TECH LIBRARY

Compressors & Compression

How they work and how to use them

Introduction

A compressor is designed to do as the name suggests, compress or "squash". It works by reducing the dynamic range of the audio fed into it. In other words, it controls the audio level to limit its maximum value gradually. Unlike severe limiting (which is the responsibility of a limiter, often used in protection scenarios for speaker processors), a compressor gently controls the dynamic range without dramatically capping it, which could introduce audible side effects.

Why do we need Compressors?

The dynamic range of any instrument or device is a measure of the ratio of the maximum possible output level to the minimum level that can be reproduced. In the case of, for example, a compact disc player, the dynamic range exceeds 90dB. However, in a typical analogue tape recorder, the higher noise floor limits the dynamic range to about 70dB.

If two pieces of equipment are connected together (to record the CD for example), the device with the lowest dynamic range determines the dynamic range for the entire system. In the example above, the dynamic range available is limited to 70dB, losing 20dB.

This limitation also applies in live sound situations, where the dynamic range from one component of the system will prove to be the limiting factor in the whole set-up. Typically this "limiting factor" will be the power amplifier or speaker system. Running the system at such a level so that the average volume is adequate will often not allow signal peaks to be handled correctly. Given a musical peak that is 6dB above the average signal level, if the system is forced to try and reproduce a level which exceeds its dynamic range, distortion will almost always occur.

How can we avoid this?

In an ideal world, all devices used to record and reproduce sound would have boundless dynamic ranges and so present no level matching problems with each other. As this is not the case, some method must be employed to limit the dynamic range of signals in situations where they will prove troublesome or incompatible.

This could be achieved just by 'turning down' the offending level until it reaches a point where the maximum possible output will not exceed the dynamic range of the system as a whole. Unfortunately, this type of action rarely produces satisfactory results, as the average level will also be attenuated, meaning the perceived volume will drop too low.

Enter the Compressor…

What is required is a device that can monitor the signal passed to it, and allow the average level to pass untouched. At a threshold level the device will begin to turn down the level by a certain amount. This amount can be varied, to offer more or less attenuation of signals above the threshold. What would also be useful would be some way of controlling the rate at which the level is turned down (slower rates will not sound as unnatural as the effect of 'reaching for the volume control and giving it a sharp twist'). Additionally, the ability to change the rate of recovery of the original gain over a period is as important.

How does a Compressor work?

A compressor functions by managing signal levels over time. Consider the diagram below, which shows a signal level versus time . It represents a burst of audio, which is initially below the threshold of the compressor - in **Region 1**. As this is in the 'safe' region, it is unaffected by the compressor. It can be thought of as the average signal level and remains unattenuated. However, in **Region 2**, the signal has risen (very sharply, hence the sudden jump) to a level above the threshold.

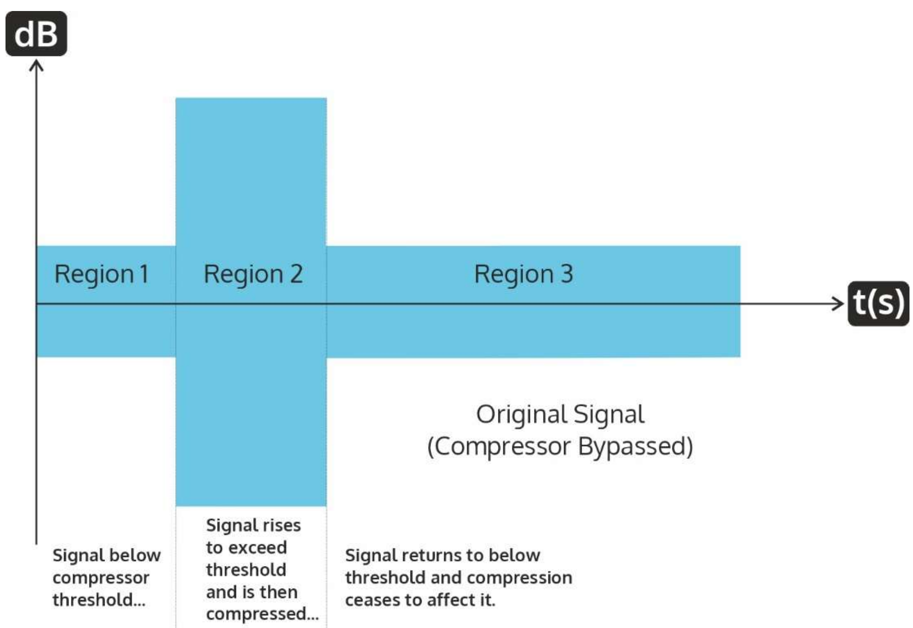

This signal is potentially too high and must be reduced in level to a more acceptable value.
The difference between the original signal and the compressed signal is set by the Ratio.

