

Courvoisier

User Manual

Revision A

1 Table of contents

1 Table of contents.....	2
2 Introduction.....	3
2.1 License.....	3
2.2 Revision History.....	3
2.3 Requirements.....	4
2.4 Compatibility.....	4
2.5 Known Limitations.....	4
2.6 Installation.....	4
3 About.....	5
3.1 Purpose.....	5
3.2 Technical Features.....	5
4 Operation.....	7
4.1 Quick Start.....	7
4.2 User Interface.....	7
4.2.1 General.....	7
4.2.2 Histogram and transfer curve.....	7
4.2.3 Histogram snapshots.....	9
4.2.4 Editing transfer curves.....	9
4.3 Mid/side and Stereo.....	10
4.3.1 The green transfer curve.....	10
4.3.2 Stereo operation and stereo link.....	12
4.4 Detection (side) chain.....	13
4.4.1 Detection algorithm.....	13
4.4.2 Filters.....	13
4.5 Main volume and dry mix.....	14
4.6 Attack and release.....	15
4.6.1 Attack.....	15
4.6.2 Adaptive release.....	16
4.6.3 User defined auto release curve.....	17
4.7 Look ahead.....	18

4.8 Miscellaneous.....	18
4.8.1 About box.....	18
4.8.2 Complex mode.....	18
4.8.3 Presets.....	19
5 Measurements.....	21
6 Credits.....	22

2 Introduction

2.1 License

Copyright (c) 2010, Tamas Zsiros

All rights reserved.

You are allowed to freely distribute this software in its original unaltered form and packaging. You are allowed to use this software for media production purposes. Selling, renting or reverse engineering this software is strictly prohibited.

This software is provided by copyright holder "as is" and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall copyright holder be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

'VST' is a Technology and Trademark by Steinberg.

2.2 Revision History

Revision	Date	Notes
A	2010-05-29	First public version

2.3 Requirements

- Win32 compatible operating system
- VST compatible host
- SS2 compatible CPU (with a minimum clock speed of 1 GHz)
- Acrobat Reader (for the user manual)

2.4 Compatibility

Courvoisier has been verified with the following hosts:

- Ableton Live 8.04
- FL Studio 9.0
- Reaper v3.52
- George Yohng's VST Wrapper
- Christian Budde's VST Plugin Analyser

2.5 Known Limitations

Curve automation is not supported in this release.

2.6 Installation

Just extract the .dll file into your VST host's VST directory. Refer to the host's documentation for details on how to do this.

3 About

3.1 Purpose

Courvoisier is designed to be a universal dynamics processor optimized for processing complex program material. The main design goal was to create a unique sounding processor with sonic possibilities beyond traditional dynamics effects (compressors, expanders etc.) with a novel user interface.

Even though Courvoisier was meant to imprint its own character over the processed program material, care has been taken to eliminate any perceivable harmonic or other forms of distortion. This enables the user to concentrate on the dynamics processing itself, and introduce any non-linearity separately.

Note: if you prefer analog style processing with added harmonics, I urge you to put some saturation / coloring plugin after or before Courvoisier. During testing I have got great results with adding an instance of FerricTDS and/or NastyVSD before or after Courvoisier. Both plugins are available for free at <http://varietyofsound.wordpress.com/>.

Courvoisier's own approach of dynamics processing comes at the price of a steeper learning curve. Various adaptive algorithms are therefore employed to ease operation and to provide good results even at default settings.

3.2 Technical Features

- Histogram representation (sliding-window statistics) of incoming program material and gain change
- Free-form dynamics processing defined by curves
- All performance-critical code optimized in assembly (utilizing SSE2 flow instructions)
- Very low harmonic distortion (THD during normal operation is below -105 dB FS)
- Two different program dependent attack algorithms
- Two release algorithms (program dependent and user defined)
- Side chain filtering (low and high pass filters)
- M/S operation with optional Side channel transfer function
- Stereo operation with adjustable stereo link
- Dry mix option for parallel processing
- Peak, RMS and Hybrid level detection options
- Look-ahead processing for minimizing distortion with extremely fast attack settings

4 Operation

4.1 Quick Start

As of this time, Courvoisier does only feature an advanced user interface. There is no “easy mode” implemented, and so it is highly recommended to read at least the User Interface chapter (4.2) of this manual.

