click the *plus* button below the track list.

Removing a Track Split

1. Click the track you want to remove in the track list.
2. Press the Delete key on your keyboard or the *minus* button below the track list.

Renaming a Track

1. Double click the track you want to rename in the track list. The entry turns into a text editor field.
2. Enter the new name of the track and press enter.


11.2.2 Restoration

There are five restoration tools integrated into the Cleaning Wizard:

- **DeClick**
Removes loud clicks and pops.
- **DeCrackle**
Removes short but frequent clicks, referred to as crackle.
- **DeClip**
Restores recordings that suffer from analog or digital clipping.
- **DeNoise**
Removes static noise like tape hiss.

- **DeHum**

Removes 50 or 60 Hz hum originating from the electric power lines



The restoration tools in the Cleaning Wizard.

You can adjust the effect of each tool by moving the sliders in the range from 0% (no effect) to 100% (full effect). Furthermore, the tools can be activated or deactivated by clicking toggle buttons to the left of the sliders.

Playing Restored Tracks


The restoration tools in the Cleaning Wizard are processed in real time during playback so that you can listen to the effect of different restoration settings immediately. You can control the playback from the transport buttons in the Restoration Page.

11.2.3 Further Editing and Processing

The *Repair* page in the *Cleaning Wizard* also lets you apply further processors to the recording using the *Processing Chain*. You can choose among all the internal processors as well external plug-ins. When you play the recording in the *Repair* page, the effects are processed in real time so that you can hear the results immediately. Please see [The Processing Chain](#)¹⁵³ for more information about how to add and set up audio processors.

11.3 Export


After you have revised the track list, set up the restoration tools and possibly added additional processing, you can proceed to the *Export* page using the *Next* button in the top right corner of the *Cleaning Wizard*. You can export your restored tracks to audio files or burn them directly to a CD. Please choose one of the two in the *Export* page:



The Export page in the Cleaning Wizard where you can burn your tracks to an audio CD or save them as files on your computer.

11.3.1 Burn a CD

If you choose to burn a CD in the *Export* page, the *Cleaning Wizard* proceeds to the burning page:




The CD burner dialog allows you to select a CD recording device and recording speed.

Insert a blank CD-R or CD-RW into the CD recorder. If you have several CD recording devices installed on your computer, make sure you choose the correct one from the *CD*

recording device drop-down list. You can choose among different recording speeds. Click the *Burn* button to start burning.

11.3.2 Export to Audio Files

If you choose to export your tracks to audio files in the *Export* page, the File Export Page appears:



The File Export Page in the Cleaning Wizard

You can choose a *Target folder* for your tracks, an *Album name* and the *File type* of the exported tracks. During the export, a directory will be created with the album name and the tracks are written to audio files with the name of the tracks within the album directory. Click the *Proceed* button in the top right corner of the wizard to start exporting your files.

12 Customizing the Workspace

Acoustica 7 has a flexible docking system which allows you to dock pane windows such as real-time analyzers or anchor editors in order to customize the workspace. You can adjust the window pane layout manually or choose from factory layouts. Furthermore, you can save and load your own custom window pane layouts.

12.1 Using Factory Layouts

The easiest way to change the window pane layout is to use one of the factory layouts delivered with Acoustica. Choosing a factory layout will rearrange the window panes for you automatically. You can choose factory layouts from the *Factory Layouts* sub menu under the *View* menu.

12.2 Docking Window Panes

Acoustica makes it easy to create your own custom window layout. Pane windows can be docked at one of the sides of an existing pane window. You can also group pane windows so that each window is represented with a tab in a header.



Docking options in Acoustica. You can move the mouse cursor over one of the dock buttons to dock a pane window.

To dock a pane window, click a header of a pane window (1) or the tab header (2) if it is grouped and keep the mouse button pressed. As soon as you move the mouse, dock buttons are shown (3) to indicate possible docking position. Move the mouse cursor over any of the available dock buttons and release the mouse button. You can simply release the mouse button anywhere else on the screen if you don't want to complete the docking operation.

