Compression Made Easy

Demystifying Compressor Controls & Parameters

Published September 2009 in Sound on Sound Magazine
By Mike Senior

Getting precisely the results you want from compressors can often be the key to a tight, modern-sounding mix. We explain what the controls are designed to achieve and how they relate to what you hear.

Few things confuse home-studio owners more when mixing than compression. It's easy to find out what compressor parameters do in the abstract, but answers to more critical (and less tangible) questions are thinner on the ground. How do you know whether or when to compress? How much compression is enough, or too much? What are the right attack and release times? This article should lend a helping hand with such questions. Instead of beginning by explaining about compressor design, as most tutorials seem to do, I'll start with some common mixing difficulties and show how the main compressor parameters provide tools to overcome them.

Dynamic Range

Possibly the single greatest challenge facing the mix engineer is finding the right balance. On the face of it, the task should be simple: you adjust the channel faders until you hear everything in the right proportion. In most styles of music, though, the chances of setting up a reliable 'static' balance like this are slimmer than a gerbil's toothpick.

Let's take the example of a lead vocal where some words are mumbled. If you set your fader so that the majority of the vocal is nicely audible in your mix, the lower-level mumbled words will start playing 'hide and seek'. If you push the level up so that the mumbled syllables come though, the rest of the vocal will eat Manhattan. No single fader-setting gives a good balance because the difference between the highest and lowest signal levels (the 'dynamic range') is too large.

Compressors remedy this by reducing a sound's dynamic range: compression will reduce the level differences between the mumbled and unmumbled words, making it easier to find a static fader setting that works. The compressor does this by turning down (or 'compressing') the louder signals so that they match the quieter signals more closely — and all it needs from you is an indication of which signals you think are too loud. Every compressor has a control for this, but it can be implemented in different ways...

Compression, Peak Reduction, Threshold & Input Gain

The simplest approach is to have a single control that makes the compressor react to more of the signal the more you twist the dial: at the minimum setting, the signal remains uncompressed; as you turn the control up, only the signal peaks are reduced in level; and as it reaches the maximum setting, all but the softest signals are sat on. This knob is sometimes called 'Compression' (on some JoeMeek and Focusrite Platinum units, for example, as well as plug-ins such as Digital Fishphones' Blockfish), but on the iconic Teletronix LA2A it was called 'Peak Reduction', a term that can be found on other hardware and software (for example Tin Brooke Tales' TLS 3127 LEA). Both terms make sense, because you get more compression (more reduction of peak levels) as you crank the knob.

Antress Painkiller takes the approach of the famous Teletronix LA2A compressor: you turn up the Peak Reduction knob to increase the amount of compression.

Cubase's Vintage Compressor uses another approach, as found in the UREI 1176, whereby an input gain control pushes the signal up against a fixed compression threshold to increase the amount of compression.

A more common approach is for compressors to have a 'Threshold' control, which works the opposite way around: you get more compression as you turn the knob down, because the threshold is the level above which the compressor considers the signal too loud. With the threshold set to maximum, very little is considered too loud, so precious little compression occurs; but set it to minimum and most things will be too loud, and the level of all but the very softest bits will be reduced.

These example compressors (Tin Brooke Tales' TLS 2095 LA and TLS 3127 LEA) sound quite different even for similar amounts of gain reduction — and you don't need to know why this is to take advantage of it.

A final control layout you may encounter is the one used on the Urei 1176LN and which now appears on many of the plug-ins it has inspired, such as Cubase's Vintage Compressor. In this design, there's a fixed signal level, above which the compressor will turn the volume down. The only way to specify the amount of compression is to adjust the input level with an input-gain control. The more you turn up this control, the more the signal exceeds the threshold, and the more compression you get.

If the action of the Input Gain control sounds rather like that of the Compression and Peak Reduction controls above, you're not wrong. They all increase the amount of compression as you turn them up. The crucial difference is that as you compress with a Compression or Peak Reduction control (or a Threshold control), the overall processed sound tends to reduce in level, while with the Input Gain approach the overall signal level gets louder. For this reason, I tend to steer newcomers to compression away from 1176LN-style processors, because the overall level increase that you get as you turn up the Input Gain control always tends to give the impression that your processing is improving the sound — even if the amount of compression is inappropriate. That aside, it pays to get comfortable with all three common control setups, so that you have the widest available choice of different compressors while mixing.

