technique mix compression

mix mastery

HOW & WHEN TO USE MIX COMPRESSION

The craft of mixing is all about controlling and balancing the relative levels of different instruments. Most of this balancing is done by riding the faders, but dynamic compression can also be a very useful and powerful tool if used appropriately. **Hugh Robjohns** explains some of the techniques and traps involved in mix compression.

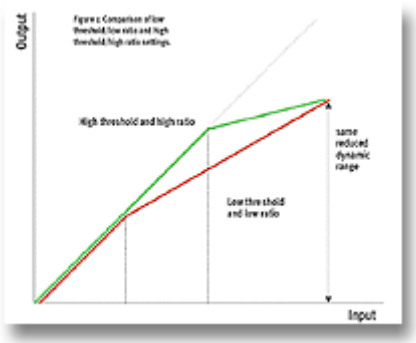

Dynamic compression is one of the most commonly used processes when it comes to mixing, probably second only to equalisation. However, there are many different kinds of compressor and even more ways of using them, each producing different end results with a variety of side effects. The purpose of this workshop is to offer some guidance on using compression in the mix, as well as to give you a few ideas for things to try on your own projects.

To start off with, let's consider the simplest kind of compression - the limiter. A limiter is just a compressor with a very high ratio (10:1 or more) and a threshold set to something close to the maximum allowable peak signal level. Limiters are great for overload protection, and possibly for getting a tiny bit of extra loudness from a mix, but not much else. Limiting always tends to sound a little unpleasant and harsh, because it has such a drastic and unnatural effect on signal dynamics. Everything below the threshold contains normal dynamic changes, whereas everything above has no dynamic change at all! In general, I would suggest you avoid using limiters during mixing (your music will see enough of them if it ever gets broadcast), but if you feel the need, try to restrict the effect to only the top 2 or 3dB of signal peaks. This will give the impression of a little extra loudness without doing too much damage.

Overall Compression

The simplest way to use a compressor on the mix is to apply it across the overall stereo mix. Ideally, this should always be done by patching it across the desk's main stereo buss insert points - if you simply plug a compressor between the desk's output and the recorder's input you can no longer perform fade-outs using the desk's master fader(s) as this will affect how hard the compressor is being driven and possibly therefore the sound or balance of the mix during the fade-out.

Whenever you are using any kind of dynamic control device on a stereo (or multi-channel) source, it is absolutely vital to always use the channel ganging mode (often called 'stereo link') so that both (or all) channels are processed in the same way. Whichever channel

contains the loudest signal peak at a given moment, the resulting gain reduction is applied equally to both (all) channels. If you don't link channels in this way, your carefully crafted stereo image will go completely awry, with musicians and sounds wandering all over the place!

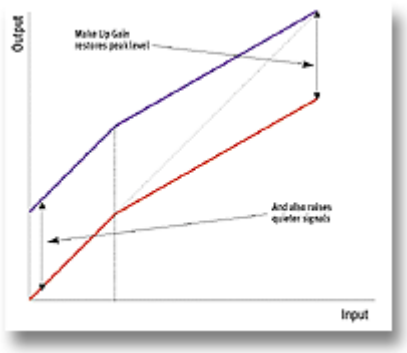

Depending on the model of compressor you are using, one channel (usually channel 1) will become the master controller for both (all) channels, but you may also have to set some of the controls on the other channels to identical settings - typically the gain make-up, threshold and ratio, but sometimes also the attack and release times. If in doubt, dig out the handbook and find out, because incorrect settings can have some very peculiar effects!

Compressors all work by pulling loud things (any signal above the threshold) down in level. The amount of reduction is determined by the ratio control in combination with the threshold, and is shown on the gain-reduction meter. Low thresholds mean the compression starts at a lower level, so there will be a lot of gain reduction. High ratios squash signals above the threshold harder, also introducing a lot of gain reduction. It is, therefore, normal to balance threshold and ratio so that low thresholds tend to be associated with gentle ratios, while high thresholds have steeper ratios (see Figure 1).

The idea of overall compression is to reduce the dynamic range of a complete mix, so the whole thing sounds that little bit louder and more impressive. However, strapping a compressor across the output of the desk and dialling in a suitable threshold and ratio will actually make it all quieter... which is not what we want at all, so we have to use the gain make-up control to bring the overall level back up again. The effect of this combined processing is to make the quieter elements of the mix louder, whilst the loud stuff stays more or less where it was (see Figure 2, right).

