BusTools v3
Installation and user manual

Contents

1 Setting up the plugins for first use

1.1 Installation

Download and install the freely-available trial/evaluation versions of the plugins. This will allow you to test the plugins prior to making any purchase decisions. trial/evaluation versions can be downloaded from the ToneBoosters.com website. Trial/evaluation plugins can have one or more of the following limitations:

- Trial/evaluation versions will not store nor save settings.
- Trial/evaluation versions will show a reminder to purchase a license.

1.1.1 Installing the trial/evaluation plugins

- The plugins will be installed in the locations given in the table below.

Operating system	Plugin folder(s)
Windows	`C:\Program Files\Common Files\VST2\ToneBoosters\`
Mac OS	`/Library/Audio/Plug-Ins/Components/` `/Library/Audio/Plug-Ins/VST/`

1.1.2 System requirements (Windows)

- Windows 7 SP1 or higher
- 32 or 64 bit host program that supports VST plugins

1.1.3 System requirements (mac OS)

- macOS 10.10 or higher
- 64-bit host program that supports 64-bit VST or Audio Unit plugins

1.1.4 Configuring your host program

After installation of the trial/evaluation plugins, you may have to inform your host program about the presence of new plugins. Most host programs require you to provide the folder where plugins are installed.

- Consult your host program manual how to configure plugin folders. On Windows, make sure you add the following VST scan path to your host program settings:
  `C:\Program Files\Common Files\VST2\`
- Refresh and/or re-start your host program to allow it to scan for new plugins on your computer.

1.2 Upgrading from trial/evaluation to full/registered versions

Existing installations of trial/evaluation plugins can be upgraded to fully functional versions with a separate registration 'key file'. The registration key file can be acquired in the online shop at www.toneboosters.com. The same key files can be used for Windows and OS X. After purchase:

- Download and extract (unzip) the registration key file(s) ('TB_PluginName.key') and place it in the exact same folder as the corresponding demo/evaluation version of the already installed VST or AU plugin(s).

- On a Windows computer, in one and the same directory, you should see the following pair of files for each registered VST plugin:

  ```
  TB_PluginName_v3.dll
  TB_PluginName.key
  ```

  On a Mac, point finder to your /Library/Audio/Plug-Ins folder. Typically, the Library folder is a 'hidden' folder so you cannot simply browse to that folder. Instead, type Command+Shift+G from the Mac desktop (or Finder > Go > Go to Folder) and type in */Library* to temporarily access the Library directory in the Finder. Then navigate to your VST and Components folders and make sure you copy key files to see the following:

  ```
  TB_PluginName_v3.vst
  TB_PluginName.key

  TB_PluginName_v3.component
  TB_PluginName.key
  ```

- Restart the host program. The plugin should now display 'registered' in the lower-right corner of the GUI, instead of 'demo'.

Please make sure to make a backup copy of this registration key file; if the registration key file is lost or damaged the plugin will automatically downgrade to a demo version. Your computer's hard drive is **NOT** a good place for a backup.

Do not rename nor edit the key file. The registration key file comes in a zip archive. Just unzip the archive and copy the resulting key file into your plugin folder. Renaming or modifying the file will cause the registration key file to become dysfunctional.

Some FREE plugins also have an associated registration key file. This registration key file is included in the evaluation download package and allows verification that the key registration system works on your computer. Please do not delete these key files as it will downgrade these free plugins to demo/evaluation versions.

1.3 Installing plugin updates

Download the latest and greatest trial/evaluation versions from the downloads page at www.toneboosters.com and overwrite your plugin files with the newer ones. **Do not delete or modify your registration key files!** The registration key files are the files that have the 'key' extension. Restart your host program and you are all set.

1.4 Disclaimers

VST is a trademark of Steinberg Media Technologies GmbH.

2 User interface common controls

2.1.1 Controlling Knobs and sliders

The various knobs and sliders on the graphical user interfaces (GUIs) of the plugins can be controlled by left-mouse clicks (for switches) or left-mouse drags (for rotary controls and sliders). The following key combinations apply that modify the behavior of the GUI elements:

Windows:

- 'Control' key + left mouse click: set the control at its default value.
- 'Shift' key + left mouse drag: fine-tuning of the control.
- 'Alt' key + left mouse click + mouse move: jump to the clicked position.
- Mouse wheel: change the value up or down.
- 'Shift' key + Mouse wheel: fine-tuning of the control.
- Left or down key: change value down.
- Up or right key: change value up.
- Double left click (if the control has a numeric entry): manual data entry.

OSX:

- 'Command' key + left mouse click: set the control at its default value.
- 'Shift' key + left mouse drag: fine-tuning of the control.
- 'Alt' key + left mouse click + mouse move: jump to the clicked position.
- Mouse wheel: change the value up or down.
- 'Shift' key + Mouse wheel: fine-tuning of the control.
- Left or down key: change value down.
- Up or right key: change value up.
- Double left click (if the control has a numeric entry): manual data entry.

2.1.2 Controlling nodes

Some plugins have nodes (round dots) in a two-dimensional space or graph that can be dragged to change parameters. These can be dragged with the mouse to change their position and the associated parameter values. The following key combinations apply:

- Left mouse click: activate node.
- Right mouse click: de-activate node.
- Drag with left-mouse knob: modify position in active mode.
- Drag with right-mouse knob: modify position in inactive mode.
- Drag with left-mouse knob and 'alt' key: lock Y while moving node
- Drag with left-mouse knob and 'control'/'command' key: lock X while moving
- 'Shift' key + drag with left-mouse knob: fine-tune position in active mode.
- 'Shift' key + drag with right-mouse knob: fine-tune position in inactive mode.
- Mouse wheel: change secondary parameter (for example filter quality).
- 'Shift' key + mouse wheel: fine-tuning of secondary parameter.

2.1.3 VU meters

VU meters will often support 'peak hold' functionality, in which the most extreme value across time is indicated by a horizontal line with the peak value displayed numerically above this line.

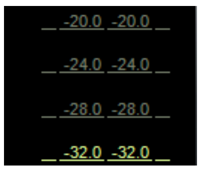

- Click on the VU meter to reset the peak hold value (if supported).
- Drag the VU meter scale to change its range (only in a limited set of plugins)

2.1.4 Bypass function

Most ToneBoosters plugins do not have a bypass switch – since virtually all VST hosts have their own bypass functions it is assumed that bypass operation is managed by the host, not by the plugin.

3 TB Isone v3

Binaural stereo loudspeaker setup and reproduction environment simulator for headphones.

3.1 Introduction

With TB Isone, a virtual stereo reproduction system and listening room can be experienced using high-quality headphones. Allowing for full control over loudspeaker cabinet type, loudspeaker distance, and room reverb, the virtual listening room can be largely customized. TB Isone can therefore be used to simulate a wide variety of loudspeakers and reproduction rooms during mixing, mastering, or to generate binaural recordings by post processing.

3.2 Features

TB Isone is a plugin that allows real-time, zero-latency binaural speaker and room simulation over headphones. Isone is best used with high-quality (full range) headphones having a flat frequency response. It features:

- Zero-latency processing, allowing for studio and live operation.
- Support of all sampling rates from 22 to 192 kHz.
- Loudspeaker designer to model speaker frequency response.
- Customizable room (volume, distance, early reflections, diffusion).
- Customizable loudspeaker azimuth angle (0 to 45 degrees).
- Customizable HRTFs (strength, head size, ear size).

3.3 The user interface

GUI section	Control	Purpose
Speaker setup designer	Loudspeaker response display	Shows the frequency response of the current loudspeaker model, both on-axis (thick line) and 45-degrees off-axis (thin line).
	Tweeter size	Sets the characteristic size of the tweeter and consequently the directivity of the loudspeaker.
	Speaker angle	The azimuth angle of the loudspeakers. Best set to 30 degrees.
	Channel mode (menu)	Allows to down-mix or solo the input channels, for example to verify mono compatibility of the audio mix at hand.

	Speaker presets (menu)	Selects a preset loudspeaker model.
	Out	Output VU meter indicating the overall output signal level. Click to reset the peak hold meter. Clipping may occur for signal peak levels above 0 dB. Reduce the SpkLev parameter to prevent clipping if necessary.
	CSC (Crosstalk Spectrum Compensation)	Enables a filter to compensate for the low-end bias of cross-talk signals.
	180	Inverts the phase of the output signals by 180 degrees.
Room designer	Room designer on/off switch	Enables/disables room acoustic simulation. When disabled, TB Isone will emulate an anechoic room.
	Size	Changes the simulated room size (volume).
	Early reflections	Changes the early reflections level.
	Diffusion	Changes the amount of diffusion of sound reflected from walls.
	SpkLev	Changes the loudness of the speaker in the room and consequently the output signal level of the plugin.
	T60	Changes the late reverb time of the room simulation.
	Room presets (menu)	Selects a preset environment (room) model.
HRTF designer	HRTF strength	Changes the strength (effect size) of the HRTF elevation cues.
	Ear size	Changes the HRTF ear size.
	Head size	Changes the HRTF head size.

3.4 Setting up and using TB Isone

3.4.1 HRTF Adjustment

1 – Recommended initial settings

The calibration of the HRTFs to each user's ears can be a somewhat tedious process, but fortunately is required only once if performed correctly. Here are some recommended settings that will work in most cases:

- Relatively small loudspeaker distance (about 2 meters).
- Room acoustic simulation disabled.
- CSC switch disabled.
- HRTF strength to relatively large values (90% or higher).
- Loudspeaker simulation disabled (flat response).

A dedicated preset 'Calibrate me!' is included to support this process.

2 – Selection of suitable audio material

The next step is to select suitable audio content to use during the adjustment process. It is recommended to use material that:

- You are very familiar with.
- Has a broad frequency spectrum.

Suitable material comprises voice recordings, snare drums, etc. Use mono content, or stereo content with very limited stereo depth, and little or no reverb.

Do not use band-limited signals such as sinusoids, or instruments covering a narrow frequency range, and alike – the human hearing system cannot localize such signals accurately.

The best situation for the HRTF calibration is when you sit in front of an actual loudspeaker setup with loudspeakers at the correct positions (-30 and +30 degrees azimuth, 0 degrees elevation).

3.4.2 Ear size adjustment

This is best performed by setting the channel mode to 'Dual Mono'. Listen closely to the test material. Ask yourself the following questions:

- Where does the sound come from?
- Does it come from above, or more from the front?
- Does it sound natural, or do I perceive unnatural timbres or frequency notches?

Rotate the ear size knob until the sound is perceived most natural, and coming most likely from the front. Wrong settings usually result in a sound perceived from above. Some people report that the adjustment process works best with their eyes closed.

3.4.3 Head size adjustment

This works best by setting the channel mode to 'Left' or 'Right'. Listen closely to the test material. Ask yourself the following questions:

- Where does the sound come from?
- Do I hear a well-defined image, or is it spatially blurred or ambiguous?

Rotate the head size knob until the sound position is most defined and natural, and is perceived at 30 degrees azimuth.

3.4.4 HTRF strength adjustment

The cue strength knob modifies the strength of the HRTF elevation cues. If this knob is set to 0, no elevation cues will be inserted in the audio and the HRTFs will have a flat frequency response. Higher values will insert more (stronger) elevation cues. Depending on your own preferences, and the audio content, the cue strength can be adjusted as desired.

Note: If the HTRF strength is set to zero, the ear size setting will not have any effect!

3.4.5 CSC – Crosstalk Spectrum Compensation

The CSC switch allows to enable or disable this compensation filter. Due to cross-feed, the signals of the right channel will not only be fed into the right ear, but also into the left ear and vice versa. Therefore the total signal power and the loudness will in most cases increase as a result of this cross-channel summation. The cross-feed signal has a low-pass character because to account for the acoustic shadow effect of the head. Consequently, without compensation such cross-feed will result in a stronger increase in signal energy at low frequencies than at high frequencies, resulting in a perceptible change in overall timbre or spectral balance. The cross feed spectrum compensation algorithm applies a correction filter to reduce this effect.

3.5 Model presets

3.5.1 Room/Environment presets

Nearfield	Typical near-field room setup with speakers at 0.75m from the listener and a T60 reverberation time of 0.3 seconds.
Midfield	Similar as above, but with a loudspeaker distance of 1.50m.
Farfield	Similar as above, but with a loudspeaker distance of 2.25m.
Even further	Simulation of loudspeakers placed far away.
Small studio	Typical simulation of a small, relatively damped (T60=0.4s) studio room.
Large studio	Typical simulation of a larger studio with a longer reverberation time (T60=0.6s).
Very small studio	Relatively dry studio with a low late reverb modal density.
Anechoic room	Simulation of an environment without reflecting surfaces.

Untreated box	Simulation of a almost square room with hard walls, resulting in substantial standing waves and flutter echoes.
Echo box	Simulation of a very large room and sound sources at a great distance with almost distinct echos
Very dry room	Simulation of a room with only very subtle room acoustics and a short reverberation time (T60-0.2s)

3.5.2 Speaker presets

This menu provides a variety of loudspeaker models that can be selected. The following table describes the characteristics of the various loudspeaker models. Please note that TB Isone does not contain or employ measured characteristics of existing loudspeakers but instead relies on analytical/theoretical models of loudspeaker cabinets, including the size, volume, driver type, resonance frequencies, enclosure type, and so on.