You can detach a window completely by releasing the mouse button outside any of the dock buttons. This is handy for multi-monitor setups where you might want to place pane windows on a separate monitor.

Restoring the Default Layout

You can always restore the default pane layout by choosing *Restore Default Layout* from the *View* menu.

12.3 Loading and Saving Custom Layouts

You can save and load your own window layouts in Acoustica. To save the current window layout, select *Save Layout to File...* from the *View* menu. A save file dialog box appears where you can edit the name of the file to save. If you save the layout file in the suggested folder, Acoustica will pick it up as a *User Layout* and make it accessible in the *User Layouts* submenu in the *View* menu.

Alternatively, you can select *Load Layout from File...* from the *View* menu to load a layout file from any folder on your computer.

13 Preferences

You can adjust preferences such as the active time code format, temporary folders, keyboard short-cuts and audio device settings in the *Preferences* dialog. Choose *Preferences...* from the *Edit* menu to open the preferences. The preferences are grouped into three different tabs:

- General
- Keyboard shortcuts
- Audio device settings

13.1 General

The *General* tab in the preferences contains the following settings:

Folder for temporary files

This is the folder where Acoustica stores temporary files. It is critical for the performance that this is on a quick hard drive or SSD and that you have sufficient free space on the volume. You can choose either to type in a folder path or use the browse button on the right side to open the OS specific folder browser.

Time format:

You can choose the time format that Acoustica uses when displaying time positions. SMPTE time code options are available if you are working with video or film as well as bar and beat oriented time formats for music production. Otherwise, we recommend the standard Hour:Min:Sec:MS format.

Close tool window after process

Choose yes if you would like processing tool windows to be closed automatically when you click the *Process* button.

Playback scroll mode

You can choose how Acoustica should scroll during playback. The following options are available:

- *No scrolling*: Acoustica won't scroll to keep the play cursor in the view
- *Page jump scrolling*: Acoustica jumps a page at a time to keep the play cursor in the view
- *Keep cursor in middle of view*: Acoustica scrolls smoothly with the cursor in the middle of the view
- *Keep cursor at end of view*: Acoustica scrolls smoothly with the cursor in the end of the view

Always play from cursor position

Acoustica will normally revert the play cursor to the last start position when you stop playback. Choose *Yes* if you prefer that the play cursor remains at the stop position.

Mouse wheel zoom

You can choose whether you want to zoom around the mouse cursor or the play cursor when you use the mouse scroll wheel.

Maximum plug-in validation time (s)

Acoustica will automatically cancel plug-in validation if a plug-in freezes for a long period of time. You can set the time out in seconds.

Language


You can choose the language of the user interface in Acoustica. Currently, English and German are the supported languages.

Check for updates when starting

Acoustica checks for new versions at start-up per default. You can disable the check by choosing *No*.

13.2 Keyboard shortcuts

You can create keyboard shortcuts for commands that you use frequently or re-assign existing keyboard shortcuts. Click the *Keyboard shortcuts* tab in the *Preferences*. You will see a list of all the available commands sorted in categories:




The Keyboard shortcuts editor in the Preferences allows you to assign your own keyboard shortcuts to the Acoustica commands.

Click the plus icon to the right of the command entry to add or change an existing keyboard shortcut. A pop-up window appears and you can click the desired key combination. Click *OK* if you wish to keep the new keyboard shortcut.

13.3 Audio Device Settings


You can set up your audio device and assign input and output channels in the *Audio device settings* of the *Preferences*:



The audio device settings let you choose a driver type, buffer size and the input and output device to use.

We recommend ASIO as driver type on Windows if available and CoreAudio on Mac. The recommended buffer size is 512 samples.

If you have an audio device with several input and output channels or you are working with surround audio, you can set up the channel routing by clicking the route button (🔗) to the right of the input / output device settings and the following window appears:



You can set up the input and output channel routing maps in Acoustica.

The channel map editor lets you assign channels on your audio interface to the speaker positions defined in Acoustica.

Some audio drivers also have a control panel where you can configure your audio interface. A button with cogs symbol (⚙️) appears to the right of the input / output device settings in this case.

