

TOKYO DAWN LABS

Feedback Compressor II – Manual

Product version: 2.0.0

Revision 1.7, last update: 03/16/2013

Editors: Fabien Schivre, Dax Linière (Puzzle Factory)



CONTENTS

Concept	4
Feed-back Compression	5
Precision, Aliasing and the solution	6
Control Scheme	7
Threshold, Peak Crest, Soft Knee and Ratio	8
Makeup Gain	8
Dry Mix	8
Output Gain	8
Sidechain HP Filter	9
Sidechain Stereo Diff	9
Attack	10
Release PEAK/RMS	10
Transfer function, gain reduction display	11
Precise/Eco Modes	11
Delta	11
Bypass	11
Controls and displays	12
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 the highest fidelity stereo program (2-buss) compression, but shines equally in classic mixing tasks.

Most modern compressors analyze the input signal to control gain reduction. This is known as a “feed-forward” topology. The TDR Feedback Compressor II, however, analyzes the output. This approach delivers unobtrusive and highly musical compression characteristics which mean the compressor is able to handle complex signals with ease (more about this in the “Feed-back compression” chapter).



The TDR Feedback Compressor II takes this traditional compression topology to new heights as a state of the art dynamics processor combining an unusual and highly flexible, yet intuitive control scheme. The compressor has been carefully tuned for intuitive and musical operation for almost every situation.

No compromises have been made in order to achieve the highest possible quality of dynamic control.

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 words. This is a proud digital processor, made with an immense amount of love and care.

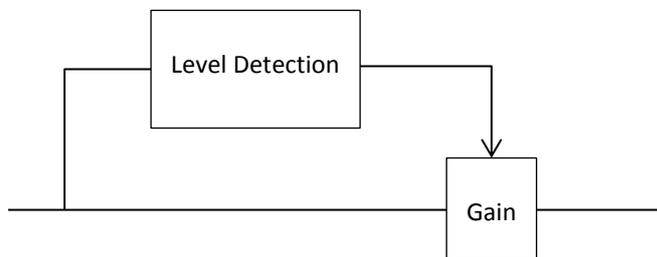
Beside the inherent “feed-back” behavior, the processor comes with an array of unusual features which offer great flexibility. The most notable are probably the unusual two-stage release timing controls, the “Peak Crest” control scheme, a generous transfer function display and metering section, as well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

The following pages cover the advanced features of the processor; it is a prerequisite that the reader is accustomed to dynamics processors and their uses.



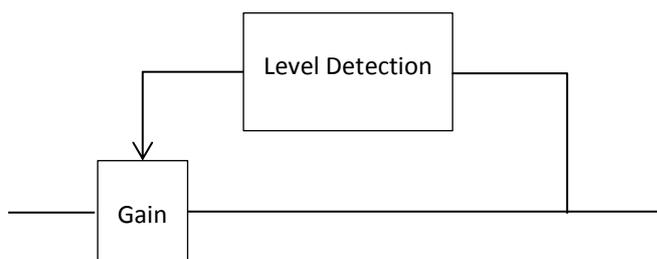
FEED-BACK COMPRESSION

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



This kind of control structure is called "feed-forward". The device analyzes the input and reduces gain accordingly. Most modern compressors are built this way.

A different behavior can be achieved by feeding the detector with the output of the compressor:



In this case, the detector (envelope follower) detects the level at the output of the compressor. That is, the compressor listens to "what it has done". This interesting topology is common in the analogue domain, mostly due its better cost efficiency and the restrictions of certain circuits.

From a technical point of view, the feedback structure has several disadvantages. Typically, it's much more difficult to control, and distortion is significantly higher and compression dependent "by design". It has difficulty achieving ratios higher than approximately 3:1, and the maximum amount of gain reduction is limited.

On the other hand, the feedback approach has some musically related advantages; Most of all a very natural compression behavior and interesting harmonic distortion distribution.

The compression timing and ratio is program dependent and seems to naturally adapt to most sources quickly. No complicated tweaking is required for the compressor to sound good in most situations. This is the route we took with the project of building the TDR Feedback Compressor.



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 delving 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. But, 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 detail 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 newly 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, resulting in an unpleasant sound.

The TDR Feedback Compressor II’s algorithm was carefully designed to avoid these issues. 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’s level is adjusted by the control signal delivered by the sidechain. Such an operation is a non-linear process and doubles the bandwidth. To handle this bandwidth, this multiplication always runs at least at twice the audio sample rate (i.e. at least 88.2kHz). One interesting design feature of TDR Feedback Compressor II is that only the compressed portion of the signal (listen via “Delta” switch) is oversampled. This leaves the original signal completely untouched - as long as no gain reduction is occurring, the plugin is 100% bit transparent.

Sidechain

The sidechain is responsible for generating a control signal for the gain cell. It consists of several complex, non-linear elements such as the threshold and timing filters. To avoid aliasing in the control signal (which would severely limit the accuracy of the compression), the sidechain 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 program material. To achieve this, the compressor runs two distinct and separate detection circuits in parallel:

A. Peak Path

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

B. RMS Path

The RMS path reacts to overall energy levels (similar to human hearing) and ignores instant events. It's best imagined as a very slow and smooth compressor that handles the constant or 'steady-state' parts of the signal.

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



The approach offers the advantages of peak compression combined with the advantages of RMS compression. Furthermore, the RMS path negates the disadvantages of the peak path and vice versa. Its design guarantees a very stable sound image and low compression artifacts over a wide range of material and settings.