Flat	Reference setup, consisting of an essentially flat frequency response, and speakers placed at + and – 30 degrees azimuth, with directionality kicking in at around 3 kHz.
HiFi speaker	Typical HiFi loudspeaker with a broad frequency response and a small boost at 60 Hz and 20 kHz.
Small monitor	Typical small, single-driver, stereo loudspeaker setup with a relatively narrow response and high directivity.
Monitor A	A model that represent popular, commercially available near-field speakers.
Monitor B	A model that represent popular, commercially available near-field speakers.
Monitor C	A model that represent popular, commercially available near-field speakers.
Portable	Typical frequency response of a portable stereo audio player with speakers placed closely together.
Laptop	Very small loudspeaker simulation producing high frequencies only.
Flatpanel	Simulation of a flatpanel TV watched from a distance.
Mono radio	Single driver, mono and band-limited loudspeaker simulation as found in mono portable radios.
Too much!	Very wide loudspeaker setup (+/-45 degrees azimuth) with significant bass and treble boost.

3.6 Starting from scratch – Build your own speakers!

TB Isone has a speaker setup designer that allows you to set the most important properties of a virtual loudspeaker setup. These include:

- The frequency response.
- The characteristic size of the smallest driver (typically the tweeter).
- The loudspeaker setup angles.
- The speaker setup designer is shown below.

The two curves in the graph indicate the frequency response of the virtual loudspeaker, both on-axis (yellow line) and 45-degrees of-axis (grey line). The curves represent the response expressed in dB as a function of frequency in Hz.

The response can be modified by dragging each of the 3 available nodes (numbered 1 to 3) with the mouse. The left-mouse button will activate a certain node; the right-mouse button disables the node.

Nodes 1 and 2 determine the bandwidth of the loudspeaker (the lowest and highest frequencies that are reproduced by the speaker). Dragging these nodes above the 0 dB line will create a resonance at that particular frequency. Node 3 can be used to generate an additional resonance or dip at any desired frequency.

The grey line (45-degrees off-axis response) will largely follow the on-axis response. The difference between these two responses is determined by the tweeter size. A large tweeter will typically result in a more directive response at high frequencies, and consequently the 45-degrees of-axis response will decrease at high frequencies. A more directive loudspeaker response will reduce the amount of wall reflections and reverberation at high frequencies.

The speaker angle represents the azimuth angle of the loudspeaker. A value of 30 degrees indicates that the left and right loudspeakers are placed at +/- degrees azimuth from the listener's point of view.

The channel mode menu allows to solo the left or right loudspeaker, or to create a mono down mix that is subsequently reproduced by both virtual loudspeakers (dual mono mode).

The preset menu contains a list of presets for the speaker setup designer that may be good starting points for tweaking.

3.7 Virtual listening setup

3.7.1 Sound localization cues

A virtual listening room is typically created by simulating the acoustical transfer from all loudspeakers to both ears. These acoustical transfer properties are often referred to as Head Related Transfer Functions (HRTFs). Such HRTFs can be measured for each individual using specialized equipment. The measured transfer functions can subsequently be used as filters to simulate a virtual sound source over headphones.

HRTFs can be decomposed into two aspects:

- Binaural cues, defined by differences between the left-ear HRTF and the right-ear HRTF. These cues comprise (1) inter-aural time differences (ITDs) resulting from the difference in path length from a source to both ears, and (2) inter-aural level differences (ILDs), resulting from the acoustical shadow effects of the head. The binaural cues predominantly determine the perceived azimuth (left-right) of a sound source and hence result from acoustical cross-talk between both ear signals.
- Monaural cues, resulting from reflections in the pinnae, shoulder and from the torso of the human body. These reflections result in specific peaks and troughs in the signal spectrum that depend on the elevation of the sound source.

When room reflections are present, the pair of transfer functions including the wall reflections for a certain sound source position to both ears is referred to as binaural room transfer function (BRTF). This is shown in the figure below.

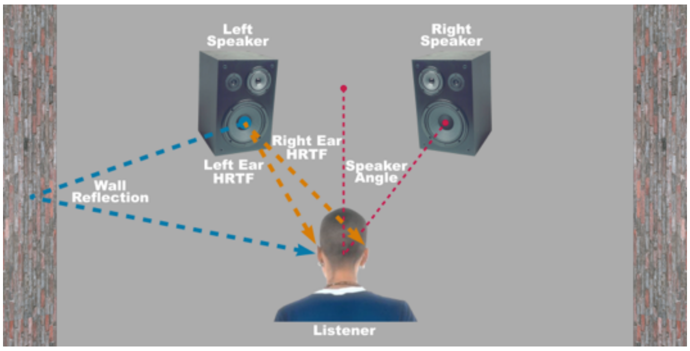

HRTFs for the left loudspeaker are indicated by orange lines. A single wall reflection is indicated by the line line. The speaker angle is between the red lines.

3.7.2 Parametric HRTF technology

Although the use of HRTFs has been shown to be very effective in numerous scientific publications, it also has well-documented shortcomings. For example, HRTFs vary from person to person as a result of differences in the head size, ear size, ear shape, and so on. Application of the wrong HRTFs results in significantly degraded sound source localization. It is therefore very important to match HRTFs to each individual listener for a convincing and accurate effect.

Isone Pro, the precursor of TB Isone, was the first VST plugin ever that provides such pseudo-personalized HRTFs. Now TB Isone provides the means to adjust the HRTFs for each individual listener, by compensating for differences in the anthropometric properties of the head and ears (pinnae).

The head size has the strongest influence on the binaural cues – inter-aural time and level differences. Hence a mismatch in head size often results in a wrong azimuth, but can also result in an ill-defined sound source position, or an unnatural sound percept.

The ear size has the strongest influence on the elevation cues – peaks and throughs in the spectrum induced by reflections in the ear. Hence a mismatch in the ear size often results in a lack of externalization, or sources erroneously perceived from above.

3.7.3 Loudspeaker simulation

Besides HRTF adjustment, TB Isone also allows to simulate a variety of virtual loudspeakers. Instead of simulating specific loudspeaker models, the approach taken in TB Isone is to simulate characteristic, common attributes of loudspeakers instead of accurate simulation of specific models.

3.7.4 Room acoustic modeling

Simulation of the acoustic environment is essential for a compelling simulation of loudspeaker listening over headphones. Music or other audio content is almost never listened to in an anechoic environment, and those who have experienced audio playback in such anechoic rooms know that this results in a very unpleasant listening experience. Moreover, your audience will listen to the content you work on in the car, in the living room, or any other echoic environment and hence it is crucial that the audio producer or engineer can estimate the effect of a room on the content he or she is working on. Nevertheless, the room simulation module in TB Isone can be switched on or off if desired.

4 TB EBULoudness v3

Loudness and true-peak meter compliant with EBU R128, ATSC A/85, and ITU-R BS.1770.

4.1 Introduction

The EBU published its Loudness Recommendation EBU R128. It tells how broadcasters can measure and normalize audio using loudness meters. TB EBULoudness and TB EBUCompact calculate k-weighted momentary loudness (LM), short-term loudness (LS), integrated loudness (LI) and loudness range (LRA) compliant with the EBU, ATSC and ITU specifications. Furthermore, true-peak levels (dBTP) are monitored as well.

Besides compliance to loudness requirements, the TB EBU Loudness plugin is also very useful tool to align the perceived loudness of different audio tracks (for example on an album). Differences in loudness (expressed as loudness units, or LU) can be directly translated into attenuation or gain expressed in dB to align the loudness of two or more tracks. Furthermore, the loudness range indicator can provide valuable information to verify the dynamic range of a track, and the potential need for dynamic range compression or expansion.

4.2 Features

- Loudness monitoring/metering compliant with ITU-R BS.1770, ATSC A/85, EBU R128, and EBU Tech report 3341.
- Loudness range (LRA) support according EBU Tech report 3342.
- EBU mode LUFS, EBU+9, EBU+18 and EBU+27 loudness scales and ITU-R BS.1770 LKFS loudness scale.
- Inter-sample (ISP) / ITU-R BS.1770 compliant 'true peak' detection
- Support of all sampling rates from 22 kHz upwards
- Stereo and 5.1 surround modes
- Includes a separate 'compact' plugin for <u>stereo content only</u> (several features are excluded)
- Virtually unlimited integration time
- Loudness history (up to a maximum of 2 hours) with hover and zoom functionality
- Ability to sync with play/pause of the DAW host (if supported by host)
- Based on the VST 2.4 specification to allow compatibility with virtually all host programs.

4.3 TB EBULoudness and TB EBUCompact

This plugin comes as a set of two plugins:

- **TB EBULoudness** is a 5.1-channel plugin assuming channel order Lf, Rf, C, LFE, Ls, Rs. It can process stereo and 5.1 content, provided that the host program is capable of running 6-channel plugins.
- **TB EBUCompact** is a stereo plugin and cannot process multi-channel content.

4.4 The user interface

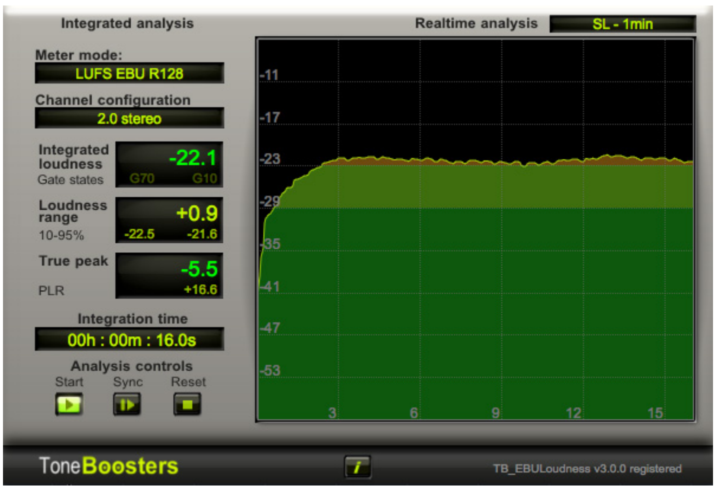

GUI section	Control	Purpose

Integrated analysis	Integrated loudness	K-weighted LI integrated loudness across the full integration time expressed in LU, LUFS or LKFS. The G10 and G70 indicators will illuminate when the relative and absolute gates are active, respectively (not for ITU-R 1770-0).
	Loudness range	K-weighted LRA loudness range across the full integration time, expressed in LU, LUFS or LKFS. The numbers below of the loudness range indicate the 10% and 95% percentiles of the short-term loudness distribution*.
	True peak	Maximum true peak (dBTP) observed since the last meter reset. The number blow the true peak value will indicate the PLR (peak-to-loudness ratio)*.
	Meter mode	Sets the display and metering modes to one of: • LU EBU R128 (2014) • LU EBU +9 • LU EBU +18 • LU EBU +27 • LKFS ATSC A/85 (2013) • LUFS EBU R128 TB-3 • LKFS ITU-R BS.1770-0 • LKFS ITU-R BS.1770-3 • LU K20 v2 (-20 LUFS) • LU K16 v2 (-16 LUFS) • LU K14 v2 (-14 LUFS) • LU K12 v2 (-12 LUFS) • LU K16 v2 d (-16 LUFS)
	Channel configuration*	Select 2.0 stereo or 5.1 surround metering configuration. For 5.1 surround, the channel order must be front left, front right, center, LFE, left surround, right surround.

Realtime analysis	Mode*	Selects the real-time analysis mode:
		• VU meters: shows momentary and short-term loudness VU meters, as well as true-peak meters for each audio channel.
		• LS (time): Shows the history of observed short-term loudness values. Time indicates the range from most recent value backward. In this mode, the following interactions are enabled:
		  ○ Hover: if one moves the mouse pointer over the plot, the loudness value corresponding to the x-coordinate of the mouse pointer is given.
		  ○ Select: by left-mouse-click and dragging, a selection of the curve can be made for a zoom / detailed view of the data.
		  ○ Left-mouse-click (without drag) to zoom out completely.

Analysis controls	Start	Start / continue the integrated loudness and loudness range measurement.
	Sync	Enable or disable the pausing of the integrated loudness and loudness range meters if the host DAW stops playback (only for hosts that support this function).
	Reset	Reset all meters.
Integration time	-	Amount of time used for integrated loudness and loudness range measurement.

** Not available in the EBUCompact meter*

4.5 Loudness standards and target loudness

Both EBU R128 (for Europe) and ATSC A/85 (for USA) are both based on loudness metering defined in ITU-R BS.1770. The target loudness, gating mechanism, the loudness units and maximum allowed true peak levels are nevertheless different, as indicated in the table below. If true peaks or the integrated loudness value are outside the valid range, the plugin will display the values in red instead of green.

Please note that the values below are taken from the 2011/2012 versions of the standards; please consult the respective documents to verify that these values are still correct.

Mode	Loudness unit	Gating	Target loudness	Maximum true peak
LUFS EBU R-128 (2014)	LUFS	Yes	-23 +/- 1 LUFS	-1 dB FS
LU EBU +9	LU	Yes	0 +/- 1 LU	-1 dB FS
LU EBU +18	LU	Yes	0 +/- 1 LU	-1 dB FS
LU EBU +27	LU	Yes	0 +/- 1 LU	-1 dB FS
LKFS ATSC A/85 (2013)	LKFS	Yes	-24 +/- 2 LKFS	-2 dB FS
LKFS ITU-R BS.1770-0	LKFS	No	-23 +/- 1 LKFS	-1 dB FS
LKFS ITU-R BS.1770-3	LKFS	Yes	-23 +/- 1 LKFS	-1 dB FS
LUFS EBU R128 TB-3	LUFS	Yes	-23 +/- 1 LUFS	-3 dB FS
LU K-20 v2	LU	Yes	-20 +/- 1 LU	-1 dB FS
LU K-16 v2(d)	LU	Yes	-16 +/- 1 LU	-1 dB FS
LU K-14 v2	LU	Yes	-14 +/- 1 LU	-1 dB FS
LU K-12 v2	LU	Yes	-12 +/- 1 LU	-1 dB FS

4.6 Setting up and measuring loudness

- Include the plugin in the last stage of the master bus as ‘insert’ plugin. Make sure that no other audio processing is performed subsequent to the loudness measurement. The loudness measurement plugin does not modify the audio signal, it only performs real-time metering.
- Specify the desired loudness measurement method. Use the ‘mode’ drop-down menu to select one of the supported loudness measurement methods / standards.
- Reset the meters by clicking on the ‘reset’ button.
- Determine whether you want to stop and start measurement of integrated loudness via the host (sync enabled) or via the plugin (sync disabled).
- Play the audio with the meters activated. Stop the host and/or plugin when the measurement period is finished.
- Read out the loudness and peak values of interest.

