

TOKYO DAWN LABS

Feedback Compressor II – Manual

Product version: 2.0.0

Manual revision 1.4, last update: 02/01/2013



CONTENTS

Concept	4
Feed-back Compression	5
Precision, Aliasing and the solution	6
Control Scheme	7
Controls and displays	8
Threshold, Peak Crest, Soft Knee and Ratio	8
Makeup Gain	8
Dry Mix	9
OUTPUT GAIN	9
Key HP Filter	9
Key Stereo Diff	9
Attack	10
Release PEAK/RMS	10
Transfer function, GAIN REDUCTION display	11
Precise/Eco Modes	11
Delta	11
Bypass	11
Technical Specifications	12
Get in touch!	12
End user license agreement	12



Thank you for choosing a Tokyo Dawn Labs product!



CONCEPT

The TDR Feedback Compressor II is a major design update of its critically acclaimed predecessor. The compressor is dedicated to highest fidelity sum/bus compression, but equally shines in classic tracking tasks.

Most modern compressors analyze the input signal to control gain reduction. This compressor however design however, listens to the output. This non-intuitive approach delivers an unobtrusive and highly musical compression characteristic that is able to handle complex signals with ease (more about this in the “Feed-back compression” chapter).



The Feedback Compressor II embeds this old compression topology into a state of the art processing structure and combines it with unusual, but highly flexible, yet intuitive control scheme. The compressor is carefully tuned for intuitive and musical operation under most circumstances.

No compromises have been made in order to achieve the highest quality dynamics compression possible.

We want to emphasize the fact that the processor neither tries to emulate any previously available device, nor does it follow popular trends like “virtual analogue”, “circuit modeling” or similar buzz terms. This is a proud digital processor, made with an immense amount of love and care.

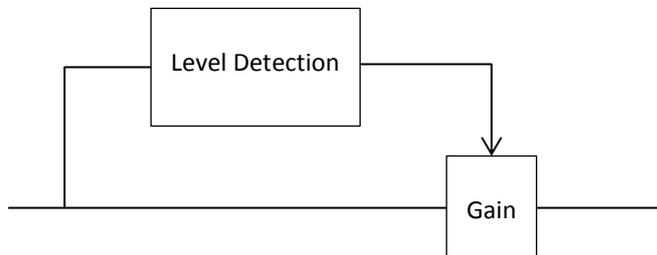
Beside the inherent “feed-back” behavior, the processor comes with an array of unusual features and great flexibility. The most notable are probably the unusual two-stage release timing controls, the “Peak Crest” control scheme, a generous transfer function display and metering section, as well as mastering optimized side-chain filters and stereo link options.

The following pages cover the advanced features of the processor; the reader is expected to have practical experience with the matter.

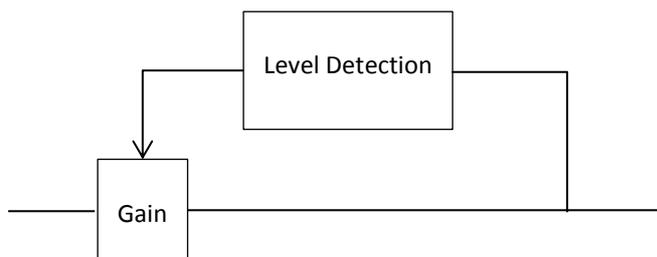


FEED-BACK COMPRESSION

Dynamic range compression can be accomplished in several ways. The most straight forward approach is the idea to detect the input level and use the signal to control an amplifier.



This kind of control structure is called "feed-forward" and works as expected. The device analyzes the input and reduces gain accordingly. Most modern compressors are built this way. However, a similar behavior can also be achieved by feeding the detector with the output of the compressor:



In this case, the detector detects to the output level of the compressor. That is, the compressor listens to what "he has done". This interesting variation is well established in the analogue domain mostly due its better cost efficiency and restrictions of certain circuits.

From the technical point of view, the feed-back structure has several disadvantages. It's much more difficult to control, distortion is significantly higher and compression dependent "by design". It has difficulty to achieve higher ratios than ~3:1, the maximum amount of gain reduction is limited too.

On the other hand, the feed-back approach has some musically related advantages: Most of all a very natural compression behavior and interesting distortion pattern. The compression timing and ratio is program dependent and seems to naturally adapt to most sources quickly. No complicated tweaking required, the compressor sounds fine in most situations. This is the route we took with this project.



PRECISION, ALIASING AND THE SOLUTION

Controlling the dynamic range of an audio signal in the digital domain is not as easy as it looks. A whole array of problems makes it very difficult to build a truly effective and musically attractive compressor in the digital domain.

