Academy of Sound Workshop #4 Compression

SETTINGS AND CHARACTERISTICS

Threshold: With manual gain riding, the level above which the signal becomes unacceptably loud is determined by the engineer's discretion: if it sounds too loud to him, he turns it down. In the case of a compressor, we have to 'tell' it when to intervene, and this level is known as the Threshold. In a conventional compressor, the Threshold is varied via a knob calibrated in dBs, and a gain reduction meter is usually included so we can see how much the gain is being modified. If the signal level falls short of the threshold, no processing takes place and the gain reduction meter reads 0dB. Signals exceeding the Threshold are reduced in level, and the amount of reduction is shown on the meter. This means the signal peaks are no longer as loud as they were, so in order to compensate, a further stage of 'make-up' gain is added after compression, to restore or 'make up' any lost gain.

Ratio: When the input signal exceeds the Threshold set by the operator, gain reduction is applied, but the actual amount of gain reduction depends on the 'Ratio' setting. You will see the Ratio expressed in the form 4:1 or similar, and the range of a typical Ratio control is variable from 1:1 (no gain reduction all) to infinity:1, which means that the output level is never allowed to rise above the Threshold setting. This latter condition is known as limiting, because the Threshold, in effect, sets a limit which the signal is not allowed to exceed. Ratio is based on dBs, so if a compression ratio of 3:1 is set, an input signal exceeding the Threshold by 3dB will cause only a 1dB increase in level at the output. In practice, most compressors have sufficient Ratio range to allow them to function as both compressors and limiters, which is why they are sometimes known by both names. The relationship between Threshold and Ratio is shown in Figure 2, but if you're not comfortable with dBs or graphs, all you need to remember is that the larger the Ratio, the more gain reduction is applied to any signal exceeding the Threshold.

Hard Knee: This is not a control or parameter, but rather a characteristic of certain designs of compressor. With a conventional compressor, nothing happens until the signal reaches the Threshold, but as soon as it does, the full quota of gain reduction is thrown at it, as determined by the Ratio control setting. This is known as hard-knee compression, because a graph of input gain against output gain will show a clear change in slope (a sharp angle) at the Threshold level, as is evident from Figure 2.

Soft Knee: Other types of compressor utilise a soft knee characteristic, where the gain reduction is brought in progressively over a range of 10dB or so. What happens is that when the signal comes within 10dB or so of the Threshold set by the user, the compressor starts to apply gain reduction, but with a very low Ratio setting, so there's very little effect. As the input level increases, the compression Ratio is automatically increased until at the Threshold level, the Ratio has increased to the amount set by the user on the Ratio control. This results in a gentler degree of control for signals that are hovering around the Threshold point, and the practical outcome is that the signal sounds less obviously processed. This attribute makes soft-knee models popular for processing complete mixes or other sounds that need subtle control. Hard knee compression can sometimes be heard working, and if a lot of gain reduction is being applied, they can sound quite heavy-handed. In some situations, it can make for an interesting sound -- take Phil Collins' or

Kate Bush's vocal sounds, for example. The dotted curve on the graph in Figure 2 (p.118) shows a typical soft-knee characteristic.

Attack: The attack time is how long a compressor takes to pull the gain down, once the input signal has reached or exceeded the Threshold level. With a fast attack setting, the signal is controlled almost immediately, whereas a slower attack time will allow the start of a transient or percussive sound to pass through unchanged, before the compressor gets its act together and does something about it. Creating a deliberate overshoot by setting an attack time of several milliseconds is a much-used way of enhancing the percussive characteristics of instruments such as guitars or drums. For most musical uses, an initial attack setting of between 1 and 20 mS is typical. However, when treating sound such as vocals, a fast attack time generally gives the best results, because it brings the level under control very quickly, producing a more natural sound.

Release: The Release sets how long it takes for the compressor's gain to come back up to normal once the input signal has fallen back below the Threshold. If the release time is too fast, the signal level may 'pump' -- in other words, you can hear the level of the signal going up and down. This is usually a bad thing, but again, it has its creative uses, especially in rock music. If the release time is too long, the gain may not have recovered by the time the next 'above Threshold' sound occurs. A good starting point for the release time is between 0.2 and 0.6 seconds.

Auto Attack/Release: Some models of compressor have an Auto mode, which adjusts the attack and release characteristics during operation to suit the dynamics of the music being processed. In the case of complex mixes or vocals where the dynamics are constantly changing, the Auto mode may do a better job than fixed manual settings.