After the loudness of a program is measured, the required corrective gain (in dB) for loudness compliance can be simply obtained by taking the target integrated loudness and subtracting the measured integrated loudness:

$$G(dB) = LI_{target} - LI_{measured}$$

For true-peak compliance, it is advised to use an ITU-R BS.1770 compatible peak limiter with true-peak detection functionality, such as TB Barricade.

5 TB Barricade v3

Mastering-grade, transparent, highly customizable peak limiter with integrated dithering and perceptual noise shaping.

5.1 Introduction

TB Barricade is a stereo, mastering-grade peak limiter which supports control over the attack and release times, look-ahead time, and includes a quantization, dithering and perceptual noise shaping module to deliver high-quality delivery signals with limited bit depths. It is especially suitable to generate pristine final delivery signals for CD, DVD, online delivery, broadcast or podcast applications.

5.2 Features

- Fixed delay (1023 samples)
- Adjustable input and output gains
- Adjustable look ahead, attack and release times
- Inter-sample (ISP) / ITU-R BS.1770 / EBU R128 compliant 'true peak' detection and limiting
- Supports both waveform and envelope limiting
- Highly transparent limiting even with very high input levels
- Peak-hold VU meters with adjustable scales (K12, K14 or K20, or digital peak)
- Peak-hold RMS meters
- Quantization, dithering and perceptual noise shaping module
- Support of all sampling rates from 22 to 192 kHz
- Based on the VST 2.4 specification to allow compatibility with virtually all host programs.

5.3 The user interface

GUI section	Control	Purpose
Limiter gain	Env	Displays the limiter envelope reduction in dB including peak hold. Click on the scale to reset the peak hold function.
	Multiband	Displays the signal amplitude reduction resulting from multi-band limiting. Click on the scale to reset the peak hold function.
Signal levels	Input gain	Gain applied to the input signal before limiting (in dB).

	Out ceiling	Maximum output level of the limiter (in dB).
Limiter dynamics	Attack	Response time constant to loudness increases (in seconds).
	Release	Response time constant to loudness decreases (in seconds).
	Lookahead	Lookahead time of the limiter to respond to overs (in seconds).
	Stereo link	Amount of linkage between the limiter operating on the left and right audio channels. Higher stereo link levels will improve the stereo image at the (potential) expense of lower overall loudness. Stereo link does not influence the waveform auto saturation operation.
	Multiband	Amount of multiband limiting. Set to 0 to exclude multiband limiting.
Output resolution	Dithering	Bit depth for final delivery output signals. Set to 'off' to exclude quantization and dithering.
	Noise shaping	Amount of perceptual noise shaping applied to the quantization errors and dithering signals. Higher values will result in lower quantization noise audibility.
Output level	VU meters	Peak (with peak hold) and RMS (with peak hold) display. Click to reset peak-hold values.
	Meter type	Select the meter scale (peak, K12, K14 or K20).
Switches	ISP	Enable true-peak / ISP limiting (for final delivery signals).
	Monitor	When enabled, the limiter operation is applied to the input signal without incorporation of the input and output gains. This allows to listen to the limiter operation without impacting loudness.
	AES17 +3dB	When enabled, the RMS readout is increased by 3.01 dB to align peak and RMS levels of sinusoidal signals.
	Meter reset	Reset all peak-hold values of the GUI VU meters.

5.4 Setting up and using TB Barricade

5.4.1 Input gain and output limit

TB Barricade limits the maximum amplitude of its input signals. The amount of limiting is determined by

- the input signal level;
- the input gain control, and
- the output ceiling control.

Input signals are first attenuated or amplified by an amount determined by the input gain control. Subsequently, the maximum amplitude is limited to the value of the out ceiling control. Values above the value of the output limit control are referred to as 'overs'.

The amount of gain (or attenuation) applied by the limiter is indicated by the limiter gain meters for the left and right channels individually.

The effect of the limiter (gain) can be evaluated without actual incorporation of the input gain by activating the 'mon' (monitor) switch.

5.4.2 Lookahead, attack and release times

Limiters need a certain amount of lookahead to allow for a smooth gain curve. If an over is detected, the limiter will already gently start attenuating the signal a few milliseconds in advance of the over. This lookahead prevents distortion and intermodulation artefacts. Depending on the audio content, values between 1 and 3 milliseconds will generally suffice.

Part of the character and transparency of TB Barricade results from its intelligent algorithm that discriminates between instantaneous (short) peaks, or overs, and long-term loudness increases that result in many consecutive overs.

- Instantaneous, sporadic overs are limited by fast reacting limiting action which is determined by the lookahead time.
- Long-term loudness increases resulting in many or consecutive overs are limited by longer-term loudness estimation. The attack and release times of this loudness analysis are set by the attack and release controls:
- A long attack time will result in a slow reaction to loudness increases, and will typically result in more loudness at the output of the limiter.
- A short release time will quickly recover the limiter from loud passages, resulting in more loudness at the expense of a (risk of) breathing/pumping artefacts.

5.4.3 Stereo link

If the limiting gain is different for both audio channels, the spatial image of the audio content may shift towards the center position. To prevent distortion of the spatial image, TB Barricade allows to link the limiter action between the channels.

- Stereo link values between 0 and 50% will gradually link loudness estimates between the channels, but allow the limiter to still process instantaneous peaks in channels individually.
- Values above 50% will gradually link instantaneous peaks across channels as well.

5.4.4 Multiband

Barricade features a fully automatic multiband limiting algorithm. Opposed to wide-band envelope limiting, this stage processes individual frequency components. For many types of content, a certain amount of multiband limiting will result in more transparent limiter behavior in situations of very high signal levels, or extreme limiting. Setting the control to 0 will switch off the multiband limiter. The amount of multiband limiting is visualized in the limiter gain VU meters. In most cases, the signal attenuation as a result of multiband limiting will not exceed 6-8 dB to ensure that the timbre of the audio content is not changed significantly.

5.4.5 VU meters and scales

TB Barricade features RMS and peak output meters. Peak meters indicate instantaneous digital peak (maximum amplitude); RMS meters indicate the average signal power with an exponentially-decaying time constant of 300 ms.

Four different output scales can be used:

- 'Digital peak': A full-scale digital signal corresponds to 0 dB on the meters.
- 'K12': A full-scale digital signal corresponds to 12 dB on the meters. This scale is typically used for broadcast applications.
- 'K14': A full-scale digital signal corresponds to 14 dB on the meters. This scale is also typically used for CD mastering.
- 'K20': A full-scale digital signal corresponds to 20 dB on the meters. This scale is typical for DVD authoring, and classical music.

The aim of these various scales is to control the amount of headroom for peaks in the audio content with respect to the RMS (or loudness) level. The proper use of these metering systems is beyond the scope of this manual. The reader is referred to other resources. The peak hold values indicated by the meters can be reset by clicking on the respective indicator.

5.4.6 AES17 RMS+3

Mathematically, a sinusoidal signal has a peak value that is 3.01 dB higher than its power (RMS). For output metering, on the other hand, it can be convenient to align peak and RMS values for sinusoidal signals. If this behaviour is intended, the RMS+3 control should be activated. This setting will increase the RMS readout by 3.01 dB and is recommended when interpreting RMS values of the various K scales.

5.4.7 ISP

The 'True peak / ISP' switch determines whether Inter-Sample Peaks (ISP) will be taken into account in the limiter (if set to 'on'). Digital-to-Analog (D/A) converters often employ up-sampling and interpolation of audio signals. During this process, new audio samples are inserted in-between current audio samples. These samples may extend the full digital scale, even if the original samples are all within the full digital scale.

When the True peak/ISP switch is on, the limiter will protect against such potential clipping problems. The use of True peak/ISP is only necessary if used as limiter operating on the master bus for generation of final output delivery signals.

The True peak/ISP implementation of TB Barricade is compliant with ITU-R BS.1770.

5.4.8 DC reject filter

To ensure a DC-free output signals, TB Barricade has a build-in DC rejection filter with a fixed -3 dB cut-off frequency of 1 Hz.

5.4.9 Output resolution

If TB Barricade is used to deliver final delivery signals (e.g. on a master bus) with a limited bit depth (for example 16 or 24 bits), dithering and perceptual noise shaping module should be enabled. Set the dithering resolution to the number of bits of the final output format (16 or 24 bits).

Quantization and dithering always results in the generation of quantization errors, or quantization noise. The audibility of this noise can be greatly reduced by the processes of perceptual noise shaping. Noise shaping changes the spectrum of the quantization noise such that it becomes less audible. The amount of perceptual noise shaping can be controlled with the noise shaping control. A value of 0 indicates no noise shaping; 100% indicates maximum noise shaping.

Dithering and noise shaping should only be enabled if TB Barricade is the **last processing plugin** to render a final output signal.

6 TB ReelBus v3

Analog tape simulation plugin carefully modeled after legendary Japanese and Swiss reel-to-reel recorders.

6.1 Introduction

TB ReelBus is an analog tape recording simulator that aims at accurate simulation of all properties related to tape, including its frequency and level dependent saturation, inter-modulation effects, bias dependencies, tape hiss, asperity noise, wow and flutter, and clipping of electronic circuitry. It is especially suitable for bus processing (including the master bus) to subtly sweeten and enhance the sound.

TB ReelBus contains several tape recorder simulations (device models), which can be adjusted individually by offsetting their tape hiss, asperity noise, amount of spectrum and saturation processing, and alike.

6.2 Features

- Very low-latency processing (4 samples, compensated for by host) as a result of analog design
- Support of all sampling rates from 44.1 up to 192 kHz
- Adjustable record level with auto level makeup option
- Accurate simulation of existing reel-to-reel recorders with different tape speeds
- Adjustable tape hiss and asperity noise levels
- Adjustable tape spectrum and tape saturation
- Adjustable wow and flutter strength
- Option to amplify bias strength for overbiasing
- Simulation of both tape saturation as well as analog circuitry clipping
- Calibrated analog VU meters
- Each and every processing element carefully modeled after analog circuits and filters
- Based on the VST 2.4 specification to allow compatibility with virtually all host programs

6.3 The user interface

GUI section	Control	Purpose
Rec level		Adjust the (simulated) recording level. The effect of this control is visualized through the VU meters. A higher rec level will create higher internal input levels.
Device adjustment	Device model	Selects the reel-to-reel device model.

	W&F	Sets the amount of wow&flutter. Set to full left to disable wow and flutter simulation.
	Overbias	Increases the high-frequency bias signal beyond its optimal operating point for the selected device model.
	Circuit clip	Increases the amount of electronic circuitry clipping. Set to zero if no circuitry non-linearities are desirable.
Pre/post emphasis	Enable emphasis	Enables / disables pre- and post emphasis. This feature enables a pre-emphasis applied to audio signals before recorded to tape, and a complimentary (inverse) post de-emphasis applied afterwards. The pre-emphasis can be configured in the spatial domain (with the mid-side control) and the spectral domain (low-high).
	Mid-Side	Amount of pre-emphasis in the spatial domain. Negative values put more emphasis on the mid component; positive values put more emphasis on the side component.
	Low-High	Amount of pre-emphasis in the frequency domain. Negative values put more emphasis on low frequencies; positive values put more emphasis on high frequencies.
Noise adjustment	Tape hiss	Adjusts the amount of tape hiss (relative to the tape hiss level of the selected device model).
	Tape hiss -30 dB	Reduces the tape hiss by an additional 30 dB.
	Asperity noise	Adjusts the amount of asperity noise (relative to the asperity noise level of the selected device model).
	Asperity noise -30 dB	Reduces the asperity noise level by an additional 30 dB.
Color adjustment	Spectrum	Adjusts the amount of spectral changes induced by the selected device model. This can be compared to the 'EQ' part of the device. Set to 0 if no spectral changes are desired.
	Saturation	Adjusts the amount of tape saturation induced by the selected device model. Set to 0 if no or very little saturation is desired.
Output gain		Sets the output gain. If the 'Auto' switch is enabled, the inverse of the 'Rec level' control will be automatically included to compensate for level changes as a result of a non-zero rec level setting.

6.4 Setting up and using TB ReelBus

6.4.1 Signal level dependencies

The input signal level can be adjusted with the large input gain control. The operation of TB ReelBus is, similar to real tape, very much signal level and frequency dependent. Higher signal levels will correspond to stronger tape saturation. The amount of tape saturation can be controlled in two ways:

- A higher input level will result in stronger saturation; and
- The 'threshold' at which saturation starts can be adjusted with the 'saturation' control: a higher value will result in stronger tape saturation effects.

6.4.2 VU meters

Similar to real analog VU meters, the VU meters of TB ReelBus do not represent digital peak values. Instead, they compute averaged signal levels with averaging time constants that are in line with those of analog VU meters. The meters are calibrated to have a reading of 0 dB for a 1 kHz sinusoid with an RMS of -20 dB FS.

6.4.3 Device models

TB ReelBus contains settings for several tape recording machines (device models). Some of these models are intended for high-quality use, while others have a more dramatic effect. As these device models are all carefully modeled according to real tape recorder units, every device model has its own tape hiss, asperity noise, saturation, spectrum, circuitry clipping and wow&flutter properties.