14 Included Plug-ins

Premium Edition Only

Acoustica Premium Edition includes an extensive collection of plug-ins that you can use in third party host applications:

- **Restoration Suite 2** — DeClick 2, DeClip 2, DeHum 2 and DeNoise 2

- **Mastering Suite** — Dynamics, Multiband Dynamics, Limit, Equalize 2 and Dither
- Extract:Dialogue
- DeClick:Dialogue
- DePlosive:Dialogue
- DeWind:Dialogue
- DeRustle:Dialogue
- DeBuzz:Dialogue
- DeBird
- Verberate 2
- Phono Filter
- Vitalize
- Convolve

You can access the plug-ins from your host application (audio editor or DAW) of choice. The plug-in is available as a 32 or 64 bit VST or AAX plug-in on the Windows platform or as a VST, AU or AAX plug-in on Mac OS X (64 bit). Some host applications will require a rescan and possibly adding the plug-in installation directory to the list of VST directories. Please consult the manual for your host application for further details.

Acoustica ARA2 Plug-In

If you use an ARA2 compatible Digital Audio Workstation (DAW), you can get the full functionality of Acoustica's clip editor tightly integrated into your DAW using the [Acoustica ARA Plug-in](#)¹⁸⁸.


Transfer Tool for Pro Tools

Additionally, a [Transfer](#)¹⁸⁸ tool in the AAX plug-in format for Pro Tools simplifies audio transfer back and forth between Pro Tools and Acoustica.

14.1 The Acoustica ARA Plug-in

Premium Edition Only

The ARA2 standard extends the VST3 and Audio Unit (AU) plug-in standards to allow a much tighter integration between the host application and plug-ins. The Acoustica ARA2 plug-in offers the complete functionality of Acoustica's clip editor inside ARA2 compatible host applications, thus eliminating the need for time consuming file transfers. You don't have to consider anything more after having made your edits with the Acoustica ARA plug-in in your host application. Your session will be played with the Acoustica edits automatically, and the clip changes in Acoustica will be stored along with your session when you save it.



The Acoustica ARA2 plug-ins running inside PreSonus Studio One.

Enabling the Acoustica ARA2 plug-in differs from host to host. You can find a short description on how to activate the Acoustica ARA2 plug-in the most common host applications below.

Using Acoustica ARA in Cubase and Nuendo

The handling of ARA plug-ins is identical in Steinberg's *Nuendo* and *Cubase* and they are referred to as *Extensions*. You will need Nuendo / Cubase version 10 or later with the most recent updates installed. To enable Acoustica ARA:

- Select the clip you want to edit in Acoustica
- Choose *Audio > Extensions > Acoustica* from the Nuendo / Cubase main menu
- The Acoustica ARA plug-in appears in the bottom window pane, and you can make your edits as you would in the stand-alone version of Acoustica.

Nuendo 12 and Cubase 12 introduces a convenient *Make Extension Permanent* command (*Audio > Extensions > Make Extension Permanent*) that you can use to render the changes in Acoustica permanently. This will refresh the clip's waveform in the session and make it possible to share the session with others that don't have Acoustica installed.


Using Acoustica ARA in Studio One

Unfortunately, Studio One only shows the convenient "Edit with..." shortcuts in the *Audio* menu for a set of hard coded ARA plug-ins. To use Acoustica ARA, you will need to locate the plug-in in the effects list:

- Press F7 to show the Effects list, if not already visible.
- Locate *Acoustica* from the *Acon Digital* sub-folder.
- Click the *Acoustica* plug-in in the list. Now, move the mouse cursor over the clip you want to edit, press the *Alt* key and keep it pressed while you release the mouse button.

Using Acoustica ARA in REAPER

ARA plug-ins are inserted as track effects in REAPER, and ARA plug-ins must be the first plug-in in the track insert chain. Click the track FX button to show the list of available plug-ins:



Click the FX button to insert track effects in REAPER. ARA plug-ins are inserted as inserts.

The Acoustica ARA plug-in will be active for all clips in the track that has the Acoustica ARA plug-in inserted. Acoustica ARA will always show the clip that is selected in REAPER.