"If the compressor threshold is set below the peak level of the kick drum, each beat will cause the compressor to reduce the overall level..."

If the concept is simple, the application is, I'm afraid, a little more involved. The first point to note is that the gain make-up brings up *all* quiet signals - tape and mixer hiss, hums, ambient noises, and reverberation - it does not discriminate between musical and unwanted sounds, so it has to be used carefully. The second point is that the amount of compression applied at any moment in time is determined by the loudest signal in the mix. This might seem obvious, but it is very important, because whichever signals are dominant are going to modulate the level of the entire mix.

For example, in most music, the greatest signal energy lies towards the bottom end of the spectrum, and so the kick drum and bass guitar tend to be pretty dominant instruments. If the compressor threshold is set below the peak level of the kick drum, each beat will cause the compressor to reduce the overall level and the result will be a kind of pumping effect. If the release time is relatively long, this will tend to punch holes in the track as each time the kick drum is hit everything will become a lot quieter for a moment! This is probably not an effect you would want, but if the release time is set much quicker, the same process creates a different effect which might help to produce a much punchier, more dynamic-sounding mix.

This effect could also be derived from a vocal line - since the vocal is likely to be quite loud,

careful setting of the compressor will allow the vocal to dip everything else slightly whilst it is present, helping it to stand out better. Between the vocal lines, when the singing stops, the lead instruments will come back up a little in level, so there will be less 'fighting' between vocals and lead guitars or keyboards.

At The Controls

The key to making overall compression work lies in setting the compressor sympathetically. Typically, a little gentle, subtle compression is all that is needed, with a very low ratio (perhaps 1.5:1 or 2:1) combined with a low threshold, so that the compression is working most of the time over most of the song's dynamic range. I would ideally be looking for between 4 and 8dB of gain reduction at most in this kind of situation, but this is a matter of personal taste and the nature of the track. If you are after a harder, more obviously compressed effect you will need to use higher ratios such as 3:1 or maybe even 5:1, but with a much higher threshold.

Setting the threshold and ratio is only half the story, though, and the control which makes the biggest difference, and which plays a large part in determining the perceived loudness, is the release time. Most compressors have a programme-dependent 'auto' mode which works just fine, giving fast releases for large but brief transient peaks, and a slower time constant for smaller peaks. However, if your machine does not have an automatic setting, or it doesn't sound very good when selected, you will have to adjust the release time by hand. Don't expect one setting to work for the entire song, because you will almost certainly have to change it as the song progresses. It is not uncommon to have to change the release time setting for the chorus, bridge, middle eight and other sections individually, to suit the changing orchestration, pace and rhythms. In general, the faster the release time, the louder the track will sound (and vice versa) but, as always, there is a compromise. Too fast a release setting will cause unpleasant pumping or breathing effects, or just make the track sound dynamically unstable! It could also introduce distortion artefacts in bass-heavy material.

When To Compress

If you know you will want to compress a track, don't mix it to a stereo master first and then try to compress it afterwards - you will be fighting against too many compromises, with sounds interacting in the wrong ways. Far better results will be achieved if you mix through the compressor, listening to its output, so that it lends a helping hand whilst allowing you to ride the faders as necessary to control the way dominant signals interact. Since unexpected interactions can have a strong impact when using overall compression, this technique is best suited for working in an 'off-line' situation where there is no problem in starting over if it all goes horribly wrong! I would also advise against making any major equalisation changes until the compressed mix is roughly right, as these will also affect the way the compressor 'sees' dominant signals. As always, changes to equalisation will almost certainly require changes to the balance, but with a compressor across the output, those changes may not be the ones you would have expected, as sounds will interact in subtly different ways!

"...the faster the release time, the louder the track will sound (and vice versa) but, as always, there is a compromise. Too fast a release setting will cause unpleasant pumping or breathing effects..."

Multi-band Compression

One of the problems I've highlighted with overall compression is that if any one sound is dominant, that sound will drive the compressor. Sometimes this is fine and just what is wanted, but more often than not it is a problem looking for a solution. A common fix is to insert an equaliser in the side-chain of the compressor. If the kick drum or bass guitar are dominating the compressor, using the equaliser to turn down the lower frequencies present in the side-chain will reduce the compressor's sensitivity to them (effectively giving a higher threshold for low frequencies compared to everything else). Now the mid-range signals will tend to dominate, which is probably a lot more useful.