It is important to note that the controls on the user interface will always be **offsets relative to the selected device model**.

6.4.4 Noise sources

TB ReelBus simulates both tape hiss (which is an additive type of noise) as well as asperity noise (which is a modulating noise that depends on the input signal level). Asperity noise is (in part) the result of inhomogeneities in the tape oxide coating, presence of dust particles and other stochastic influences. Both the tape hiss and asperity noise have been carefully modeled for each device model and will therefore change if a different device model is selected.

6.4.5 Color adjustment

The changes in timbre and dynamics of TB ReelBus can be separated into a static 'spectrum' part which is can be understood as an equalizer that imposes a certain tape frequency characteristic onto the audio signal. The amount of this effect can be changed with the 'spectrum' control.

Besides such static characteristic, TB ReelBus also simulates the frequency and level dependent saturation and spectrum of tape. The amount of this effect can be controlled by the 'saturation' control. Higher values will give a more pronounced effect.

Both the spectrum and saturation controls are offsets relative to the selected device model.

6.4.6 Wow and flutter

Wow and flutter are the result of small variations in tape speed, that cause changes in pitch / frequency. The wow and flutter simulation can be switched off by setting the control to full left. Full right gives twice the amount of wow and flutter of the selected device model.

6.4.7 Bias and overbias

Tape recorders add a very high frequency bias signal to the incoming audio before recording the combined signal to tape. The bias signal has a frequency that is typically around 100 kHz or higher, and improves the response of the tape. Low bias levels will give a 'brighter' timbre as the high-frequency response is more neutral, at the expense of more harmonic distortion at lower frequencies. Hence the bias control provides a trade-off between high-frequency response and low-frequency distortion.

6.4.8 Circuit clip

Some tape recorders cannot handle hot input signal levels accurately, and tend to demonstrate analog clipping. This clipping behavior is also carefully modeled for each device model. A setting of 12h will result in the same analog clipping level threshold as the original device; higher values will result in more clipping effect. If no, or very little effect of analog circuitry clipping is desired, set this control to full left.

6.4.9 Pre-emphasis and post de-emphasis

The timbre of the tape simulation can be modified by enabling a pre-emphasis and a complimentary de-emphasis. The pre-emphasis can increase/decrease the level of the side signal relative to the mid signal (Mid-Side control) or increase/decrease the level of the high frequencies relative to the low frequencies (Low-High control). This process is applied on the signals before tape simulation takes place. The inverse process (post de-emphasis) is automatically applied afterwards. Increasing the input level of the high frequencies will result in more aggressive processing of these high frequencies and vice versa. The same is true for the relative levels of mid and side.

6.5 Bounce tracks with TB ReelBus

TB ReelBus includes accurate simulation of tape hiss. This can have consequences for bouncing tracks in some hosts.

6.5.1 Bouncing tracks in Apple Pro Logic

The small amount of tape hiss generated by TB ReelBus can lead to very long or infinite bouncing behavior when the 'Include Audio Tail' option is enabled (the tail will never end due to the hiss simulation). To resolve this, either:

- Disable ‘Include Audio Tail’ in the bounce dialog; or
- Reduce the tape hiss level on the plugin; or
- Insert a noise gate after TB ReelBus and adjust it such that the gate will close and remove the tape hiss at the end of the track.

7 TB FlX v3

Dynamic equalizer - blending flexible dynamics processing and equalization in one optimized plugin.

7.1 Introduction

7.1.1 Equalizer section

TB FlX (“*Flex*”) combines equalization and dynamics processing in one go. It works just as most equalizers; it has 6 filter sections with lots of controls to modify their effect on the spectrum. More than 30 filter types are currently supported, which include classic analog peaking and shelving filters and resonating low- and high-pass filters. Besides these conventional filter types, some not-so-common or entirely novel filters are available as well:

- Bell-shape filters that have a flatter top than analog filters, to give a more natural sound;
- Non-resonating shelving filters to allow for steeper filter slopes;
- Gaussian filters, because these filters have the shortest possible group delay;
- Gammatone filters, because they closely mimic the frequency analysis of our hearing system;
- Linear and logarithmically-spaced harmonic filters, for creative effects;
- Brick-wall highpass, lowpass, and bandpass filters;
- Analog resonating highpass and lowpass filters (order between 1 and 16);
- Analog bandpass filter;
- Spectral balance filter;
- and several more.

In the unlikely case that you want to create a filter shape that is different from any of the included ones, TB FlX supports a so-called ‘auto node link’ mode. In this mode, the filter shapes will be automatically constructed such, that their combined effect will give a smooth, interpolated curve through all nodes that were configured as ‘auto node link’ filter.

Each filter section has its own ‘amount’ control to modify how much of that filter is actually applied to the audio signal. Furthermore, the filter can be applied in stereo, left only, right only, mid only or side only channels.

7.1.2 Linear phase or minimum phase?

We have closely followed the sometimes intense debates on minimum phase and linear phase equalizers, and the pros and cons of each of them. We understand that you appreciate the surgical editing precision of linear phase, and the snappy, accurate time response without pre-ringing of minimum phase. Wouldn’t it be great if these benefits could be combined in one algorithm, without their cons?

TB FlX has the potential to solve the minimum vs linear phase debacle for once and for all by introducing ‘auto phase’. The phase response of the equalizer is dynamically adjusted to give the best of both worlds, continuously adapted to the input signal and the desired equalizer curve.

7.1.3 Dynamics processing without limits

Each filter section also has a dedicated dynamics processing (compressor) stage. This means that a filter response can become more or less active depending on the input signal level. The compressor input/output curve is configured by dragging nodes in an input/output graph. This way, the input/output curve can be configured beyond what is possible with conventional compressors using threshold and ratio controls. Expansion, downward compression, upward compression, positive or negative ratios, parallel compression, New-York style compression (via the ‘amount’ control) are all easy to set up, and can be configured for each filter section independently. Of course you can set the attack and release times, and an auto-release (AR) function is supplied as well.

The level detector of each compressor section also has some innovative features. Besides being able to respond to stereo level changes, the level detector side chain can also operate on mid only, side only, left or right only, and even on the side-over-mid ratio!

7.1.4 Putting it all together

All user-accessible and internal parameters are interpolated automatically with maximum precision, using 4 times oversampling. This gives ultra-smooth, zipper-noise free behaviour for no-compromise, professional-grade output quality. And did we mention the integrated DC-reject filter (DCF)?

7.2 Features

- More than 100 parameters to shape the sound in a clean and simple interface.
- 6 filter sections with many controls to modify their behavior;
- More than 30 filter types, including a unique 'auto node link' filter type
- 3-node dynamics processing editor for each filter section
- Manual and auto-release (AR) option
- Integrated output spectrum analyzer with zoom functionality
- Unique, innovative auto-phase filter mode for high-resolution transient response
- Based on the VST 2.4 specification to allow compatibility with virtually all host programs.

7.3 The user interface

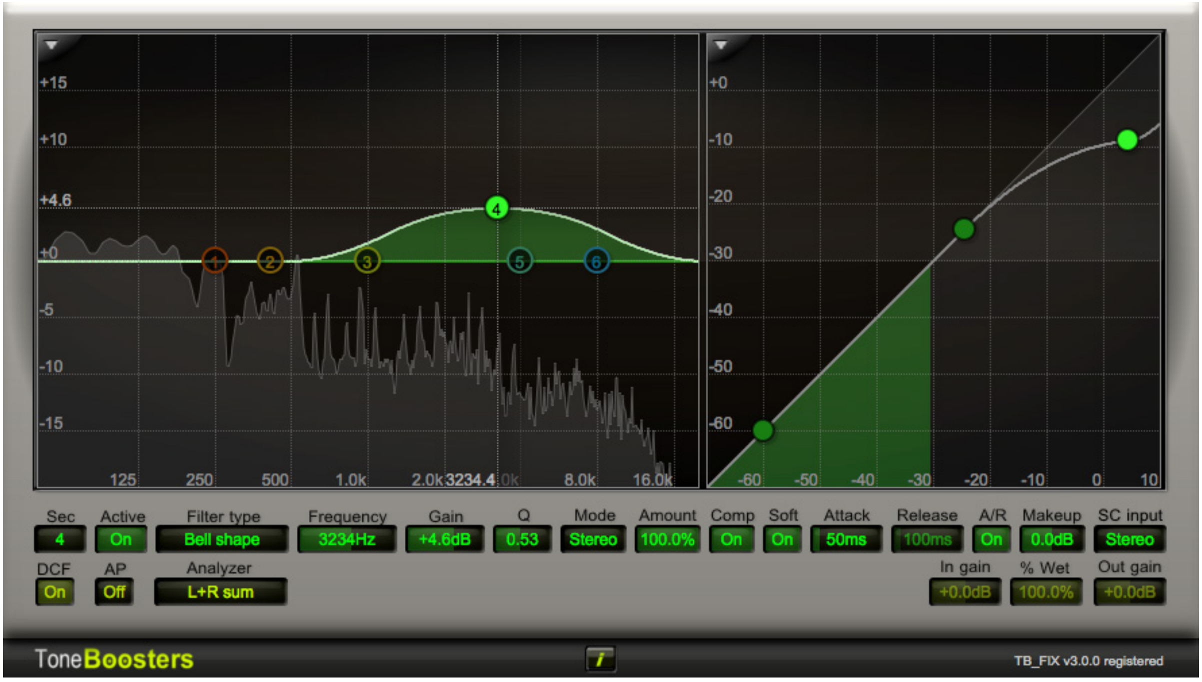

GUI section	Control	Purpose
Spectrum editor	Nodes / sections	Each node represents one equalizer/dynamics processor section. There are 6 of these sections available. Each node can be placed anywhere in the grid. The x coordinate determines the (center) frequency of the equalizer/dynamics processing section; the height determines the gain. • Left-click a node to activate a section, and to activate its dynamics editor. • Right-click a node to de-activate it. • Click elsewhere in the editor to zoom out • Drag the mouse to zoom into an area for microscopic editing. When a section is selected, the filter response corresponding to that section is highlighted with a color depending on the section index. The white line indicates the overall filter curve for all active filter sections simultaneously.
	Sect	Selects the equalizer/dynamics section (off, or 1 to 6).

	Filter type	Select the filter type of the current section (low pass, high pass, bell shape, etc)
	Frequency	Sets the frequency of the current section
	Gain	Sets the gain (in dB) of the current section
	Quality (Q)	Sets the quality factor of the current section section. A higher Q value means a narrow bandwidth, or a higher resonance (depending on the filter type).
	Mode	Determines whether the section applies its processing in stereo, left only, right only, mid only or side only channels.
	Amount	Sets the amount of processing for the current section. 0% means that the filter is not being applied; 100% indicates full processing.
Dynamics editor	Nodes	Each equalizer/filter section has a dedicated compressor input/output curve. This curve determines the compressor gain for a given input level, and the curve can be modified with 3 nodes. • Left-click a node to activate it. • Right-click a node to de-activate it. • Drag the mouse to zoom into an area for microscopic editing. • Click elsewhere (not on a node) in the editor to zoom out. The detected input level will be shown as a highlighted area under the compressor input/output curve.
	Comp	Sets the dynamics (compressor) functionality on or off
	Soft	Enables or disables smooth / soft curves rather than hard knees.
	Attack	Sets the attack time of the equalizer/compressor section.
	Release	Sets the release time of the equalizer/compressor section. The value is ignored when A/R (Auto Release) is enabled.
	A/R	Enables or disables the Auto Release (A/R) mode.
	Make up	Sets the make-up gain (in dB) of the dynamics editor.
	SC Input	Selects what signals are used for level detection (side chain input). The dynamics processor can detect stereo levels, but also only operate on mid, side, left, right, or the side-to-mid ratio. FLX4 has the additional option to use external input 3+4 for level detection (external sidechain).
Generic settings	DCF	Enables or disables a DC reject filter. If enabled, frequency below 5 Hz will be removed from the plugin's output.
	In gain	Sets the input gain (in dB).
	Auto phase	Enables or disables the Auto phase feature of the plugin.
	Out gain	Sets the output gain (in dB).

7.4 Setting up and using TB FlX

7.4.1 Spectrum editor

The spectrum editor works very similar to any equalizer plugin. It supports up to 6 filter sections, with individual controls to set the filter type, the filter gain, its frequency, and bandwidth/quality (Q) factor. Just activate (left mouse click) or de-activate (right mouse click) a section and drag it to the frequency/gain combination that is desired.

- You can freely re-order the nodes.
- You can draw a rectangle in the editor to zoom in.
- Click anywhere but on a node to zoom out, and/or to bring up the generic settings display.
- The highlighted area displays the filter curve of the selected section.

- The white line shows the overall equalizer curve in real time.

In the upper-left corner of the frequency editor there is a small drop-down menu for quick initialization / reset of the editor.