Gain-compensation Controls

Irrespective of which compressor you choose, you'll almost always find that squishing a signal's dynamic range to taste <u>will change its apparent overall level</u>. You could use the channel fader to compensate for this, but because of the large level changes that compressors can bring about, this is rarely a good solution in the real world. <u>Almost all compressors include a simple output gain control, usually called Output Gain, or Make-up Gain (or simply Gain, or Make-up), but whatever it's called, all it does is allow you to reinstate the compressed signal to roughly its former level in the mix.</u>

Having said that, there are a few 'one-knob' compressor designs with only a single Compression control. What you'll typically find if you use one of these is that the designers have implemented some kind of automatic Make-up Gain function behind the scenes, keeping the subjective level of the audio consistent no matter how much compression you've dialled in. This does make the compressor simpler to control, but these designs almost always make compressed signals feel louder than uncompressed ones which, once again, can encourage inexperienced users to dial in more compression than is necessary.

That's a lot of explanation for only two controls, but I make no apologies for that, because they can actually deal with a lot of compression tasks on their own. Furthermore, if you find compression confusing, these controls make it possible to make useful headway with any compression presets in your recording software. So before we discuss any other compression parameters or controls, let's look at how you can make the most out of what we've already covered.

Balancing, Multing & Compressor Choice

First, let me repeat myself: <u>concentrate on the balance of the tracks in your mix</u>. If the tracks balance fine as they are, no-one will arrest you for leaving them alone! <u>The trick is to wait until you spot a fader that you can't really find a suitable level for (the sound may disappear in some places, or have sections that feel too loud): that's where you may need to compress</u>. In the first instance, though, see if you can solve the level problem by splitting the audio onto two different tracks and balancing them separately. This is a common technique often referred to as 'multing'. It's easily done in most DAWs, and can head off a lot of rookie compression mistakes. Again, you may find that you don't need any compression at all to find a balance that works.

Multing can solve a lot of problems on its own, but quickly gets very fiddly if you try to use it to deal with lots of short-term balance problems (lots of single notes or words that are too loud or quiet), and this is where the automatic processing offered by a compressor can begin to complement multing. For example, you might mult out a guitar solo from the main guitar track to give it a higher fader level, but still compress that solo so that a few over-zealous notes don't pop out too far. So try multing to solve balance problems at first, but don't be afraid to let compression take over when it suits the job better.

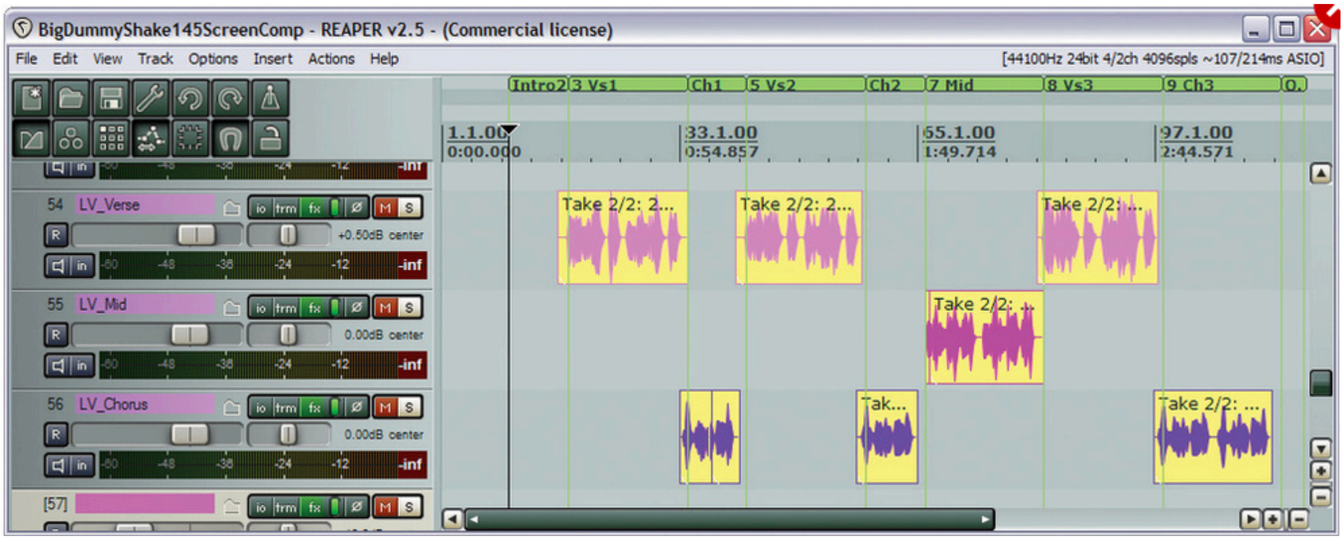

Here you can see a single lead vocal multed across three tracks to allow for different vocal processing and levels for different sections of a song. In some cases, multing tracks can allow you to avoid compressing at all, but even when it doesn't, it can still make it easier to improve your compression results, because you can better adapt each multed track's compression to the context of its section of the song.