The most significant restrictions are due to the discrete (i.e. “stepped”) nature of digital signal storage. Without stepping too deep into the mathematics, we have to keep in mind that the data saved in audio-files is not the actual analogue audio signal. It is an intermediate format. The true signal is only reconstructed after the anti-alias filter in the DA converter. So, if one wants to control the amplitude of music accurately, it is absolutely essential to know the actual waveform!

Another huge problem in digital dynamics control is a phenomenon called aliasing or Moiré Images. Aliasing appears all in discrete (i.e. “stepped”) systems as soon the frequency of a signal exceeds the Nyquist rate (half the sample-rate). The evil details here is that contrary to analogue systems, signals exceeding the bandwidth aren’t gradually “faded out”, instead, they mirror at the Nyquist frequency and move back into the audible range at full energy (and mirror again at DC, and at Nyquist again, and again..).

Compressors are non-linear systems. All non-linear systems add harmonic (and non-harmonic) content to the processed signal and thus extend the bandwidth significantly. The more aggressive the non-linearity, the stronger and higher the new generated artifacts. If the sample-rate is too low to handle the extended bandwidth, all these harmonics will alias and lose their harmonic relation to the fundamental.

The Feedback Compressor II’s algorithm was carefully designed to avoid these issues generously. To achieve this in an efficient manner, the algorithm is split in two parts both running at higher rates than the original signal:

Gain Cell

The gain cell is a central part of the compressor and the place where audio signal gets amplified (i.e. multiplied) by the control signal delivered by the side-chain. Such a multiplication is a non-linear process and doubles the bandwidth. To handle this bandwidth, this multiplication always runs at least at 88.2kHz. Only the compression itself (listen via “Delta” switch) is oversampled, not the original signal.

Side-Chain

The side-chain is responsible for generating a control signal for the gain cell. It consists of several relatively complicated non-linear elements such as the threshold and the timing filters. To avoid aliasing in the control signal (which would severely limit the accuracy of the compression), the side-chain runs at a very high rate of more than 352.8kHz.

In both cases, the resampling is done via very high quality linear phase time domain convolution.

Note that these quality improvements come at a price: CPU cycles. But from the sonic point of view, they are well worth the costs.



CONTROL SCHEME

The Feedback compressor II is a relatively complex compressor which allows access to a wide range of advanced parameters. To better understand all the controls and full range of options, it is important to have an idea of the things going on under the hood.

The compressor is designed to handle complex material. To achieve this, the compressor runs two distinct detection circuits in parallel:

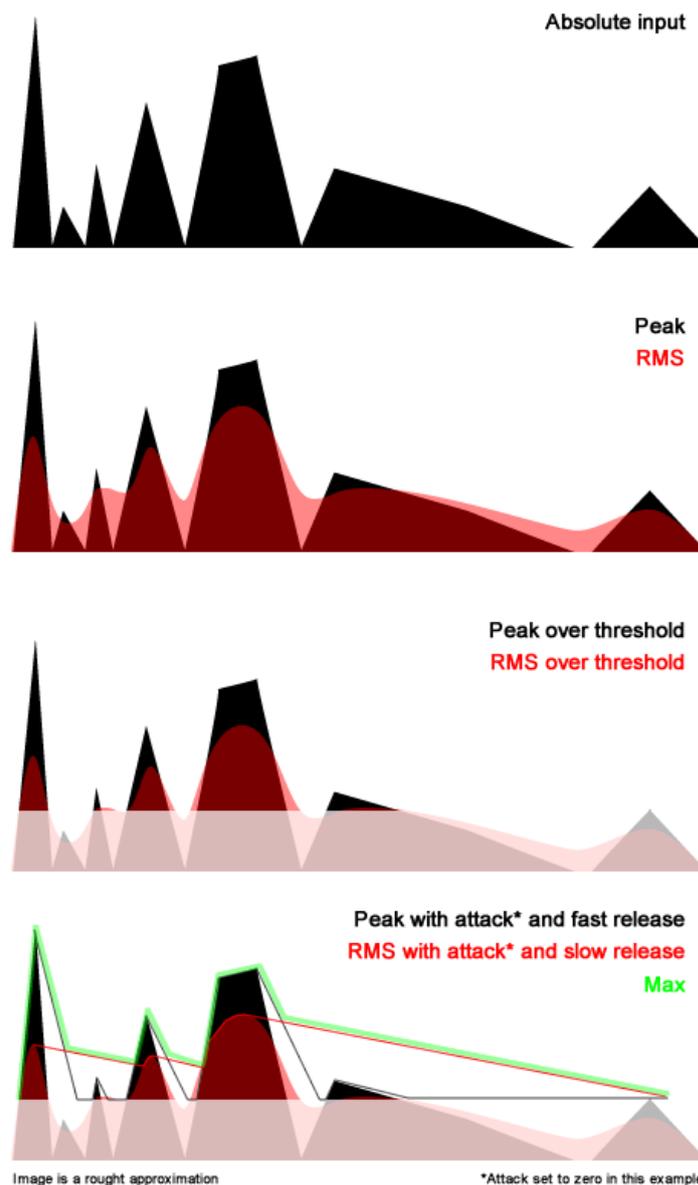
A. Peak Path

The peak detection/compression path follows the (true!) audio waveform closely.