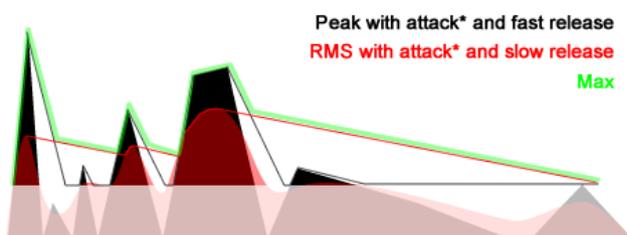
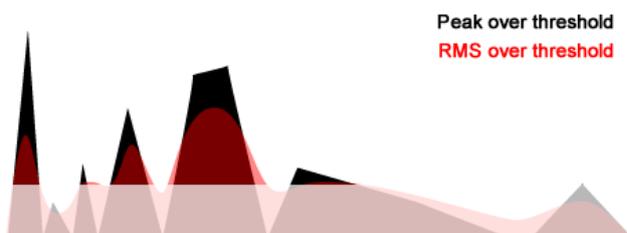
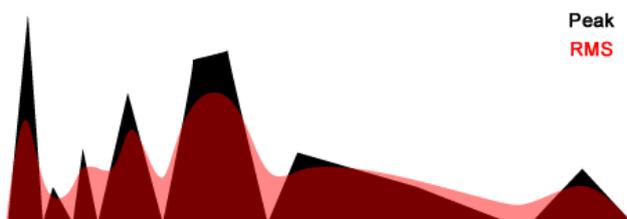


Image is a rough approximation

*Attack set to zero in this example



THRESHOLD, PEAK CREST, SOFT KNEE AND RATIO



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

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 events (transients) 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, but also ‘deeper’ than the threshold.

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 strongly” the signal will be compressed as it exceeds the 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 makeup 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 complicated routings or DAW latency compensation issues.

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

OUTPUT GAIN



Out Gain controls the output of the compressor without affecting the balance between compressed and dry mix signals.



SIDECCHAIN HP FILTER



The **Side-chain High-Pass Filter** section allows control over the compressor's sensitivity to 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 sidechain 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 crossover point can be set between 30Hz to 400Hz.

SIDECCHAIN STEREO DIFF



Side-Chain Stereo Difference controls the way compression of left and right channels is linked.

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 compressors 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 used to link the two mono channels 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.

This control differs radically from the stereo link/unlink option found on most stereo compressors. The Feedback Compressor II is never allowed to shift the stereo center! (which could occur if unlinked operation was permitted.)



ATTACK



The Attack time controls how long the compressor takes to reach full gain reduction. Fast attack times respond quickly to level changes in the sidechain. Slow attack times on the other hand respond more slowly and smoothly, letting small events (transients) pass through the compressor without gain reduction.

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

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 virtual 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 after sustained gain reduction.

The following image shows a test-signal before and after compression. It consists of one short event (such as a snare hit), 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 transient recovers much faster than the reduction caused by the second over-threshold event:



Dual 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.

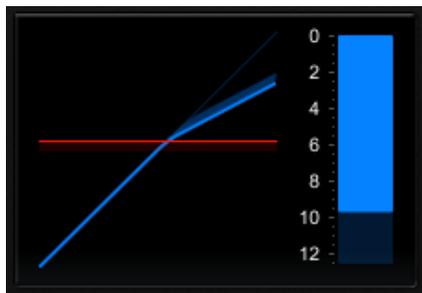


We recommend a **Release Peak** value between 25-200ms and **Release RMS** values above 160ms for complex content such as mixes. Faster settings are useful for particularly dynamic material like solo drums, percussion and acapellas.

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 reverts to a single-stage release and behaves much like a common peak compressor.

To “disable” the Peak path of the compressor, set **Peak Crest** fully clockwise and set **Release Peak** fully counterclockwise.

TRANSFER FUNCTION, GAIN REDUCTION DISPLAY



The Transfer function graph 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 in dark red, always below the main threshold line.

The blue bar on the right shows the current gain reduction applied to the signal. An additional slowly decaying dark blue bar shows the maximum reduction over a bigger time-frame. The scale range can be switched by clicking the scale or meter area (left click to increase range, right click to decrease). A red warning marker appears at the bottom of the reduction bar in case the reduction amount exceeds the current scale.

PRECISE/ECO MODES



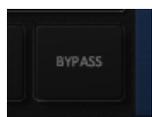
This switch allows the user to select a more efficient processing mode to save a few CPU cycles for less critical situations. This **ECO** mode is not meant as creative parameter! No matter what you mean to hear, the process is the same, just at lower quality! 😊

DELTA



This allows the user to listen to the difference between the original and compressed signals. 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 (gapped) to enable better comparisons.



CONTROLS AND DISPLAYS

All controls within the plugin feature acceleration. This is to say fast movements of the mouse make large changes and slow, small movements allow for fine tuning of settings.

You can also double-click a control to reset it to its default value.

TECHNICAL SPECIFICATIONS

Available binaries:	32bit and 64bit VST/AU
Input / Output resolution:	32bit floating point
Internal resolution:	64bit floating point
Latency:	184 samples at 48kHz or below, 8 samples for higher sample rates
Supported sample-rates:	From 44.1kHz to 192kHz
Supported channel configuration:	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.