For instance, if the input jumps from 0 dB to +10 dB with the threshold at 0 dB, and the output is limited to a rise of 5 dB, the ratio would be 2:1 (10/5).

The diagram below shows the same signal after passing through the compressor.

Important in **Region 2** is the red shaded area. This represents the time to fully apply the gain reduction set with the ration control, or the Attack time. The faster the attack time, the less signal will 'escape' through the compressor before it reacts and applies gain reduction. The setting of this control will be discussed in the next section. Its value can have a profound effect on the transparency of the compressor ; that is, how natural it sounds and how noticeable its action appears.

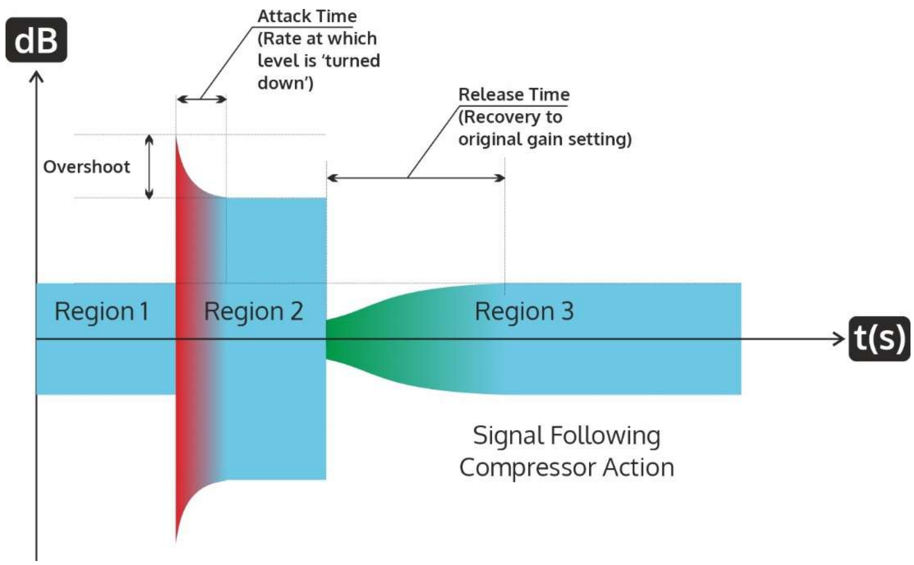

The last region of interest is **Region 3** where the signal returns once again to below the threshold. The green shaded area represents the recovery time for the compressor to stop applying gain reduction and effectively 'recognise' that the signal is to pass through unaffected. This represents the Release time and again must be carefully set so as to prevent audible artefacts.

Note the level difference marked as **'Overshoot'**. This phenomenon cannot be avoided because the compressor has no way of knowing when the signal will reach the threshold, and so cannot 'prepare in advance' for this happening. Consequently, there will always be a brief period before it can react, even at the minimum attack setting (fastest rate of gain reduction application).

Editing & Interaction of the Compressor Parameters in D-Net

The compressor comes first on all input processing signal paths on NST Audio products, ahead of the input gain control. So, adjusting the input fader will NOT affect the level being fed into the compressor. This is so that the input fader can be used to apply make-up gain if necessary.

In this example, the live signal shown by the dynamic green area below the white line is approximately 10dB below the threshold, currently set to +10dB.

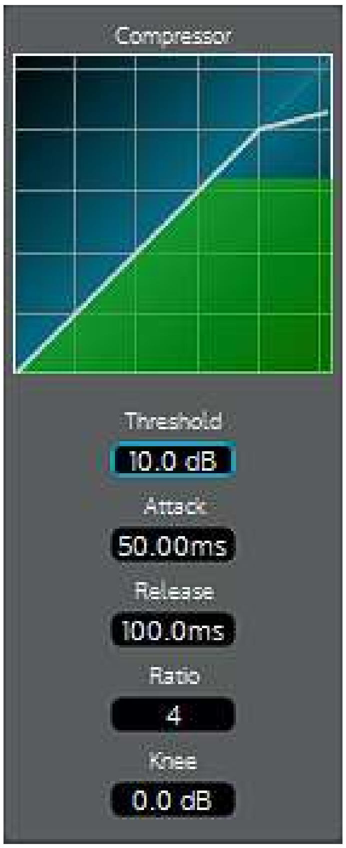

Dragging the threshold parameter down or entering a lower value, will move the point at which compression starts to nearer the signal level. If this value is reduced to 0dB, the signal will now be on the threshold of compression.