B. RMS Path

The RMS path responds reacts to energy levels and ignores instant events. It's best imagined as a very slow and smooth compressor that handles the constant/steady-state parts of the signal.

The path generating the highest reduction takes control of the whole compression action:



This approach practically offers the advantages of peak compression combined with the advantages of RMS compression. Even more, the RMS path cancels out the disadvantages of the peak path and vice versa. This clever design guarantees a very stable sound image and low compression artifacts over a wide range of material and settings.

CONTROLS AND DISPLAYS

All knobs and range display controls have a velocity based reaction. Small mouse movements only trigger small changes while faster movements create much greater changes. Double-click to re-set the control to its default value.

THRESHOLD, PEAK CREST, SOFT KNEE AND RATIO



Threshold defines the level above the compressor begins to compress the signal. The lower the threshold, the deeper the “grab”.

Peak Crest adjusts the threshold of the peak path relative to the main Threshold. A Peak Crest of zero sets both paths to equal thresholds, a high Peak Crest increases the threshold of the peak path in order to compress fast event only for higher levels and nearly disables the peak path.

Soft Knee adjusts the threshold transition “softness”. Lower knee values create a sharp transition at the threshold point, while higher knee values will compress more gradually.

Peak Ratio In default mode, both peak and RMS paths use the same compression ratio. In “LIM” mode, the peak path always uses a fixed ratio of 7:1, while the RMS path ratio is still controlled by the ordinary Ratio control.

Ratio defines “how strong” the signal will be compressed above threshold. The range is fully continuous between 1:1 and 7:1.

[i] Soft knee reduces the effective threshold.

[i] The threshold affects compression based on the compressor's output level, not the input level. The higher the compression, the lower the detection signal will be. The threshold effectively “rides” on the signal instead of accurately tracking its overloads.

MAKEUP GAIN



Make-up gain is used to compensate the gain-reduction introduced by the compressor. The make-up control can also attenuate the compressed signal up to 60dB which is useful for parallel compression techniques.



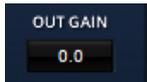
DRY MIX



Dry Mix blends the original signal into the processed (i.e. compressed) part. It enables the use of parallel compression techniques (upward compression, “NY compression”) without the complicated routings and latency compensation inside the DAW.

Dry Mix and Out-Gain form a 2 channel mixer. Dry mix is not a 0-100% Dry/Wet control.

OUTPUT GAIN



Out Gain helps to control the output of the compressor without affecting the balance between Makeup-Gain and Dry-Mix.

KEY HP FILTER



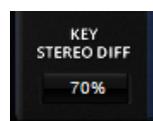
The Key High-Pass Filter section allows to control the compressors sensitivity for low frequencies. That is, the compression threshold becomes frequency dependent and thus less responsive to low frequency content. Three different high pass filter slopes are supported:

- 3dB/Octave
- 6dB/Octave
- 12dB/Octave

The 3dB/Oct filter deserves special attention. It consists of a special filter network which approximates the inverse pink noise distribution. This is probably the most reasonable side-chain pre-processing, because it basically helps the compressor to take into account the typical frequency distribution of music signals: The 3dB/Octave slope of pink noise. This mode performs particularly well on complex material.

The filter’s -3dB point can be set between 30Hz to 400Hz.

KEY STEREO DIFF



Key Stereo Diff to controls the stereo way stereo compression is linked inside the compressor.

Full right (100%) means that the compressor fully preserves the original stereo image under all circumstances. Technically, the channel with the highest level controls both sides equally. This is how most modern compressor are linked for stereo operation.

Full left (0%) is equivalent to a sum based stereo link. That is, both sides are mixed and the sum of them controls both sides equally. This is how several old compressors use to link the both sides for stereo operation. While this mode also perfectly preserves the position of the stereo center, the stereo width is allowed to change depending on compression and material. Technically, the compressor ignores the stereo difference information.

The control differs radically from the stereo link/unlink option found on several stereo compressors. The Feedback Compressor II is never allowed to shift the stereo center! (which would be the case if unlinked operation would be allowed).



ATTACK



Attack time controls the speed of gain reduction, i.e. how long the compressor takes to reach full gain reduction. Fast attack times quickly respond to level changes in the side-chain. Slow attack times on the other hand respond smoothly and let small event pass through the compressor without reduction.

Note that the Feedback Compressor II offers particularly fast reaction times many smaller than a single sample! However, too fast attack times can also easily create distort low frequency content, so be careful.