Using Acoustica ARA in Logic Pro X

ARA plug-ins are inserted as track effects in Logic Pro X, and ARA plug-ins must be the first plug-in in the track insert chain. To enable Acoustica ARA:

- Make sure the *Mixer* is visible by pressing the X key.
- Click the *Audio FX* drop-down list for the track that you want to make edits in.
- Navigate to *Audio Units > Acon Digital* and choose *Acoustica (ARA)*. Please make sure that you choose the ARA version indicated with ARA in brackets.
- Select the clip you want to edit. You will need to start and stop playback in *Logic Pro X* to get Acoustica to show the clip content.

Please note: Even if Acoustica 7.4 runs natively on Macs with Apple Silicon processors, Logic Pro X only supports ARA plug-ins when running under Rosetta 2. Please check if there's a Logic Pro X update available if you want to use the Acoustica ARA plug-in in native mode on a Mac with an Apple Silicon processor.

Using Acoustica ARA in Cakewalk

ARA plug-ins are referred to as *Region FX* in Cakewalk. To enable Acoustica ARA:


- Right-click the clip you want to edit with Acoustica and the context menu appears.
- Choose *Region FX > Acoustica > Create Region FX*
- Click the clip in the Cakewalk session

14.2 The Pro Tools Transfer Tool

Premium Edition Only

Acoustica Premium Edition is delivered with an AAX plug-in that simplifies audio transfers from Pro Tools to Acoustica and back again after processing.

1. Select the audio in Pro Tools that you want to transfer to Acoustica.
2. Choose **AudioSuite > Other > Transfer to Acoustica** from the main menu in Pro Tools.
3. The Transfer tool opens. If you have Pro Tools 12 or later, Acoustica automatically opens. In earlier Pro Tools versions, click the *Transfer* button.
4. Do your editing in the clip editor in Acoustica as you are used to. You can use all the editing tools available and even change the duration of the audio.
5. Click the accept edit button which is the very first button with a check mark icon in the Acoustica toolbar:



6. Pro Tools 12 or later renders the audio automatically back into your Pro Tools session. If you have an earlier version of Pro Tools, click the *Render* button.

Index

- A -

| | | | |
|--|-----|------------------------------|-----|
| Accelerators | 181 | <i>DeNoise Light</i> | 133 |
| Accessing plug-ins | 153 | <i>manual click removal</i> | 123 |
| Acoustica ARA2-Plug-in | 185 | <i>noise analysis</i> | 129 |
| Acoustica Plug-in Pack | 184 | <i>Vitalize</i> | 149 |
| Add track | 157 | Audio scrubbing | 22 |
| Adding clips | 160 | Audio tracks | 157 |
| Adding source files to the Batch Processor | 169 | Audit trail | 30 |
| Aggressive hum removal | 126 | Authorize Acoustica | 6 |
| Analysis tools | 37 | Automatic cropping | 29 |
| Analyze noise | 129 | Automatic de-noising | 129 |
| Analyzers | | Automatic editing | 29 |
| <i>level meter</i> | 48 | Automatic noise reduction | 29 |
| <i>Phase correlation meter</i> | 52 | Automatic recording | 34 |
| <i>Spectrum analyzer</i> | 52 | Automatic silence removal | 29 |
| <i>time display</i> | 53 | Automatic Splice De-click | 21 |
| ARA2 Plug-in | 185 | Automatic track splitting | 29 |
| AU | 151 | Automation curves | 163 |
| Audio | | Aux Bus | 165 |
| <i>recording</i> | 32 | Auxiliary Bus | 165 |
| Audio analysis | 35 | | |
| <i>frequency domain</i> | 36 | - B - | |
| <i>time domain</i> | 36 | Back to front pan curve | 163 |
| Audio CDs | | Batch Processor | 169 |
| <i>creating</i> | 166 | <i>adding source files</i> | 169 |
| <i>importing audio tracks from</i> | 168 | <i>Destination page</i> | 170 |
| <i>working with</i> | 166 | <i>source page</i> | 169 |
| Audio device settings | 182 | Birds | 148 |
| Audio editing window | 11 | Brickwall limiter | 73 |
| Audio effects | 103 | Burn audio CD | 177 |
| Audio files | | Buzz noise | 140 |
| <i>inserting</i> | 13 | | |
| <i>loading</i> | 12 | - C - | |
| <i>saving</i> | 14 | Caption Editing | 27 |
| Audio processing | 60 | CD project | 11 |
| Audio restoration | 123 | Channel Mixer | 100 |
| <i>about de-noising</i> | 128 | Channels | |
| <i>automatic de-noising</i> | 129 | <i>selecting</i> | 16 |
| <i>DeClick</i> | 124 | Chorus | 113 |
| <i>DeClip</i> | 123 | Cleaning Wizard | 171 |
| <i>DeHum</i> | 126 | <i>burn audio CD</i> | 177 |
| <i>DeNoise</i> | 129 | <i>export</i> | 176 |
| | | <i>export to audio files</i> | 178 |
| | | <i>import page</i> | 171 |
| | | <i>importing files</i> | 173 |
| | | <i>processing chain</i> | 176 |
| | | <i>recording</i> | 172 |