4.2 User Interface

4.2.1 General

All controls can be manipulated with the mouse by clicking and dragging. All linear controls (sliders and filters) can be re-set to their respective default values by CTRL-clicking.

4.2.2 Histogram and transfer curve




Figure 1: The input histogram

The histogram part of the GUI is where most of the action happens in Courvoisier. A level histogram of the incoming audio material is plotted with **red** whenever the plugin is active.

Note: a histogram is basically a sliding-window statistics about the various levels present in the signal. The scale ranges from -60 dB FS to 0 dB FS (right to left), the range where most of the - so to say - musically interesting stuff happens. You usually see a nice bell shaped curve here (unless the input is mixed in some sort of esoteric way).

The transfer function is also drawn here in the middle of the screen (black line with orange control points). For now it is totally flat, and so Courvoisier does not change the dynamics of the input material at all.

On our next example the transfer curve has been changed from its default flat shape. The curve has a dip around -19 dB FS, and therefore Courvoisier will try to make levels around -19 dB FS quieter. The curve also features a boost around -39 dB FS. Courvoisier will boost these - relatively low - levels accordingly.




Figure 2: The transfer curve and the gain change histogram

The transfer curve has the vertical scale of +/- 12 dB, where 0 dB is in the middle, and +12 dB is where the curve is touching the top of the screen.

The histogram of the resulting level change (representing the difference between the input's and output's histogram) is now also visible (blue plot). According to the statistics the processed output now contains more levels around -15 dB and less around -26 dB.

4.2.3 Histogram snapshots




Figure 3: Blue and red histogram snapshots

It is possible to create snapshots of the input histogram anytime by pressing the “RED SAMPLE” or “BLUE SAMPLE” buttons.



Figure 4: Sample buttons

Snapshots can be used for referencing purposes and are stored into the actual preset. A snapshot can be stored and cleared by clicking on the respective button.

4.2.4 Editing transfer curves

Transfer curves in Courvoisier are defined by control points. One can create control points by simply clicking on the graph. To delete a control point just click and drag the control point out from the screen, either at the top or at the bottom.

4.3 Mid/side and Stereo

Courvoisier offers two modes of operation when it comes to dealing with two-channel input material: mid/side [M/S] and stereo.

Note: in M/S operation the original 2-channel program material is converted from left/right to mid/side, where the “mid” channel contains only the mono part of the signal (sounds coming from center soundstage) and the “side” channel contains everything else (sounds coming left or right of the soundstage).

Dynamics processing in a mid/side configuration is a topic way beyond the scope of this manual. For further information please refer to the vast resources available on the internet. A good place to start is http://en.wikipedia.org/wiki/Stereophonic_sound and <http://www.bluecataudio.com/Vault/Doc/MidSideProcessing.pdf> or just search for ‘mid/side processing’.




Figure 5: Stereo - M/S selector

In stereo operation the black transfer curve is used to define gain map for both channels.

In MS mode there is a possibility to use the same black transfer curve for both mid and side channels, but as side channels are usually lower in volume, it is more practical to use a second transfer curve. Also, with a separate transfer curve one can perform level dependent widening/narrowing of the soundstage.

4.3.1 The green transfer curve

Courvoisier provides a **green** transfer curve for side channel manipulation. The **green** curve is accessed by pressing the SIDE button on the user interface:




Figure 6: The side curve

The **green** side curve works just like the black one, but only effects the side channel when:

1. Mid/side mode is enabled



Figure 7: Mid/side mode enabled

and

2. the 'MS Curve [%]' slider is set to something else than 0%



Figure 8: MS Curve at 49%

This slider determines which transfer curve should Courvoisier use for the side channel in M/S mode. At 0% the black transfer curve is used for the Side channel, while at 100% the **green** side curve solely determines what should happen to the side (and any in-between values will mix between the two).

Note: when the **green** transfer curve is displayed (SIDE button pressed) the histogram will show the analysis for the side channel.

The example provided on Figure 6 shows a typical setup of spatial widening in Courvoisier in M/S mode.

4.3.2 Stereo operation and stereo link

The opposite of Mid/Side mode is the normal stereo operation (STEREO button pressed). This will cause Courvoisier to perform level detection on both left and right channels, and then it will mix these two detected levels according to the 'Stereo Link' slider. 0% of stereo link will result in totally independent processing on the left and right channels, something which might sound just great and wide, or which might just destroy the stereo imaging of the recording. By setting this slider to 100% the same identical processing will take place for both channels, and this should leave your stereo imaging intact, but in many cases one can experience slight or even moderate decrease in soundstage width.

