technique | advanced compression

advanced compression techniques

Part 1: Paul White looks at different gain-control elements and their effect on audio.

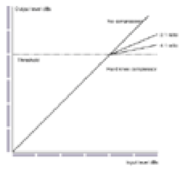

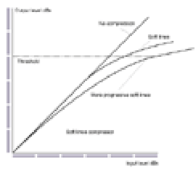

Compression is a subject that has been covered on numerous occasions in past issues of *Sound On Sound*, but it is worth revisiting, both because of the importance of compression in contemporary music production and because many people are unsure as to the best way to use their compressors. However, in order to avoid retreading old ground, I'll only be giving a very quick overview of the principles of compression before moving on to some of the more advanced concepts — if you'd like a more in-depth discussion of basic compression, then refer back to my article on compression in *SOS* April 1997.

In A Nutshell

Most of you probably know that a compressor is a device for automatically controlling the level of an audio signal. More specifically, a compressor 'turns down' the audio when the level exceeds a threshold set by the user. The amount by which the gain is turned down depends on the ratio of the compressor — for example, if a ratio of 5:1 is set, an input signal exceeding the threshold by 5dB will be output with a level of only 1dB over the threshold. Once the signal falls back below the threshold level, the gain returns to normal. It's exactly the same as manually turning the level down with a fader whenever it gets too loud, but it's much faster to respond than any human and it's totally automatic.

To make the effect of compression smooth and natural-sounding, compressors often allow attack and release time parameters to be set by the user, but just occasionally these are fully automated. The attack time determines how long the compressor takes to reduce the gain once the input signal has passed the threshold, while the release time determines how long the gain takes to return to normal after the input signal has fallen back below the threshold. If the attack and release are too fast, rapid changes in gain cause an effect known as 'pumping'. All that pumping means is that the compressor action is clearly audible rather than subtle. Because compressors work by reducing level, most models have an output control called 'gain make up' or something similar. This control is simply used to restore the peak level of the compressed output signal to that of the uncompressed input signal. In effect, this means that compression makes low-level signals louder if the peak level is returned to its former value.

The last concept to explore before moving on to the more advanced stuff is that of the 'knee'. A basic compressor does nothing to the input signal until it reaches the threshold, then the full amount of gain reduction is applied as fast as the attack time will let it. This is good for assertive level control, but can be a little too obvious when a lot of compression is being applied to critical sounds within a mix — or to complete mixes for that matter. A gentler-sounding compression can be achieved by using a so-called soft-knee compressor, where the compression ratio increases gradually as the signal approaches the threshold. Once the signal passes the threshold, the full ratio as set by the user is applied, but, because some compression is applied to signals approaching the threshold, the transition from no gain reduction to full gain reduction is far smoother. Figure 1 shows graphs of input level versus output level for both hard-knee and soft-knee compressors.

So, if soft-knee compressors are so smooth and cuddly, why don't we use them all the time? Firstly, it's sometimes nice to use compression as an effect, in which case a fairly hard compression tends to work best. A little deliberate gain pumping can give the impression of loudness and hard-knee compressors pump more readily than soft-knee types. The second reason is that, at higher ratio settings, the hard-knee compressor provides firmer gain control, so if a signal is varying in level to an excessive degree, a soft-knee compressor might not produce the required degree of levelling. The choice of which to use has to be made by ear, especially as every soft knee compressor behaves differently. Some have a relatively small knee, where the ratio increases over an input range of just a few dBs, whereas some start compressing at very low signal levels and then gradually increase the ratio over a range of 20 or 30dB. In fact some of these compressors are not so much soft knee as soft leg!

Bending The Law

As if hard-knee and soft-knee compressors didn't confuse the picture enough, there are other 'control law' effects to consider. In a theoretically perfect compressor, once gain reduction is applied (in other words, once the input is above the threshold), the response is reasonably linear, so no matter by how much the input exceeds the threshold, the output level increase will always be the fraction of that amount determined by the ratio control. Both hard-knee and soft-knee compressors settle down into this type of linear response above the threshold. However, there are some compressor types that don't exhibit a linear response above the threshold, and it's not uncommon for the amount of gain reduction actually to reduce at very high signal levels. In effect, this means that at very high signal levels the compression ratio tends to fall to a lower value, as shown in the graph of Figure 2.

Compressors that use lamps and photocells are notoriously non-linear but, rather than this being deemed a fault, it is acknowledged as one of the factors that gives them their distinctive sound. Compressors using valves within the gain-control circuitry may also be non-linear. It's not important to know a lot about the technicalities of such non-linearities — just be aware that this factor contributes to audible differences between models of compressor that might otherwise appear to have the same broad technical specification. As is so often the case in audio, 'theoretically perfect' doesn't always equate to the most musical sound.