RELEASE PEAK/RMS



The timing section differs slightly from the typical layout found on most compressors. The TDR Feedback Compressor can dynamically select the most suitable release “path” depending on the program material. Two LEDs indicate the currently active release path and controls.

Release Peak defines how quickly gain recovers after short overloads (i.e. fast transients). **Release RMS** defines how fast gain recovers sustained content.

The following image shows a test-signal before and after compression. It consists of one short event (such as a snare), a long term event and some “background noise” in order to make the effect more visible. As you can see, the reduction triggered by the short overload recovers much faster than the reduction caused by the second overload:



Two release paths offer a wide range of musically useful options. For example, single drum peaks can be allowed to recover quickly to avoid “dulling” and “breathing” side-effects. At the same time, sustained content like bass-lines or synth pads can recover slower and thus strongly reduce typical side-effects like “pumping” and distortion.

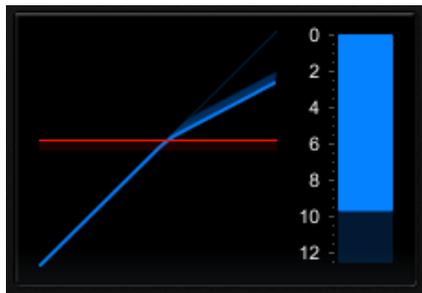


We recommend **Release Peak** value between 25-200ms and **Release RMS** to values above 160ms for complex content such as mixes. Faster settings are useful for particularly dynamic material like drum solos, bongos and acapellas.

Because the peak path is much more sensitive than the RMS path, setting the **Release Peak** to a higher value than **Release RMS** while keeping the **Peak Crest** low practically disables the RMS part of the compressor. In other words, the algorithm “falls back” to a single stage release and behave much like a common one-stage release peak compressor.

To practically disable the Peak path of the compressor, set **Peak Crest** full right and set **Release Peak** full left.

TRANSFER FUNCTION, GAIN REDUCTION DISPLAY

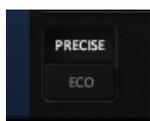


“Transfer Function” shows the current in/out gain transfer function of the compressor. The solid line shows the compressor’s overall transfer function, the half transparent line above represents the transfer function of the peak path (which is always **Peak-Crest** dB above the RMS path’s transfer function). The transfer function display covers a range between 0dB and -65dB.

The main threshold is indicated by the red line. The soft knee region appears dark red.

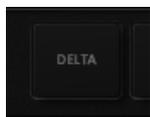
“GR” shows the current gain reduction applied to the signal. An additional slowly decaying reduction amount is shown in dark blue. The scale range can be switched by clicking the scale or meter area. A red marker appears at the bottom of the reduction bar in case the reduction amount doesn’t fit inside the current scale.

PRECISE/ECO MODES



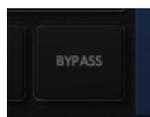
The switch allows to select an more efficient processing mode to save a few CPU cycles in critical situation. This **ECO** mode is not meant as creative parameter! No matter what you mean to hear, the process is the same, just at lower quality! ☺

DELTA



Listen to the difference between the original signal and the compressed signal. This is best described as “*what the compressors actually does*” and is very useful to get a better understanding how different settings affect the original signal.

BYPASS



Bypasses the whole processor. Processor latency is accurately compensated and the actual processing is never interrupted for better comparison.



TECHNICAL SPECIFICATIONS

VST version:	2.4
Available binaries:	32bit and 64bit
Input / Output resolution:	32bit floating point
Latency:	184 samples at 48kHz or below, 8 samples for higher sample-rates
Supported sample-rates:	From 44.1kHz to 192kHz
Supported channels:	Stereo

GET IN TOUCH!

Ideas? Bugs? Improvements? We want to hear your feedback! You can easily reach us via one of the websites below.

Check out the Tokyo Dawn Labs website for feedback, news, updates and downloads:

<http://www.tokyodawn.net/tokyo-dawn-labs/>

You can also directly head to the TDR Feedback Compressor page:

<http://www.tokyodawn.net/tdr-feedback-compressor-2/>

Tokyo Dawn Labs is closely affiliated to Tokyo Dawn Records, so make sure to check out their artists and releases!

<http://www.tokyodawn.net>

END USER LICENSE AGREEMENT

End User License Agreement

This software is provided free of charge but Tokyo Dawn Records retains copyright.

You are not allowed to redistribute this software without explicit permission from Tokyo Dawn Records.

You are not allowed to sell or to rent this software. You are not allowed to reverse engineer this software.

You are allowed to use this software for educational and artistic application including commercial music production.

This software is provided 'as-is', without any express or implied warranty. In no event will Tokyo Dawn Labs be held liable for any damages arising from the use of this software.