| | | | |
|------------------------------|----------|---|------------|
| Cleaning Wizard | 171 | <i>waveform visualization</i> | 10 |
| <i>restoration</i> | 175 | Digital drop-outs | 124 |
| <i>track splitting</i> | 175 | Distortion | 123 |
| Clicks | 124 | Dither | 75 |
| Clip fades | 162 | Dithering | 75 |
| Clipboard | 21 | Drag and drop | 20 |
| Clipping | 123 | <i>copying</i> | 20 |
| Clips | 160 | <i>moving</i> | 20 |
| <i>add fades</i> | 162 | Dynamic processing | 62 |
| <i>crossfading</i> | 161 | Dynamics | 62 |
| <i>grouping</i> | 161 | - E - | |
| <i>looping</i> | 161 | EBU R-128 | 47, 50, 99 |
| <i>moving</i> | 161 | Echo | 111 |
| <i>stretching</i> | 161 | Edit History | 30 |
| Compressor | 62 | Editing | |
| Convert channel format | 77 | <i>basic</i> | 11 |
| Convert sample rate | 77 | <i>using drag and drop</i> | 20 |
| Convolution Reverb | 109 | <i>using the clipboard</i> | 21 |
| Copy | 21 | Editing states | 30 |
| Crackle | 124 | Efecst | |
| Creating multitrack sessions | 156 | <i>Chorus</i> | 113 |
| Crop | 21, 29 | Effects | 103 |
| Crossfading clips | 161 | <i>Convolve</i> | 109 |
| Customizing the workspace | 178 | <i>Echo</i> | 111 |
| Cut | 21 | <i>Harmonize</i> | 120 |
| - D - | | <i>Modulate</i> | 116 |
| DC offset removal | 97 | <i>Reverb</i> | 107 |
| DeBird | 148 | <i>reverse</i> | 121 |
| Decibel (dB) | 10 | <i>Time Stretch</i> | 118 |
| DeClick | 124 | <i>Transpose</i> | 119 |
| DeClick for dialogue | 144 | <i>Verberate</i> | 103 |
| DeClip | 123 | Emphasis filtering | 96 |
| Deemphasis filtering | 96 | Enhancement tools | 123 |
| De-essing | 146 | EQ | 80, 92 |
| DeHum | 126 | Equalize | 80 |
| Demo version | 6 | Equalize Light | 92 |
| DeNoise | 128, 129 | Equalize preferences | 88 |
| DeNoise Light | 133 | Expand to Zero Crossings | 21 |
| DePlosive | 142 | Expander | 62 |
| Dialogue | | Export | 176 |
| <i>declick</i> | 144 | Export to audio files | 178 |
| Digital Audio | 8 | Exporting audio from multitrack sessions | 163 |
| <i>decibel (dB)</i> | 10 | Extract Dialogue | 135 |
| <i>quantisation</i> | 9 | | |
| <i>sampling</i> | 9 | | |

