Compressors – how, why and when.

A compressor is a type of audio tool that falls within the group known as dynamic processors. These also include de-essers, limiters, and gates, but the compressor is the daddy of them all.

Compressors are used in recording studios to control level fluctuations that may occur in the audio signal. For example, a singer's voice may be louder or quieter at different parts of a song – especially if they have less than perfect microphone technique! In the past this problem may have been dealt with by physically adjusting the fader level on the mixing desk to compensate for the discrepancies, a process known as 'riding the fader'. A compressor will accomplish this more efficiently – it has a much faster reaction time – and automatically, leaving the engineer with two free hands to do the other hundreds of things that need doing. Because a compressor affects the overall dynamic range of the sound, effectively 'squeezing' the signal, it can also increase the average energy of a recording giving a better level before peaking and can (subjectively) give a tighter, more consistent sound. These days we are all used to hearing compression as it is commonly used at the mastering stage to maximise the energy in a mix, making the song sound as loud as possible.

How does it work?

In order to understand how a compressor works, we need to know that it has two main components, a voltage controlled amplifier (VCA) and a side chain.
The VCA is just like a normal amplifier except that it's gain (the amount it makes things louder) is controlled by the voltage level of a control signal instead of by a knob on the front.

Compressor—block diagram

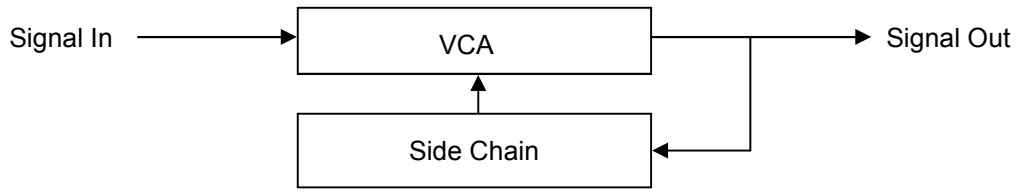

The side chain monitors the output gain of the amplifier and uses this to generate the control signal for the VCA. With the compressor switched to bypass, the side chain will merely 'follow' the level of the signal passing through the VCA and will have no effect. Switch the compressor in and the fun begins! When the output of the VCA exceeds a preset threshold level, the side chain generates a control voltage that will reduce the gain of the amplifier according to a preset ratio. The threshold and ratio settings are selected by the user along with the attack and release characteristics of the compressor—more about them later!

The controls.

Let's have a look a some typical controls that you will find on the front of your compressor. Understanding what they do is key to the effective use of a compressor and its associated techniques. The actual controls and their values may vary slightly so I'll include the standard controls with an average range of values.

Threshold: -30dB to infinity. This sets the point at which gain reduction (compression) commences. It is important to realise that it is the signal entering the side chain circuit that is governed by the threshold, not the input signal.

Ratio: 1:1 to 20:1. This controls the amount of compression. A ratio of 5:1 means that a signal that would be 5dB above the threshold is reduced to 1dB above by the compressor.

Attack: 50uS to 500mS. (50 micro seconds to 500 milli seconds) This controls the response speed of the compressor. 50uS will react to the fastest of audio signals, a slower attack time will allow transient peaks to escape compression and may give a more natural percussion to a sound.

Release: 50mS to 5S. (50 milli seconds to 5 seconds). This controls the time taken for the compressor to recover from gain reduction. A short release time will stop the effects of the compressor as soon as the signal falls below the threshold. A long release time will continue to compress the signal at the selected ratio. Too long a release time may cause a 'pumping' effect in the output signal as the sound is never allowed to die away as it would naturally. This control is sometimes labelled 'sustain', especially on compressors designed for guitars.

Output Gain: -20 to + 20dB. This is used to make up for any signal level that is lost during compression and is sometimes called 'make up gain'.

Other controls commonly found on a compressor include: a switch for bypassing the unit; a stereo link button that allows you to control both sides of a dual compressor with one set of controls for easy stereo compression; Side chain listen switch that lets you listen to the signal entering the side chain. This is sometimes called 'key listen' and is a useful feature if you use compression techniques such as ducking or de-essing. You will normally also find some kind of level meter or at least a peak indicator.

Compressors and limiters

One way to better understand what a compressor is doing to your signal is to look at the difference between a compressor and a limiter. A limiter (or limiting amplifier) is an amplifier whose output is practically constant above a preset input level. In other words, no matter how much more you put into the amp it won't get any louder. Unlike a compressor, this characteristic does not affect the whole dynamic range of the signal but just 'limits' the peak levels to prevent tape or power amp clipping. A limiter will have a ratio of at least 10:1 and typically as much as 50:1 with an input threshold set *just under* tape saturation level (in the old days), max digital input level or clipping level on an amplifier or broadcasting system. A compressor set with a high ratio and a very fast attack time will work as a limiter but will affect the whole dynamic range of the signal and so raise the level of the noise floor creating added hiss.