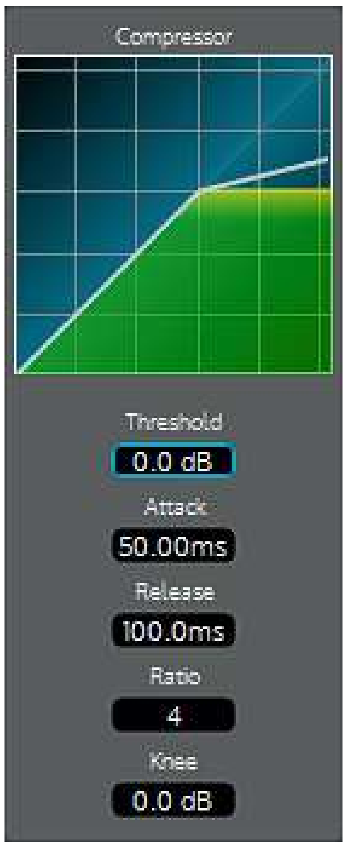

Reducing the threshold yet further, to -10dB will now be quite heavily into compression, evidenced by the signal changing colour as it exceed the threshold and is being reduced in level.

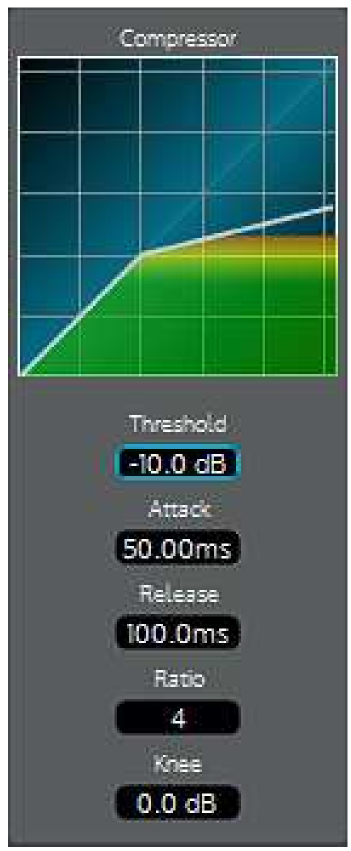

Make-up Gain

As mentioned earlier, the compressor is the first thing in the signal path in NST products, and the input fader will NOT affect the level on the compressor graph. The input fader is post-compressor and can be used to add make-up gain if required.

The best method of setting the make-up gain is to use the 'UNDO' feature in D-Net and swap between the chosen settings and a ratio of 1 (so bypassed) and compare, whilst adding gain when the compressor is set to a ration above 1. The human ear is a much better gauge than a digital meter. UNDO and REDO are quickly accessed using the Ctrl+Z and Ctrl+Y shortcuts (Windows) or Cmd + Z and Shift + Cmd + Z (mac).

Setting the Attack and Release times.

Using too fast attack and release times on low frequency program (such as a bass guitar) will cause the compressor to respond to individual cycles of the signal, rather than the overall envelope. This will result in obvious distortion, which might be described as sounding like clicking superimposed on the original signal.