With the stereo link slider one can easily find the right compromise between soundstage width and integrity.



Figure 9: Stereo link at 25%

Please note that the stereo link slider is only available in stereo mode of operation.

4.4 Detection (side) chain

4.4.1 Detection algorithm

Every dynamics processor needs an algorithm to determine the level of the incoming audio, so that it can make decisions based on that. Courvoisier has three such algorithms: Peak, RMS and Hybrid.




Figure 10: Peak-RMS-Hybrid switch

In fact all three settings use an RMS filter for the level detection, but peak's window size is so small that it 'feels' more like a "real" peak detector. Using some amount of RMS filtering ensures very low distortion with even the fastest time constants.

4.4.2 Filters

There is a possibility to apply further filtering (that is, additional to the filtering done by the RMS averaging) in the level detector circuit (side chain).




Figure 11: High and low cutoff filters

'Low Cutoff' defines a low cut filter, and can be used to get rid of the nasty effect when some high-powered sub bass modulates the whole action of the processor. The 'High Cutoff' control sets a very gentle high cut filter, and can be used - together with low cutoff to shift the "attention" of the processor to the mid frequencies (which usually convey the most information towards the listener).

Note: to set the filters just click and drag on the GUI.

4.5 Main volume and dry mix




Figure 12: Main volume and dry mix controls

Main Volume can be used to apply a +/- 6 dB boost/attenuation to the processed material. This could be used to compensate any volume differences between the processed and the unprocessed audio, and so make it possible to make comparisons at the same volume.

Note: it is very important to judge dynamics processing by comparing at the same volume otherwise one might be fooled by the "what sounds louder will automatically sound better as well" effect. For more information on this topic feel free to type 'Fletcher Munson' into your favorite internet search engine.

Dry Mix can be used to mix processed and unprocessed material. One typical use of this feature is to set up Courvoisier to perform drastic dynamics mangling, and then dial in the right amount by mixing back the unprocessed material. Useful keywords for your search engine are: 'parallel', 'processing' and optionally 'compression'.

4.6 Attack and release

A dynamics processor first determines how much it should boost or attenuate at a given point in time (this decision is based on the level detected in the side chain and the level target set by the transfer curve). The resulting target gain then needs to be applied to the input material. In order not to have abrupt changes in audio levels, a processor usually employs some kind of rate limiting to the applied gain. In plain English this means that the gain can only increase or decrease by some predefined speed. In general, Attack defines the speed by which gain can decrease and **Release** is the speed by which gain can *increase*.

It is important to have different speeds (defined by time constants) for gain increase and decrease, because in music the sounds we hear have usually faster attack times and slower release times and it is fair to assume that the better we replicate this phenomenon in our processing, the more transparent and natural end result we will get.

4.6.1 Attack

Attack is primarily set by the Attack Slider.



Figure 13: Attack slider set to 18 milliseconds

The attack slider is however only one input parameter to the complex adaptive attack algorithms which shape the attack speed continuously so that it best matches the input material.

Courvoisier provides two different adaptive attack algorithms:



Figure 14: Adaptive attack set to 'ATTACK ONE'

Both of them are different implementations of the same algorithm, with 'ATTACK TWO' being an early (and faulty) piece of code that can sound actually quite nice :-).

The bottom line is: because the attack speed is influenced by the adaptive process, the slider should not be considered as an 'exact' control but rather an approximate indicator. The default value of 18 ms will work pretty well for most input material, but this should deter no-one from experimenting a little with different values.

4.6.2 Adaptive release

Release time is set internally to 110 ms that - again - works very well with most program material, due to the fact that:

- A) 110 ms is the approximate release time of the human hearing (our ears function as compressors when encountering loud noises)
- B) the adaptive release algorithm changes the release based on the input material anyway

The 'Auto Release' slider can be used to dial in the right amount of adaptiveness into the release speed: just listen to the processed output and increase (or decrease) the auto release until it sounds right.



Figure 15: Auto release slider set to 40%

4.6.3 User defined auto release curve

There is also a third option for the absolute geek-hearted engineer/producer: by pressing the 'REL CURVE' button one can set up a red curve, that multiplies the release speed based on the detected input level.