Compressor vs. Limiter diagram

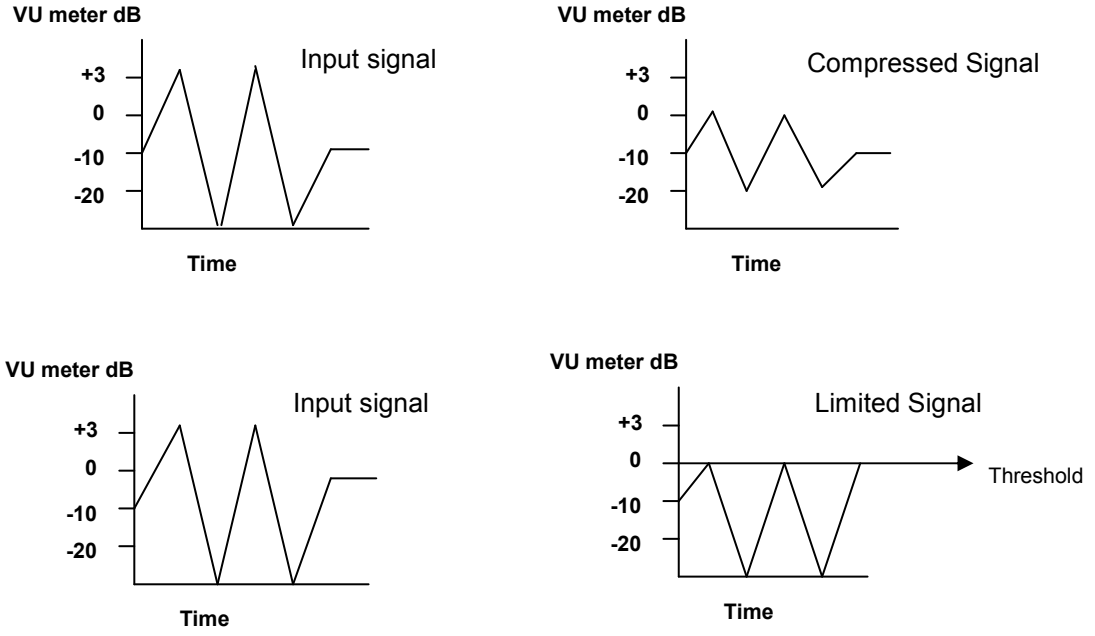

This principle can be easily observed by looking at the waveform display on your DAW. Take a dynamic sound such as drums, apply some extreme compression and you can see the loud peaks get quieter and the quiet parts get louder. Try it, you might be surprised!

The knee

Compressors come with two varieties of knee, hard and soft. The term knee describes the type of curve that the compressor responds with when it starts its gain reduction. Hard knee will give instant compression at the selected ratio whenever the threshold is exceeded. Soft knee will give a progressive compression over a range of about 10dB up to the nominal threshold. The soft knee gives a less processed sound which can be useful when subtle gain reduction is required such as a stereo mix. Hard knee produces a more aggressive reduction that can work well for more percussive sources such as drums.

Comparison of Knee characteristics

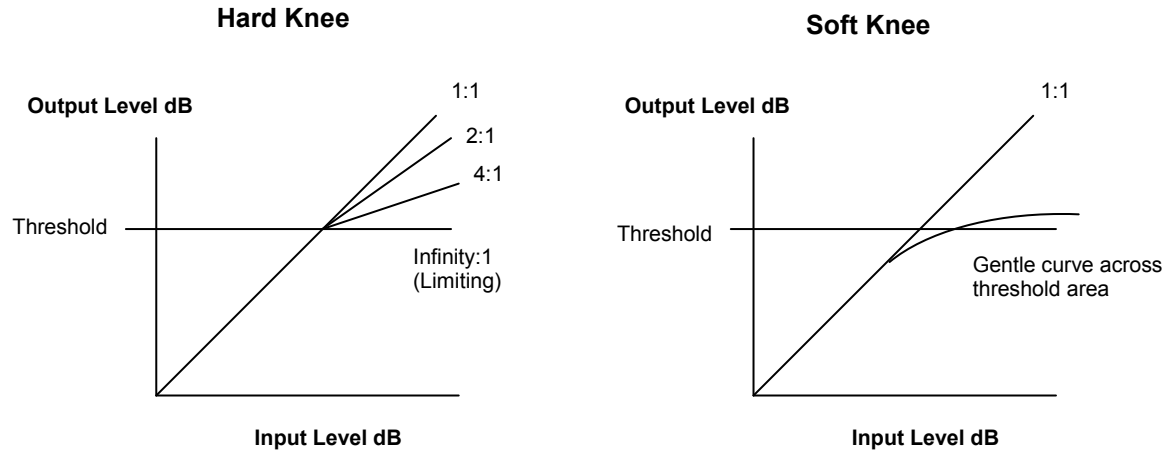

A little compression at the recording stage can maximise your levels and prevent excessive signal peaks from overloading the input and causing distortion or clipping. However in these days of multiple plug-ins more can easily be added at a later stage. Heavy compression (often used for kick drums or slap bass) may cause high frequency loss due to the compressor's attenuation of peaks in the low frequencies. A normal compressor is not frequency selective so any gain reduction will affect the whole frequency spectrum. This is one reason why its always good to get a balanced sound before you start adding the huge plug-in 'mastering suite' across your mix! To overcome this drawback, a slower attack time will allow some high frequencies (such as cymbals) through the compressor or an EQ unit could be patched in after the compressor.