7.4.2 Filter types

Filter type	Purpose / description
Auto node link	The auto node link filter type automatically adjusts its filter response to construct a smooth filter characteristic through all nodes that are configured as 'auto node link'. This way, you can create many different equalizer curves by just placing 2 or more filter nodes anywhere in the editor. The constructed filter type will be shown by a highlighted area.
LSF - no res	Non-resonating low-shelf filter (LSF). The steepness of the transition is determined by the Q factor, but the filter will not resonate as analog shelving filters do.
Bell shape	Bell-shaped filter with a peak that is more flat than analog filters for a more neutral sound.
HSF – no res	Non-resonating high-shelf filter (HSF). The steepness of the transition is determined by the Q factor, but the filter will not resonate as analog shelving filters do.
Rectangle	Rectangular filter shape to boost or attenuate a very specific frequency range.
Gammatone	Gammatone asymmetric filter. This filter is often used in perceptual models to mimic the behavior of our hearing system. Now you can use its characteristic as an equalizer curve.
Gauss	Gaussian filter shape. Gaussian filters have the shortest possible group delay of all filter types.
Harmonic lin	Combination of 8 linearly-spaced peaking filters. The decay of the individual 'harmonics' can be set with the Q factor control. This filter can give some very artistically interesting effects, especially if the frequency is changed over time.
Harmonic log	Combination of 8 logarithmically-spaced peaking filters.
Analog bell	Typical analog peaking filter.
Analog LPFx	Analog low-pass filter (LPF) of order 'x'. Higher orders will give a steeper cut-off. Resonance can be set with the quality (Q) factor.
Brickwall LPF	Brick-wall low-pass filter (LPF). Removes everything above its frequency with a very steep slope. The transition between pass and stop bands can be modified with the quality (Q) factor.
Pink -3dB/oct	Applies a -3dB/oct filter above the frequency. Useful to generate pink noise from white noise.
Spectral balance	This filter is symmetric along its frequency – when a certain attenuation is applied below the frequency, the same gain is applied above its frequency. Very useful to change the spectral balance.
Analog LSF	Analog low-shelving filter (LSF). This shelving filter will show the typical analog resonance when the quality factor (Q) is increased.
Analog HSF	Analog high-shelving filter (HSF). This shelving filter will show the typical analog resonance when the quality factor (Q) is increased.
Analog HPFx	Analog high-pass filter (HPF) of order 'x'. Higher orders will give a steeper cut-off. Resonance can be set with the quality (Q) factor.
Brickwall HPF	Brick-wall high-pass filter (HPF). Removes everything below its frequency with a very steep slope. The transition between pass and stop bands can be modified with the quality (Q) factor.
Analog BPF	Analog band-pass filter (BPF) with -6dB/oct slopes.
Brickwall BPF	Brickwall band-pass filter (BPF). Will remove everything excerpt for a narrow frequency range around the frequency parameter.
Broadband gain	Applies an overall gain independent of frequency. When the compressor is activated, a frequency curve will be shown that indicates the filter response of the compressor level detector (side chain). This allows you to precisely select which frequencies the compressor should respond to.

7.4.3 Compressor editor

When one of the nodes in the frequency editor is selected, a corresponding compressor editor will be activated. This editor shows the compressor input/output curve for the selected node in the frequency editor.

Similar to the spectrum editor, dragging a rectangle with the left-mouse button will zoom in; a left-mouse click anywhere in the editor but on a node will zoom out.

Nodes can be placed anywhere in the compressor editor. The line between the nodes will indicate the compressor input/output curve. Some examples are given below.

Upward compression. In this case, compression is applied to low input levels. Low input levels are brought up in level, while high input levels are not modified, other than a static gain. This type of compression is useful when soft parts of a signal need to be louder without modifying loud parts and transients.

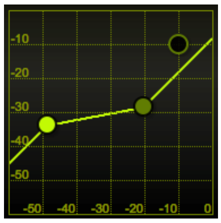

Downward compression. In this case, signals with low levels are not modified, while high input levels are decreased in level. This method is especially useful to change the character of a signal, for example to change the punchiness of a percussive sound.

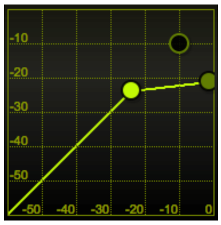

Soft-knee (downward) compression. In this downward compression configuration, the curve is smoothed instead of having a hard knee. This mode of operation is often somewhat more transparent than hard-knee compression.

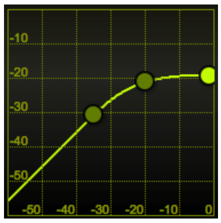

Expansion. In this configuration, high input levels are further increased in level, while low input levels are not modified. This mode allows to increase the dynamic range of transients.

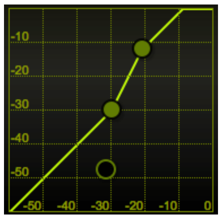

Negative ratio compression. In this configuration, a sound that decreases in input level will become louder at the output.

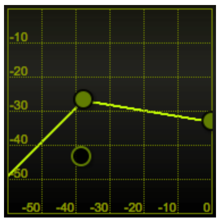

Compressor settings can be initialized efficiently by using one of the preset curves (accessible via the small drop-down menu in the upper-left corner). This menu also allows you to copy curves from one equalizer/dynamics processor section to another.

7.4.4 Auto phase option

The auto-phase option activates a novel method to modify the phase response of an equalizer. Most equalizers available today have a linear-phase or minimum-phase response, the latter having a lower delay than the first.

When the ‘auto phase’ option is activated, TB FlX activates a novel method to construct the phase response, which aims at combining the best features of linear phase and minimum phase. Depending on the input signal, and the desired frequency response, TB FlX will fully automatically modify its phase response to anything from linear phase to (close to) minimum phase, to give the best possible sound quality.

7.4.5 FlX vs FlX4 - external side chain

FlX comes as a set of two plugins:

- FlX – the default plugin with stereo in, stereo out plugin. This version of the plugin does not support external side chains.
- FlX4 – this version has 4 inputs and stereo out. The 2 additional inputs can be used as external side chains by selecting Ext 3+4 in the SC input control of the dynamics editor. Please consult the manual of your host program whether external side chains are supported and how to enable them.

8 TB Dither v3

World's first quantization and noise shaping plugin that allows the design of your own noise shaping curve – as easy as working with an EQ!

8.1 Introduction

TB Dither is a plugin designed to modify the bit depth of audio signals, for example when authoring a CD or for archival purposes, with minimum quality degradation. Such process typically involves dithering, quantization, and noise shaping. TB Dither supports industry-standard dithering noise types such as RPDF (rectangular probability density function, 1 LSB wide) and TPDF (triangular probability density function, 2 LSBs wide). A GPDF (Gaussian probability density function) is provided as well.

TB Dither's uniqueness lies in the flexibility to shape and minimize the audibility of noise inherently introduced by bit depth reduction. Instead of providing a very limited set of a few, fixed noise shaping curves, TB Dither allows you to design the spectrum of the quantization noise using familiar tools such as low-shelf, high-shelf and peaking filters, just as any equalizer! This provides an unprecedented ability to adjust quantization noise spectra according to the audio content, and envisioned reproduction system(s). If you can work with an EQ, you can work with TB Dither!

To get started, no less than 7 different noise shaping curves are provided and can be recalled from a menu, ranging from threshold-in-quiet curves, inverse dB(A) weighting, inverse ITU-R 468 curves, and several more.

8.2 Audibility of sample rate and bit depth reduction

According to double-blind tests, the only audible effect when converting high-resolution audio to a sample rate of 44.1 kHz and 16 bits is the injected (dithering and quantization) noise. TB Dither resolves this by decreasing the quantization noise level in the frequency range the human ear is most sensitive to, and thereby increasing the dynamic range for those frequencies beyond the 16-bit limit.

8.3 Dithering and information theory

When working with TB Dither, there are a couple of things one should know about quantization and the resulting quantization noise. Bit depth reduction will always introduce errors. If appropriate dithering is applied, these errors will manifest themselves as spectrally flat (or white) noise. With additional noise shaping techniques, we can change the spectrum of that quantization noise to make it less audible. A very common approach is to reduce the noise level in the frequency region the human hearing system is most sensitive to (1 kHz to 10 kHz, approximately). Such noise shaping is, however, subject to information-theoretic limitations, the most important one being that the total amount of noise cannot be reduced, only be increased. In other words, decreasing the noise level between 1 and 10 kHz will always result in an increase of noise below 1 kHz and/or above 10 kHz!

Fortunately, as end-user of TB Dither, you do not have to worry about information-theoretic principles. The way it works is as follows. Use the four noise shaping sections to construct a desired noise spectrum, in the same way you operate an equalizer. TB Dither will, in real time, automatically find the best-possible realization of that curve by shifting the curve up and down such that the overall quantization noise level is the lowest possible noise level for that specified curve. An example is given in this screenshot. The colored areas indicate the targeted spectrum, while the white curve shows the resulting noise spectrum, compared to applying no noise shaping at all. In this example, the quantization noise has been reduced by up to 20 dB for frequencies below approximately 10 kHz, at the expense of an increase of noise in the frequency region humans are not so sensitive to.

Dithering and noise shaping should always be the **last processing plugin** to render a final output signal. This means that TB Dither must be applied **post fader** (or pre-fader with a fader gain of exactly 0 dB).

Any **level normalization** process applied subsequent to dithering and noise shaping will eliminate the effect of dithering and noise shaping, and should therefore always be disabled.

8.4 Features

- Zero-latency processing
- Support of all sampling rates from 44.1 to 192 kHz
- Supports industry-standard RPDF and TPDF dithering noise
- Unrivalled flexibility to shape dithering and quantization noise
- Supports any bit depth between 8 and 24 bits
- Dedicated modes to listen to quantization noise only
- Advanced noise-shaping overload protection algorithm

8.5 User interface

GUI section	Control	Purpose
General	Resolution	Sets the number of bits that audio is quantized to. Can be any integer value between 8 and 24 bits.
	Quantizer mode	Sets the type of quantization: • Rounding (no dithering, no noise shaping) • RPDF dither (industry-standard 1-LSB wide RPDF dither; no noise shaping) • TPDF dither (industry-standard 2-LSB wide TPDF dither; no noise shaping) • GPDF dither(Gaussian-noise dither; no noise shaping) • Noise shaping (dither + user-controlled noise shaping)
	Noise shaping	Sets the amount of noise shaping between 0 and 100% relative to the curve specified in the noise shaping editor
	Output mode	Sets the output mode: • Normal (output is input with dither and noise shaping applied); • Muted input (output with silent input signal); • Output-input (difference between input and output to audition the effect of quantization, dithering, and noise shaping).
Noise shaping editor	Section	Identify the noise shaping equalizer section (1 to 4). The noise shaping editor is only active if the quantizer mode is set to 'Noise shaping'.
	Type	Set the noise shaping equalizer section type: • LSF (low shelving filter) • PEQ (parametric equalizer) • HSF (high shelving filter)
	Frequency	Sets the (center) frequency of the active noise shaping equalizer section.
	Gain	Sets the gain of the active noise shaping equalizer section.

Q factor	Sets the Q factor (or inverse of the bandwidth) of the active noise shaping equalizer section.
Nodes	Nodes can be dragged in the noise spectrum graph to modify the spectrum of the quantization noise. The white line will always indicate the overall noise shaping curve subject to information theoretic limitations. • Left-click a node to active a noise-shaping equalizer section; • Right-click a node to de-active the corresponding equalizer section.
Drop-down menu	Several noise-shaping presets are provided via the drop-down menu indicated in the upper-left corner of the noise shaping editor.

8.6 Typical workflow for dithering and noise shaping

8.6.1 Determine the desired bit depth

For CDs the bit depth is 16 bits, while for DVD audio, bit depths of 24 bits are typically used. The desired bit depth depends on the application at hand.

8.6.2 Insert TB Dither as the very last plugin in the processing chain

Dithering and noise shaping must always be the very last step in the effects chain, preferably even post master fader. Dithering and noise shaping processes depend on the exact quantization levels that are used during the final export. Any level adjustment, filter, or other effect being applied in-between dithering and export will completely eliminate the positive effects of dithering and noise shaping. This also implies that peak limiting must be applied **prior to dithering**, and that any level normalization applied by the host **must be disabled**.

8.6.3 Choose the dithering and noise shaping settings

TPDF and GPDF dithering

For 24-bit audio, the dynamic range provided by 24 bits is in many cases sufficient to simply use spectrally white noise for dithering without noise shaping. That noise can be produced by the TPDF (triangular probability distribution function) noise, or GPDF (Gaussian probability distribution function) dithering setting. These two dither noise types are specifically crafted to have the following two properties:

- The quantization noise due to bit-depth reduction is not correlated with the audio signal; and
- The 2nd moment (power) of the quantization noise due to bit-depth reduction is not correlated with the audio signal.

In practice, this means that the effect of quantization and dithering is a steady, continuous, low-level white noise signal that is independent of the input audio signal.

RPDF dithering

RPDF (rectangular probability distribution function) dithering noise has a lower noise level than TPDF and GPDF noises. However, the drawback of this particular dithering noise type is that, although the quantization noise is not correlated with the audio signal, its 2nd moment still is. In other words, the quantization/dithering noise level will fluctuate with the audio signal itself, which can be an undesirable property.

Noise shaping

The goal of dithering and noise shaping is to modify the spectral properties of the quantization noise introduced by bit-depth reduction in such a way that it becomes less audible. Although due to information-theoretic constraints the total amount of noise cannot be reduced, one can exchange a lower noise level in one frequency region for higher noise levels in other regions. This trade-off is provided by the noise shaping editor, which basically works like an equalizer operating on the quantization noise only. A reduction of quantization noise in a specific frequency range will automatically result in an increase for neighboring frequencies, and vice versa. Carefully crafted noise shaping can result in 16-bit audio that has the same noise audibility as 18 bit audio when using TPDF (white) noise, and can give quite good quality for bit depths as low as 8 bits/sample.

The most common technique is to increase the quantization noise level above 16 kHz to allow a lower level in the 4-8 kHz range. However you are encouraged to experiment with different noise shaping characteristics.

Noise shaping is more effective at higher sampling rates. 44.1 kHz is the **minimum sampling rate** for noise shaping to work properly, but 48 kHz or 96 kHz will make the process much more effective.

There is no benefit from using multiple dithering/noise shaping algorithms on the same audio signal; in fact, it is better to avoid this from happening. If TB Dither is used, make sure that all other processes do **not** apply dithering and/or noise shaping (either in a plugin such as a peak limiter, or during export by the plugin host).