- F -

| | |
|-----------------------|-----|
| Fade | |
| <i>quick fade in</i> | 103 |
| <i>quick fade out</i> | 103 |
| Fade in | 162 |
| Fade out | 162 |
| Fades | 102 |
| Frequency analysis | 42 |
| Frequency domain | 36 |
| Fundamental frequency | 126 |

- G -

| | |
|---------------------|-----|
| Gate | 62 |
| Generating | |
| <i>test signals</i> | 77 |
| Grouping clips | 161 |

- H -

| | |
|-----------------|-----|
| Harmonics | 126 |
| Harmonize | 120 |
| Harshness | 146 |
| Histogram | |
| <i>spectral</i> | 44 |
| Hum | 126 |

- I -

| | |
|-------------------------------------|-----|
| Import page | 171 |
| Import stems to multitrack session | 165 |
| Impulse response | 109 |
| Insert | |
| <i>silence</i> | 77 |
| <i>test signal</i> | 77 |
| Inserting an audio file | 13 |
| Interpolate | 123 |
| Introduction | 5 |
| <i>Requirements</i> | 8 |
| <i>What is new in Acoustica 7.4</i> | 6 |
| Invert phase | 98 |
| Isolate dialogue | 135 |
| ITU-R BS.1770 | 99 |

- K -

| | |
|--------------------|-----|
| Keyboard shortcuts | 181 |
|--------------------|-----|

- L -

| | |
|---------------------------------------|--------|
| Labels | 24 |
| Lavalier noise | 138 |
| Left to right pan curve | 163 |
| Level meter | 48 |
| License key | 6 |
| Limit | 73 |
| Limiter | 73 |
| Linear phase | 90 |
| Link thresholds | 123 |
| Loading a workspace file | 30 |
| Loading an audio file | 12 |
| Loading multitrack sessions | 163 |
| Looped playback | 14 |
| Looping clips | 161 |
| Loops | |
| <i>adding loops for MIDI samplers</i> | 25 |
| Loudness | 50 |
| Loudness range | 50 |
| Lower threshold | 123 |
| LUFs | 47, 50 |

- M -

| | |
|------------------------------|-----|
| Manual click removal | 123 |
| MIDI samplers | |
| <i>adding loops</i> | 25 |
| Minimum phase | 90 |
| Mix paste | 21 |
| Mixed phase | 90 |
| Modulate | 116 |
| Monitoring | 32 |
| Morlet wavelet | 40 |
| Mouth declick | 144 |
| Moving clips | 161 |
| Multiband compressor | 69 |
| Multiband dynamic processing | 69 |
| Multiband Dynamics | 69 |
| Multiband expander | 69 |
| Multiply | 113 |
| Multitrack sessions | 155 |
| <i>creating</i> | 156 |
| Mute | 99 |