With very fast attack and release times, the compressor may attempt to respond to individual wave cycles at low frequencies (as these are the longest cycles) which will lead to distortion as the wave-form is compressed out of shape. Some compressors have a 'hold' function to delay the release cycle or you can make sure the release is set to longer than the lowest frequency cycle—50mS should be safe to 20Hz.

Other compression techniques.

Sibilance caused by air in the teeth during speech or singing is a common issue in recordings. This effect can occur from 5kHz up to 10kHz. To reduce this unpleasant sounding phenomenon an EQ unit can be inserted into the side chain with the sibilant frequencies boosted and the normal signal sent through the compressor as usual. Because of the EQ boost, sibilant frequencies will be compressed before the rest of the signal. This technique is commonly known as de-essing. Many modern compressors will have a dedicated de-essing function where the side chain is set up for you and all that is required is a tweak of the required frequency and/or level.

De-essing with compressor and EQ.

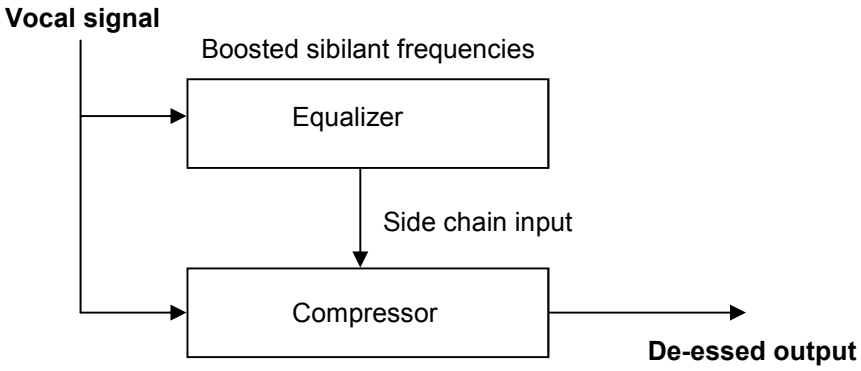

You may have noticed that many DJ's manage to fade the music down while they are talking so that their speech is more clearly heard. This is known as ducking and it is achieved by sending the voice signal through the side chain and the music through the compressor. When the voice exceeds the threshold, the music is compressed. Ducking is used extensively in radio and television for voiceovers and with a little imagination can produce some interesting and unusual effects. Experimenting with ducking is also a great way to understand how the side chain interacts with the VCA in a compressor and demonstrates that it is only the loudest part of a signal that triggers compression, even though the resulting gain reduction will affect the whole sound.

Ducking with a compressor.

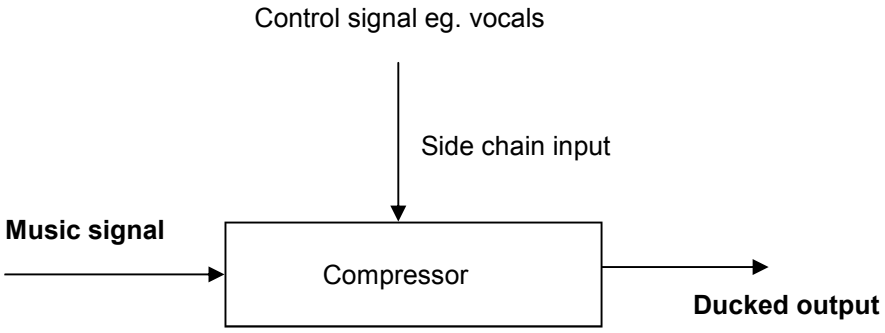

Here's a simple experiment to try. Feed a drum part into the side chain with its reverb going into the compressor. Whenever the drums stop, the reverb tail can be heard but while the drums are playing, the reverb is ducked and lost in the mix. Once again, many compressor have a dedicated ducking function that makes the whole process simpler.

It is very tempting to now include a list of recommended compressor settings for various common recording tasks such as vocals, drums, mix etc.. I believe that if you take some time to understand how the compressor works and experiment, the experience will be rewarded with an ability to quickly dial in just the right settings to get the sound you want.

Many engineers will tell you that, used correctly, the compressor will be inaudible. Why use a compressor if you can't hear it? I hear you ask. In this case the compressor is being used to counteract some of the artefacts of the recording process that result in an unnatural colouring and return the sound to one that is as natural as possible. On another occasion an engineer may use the effect of deliberate over-compression to create a unique and interesting sound that is most audible. There really are no rules and ballpark settings are just that—ballpark and approximate. The secret is—as ever—to understand your equipment so you can get the best out of it.