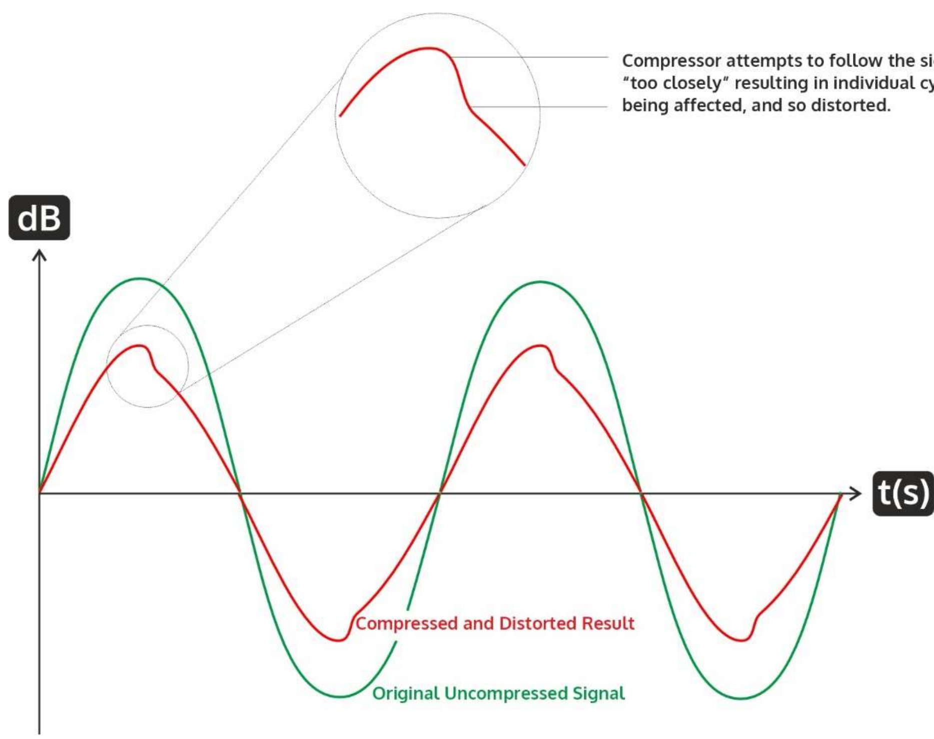

The compressor release time has been deliberately restricted to a minimum of twice the attack time to prevent excessive distortion on low frequency signals, even with fast attack times and high ratios. None the less, it is still possible to introduce some distortion if care is not taken with the settings. The best way to ensure that the signal is not being excessively distorted is to make good use of the undo and redo feature constantly comparing the original signal with the compressed version.

The Compressor 'Knee' Control.

One of the most misunderstood parts of a compressor is the parameter usually labelled the 'Knee'. This may be a fully variable control, or a switchable parameter, normally with 'Hard' and 'Soft' settings. The knee control permits a softening of the compressor action, which can prove to be especially useful at high compression ratios.

Without the inclusion of a knee control, at the threshold of gain reduction, a sudden transition occurs between unity gain and the ratio by which the compressor attenuates. When using high compression ratios, the use of a hard knee can result in a very unnatural sound.

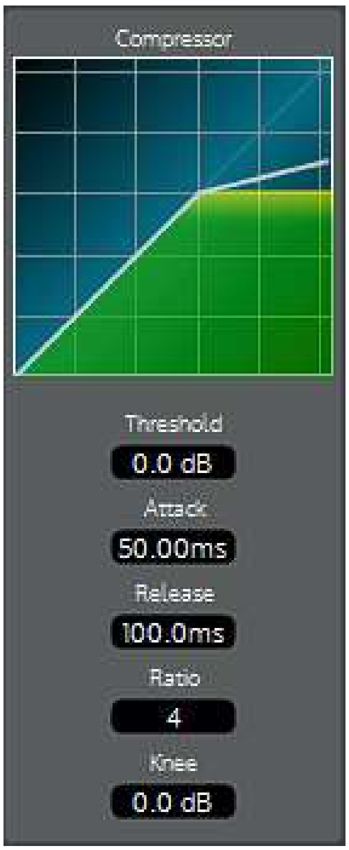

Consider a compressor with a ratio of 4:1. The in-out transfer function as seen in D-Net would be like this. The very sharp introduction of gain reduction is obvious at the threshold point (0dB). The 'Knee ' parameter is set to 0dB in this example. This corresponds to a 'hard' setting.

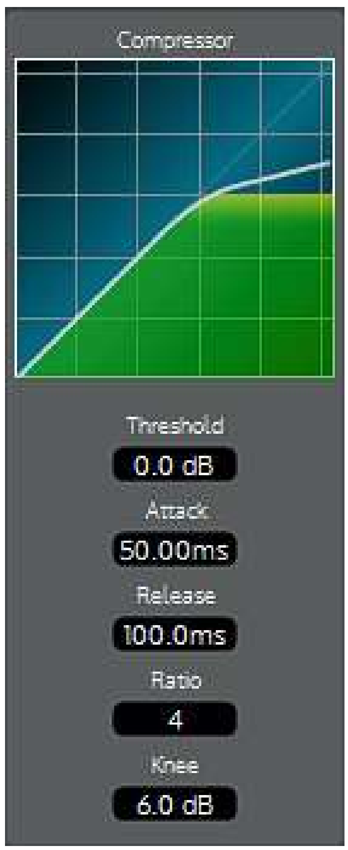

Increasing the 'Knee' to its maximum of 6dB, spreads the onset of the compression over a wider area, (3dB above and 3dB below the threshold), reducing the severity of the compressor. This shows on the graph as a "smoothing" across the threshold, rather than a sharp bend.

In Conclusion

Understanding how compressor controls interact and utilizing features like undo/redo in D-Net can unlock the full potential of this powerful dynamics module. Even applying subtle compression to highly dynamic input signals can elevate a performance from "okay" to truly outstanding.