- N -

| | |
|----------------|-----|
| Noise analysis | 129 |
|----------------|-----|

| | | | |
|------------------------------|----------|---------------------------------------|----------|
| Noise reduction | 128, 135 | Recording | 32 |
| Noise shaping | 75 | Redo | 30 |
| Normalize | 99 | Reduce | |
| Notch filters | 126 | <i>bird noise</i> | 148 |
| - O - | | <i>buzz</i> | 140 |
| Odd harmonics | 126 | <i>hum</i> | 140 |
| Order Acoustica | 6 | <i>lavalier noise</i> | 138 |
| Output level | 153 | <i>lip smacks</i> | 144 |
| Overwrite paste | 21 | <i>mouth clicks</i> | 144 |
| - P - | | <i>plosives</i> | 142 |
| Pan curve | 163 | <i>pops</i> | 142 |
| Paste | 21 | <i>RF noise</i> | 140 |
| Phase | | <i>rustle</i> | 138 |
| <i>inversion</i> | 98 | <i>wind noise</i> | 136 |
| <i>rotation</i> | 97 | Region | |
| Phase correlation meter | 52 | <i>selecting</i> | 15 |
| Phase modes | 90 | Region markers | 24 |
| Phono filter | 96 | Remix | 79 |
| Pitch shifting | 119 | Remove DC offset | 97 |
| Playing | 14 | Remove track | 157 |
| Playing in loop | 14 | Repair page | 174 |
| Plosives | 142 | Requirements | 8 |
| Plug-in manager | 151 | Re-recording | 33 |
| Plug-in Pack | 184 | Restoration | 175 |
| Plug-Ins | 151 | Retouch tool | 58 |
| Polarity | | Reverb | 103, 109 |
| <i>inversion</i> | 98 | Reverb (Standard Edition) | 107 |
| Preferences | 180 | Reverse | 121 |
| <i>audio device settings</i> | 182 | RF noise | 140 |
| <i>general</i> | 180 | RIAA | 96 |
| <i>keyboard shortcuts</i> | 181 | Rich Text Format (RTF) | 27 |
| Pre-roll time | 33 | RMS | 47 |
| Pro Tools integration | 188 | <i>AES17-1998</i> | 99 |
| Processing chain | 153, 176 | <i>mathematical</i> | 99 |
| Processing tools | 61 | Rotate phase | 97 |
| Purchase Acoustica | 6 | Rustle reduction | 138 |
| - Q - | | - S - | |
| Quantization | 9 | Sample peak | 47 |
| Quick fade in | 103 | Sample peak level | 99 |
| Quick fade out | 103 | Sample rate conversion (SRC) | 77 |
| - R - | | Samplers | |
| Realtime analyzers | 48 | <i>adding loops for MIDI samplers</i> | 25 |
| | | Sampling | 9 |
| | | Saveing a workspace file | 30 |
| | | Saving an audio file | 14 |

| | | | |
|-------------------------------|---------|--------------------------------|---------|
| Saving multitrack sessions | 163 | Timer record | 34 |
| Scanning plug-ins | 151 | Tools | 61 |
| Scrollbar | 18 | Track splitting | 29, 175 |
| Scrolling | 18 | Tracks | 157 |
| Scrubbing | | Transcription | 27 |
| <i>audio</i> | 22 | Transfer | 188 |
| Selecting a region | 15 | Transpose | 119 |
| Selecting the active channels | 16 | True peak | 47 |
| Selection | 14 | True peak level | 99 |
| <i>playing</i> | 14 | | |
| <i>playing in loop</i> | 14 | - U - | |
| Selections | | Undo | 30 |
| <i>time-frequency</i> | 54 | Upper threshold | 123 |
| Send knob | 165 | | |
| Send level curve | 163 | - V - | |
| Sibalance | 146 | Verberate | 103 |
| Side chain | | Vitalize | 149 |
| <i>filtering</i> | 62 | Volume | 98 |
| Signal generator | 77 | <i>adjusting</i> | 100 |
| Silence | 99 | <i>applying a volume curve</i> | 101 |
| Silence removal | 29 | <i>channel mixer</i> | 100 |
| Sinusoidal resynthesis | 126 | <i>fade ins and fade outs</i> | 102 |
| Spectral editing | 53 | <i>normalize</i> | 99 |
| Spectral histogram | 44 | Volume curve | 163 |
| Spectrogram | 36, 37 | Volume manipulation | 98 |
| Spectrogram view settings | 56 | VST | 151 |
| Spectrum | 42 | VST3 | 151 |
| Spectrum analyzer | 52 | | |
| Spleeter for Acoustica | 79 | - W - | |
| SRC indicator | 153 | Waveform visualization | 10, 11 |
| Statistics | 47 | Wavelet | 36, 40 |
| Stem separation | 79, 165 | Wind noise reduction | 136 |
| Stretching clips | 161 | Workspace | 11 |
| SubRip | 27 | <i>customizing</i> | 178 |
| Subtitles | 27 | Workspace files | 30 |
| | | | |
| - T - | | - Z - | |
| Tame resonances | 146 | Zero crossing rate | 47 |
| Threshold | 146 | Zooming in or out | 18 |
| Threshold levels | 123 | | |
| Time display | 53 | | |
| Time domain | 36 | | |
| Time Stretch | 118 | | |
| Time stretching clips | 161 | | |
| Time-frequency domain | 37, 40 | | |
| Time-frequency selections | 54 | | |