You can audition the effect of noise shaping very easily by temporarily making the following adjustments:

- Set the bit depth to a very low number, such as 8 bits/sample;
- Set the output mode to 'Output-input' so you can listen to the effect of quantization without the input audio.

8.6.4 Export

If all noise shaping parameters are tuned correctly, export the audio signal. The export bit depth of the host must be set to the exact same value as used in TB Dither (e.g., export as 16-bit PCM if TB Dither was set to 16 bits).

9 TB BusCompressor v3

High-quality, transparent dynamics processor with adjustable knee and auto-release functionality suitable for single tracks as well as complex mixes.

9.1 Introduction

TB BusCompressor is a very transparent, musical, all-round dynamics processor designed to be able to handle everything from single tracks to complex, full mixes. Even with ultra-short attack and release settings, harmonic distortion is extremely low (often better than -150 dB re FS*), and CPU load is typically below 0.5% (depending on hardware).

TB BusCompressor has the unique feature to set the compressor hold time in cycles rather than in seconds. This dramatically reduces intermodulation distortion even with ultra-fast attack and/or release settings. Expression of the hold time in cycles creates longer hold times at low frequencies (at which one cycle has a long duration) while still having a very fast response at high frequencies.

Another unique feature is to adjust the compressor sensitivity to noisy (as opposed to) harmonic signal components. TB BusCompressor's advanced signal analysis toolset includes the separation of tonal/harmonic and noisy/percussive signals. Therefore, you can control the relative amount of these signal types that the compressor responds to. For example, in a certain situation you might want to compress harmonic instruments present in a mix more than the (noisy) snare drums. The noise control of TB BusCompressor changes the amount of noisy components that the compressor is responding to. A second application for this feature is the compression of vocals. By changing the sensitivity to noisy components, fricatives and sibilants will (relatively) be more compressed, reducing the need for additional de-essing.

9.2 User interface

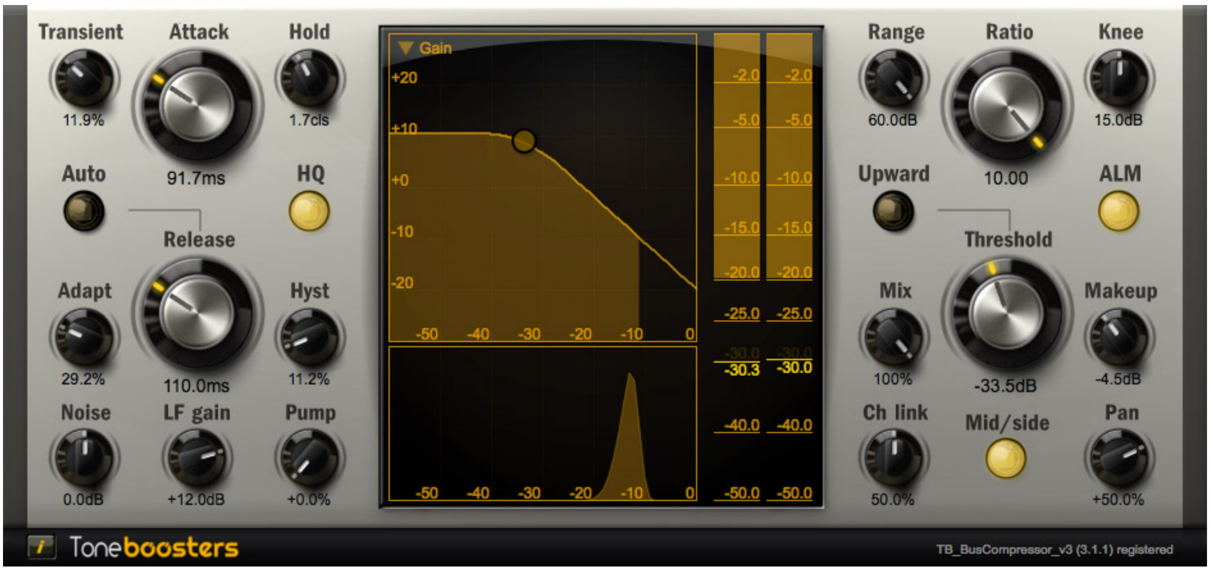

GUI section	Control	Purpose
Attack	Attack (time)	Sets the time to respond to increases in input level. A short attack time will not let strong transients through; a longer attack time will cause a slower decrease in gain when the input level increases.
	Transient	Sets the amount of additional compression applied to transients. A higher value will result in stronger compression of transients. Only transients with a level above the threshold will be affected.

	Hold	Sets the hold time in cycles. When set to +1, at most 1 cycle (or less) is used to hold the gain. A typical setting of 2 cycles should sound great on many sources and prevents intermodulation distortion.
	High Quality (HQ)	Enables the high-quality mode. Engage the High Quality (HQ) mode to increase the oversampling factor of TB BusCompressor for sub-sample accuracy. Rest assured that even with the HQ mode disabled, oversampling will still occur in TB BusCompressor, but enabling the HQ mode will shift the oversampling parameters to the next gear for even more accurate timing.
Release	Release (time)	Sets the minimum time to respond to decreases in input level.
	Adaptive release	Adaptive release increases the release time if the signal is not quickly dropping in level, ensuring that the gain riding behavior of the compressor more closely matches the signal envelope.
	Hysteresis	Hysteresis makes the release time history dependent. If signals in the past were of relatively low level, and the compressor is merely reacting to a short transient, its release will be short to quickly recover from the short transient. If, on the other hand, the signal was consistently loud previously, the compressor will react with a slower release.
	Auto	Clicking the “Auto” button will engage the automatic (content-dependent) release mode. The release time, hysteresis and adaptive release controls will become inactive if the auto release mode is enabled.
	Noise	Sets the relative sensitivity to noisy signal components in the input (as opposed to harmonic components). A higher value will cause the compressor to react relatively stronger to sibilants, percussion instruments, noise-like signals, snare drums, and alike.
	LF Gain	Low-Frequency gain sets the (relative) sensitivity to low frequency content. Turning this knob will boost or reduce the low frequencies in the side chain only (adjusting the compressor sensitivity).
	Pump	The “pump” control changes the behavior of the compressor. With a setting of 0, the compressor will typically work in a very transparent manner with minimum amount of pumping. For electronic music, however, pumping might be a desirable effect. Increasing the value of this control will result in a stronger pumping behavior, especially if the LF Gain is set to positive values.
Ratio	Ratio	Sets the amount of compression. A ratio of 4 indicates that 4dB above threshold will be reduced to 1dB above threshold.
	Range	Sets the maximum gain / attenuation that can be applied. If the range is set to 20 dB for example, the compressor gain (or attenuation) is limited to -20 to +20 dB.
	Knee	Sets the soft knee for a smoother compression behavior near the threshold point : A soft knee applies the ratio exponentially as the signal approaches the threshold point. With the right setting, it gives a more transparent sound. For instance, using a 6dB knee and a -12dB threshold, subtle compression will begin at -18dB (6dB below the threshold) and will gradually become stronger until the maximum compression is obtained at -6dB.
Threshold	Threshold	Sets the input level below or above which the compressor becomes active.
	ALM	Assisted Level Makeup (ALM) provides support in levelling compressor output with input by adjusting the compressor input-output curve depending on the compressor settings. In many cases, these automatic adjustments should reduce the need to use the manual make-up level. ALM has three levels: • Off: ALM is disabled. • Green: ALM is set to ‘normal’ which aims at keeping the loudness constant when threshold, ratio, and dry/wet controls are changed; • Yellow: ALM is set to ‘boost’ which will usually give a boost in loudness.
	Upward	Engages the upward compression mode.

		When this setting is enabled, the quieter passages (below the specified threshold) will be boosted while leaving louder passages unchanged (upward compression). If disabled, louder passages (above the threshold) will be reduced in level, while leaving quieter passages unchanged (downward compression).
	Mix	The dry/wet mix control allows New York style / parallel compression inside the compressor itself. A value of 75% indicates that the output consists of 75% of compressed signal, and a remaining 25% of (unmodified) input signal. Uniquely to TB compressors, the effective input-output curve is visualized accordingly.
	Makeup	Sets the makeup level of the compressor output.
	M/S	Mid/side mode. TB BusCompressor can operate in left/right or mid/side mode. In left/right mode, the left and right channels are compressed. In mid/side mode, on the other hand, the mid (left+right) and side (left-right) channels are compressed.
	Ch link	Determines the amount of linking between left/right compression (in left/right mode) or mid/side (in mid/side mode). A value of 100% results in full coupling of the compression in both channels; a value of 0% gives fully independent compression operation in both channels. The setting of this control will be stored independently for M/S mode disabled and enabled.
	Pan	Adjusts the side-chain input level balance. In left/right mode, this knob works like a left/right pan knob on the side chain, determining the left/right balance adjustment going into the side chain level detector. The same applies for mid/side in the mide/side mode. The pan control will remember its setting independently for mid/side and left/right modes.
Display	Compressor curve	The display at the center of the GUI gives a graphical representation of the compression curve and the current input level. The handle attached to the curve can be used to adjust a few basic compressor settings: • Drag handle to change the threshold; • Use the mousewheel to change the ratio; • right-button click on the handle to engage bypass; Apply a left-button click on the handle to disable bypass.
	Display selector	Click on the small downward triangle in the upper-left corner of the display to change the display mode. • The "I/O" mode presents compressor *output level* as a function of input level; • The "Gain" mode presents the compressor *gain* as a function of input level.
	Histogram	The lower half of the display shows a real-time histogram of the input levels. This may provide guidance for threshold adjustment. The height of the curve represents how often a certain input level was observed in the last 30 seconds (approximately).

10 TB Sibalance v3

10.1 Introduction

De-essers can be an evil necessity. Vocal recordings may be too sibilant requiring de-essing (or excess sibilance removal), but most de-essers come with very clearly audible drawbacks as well. After de-essing, vocals may sound muffled, the ‘s’ may sound more like an ‘f’, or even worse, the operation of a de-esser manifests itself as a clearly-audible time-varying filter. TB Sibalance provides very powerful tools to reduce excess sibilance in a minimally invasive way. In contrast to conventional de-essers, TB Sibalance uses so-called ‘matched filter’ technology to only process those frequencies that are causing excess sibilance, while leaving all other frequency components untouched. The result of TB Sibalance will therefore sound cleaner and more transparent than that obtained with other de-essers.

10.2 Features

10.2.1 De-essing like a compressor

You may see a very familiar input/output curve in the screenshot that looks like a compressor. In this case, the input/output curve does not relate to level, but to (excess) sibilance. Sibilance is a property of audio that is largely independent of level; signals sound sibilant if there is a relatively large amount of signal energy present in the sibilant range (typically 4-11 kHz) compared to the overall level. Nevertheless, TB Sibalance allows control of sibilance by means of a threshold, a ratio, a soft knee, and a range parameter; much like a compressor. Of course, a dry/wet control is included as well.

10.2.2 Algorithm fusion

TB Sibalance has three algorithms: (1) a broad-band de-esser, (2) a band-limited de-esser, and (3) a matched-filter de-esser. The latter will create a filter dynamically that only reduces sibilant frequencies while leaving everything else untouched. In contrast to many other de-essers, these algorithms can be fused on a continuous scale. Do you want 60% of de-essing using a band-limited de-esser, and the remaining 40% using matched filter technology? Just set the algorithm slider to 1.6 and it is all set. Moving the algorithm slider from 0, to 1, to a value of two seamlessly fuses the broad-band de-esser, the single-band de-esser, and the matched-filter de-esser.

10.2.3 Tonal component sensitivity

In many practical situations, vocal sibilance consists of noise-like signals that one would like to suppress. Tonal or voiced signals, on the other hand, are often better left untouched. Conventional de-essers cannot discriminate between noise-like and harmonic signals – they simply measure energies. With TB Sibalance, the relative contribution of tonal components to the measured sibilance can be adjusted so that the de-esser works much more accurately.

10.2.4 Mid/side processing and high-quality modes

If a bus signal or full track needs de-essing, great results can be obtained by de-essing the mid channel only (this is typically where the vocals are), while leaving the side signals untouched. Mid/side mode of operation is available on TB Sibalance, as well as a control to engage a high-quality mode.

10.2.5 Signal level dependencies

It can be very desirable to reduce excess sibilance for relatively loud parts of a track, while leaving less loud elements untouched. The ‘level threshold’ control of TB Sibalance influences how sibilant low-level signals are. Basically, if the input signals approach the threshold set for level, the measured sibilance will gradually be reduced, and hence the amount of de-essing will become more subtle or even absent.

10.2.6 Processing of full mixes

Ideally, a sibilance tool is sufficiently flexible to also process full mixes, for example to catch excess esses in a mix, or simply to reduce the mix’s harshness. This is why you’ll see controls to change the ‘voiced’ frequency range analysis; for mix processing, these can be set to cover almost the complete audible frequency range.