Inside Vintage Compressors

In the days before dedicated VCAs (Voltage Controlled Amplifiers), the most common gain elements used in compressors were valves, FETs (field effect transistors) and photo-resistive devices. While these gain-control elements are not nearly as accurate as VCAs, each gives a particular sonic character to the gain reduction which has been deemed musically useful for many types of music. Though some manufacturers, for example Aphex, now concentrate on the sonic purity of highly sophisticated VCA chips with ultra-low distortion and linear control responses, ultimately the choice of a compressor that flatters or one that controls dynamics as transparently as possible is an artistic one.

Valve gain elements tend to add a certain amount of distortion — I find that this distortion sounds subjectively similar to that produced through soft limiting. Transients that are not caught by the compressor still tend to be softened by the non-linear characteristics of the valve circuitry, which is one of the reasons why valve designs are often felt to be more musical than their more accurate VCA counterparts. However, their control law is also somewhat non-linear, so some modern hybrid compressors use VCAs for the gain-control elements while using valves for the amplification stages. This can provide better control of the compression while emphasising the valves' soft-limiting sound, and such units can have a warm, musical sound when well-designed.

The Field Effect Transistor (FET) is often used as a gain-control element within more cost-effective solid-state designs. FETs have similar non-linear transfer functions to valves, and so tend to distort the signal in a similar way. LA Audio have made a number of compressors which are known for their vintage FET sound.

Photo-resistive gain-control devices are particularly interesting because they actually add very little direct distortion to the signal. Being purely resistive, they can be hooked into a circuit much like any other potentiometer. However, while they don't add distortion to the signal being processed, the non-linear control law of a combined photo-resistive device and light source gave them a unique sonic signature. The very first such compressors were designed before the invention of the LED and so used regular filament light bulbs. Compared to the dynamics of audio, light bulbs have a pretty slow turn-on/turn-off rate, so big attack overshoots were typical. Modern photo-electric compressors use LEDs along with compensation circuitry to speed up the gain change response, but there are still non-linearities in some designs that produce a musically interesting result. Modern opto-compressors are made by companies such as Joemeek and Focusrite (their Platinum range). You can approximate the sound of a vintage opto compressor using a regular VCA or FET compressor by setting a fairly long attack time (around 100mS) and a faster release than normal.

Digital compressors and plug-ins can be designed to emulate all types of hardware compressor, though how good they sound depends on the designer's understanding of the distortion and control law mechanisms of the original. Even today, experts argue over which aspects of a valve's electrical characteristics have the greatest influence on 'musicality', so do try out the real thing if you get the opportunity, rather than relying entirely on plug-ins just by default — you might be surprised how different real analogue circuitry sounds!

Dynamic Damage

Basic compressors are little more than automated faders, but sometimes their action is at odds with the way audio behaves. It's well-known that you need a lot more energy to make a loud bass sound than a loud high-pitched sound, so it comes as no surprise in pop music to discover that most of the sound energy in a mix comes from the kick drum and the bass guitar or bass synth. When you compress a mix, it stands to reason that the compressor will respond mainly to the levels of these instruments so that, whenever a loud kick drum comes along, the level of the whole mix will be reduced for a few moments. Unless the amount of compression is quite modest, this can lead to an audible pumping of the high frequencies in a mix as they are reduced in level needlessly. Setting an attack time long enough to allow high-frequency transients to pass before gain reduction occurs can help in some cases, but this isn't always successful. Furthermore, there are occasions on which a fast attack time is necessary to achieve the right overall effect.

As we shall see later, the best solution to this problem is to use a multi-band compressor, but the designers of conventional compressors have also come up with some ingenious solutions to lessen this problem. For example, some designs use circuitry that allows a small amount of high-frequency signal to bypass the compression process so that, when a loud bass sound causes a drop in the overall level, the high end doesn't get killed. Once again, the technicalities aren't as important as the results, and what I'd like to get over to you is that, when trying out any compressor, you ought to listen to the way the high end changes when heavy compression is being triggered by low-frequency sounds. The variety is enormous — some compressors sound quite dull and choked while others maintain the high end very effectively.

Peak Or RMS?