Peak/RMS operation: Every compressor uses a circuit known as a side chain, and the side chain's job in life is to measure how big the signal is, so that it knows when it needs compressing. This information is then used to control the gain circuit, which may be based around a Voltage-controlled Amplifier (VCA), a Field Effect Transistor (FET) or even a valve. The compressor will behave differently, depending on whether the side chain responds to average signal levels or to absolute signal peaks.

An RMS level detector works rather like the human ear, which pays less attention to short-duration, loud sounds than to longer sounds of the same level. Though RMS offers the closest approximation to the way in which our ears respond to sound, many American engineers prefer to work with Peak, possibly because it provides a greater degree of control. And though RMS provides a very natural-sounding dynamic control, short signal peaks will get through unnoticed, even if a fast attack time is set, which means the engineer has less control over the absolute peak signal levels. This can be a problem when making digital recordings, as clipping is to be avoided at all costs. The difference between Peak and RMS sensing tends to show up most on music that contains percussive sounds, where the Peak type of compressor will more accurately track the peak levels of the individual drum beats.

Another way to look at it is to say that the greater the difference between a signal's peak and average level, the more apparent the difference between RMS and peak compression/limiting will be. On a sustained pad sound with no peaks, there should be no appreciable difference. Peak sensing can sometimes sound over-controlled, unless the amount of compression used is slight. It's really down to personal choice, and all judgements should be based on listening tests.

Hold Time: A compressor's side chain follows the envelope of the signal being fed into it, but if the attack and release times are set to their fastest positions, it is likely that the compressor will attempt to respond not to the envelope of the input signal but to individual cycles of the input waveform. This is particularly significant when the input signal is from a bass instrument, as the individual cycles are relatively long, compared to higher frequencies. If compression of the individual waveform cycles is allowed to occur, very bad distortion is audible, as the waveform itself gets reshaped by the compression process.

We could simply increase the release time of the compressor so that it becomes too slow to react to individual cycles, but sometimes it's useful to be able to set a very fast release time. A better option is to use the Hold time control, if you have one. Hold introduces a slight delay before the release phase is initiated, which prevents the envelope shaper from going into release mode until the Hold time has elapsed. If the Hold time is set longer than the duration of a single cycle of the lowest audible frequency, the compressor will be forced to wait long enough for the next cycle to come along, thus avoiding distortion. A Hold time of 50ms will prevent this distortion mechanism causing problems down to 20Hz. If your compressor doesn't have a separate Hold time control, it may still have a built-in, preset amount of Hold time. A 50ms hold time isn't going to adversely affect any other aspect of the compressor's operation, and leaves the user with one less control to worry about.

Stereo Link: When processing stereo signals, it is important that both channels are treated equally, for the stereo image will wander if one channel receives more compression than the other. For example, if a loud sound occurs only in the left channel, then the left channel gain will be reduced, and everything else present in the left channel will also be turned down in the mix. This will result in an apparent movement towards the right channel, which is not undergoing so much gain reduction.

The Stereo Link switch of a dual-channel compressor simply forces both channels to work together, based either on an average of the two input signals, or whichever is the highest in level at any one time. Of course, both channels must be set up exactly the same for this to work properly, but that's taken care of by the compressor. When the two channels are switched to stereo, one set of controls usually becomes the master for both channels -- though some manufacturers opt for averaging the two channel's control settings, or for reacting to whichever channel's controls are set to the highest value.

ALL IN THE EAR

You may have noticed, or at least read about, the fact that different makes of compressor sound different. But if all they're really doing is changing level, shouldn't they all sound exactly the same? As we've already learned, part of the reason is related to the shape of the attack and release curves of the compressor, and of course peak sensing will produce different results to RMS, but at least as important is the way in which a compressor distorts the signal. Technically perhaps, the best compressor is one that doesn't add any distortion, but most engineers seem to like the 'warm' sound of the older valve designs which, on paper, are blighted by high distortion levels. The truth is that low levels of distortion have a profound effect on the way in which we perceive sound, which is the principle on which aural exciters work. A very small amount of even-harmonic distortion can tighten up bass sounds, while making the top end seem brighter and cleaner.

The best-sounding contemporary compressor designs include valve models with a degree of distortion built in, while others use FETs, which mimic the behaviour of valve circuits.

Academy of Sound Workshop #4 Compression Settings

As digital recorders and mixers are introduced into the signal chain, more people are becoming interested in equipment that can put the warmth back into what they perceive as an over-clinical sound.