10.3 User interface

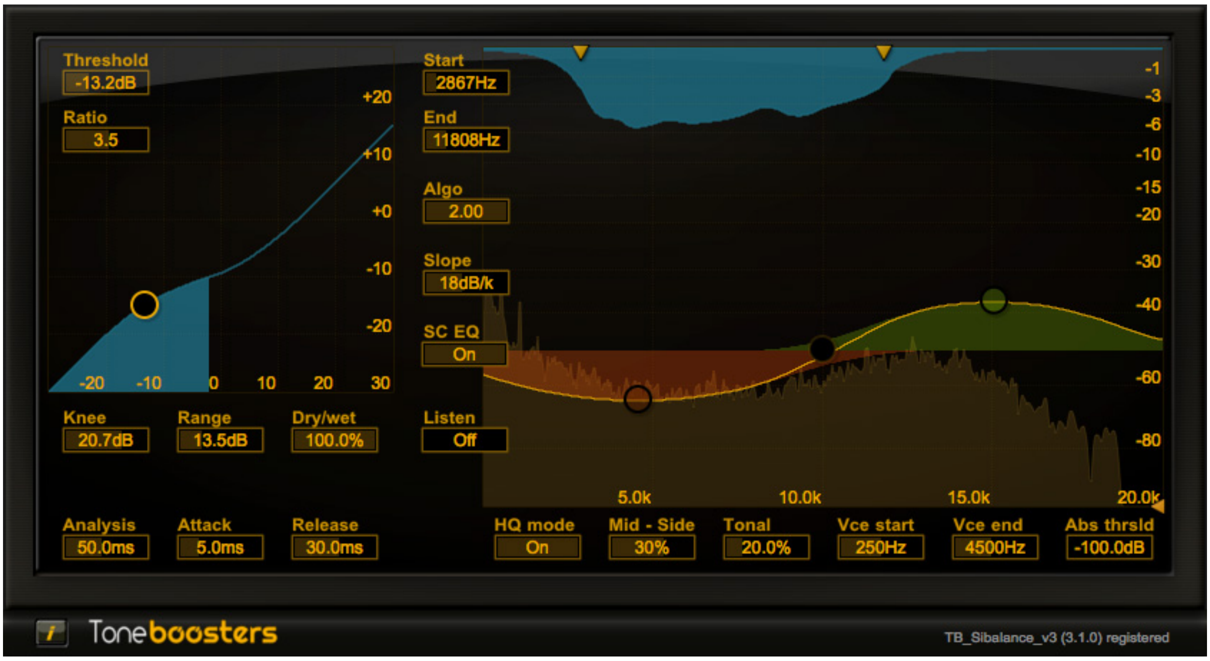

GUI section	Control	Purpose
Sibilance input/output function	Input/output function	The sibilance input/output function sets the output sibilance as a function of the input sibilance. • The handle in the graph sets the threshold sibilance at which the de-esser becomes active. • Use the mouse wheel to set the ratio – a higher ratio will result in a larger reduction of sibilance. • The blue fills in the input/output function graph will indicate the currently detected sibilance level.
	Threshold	Sets the sibilance level at which the de-esser should become active in reducing sibilance.
	Ratio	Sets the ratio of the sibilance input-output function. A ratio of two indicates that a sibilance level that has a value of X above the threshold, will be reduced to X/2 above the threshold. In other words, a larger ratio will result in a stronger reduction of sibilance.
	Knee	Sets the size (in dB) of the transition around the threshold. Increasing the knee value will create a softer knee around the threshold value.
	Range	Sets the maximum reduction of sibilance in dB. The de-esser will never apply a reduction larger than this value.
	Dry/wet	Sets the amount of dry/wet mixing. A value of 100% indicates that only the wet signal (processed signal) is sent to the plugin’s output.
	Analysis	Sets the analysis time for determining excess sibilance. A larger value will cause the de-esser to have a smoother (but slower) response.
	Attack	Sets the time (in milliseconds) to react to increases in sibilance level.
	Release	Sets the time (in milliseconds) to react to decreases in sibilance level.
Spectrum analyzer	Spectrum analyzer	Shows the real-time spectrum from which excess sibilance is detected in yellow.

	Gain visualization	The de-esser attenuation as a function of frequency is visualized in blue.
	Start	Sets the de-essing start (minimum) frequency in Hz.
	End	Sets the de-essing end (maximum) frequency in Hz.
	Algo	Sets the algorithm used for calculating the de-esser attenuation function. • A value of 0 results in a broad-band de-esser, e.g., *all frequencies* will be attenuated by the same amount if excess sibilance is present. • A value of 1 results in a single-band de-esser, e.g., all frequencies between the start and end frequency will be attenuated by the same aount of excess sibilance is present. • A value of 2 results in a matched-filter de-esser, which targets specific frequencies only in-between the start and end frequency range that are causing excess sibilance. • Any value in-between 0 and 1 will give a response in-between a broad-band and single-band de-esser. • Any value in-between 1 and 2 will give a response in-between a single-band and matched-filter de-esser.
	Slope	Sets the filter slopes in dB per kHz. Lower values will give a smoother frequency response in the gain function; higher values will allow more surgical processing in the frequency domain.
	SC EQ	Enables or disables a side-chain equalizer. If enabled, three equalizer sections will appear that allow modification of the results shown by the spectrum analyzer. • Drag the handles to change their frequency and gain values. • Left click or right-click the handles to activate /de-active an equalizer section. • Use the mouse wheel to modify the Q factor / bandwidth of the equalizer section.
	Listen	If enabled, the difference between original and processed signal will be produced at the output. If no de-essing takes place, the output will therefore be silent.
	HQ mode	Enables the HQ (high-quality) mode. This mode will run the algorithm at a higher sampling rate internally.
	Mid/side	Sets the amount of de-essing for mid and side. • A value of -100% will apply de-essing on mid only • A value of 0% will apply de-essing in stereo • A value of +100% will apply de-essing on side only.
	Tonal	Sets the contribution of tonal components in detecting excess sibilance. • A value of 0% will set the de-esser sensitivity to tonal signals to its minimum. Excess sibilance will mainly be detected for noise-like signals. • A value of 100% will set the de-esser sensitivity to tonal signals to its maximum. Excess sibilance will be detected for both noise-like as well as tonal / harmonic signals.
	Vce start	Sets the start frequency for voiced signal detection. The level within the voiced signal range determines if signals in the sibilant range are excess sibilance or not.
	Vce end	Sets the end frequency for voiced signal detection.
	Abs thrsld	Sets the absolute threshold for excess sibilance. If the spectrum analyzer indicates levels below this value, the resulting signal will gradually not be classified as sibilant.

10.4 Understanding excess sibilance

10.4.1 Voiced and sibilance frequency ranges

The goal of TB Sibalance is to reduce or remove excess sibilance, or said differently, sibilant sounds such as 'ess' that are too loud are to be reduced in level. It is important to realize that the phrase 'too loud', or excess sibilance, is defined within its context. This context dependency is explained schematically below.

Let us start with showing a spectrum of an audio signal. In the figure below you will see the power spectrum level of a sound as a function of frequency. We can identify two frequency ranges that are not necessarily mutually exclusive (they are allowed to overlap in frequency):

- A voiced frequency range, typically around 200 – 4000 Hz, which is the frequency range in which voiced parts of speech (such as 'a', 'e', 'i', and alike) are predominantly present, and
- A sibilance frequency range, typically around 5000-11000 Hz, which is the frequency range in which sibilant sounds (such as 's', 't', and alike), and excess sibilance often occurs.

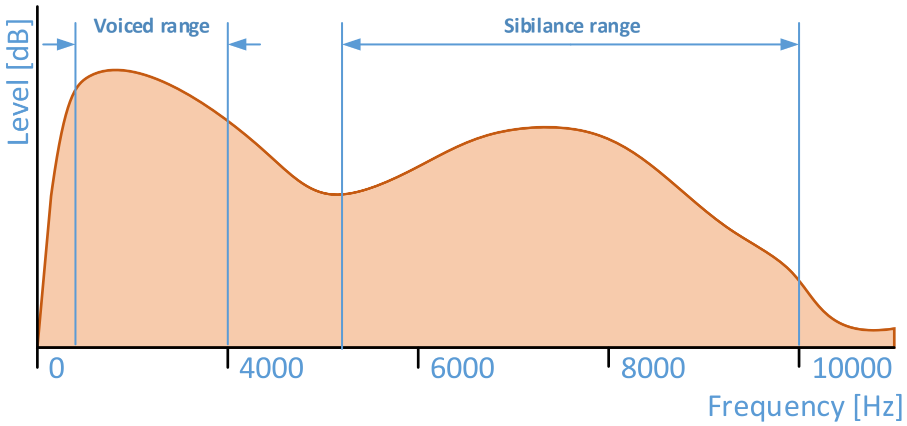

10.4.2 Sibilance level

Let us consider an example in which excess sibilance occurs. The figure below one can clearly observe that the power spectrum level within the sibilance range is much higher than the (average) power spectrum level in the voiced range. Schematically, we could therefore consider the power in the sibilance range above the voiced level as *excess* sibilance, as it clearly stands out in the context of the overall power spectrum, and with respect to the voiced level. Said differently, the sibilance level is *positive*. If, on the other hand, the power spectrum level in the sibilance range would be below the voiced level, the sibilance level is *negative*.

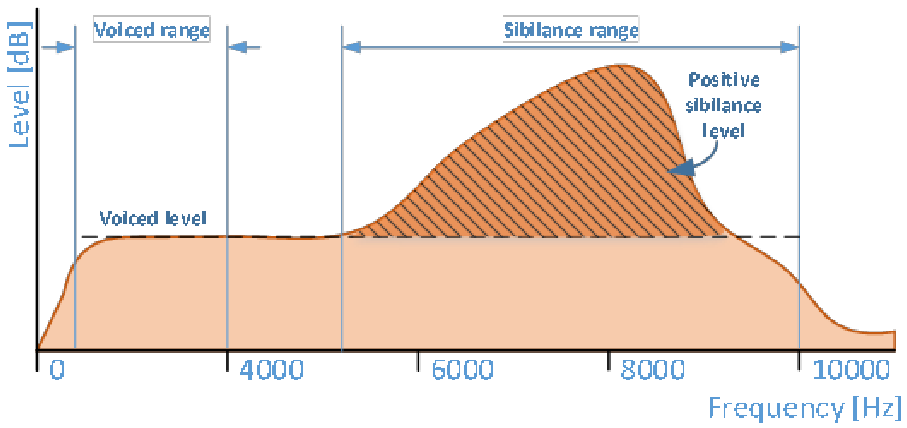

10.4.3 Absolute threshold

As we have seen in the previous section, a positive sibilance level manifests itself as power spectrum levels within the sibilance frequency range that stick out in their context. In practice, however, this method of determining the sibilance level is not sufficient. For example, it can happen that very soft sounds (such as background signals, or whispered voices) show a power spectrum in which the sibilance frequency region sticks out, while one does not want to process it because the signals are very soft in level, and are therefore not being perceived as having excess sibilance. Such level dependencies can be accomplished by means of the 'absolute threshold' parameter as shown below.

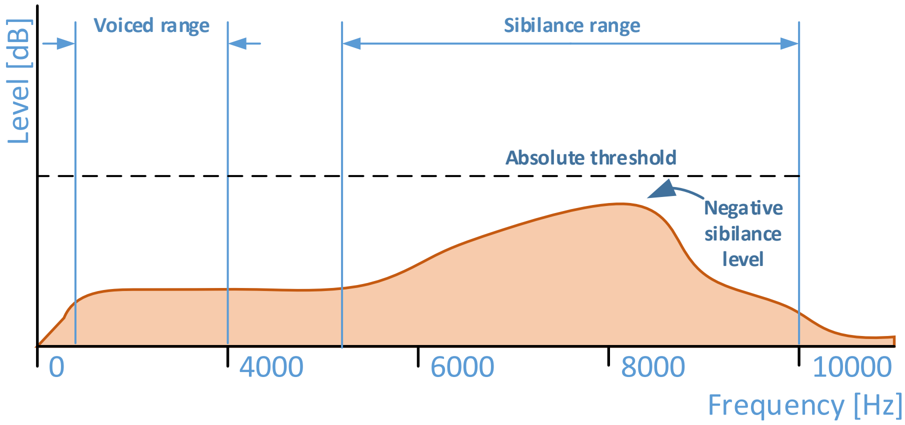

In this example, the power spectrum level in the sibilance range sticks out with respect to the voiced level, but is below the absolute threshold level set by the user. As a result, the sibilance level is negative.

In these exemplary figures, the absolute threshold is visualized as a hard decision, while in TB Sibalance, the absolute threshold function results in a *gradual decrease* in detected excess sibilance when the power spectrum is near the absolute threshold parameter setting.

10.4.4 Tonal and noise sensitivity

Excess sibilance is often very noise like, or said differently, the power spectrum does usually not show strong harmonics. Conventional de-essers can unfortunately not discriminate between harmonic and noise-like power within the sibilance frequency range, and therefore have a tendency to respond to tonal signals as well (e.g. by determining tonal signals as being excess sibilant). Such dependency may not always be desirable, and therefore TB Sibalance includes sophisticated algorithms to separate harmonic and noise-like signals.

- Noise-like power within the sibilance frequency range is always contributing to the estimation of the sibilance level;
- Harmonic or tonal signals within the sibilance frequency range have a *user-configurable* contribution to the sibilance level between 0 and 100%. If set to 0%, harmonic or tonal signals are suppressed in the detection algorithm. If set to 100%, on the other hand, the sibilance level works in the same way as conventional de-essers, without specific sensitivities to noise or tonal signals.

10.4.5 Sibilance level summary

In summary, *the sibilance level* is determined as power spectrum levels within the sibilance frequency range relative to:

- the (average) level of noise-like components in the voiced frequency range, *and*
- the absolute threshold level set by the user.

It should be noted that the examples and description in these sections are very schematic and exemplary of nature to outline the concept of sibilance; the actual algorithms in TB Sibalance are much more sophisticated than visualized here.

10.5 Reducing sibilance

10.5.1 Sibilance input/output graph

The sibilance input/output graph provides a wide range of controls to modify sibilance levels. The input/output graph shows some similarities with input/output graphs shown on the ToneBoosters compressors. However, with TB Sibalance, the input/output graph shows the input sibilance level along the horizontal axis, and the desired (or output) sibilance level along the vertical axis. The units are in Decibels.