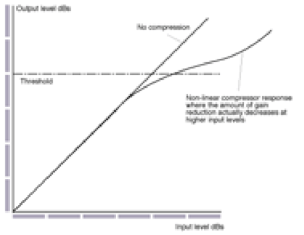

Going back to the 'compressor as a fader' analogy, the side-chain of the compressor is that part of the circuitry that listens to the incoming signal to see if it needs turning down or not. Most often, compressor side-chains are designed to respond pretty much like the human ear, which means that short duration sounds aren't perceived as being as loud as longer sounds of exactly the same level. This is called an RMS response (an abbreviation for 'Root Mean Square'), a mathematical means of determining average signal levels. The implications of using a compressor with an RMS control law are that the compression will sound natural, but short duration, high amplitude sounds may pass through at a higher level than you expect. One solution when feeding digital systems that can't tolerate overload is to use a fast acting peak limiter after the compressor.

Some compressors offer switchable RMS/Peak operation, and in Peak mode, the gain control responds more accurately to brief signal peaks than in the RMS 'averaging' mode. This ensures peaks are more accurately controlled, but at the same time introduces a greater risk that the broadband audio will be squashed unacceptably whenever a loud, short transient sound occurs. For this reason, it may be most effective to use Peak compression when treating individual drum and percussion sounds prior to mixing.

The Future Of Audio

Just like our hypothetical engineer controlling levels with a fader, a compressor can't take action until it 'hears' something that's too loud. To put it another way, a basic compressor's level corrections inevitably come slightly late. If a compressor is set to have a very fast attack time, the signal level can be brought under control before it overshoots, but even then the rise of an attacking sound will be distorted slightly by the compressor action — though, fortunately, very short periods of distortion during transient sounds are not generally audible.

One way of getting around the 'too late' problem is to use a so-called 'look-ahead' compressor, where the side-chain is allowed to see the input signal a fraction of a second before it arrives at the gain-control stage. To do this in real time would require circuitry that could see into the future, so a more practical ploy is to delay the audio passing through the gain-control stage by just a few milliseconds while the audio feeding the side chain remains undelayed. In normal situations, a delay of three or four milliseconds is insignificant to a signal, but you should be aware that any hardware look-ahead compressor will introduce a tiny delay. Figure 3 shows the block diagram of a look-ahead compressor.

Software plug-ins used to process audio that's already been recorded fare somewhat better than hardware, because they can often get a chance to read an audio file slightly in advance of playback, therefore enabling them to work without introducing any delay. It's for this reason that look-ahead functionality is far more common in software compressors than in hardware. Many traditionalists don't like look-ahead compressors, because they don't give the same result as the analogue compressors they are derived from, but in situations where transients with extremely fast attack times are present, using a look-ahead compressor may be the only way of bringing peaks under control fast enough.

Within Limits

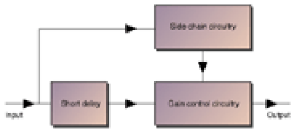

By delaying the audio slightly prior to the gain control stage, the side-chain has extra time to react to signal peaks, the advantage being that overshoots can be prevented.

We are often told that a limiter is simply a compressor with an infinitely high ratio, so that once a signal reaches the threshold, it is prevented from exceeding it. This is pretty much true, but in this digital age where even very short periods of clipping may not be acceptable, a regular compressor is unlikely to be able to act fast enough to function as an effective limiter — fast transients can pass through a system before your compressor is able to react, and this can result in clipping at your A-D conversion stage. In the days of analogue tape, this didn't matter so much, as short periods of analogue overload tended to be inaudible, but some digital systems can't cope with any clipping at all, however brief. In such situations, a dedicated, fast-acting limiter is the best bet.

In order to control signal peaks without affecting a sound's subjective level, some digital limiters may be programmed to allow a certain number of samples to clip before the level is reduced. In situations where the equipment next in line doesn't object to short periods of clipping, this can actually make the material seem much louder, though, as a rule of thumb, the period of clipping should be less than 1mS, which is equivalent to 44 consecutive samples at the sampling rate of CD-quality audio. However, if frequent clipping is expected, then the maximum length of clipped signal should be reduced to below 10 samples, as research indicates that repeated clipping within a short space of time is more audible than widely spaced instances of clipping. Some limiters emulate analogue soft clipping, where the top few dBs of any peaks are rounded off rather than clipped. Soft clipping can also help preserve the impression of loudness, though the effect can be audibly unpleasant if the signal is forced into limiting for more than very brief periods of time.

That's all for now, but in part two, next month, I'll be covering the many uses of the powerful dynamics-processing tool that is multi-band compression. SOS

Glossary

http://www.sospubs.co.uk/sos/regular_htm/glossary.htm

SOUND ON SOUND
Europe's No1 Hi-Tech Music Recording Magazine Sound On Sound

Go To SOS Website Homepage

Media House, Trafalgar Way, Bar Hill, Cambridge CB3 8SQ, UK.
Telephone: +44 (0)1954 789888 Fax: +44 (0)1954 789895
Email: info@sospubs.co.uk Website: www.sospubs.co.uk