Compressing In Sections

An alternative to multi-band compression is to compress separate sections of similar instruments at group level - say, the vocals, lead instruments, bass and drums - and then combine these in the usual way after compression. The big advantage of this technique is that dominant instruments in different sections can no longer interfere with each other (although, as we've seen, this can also be useful in some circumstances), and more compression can be applied. This last point is not entirely obvious, but is true, and is probably due to the fact that in the combined mix there is no one fixed reference point for the ear to measure changes in gain, making it far more tolerant of heavier compression - the same argument also applies to the effectiveness of multi-band compressors. The main disadvantage of this approach is that you need more compressors, of course, but personally I prefer working in this way whenever possible, as I feel it gives more control and there are fewer surprises, making it far more appropriate for live working.

I would almost always compress vocals fairly hard (depending on the voices and the nature of the song), because the vocal line usually needs to remain clearly audible above the rest of the orchestration. Reducing its dynamic range to something relatively small (in keeping with the performance) is a good way to achieve that, but be wary of squashing the vocal too much as pulling up the overall level with the make-up gain could emphasise breaths and lip noises disproportionately. As always, the more highly skilled the performer, the more compression can be applied. A typical starting point would be a 3:1 ratio with a moderately low threshold and fairly fast release of perhaps 160mS.

Lead instruments can usually be compressed in a group because they rarely all play at the same time, although they will usually need different release settings unless the auto-release mode is used. The amount of compression needed is usually fairly small, say 2:1 with a fairly high threshold, just to keep a lid on the more enthusiastic dynamic sections.

Compressing drums requires a lot of care, because it is very easy to destroy the attack and crispness in the sound. This is one case where limiters actually work quite well to catch the transient peaks and allow a little increase in overall loudness. Slowing the attack time is another handy tip, as this will let the initial transient through unmolested, whilst allowing the compression to pull up the later resonance of the drums to give a fatter overall sound.

Bass guitars almost always need fairly hard compression, which is best done in isolation. A ratio of up to 5:1 is typical, with the threshold adjusted so that the majority of notes lie within a 4 to 6dB window. This is a particularly important technique for dance tracks where the bass is a driving line which needs to remain very consistent. Try a slowish attack (say 10mS) to allow the transient pluck through nicely, with a long release of about 250mS or so - beware of using too fast a release, as this can introduce pronounced harmonic distortion, with the compressor tracking the low-frequency waveform of the bass signal!

A more elaborate but far more effective solution is to employ a multi-band compressor (a further alternative, which allows an even greater degree of control, is to use several different compressors, patched into group insert points on your mixer, to compress different sections of the mix independently - see the 'Compressing In Sections' box for more details). The multi-band compressor splits the input signal into three or more frequency bands, each being processed by a separate compressor before being recombined. The advantage is that bass signals can be controlled and squashed as necessary, as can the mid-range

signals, but loud peaks in any one section will not affect the levels of the others.

As with single-band compressors, the best results will always be obtained if you mix through a multi-band compressor, rather than applying it afterwards, although this is less of an issue than with single-band overall compression, because there is going to be far less detrimental interaction going on anyway.

The multi-band approach allows rather more compression to be applied, resulting in a far louder and more dense-sounding track. You can hear the effect clearly if you compare a single or album track direct from the CD player with the same thing as broadcast by any of the national or local pop music radio stations (the effect is most pronounced with music from the '60s and '70s). Most broadcasters use very elaborate five- or even seven-band compressors on their transmitted outputs to make sure that the music is as compressed (and therefore as loud) in each band as is technically possible! From what has already been described, it is clear that this kind of severe processing is going to alter the sound of the original mix, and so a lot of recent pop music is very heavily compressed at the mixdown and mastering stages, so that what the consumer hears at home on their CD sounds more similar to the radio broadcasts (the idea being that the radio stations' multi-band compressors can't squash something that already has no dynamic range!). SOS

SOUND ON SOUND

Europe's No1 Hi-Tech Music Recording Magazine

Sound On Sound
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

Go To SOS Website Homepage