Figure 16: User defined auto release curve

By this it is possible to create an own "adaptive release" algorithm. The curve defines a multiplier that modulates the release speed (calculated by the default adaptive release algorithm described in chapter 4.6.2). In other words the user-defined adaptive release does not override the built-in one, only changes it based on the detected level at any given time,

A completely horizontal line will not change the release behavior in any way, values towards the top of the screen slows down the release, while values towards the bottom make it faster. The complete range of the release curve multiplier goes from 0.1 (bottom) to 10 (top).

4.7 Look ahead



Figure 17: Look ahead enabled

Look ahead is an old digital compressor trick that makes the processor respond to transients earlier than they actually happen: something that a normal, non-relativistic universe would not make possible at all. Courvoisier does this (not by speeding up to relativistic speeds, but) by simply delaying the audio compared to the detection circuit, thereby "seeing into the future". Be aware: enabling this function introduces a 1 millisecond delay into the processing, something that your host might compensate against... or not. Fortunately, look ahead is only needed with very fast (1-4 ms) Attack times, so if you stick with anything higher, just be on the safe side and switch it off.

4.8 Miscellaneous

4.8.1 About box

To display credits, revision information and some other info just click on the 'Courvoisier' logo on the user interface.

4.8.2 Complex mode

There is a mysterious :-) switch 'hidden' at the bottom of the about box that says 'Complex Mode'.

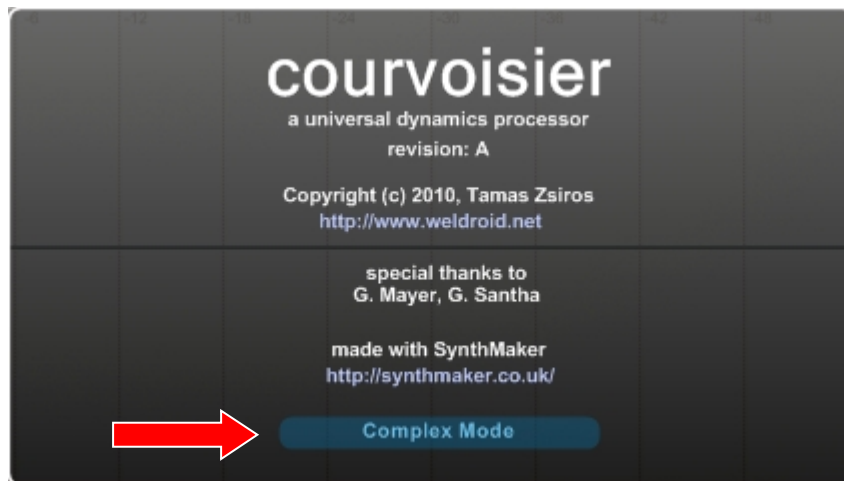


Figure 18: 'Complex Mode' switch

Complex mode started its life as plain old oversampling (it is said that oversampling generally improves the sound quality of dynamics processors), but during initial testing it became clear that Courvoisier benefits very little from oversampling when it is done properly. (To cut a long story short, a DSP algorithm needs to do some things differently when oversampling is applied.)

Then came - by mistake - a version of Courvoisier where oversampling was only halfway implemented, and it turned out to sound quite interesting due to the tricks it plays with the adaptive algorithms, which were not fully adapted to the internally increased sample rate.

Complex mode is now included for those who like to experiment a little and can be switched on from the about box (its value is saved with the preset as well).

4.8.3 Presets

Presets are managed via the actual host's preset management function. Presets using the 'Complex Mode' are marked with 'C' at the end of the preset name.

Presets can be a quite good way to get a showcase of Courvoisier's capabilities. Every preset included with Courvoisier has the input histogram snapshot of the original input material as well, and it must be noted that the presets will only work properly if the actual input's histogram is somewhat similar to the original histogram.

Red snapshots are always mono histograms, and blue snapshots are always side channel histograms.

5 Measurements




Figure 18: Harmonic distortion, Complex mode off




Figure 19: Harmonic distortion, Complex mode on

6 Credits

Tamas Zsiros: original concept, GUI, coding, testing and user manual.

<http://www.weldroid.net>

<http://weldroid.blogspot.com>

Gabor Mayer: feedback and testing.

<http://wolandroid.blogspot.com>

Gergo Santha: feedback and testing.

<http://phonland.blogspot.com>

Special thanks to Christian Budde for the VST Plugin Analyzer and
H. L. Goldberg for the inspiration.