- The reduction in sibilance can be thought of as the difference between a (dashed) line that connects equal input and output sibilance, and the actual input/output curve shown by the solid line. That amount is set by the ratio parameter. A higher ratio will result in a stronger reduction of sibilance.
- The threshold determines the input sibilance level at which reduction of sibilance starts to take effect. In other words, the threshold value allows you to determine what sibilance level is *excess* sibilance (above the threshold), and what is not considered excess (below the threshold).
- The currently detected sibilance level is shown by the filled polygon.
- A knee parameter determines the range of the transition from no sibilance reduction (below the threshold) to de-essing (above the threshold). A larger knee value results in a softer knee around the threshold.

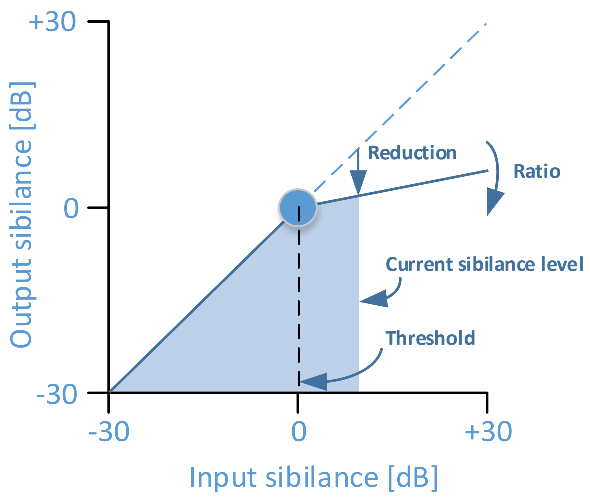

10.5.2 Set a maximum reduction in sibilance

The range parameter sets the maximum change (reduction) in sibilance that is allowed. For example, if the range parameter is set to 10 dB, the maximum difference between input and output sibilance will be exactly 10 dB.

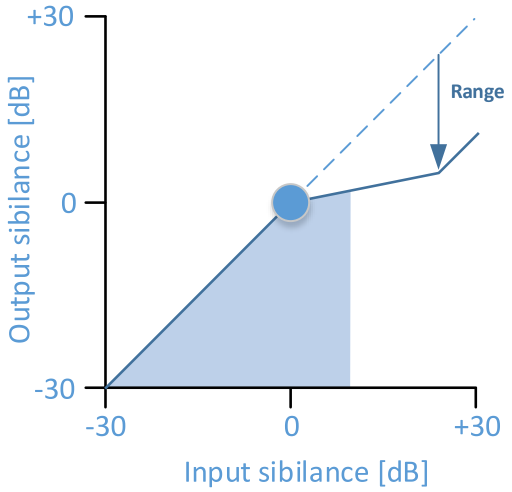

10.6 Algorithm tuning

10.6.1 Broadband, single band, or matched filter

Now that we have defined sibilance levels, and have determined the amount of sibilance reduction we would like to apply through the input/output curve, we are ready to set the method or algorithm for applying this reduction. TB Sibalance supports 3 algorithms which can be blended seamlessly:

- Algorithm 0: broad-band attenuation. With this algorithm, the signal is attenuated and all frequencies are treated equally.
- Algorithm 1: single-band attenuation. With this algorithm, the frequencies within the sibilance range are attenuated by the same amount, while frequencies outside the sibilance range are not attenuated.
- Algorithm 2: matched filter. With this algorithm, only specific frequencies within the sibilance range will be attenuated, namely those that were responsible for the (high) sibilance level. Usually this algorithm gives the most transparent results.

The difference in attenuation (or negative gain) for the single-band and matched-filter algorithms is shown below. The single-band algorithm attenuates the full sibilance range, alike most conventional de-essers. The matched-filter algorithm, on the other hand, applies a more surgical cut of frequencies that are most offensive in terms of sibilance level, while leaving the remaining signals untouched.

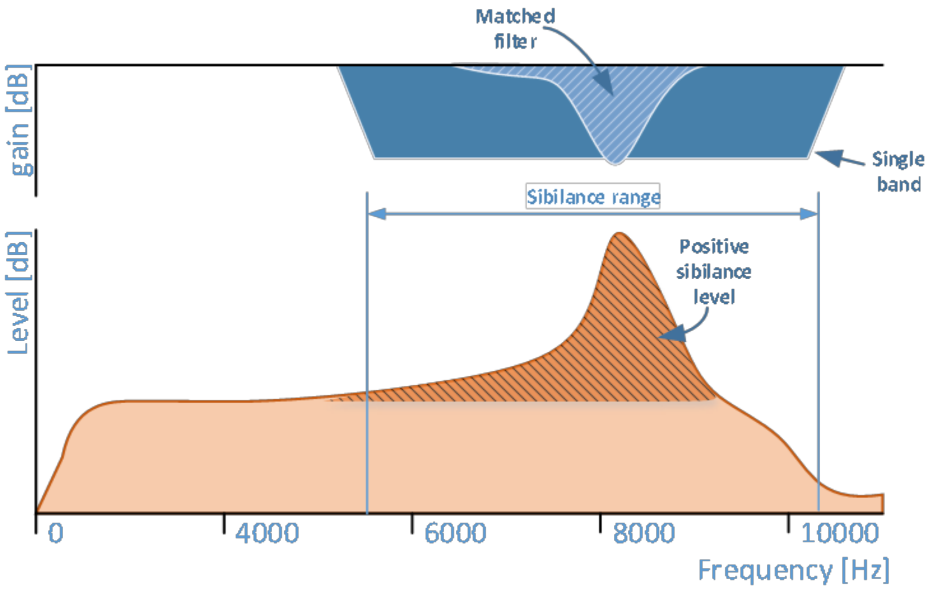

The algorithm selection control can be set to intermediate values as well. For example, a value of 1.5 will give an attenuation behavior that is in-between a single-band and matched-filter algorithm.

10.6.2 Filter slope

The steepness of the sibilance attenuation filter can be changed with the 'slope' parameter.

- For broad-band de-essing (algorithm 0), the slope parameter has no effect.
- When using the single-band algorithm (algorithm 1), the slope parameter determines the steepness of the filter around the sibilance range. Larger values will cause steeper filter responses.
- When using the matched-filter algorithm (algorithm 2), the slope parameter determines the specificity of the filter – larger values will cause more surgical cuts in very specific frequency ranges, while smaller values will result in a more global reduction of sibilance.

10.6.3 Mid, stereo, side processing

Excess sibilance may exist in certain regions of the spatial image only. One example is a complex stereo track in which sibilant vocals sit in the middle of that mix. In such cases, the mid-side parameter can help to mainly process the vocals in the complex mix, while leaving other elements largely untouched.

- Setting the mid-side parameter to 0% will apply de-essing to left, right, mid, side equally.
- Setting the mid-side parameter to mid (-100%) will apply de-essing to mid only.
- Setting the mid-side parameter to side (+100%) will apply de-essing to side only.
- Values in between will apply partial de-essing to mid, stereo or side depending on the exact value.

10.6.4 Attack and release

For certain languages, the onset of sibilant sounds is very important for language intelligibility. If a pass-through of such onsets is desirable, and de-essing should only start a short time after such onset, increase the attack parameter. This will cause the de-essing to kick in later and allow pass-through of onsets of sibilant sounds.

Similarly, the release parameter determines how fast the de-esser recovers from attenuating sibilant sounds. A longer value tends to give a smoother behavior, but a too slow value may cause the de-esser to recover too slowly before a non-sibilant syllable starts. For very fast talkers, a shorter release time may therefore be beneficial.

10.6.5 Side-chain equalizer (SC EQ)

If more precise control is required to determine what frequencies TB Sibalance is responding to, one can use the side-chain equalizer (SC EQ) to shape the spectrum before it is analyzed for sibilance levels.

The side-chain equalizer has no effect on the signal at the output of the plugin; it cannot be used as a signal equalizer. It only shapes the signal used to analyze sibilance levels.

Enable the side-chain equalizer by clicking on the SC EQ button. Three different equalizer section handles will appear which, combined, determine an equalizer curve that is applied on the signal before analyzing sibilance levels. Drag the handles around to change the sensitivity to detect sibilance at certain frequencies.

- Like with all TB plugins using handles, left or right-click on the handle enables or disables the section.
- Use the mouse wheel to change the Q (or bandwidth) of each section.
- The integrated spectrum analyzer will always show the modified spectrum (e.g., which includes the side-chain equalizer curve).

10.7 Excess sibilance in signals other than vocals

Signals other than vocals may benefit from processing with TB Sibalance as well. For example, full mixes, drum sounds, or an instrument may exhibit some harshness caused by excess energy in the 4-8 kHz range. Some aspects that might be good to remember when using TB Sibalance on signals other than vocals:

- As explained in the previous sections, the sibilance level does not only depend on spectral power in the sibilance range, but also on its context, determined by the voiced level. The frequency range for voiced level analysis can be modified. For full-mix processing, for example, it can be beneficial to extend the voiced range to a much wider set of frequencies, for example 100 Hz – 10 kHz. In this way, the full context is taken into account to determine sibilance context, not just the 300-3000 Hz range.
- The tonal parameter may help in restricting the level of tonal components. Setting it to higher values will cause an increasing sensitivity to harmonic signals such as instruments present in a mix. On the other hand, a low value is beneficial to process rhythm instruments while leaving harmonic instruments untouched.
- Mixing dry and wet by setting the dry/wet parameter to a value less than 100% can help in creating a more neutral or transparent behavior.
- Additionally, the range parameter may be of use to set a limit of how much attenuation is applied.
- Use the side-chain equalizer to modify the sensitivity in specific frequency ranges.
- Use the ‘listen’ option to evaluate what signals are being removed or processed with the current settings.

11 TB VoicePitcher v3

11.1 Introduction

TB VoicePitcher is a plugin to change the overall pitch of dialog and vocals, including singing voices. Because its algorithms are specifically designed for the human voice, the results with TB VoicePitcher will typically sound better, cleaner, and more natural than with other general-purpose pitch shifting algorithms*. Pitch shifting can be performed in real time, with a small latency that is compensated for by the DAW host (if supported). Furthermore, spectral (formant) corrections to improve the timbre of pitched vocals can be applied as well.

**Because TB VoicePitcher is specifically designed to process the human voice, results may vary for other content such as polyphonic music or rhythm tracks.*

11.2 User interface

GUI section	Control	Purpose
Pitch	-	Determines the amount of pitch change, in semitones. The range is from -12 to +12 semitones, covering a range of 2 octaves.
Formant correction	Amount (%)	Sets the amount of formant correction. Setting this control to 0% disables formant correction. Increasing its value will change the spectral envelope to better match that of the input, giving a more natural timbre of the pitched output, especially when pitching up speech or singing voice. Adjust this control to your liking to optimize the timbre of the output of TB VoicePitcher.
	Bandwidth (Hz)	This control sets the bandwidth for formant correction. A small value will result in a very accurate, detailed formant correction, while a larger value will result in a more global, less detailed formant correction. Adjust this control to your liking to optimize the timbre of the output of TB VoicePitcher.
	Wet mix	Sets the amount of dry and wet in the output. A value of 0% will only produce unpitched results; a value of 100% will only produce pitched results.
	Threshold	Sets the threshold input level at which the plugin will automatically enable/disable itself. If the signal level drops below the value set by this control, the plugin will switch itself in a hibernation mode to reduce CPU load. Whether the plugin is active or inactive is visualized by the 'Active' LED.
	Active	Indicates whether the plugin is active or inactive. The latter will occur if (1) the signal input level is below the level set by the threshold control, or (2) the dry/wet control is set to 0%.
	Stereo	Sets the channel mode of the plugin to stereo or mono. For monophonic inputs (even when supplied on a stereo track in your DAW), running the plugin in mono mode will typically reduce CPU load.

11.3 Pitch and formants

11.3.1 Introduction

When describing the pitch and timbre of vocals, the terms pitch and formants are often used. Pitch is the fundamental frequency of vibration of the vocal folds. They vibrate quasi-periodically only for voiced phonemes. The rate of vibration gives rise to a perceived pitch, corresponding to a specific note on a keyboard for example.

The sound produced by the vibrating vocal folds is changed due to the frequency shaping by the vocal tract, which is everything from nasal tract, tongue, teeth, lips, mouth. The particular configuration for every phoneme creates resonances at specific frequencies that we call 'formants'. Such formants allow us to distinguish between different (voiced) phonemes that can all have the same pitch, or alternatively, recognize phonemes independently from their pitch.

The human voice organs allow us to change pitch from formant frequencies rather independently; the same word can be said or sung with a different pitch. With such a pitch change, the formants will typically remain (almost) constant. This is in contrast to what happens when we change the playback speed of a voice recording; in that case, both the pitch and the formants will shift in frequency by the exact same amount. The result of such playback speed change is therefore often very unnatural, especially when the playback speed is faster than intended. Additionally, the duration will obviously be different.

11.3.2 Pitch change without affecting the duration

In contrast to what happens when changing the playback speed of a digital recording, changing the pitch with TB VoicePitcher will never change the tempo or duration of the content. Even with extreme settings of pitch shifting by an octave, the output will stay in-sync with the input.

11.3.3 Pitch change without affecting formants

If formant correction is disabled (by setting its control to 0%), both the pitch and the formants will be modified by the exact same amount, determined by the pitch control. Locking pitch and formant shifts by the same amount can give somewhat funny or unnatural results if voices are pitched up by a large amount. Therefore, a formant correction algorithm is included, which applies a correction to the formant frequencies present in the vocal utterance. This algorithm will analyze formant resonances present in the input, and repair (or correct) them in the output if these resonances have been modified. This is mostly helpful when pitching vocals up.