Which compressor should you choose? At the risk of uttering studio heresy, there are more important things to worry about when starting out than the model of compressor you use. You might as well use whichever one comes to hand, albeit bearing in mind the advice I offered above regarding 1176LN-style or one-knob designs. It's also worth finding one that has a gain-reduction meter of some type, because this helps you to see when and how hard the compressor is turning down the level of louder signals.

Gain-reduction displays are typically in the form of VU-style moving-coil meters or LED bargraphs, and sometimes the compressor's normal level-meter can be switched to show gain-reduction instead.

Starting To Mix With Compression

Now insert your chosen compressor into the channel in question, and if there are presets available, select something likely-looking — again, there's no need to give it too much thought for now, just go with your gut. To start with, pile on a fair dollop of compression using the Threshold control, so that the gain-reduction meter (usually calibrated in decibels) shows at least 6dB of compression occurring on signal peaks. Once this is set up, adjust the Make-up Gain control to compensate roughly for any overall change in signal level caused by the compression. (For the sake of discussion I'll refer to Threshold and Make-up Gain controls, but the same principles apply with a different control set.)

Now's the time to return to the main question: can you now find a level for the channel fader that delivers a better mix balance? There are a lot of possible answers to this question, so let's look at how you deal with each in practice. Clearly, if your compression solves your balance problem, the job is done, but even if you think that this is the case, it makes sense to try turning the threshold back up a little and seeing how little compression you can get away with. Pushing your channel compressors too hard is a common mistake that can slowly suck the life out of a mix if it's duplicated across all your tracks, so it pays in the long run to be a little wary.

With its pretty simple interface, Tin Brooke Tales' freeware TLS 3127 LEA plug-in is a good starting point. [for Windows]

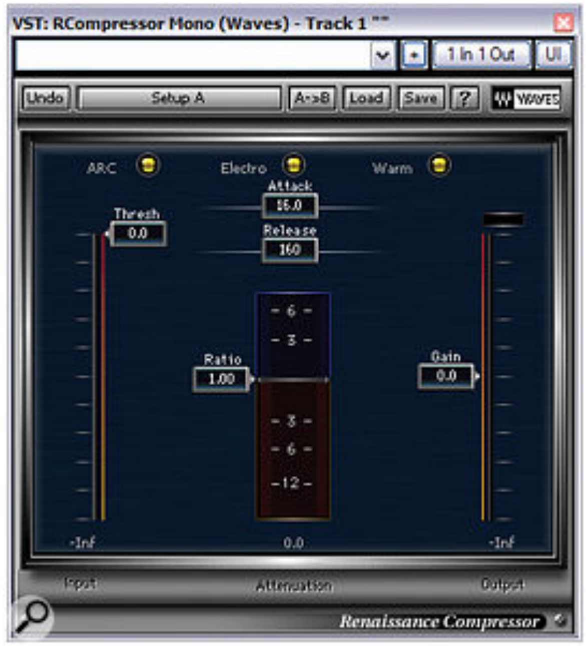

Waves' Renaissance Compressor takes the approach of SSL's popular bus compressor, where you get more compression as you bring the threshold down.

If the balance problems can't be solved, try rolling the threshold down further, to see if that makes it easier to find a decent fader level. Feel free to completely max out the control if you like, even it if makes the result sound rather unnatural for the moment: the important thing is to keep concentrating on the balance, and whether the compression can deliver the static fader-level you're after. Again, if you can get the balance you're happy with, and you find any side-effects of the compression appealing (as they often can be), then consider yourself a hero, and turn your attention to the rest of the instruments.

On the other hand, although you may find an appropriate balance through heavy compression, you could find that the processing isn't doing nice things to the instrument's sound. Perhaps it's making the performance feel lumpy and unmusical, or altering the tonality in some unpleasant way. The remedy? Just switch to a new compressor or preset and have another crack. Different compressors and presets can respond very differently for similar settings of our two main compressor controls — and you don't really need to understand why this is to reap the benefits. Try out a few different ones and choose the one that does the job best. With a little experience, you'll soon build a shortlist of personal favourites for different instruments. In fact, if you're working with a computer sequencer, you may well be able to set up several different compressors in the signal path for comparison purposes, bypassing them as necessary.

When Compression Is Not The Answer

There are almost always cases where no matter which compressor you use, or how you set the threshold, you can't find a good fader setting for the track in the mix, even if you've already done some sensible multing. This is the point at which a lot of inexperienced engineers throw in the towel and simply settle for a compromise between dodgy balance and unmusical processing side-effects. What you need to realise, though, is that your mix is probably trying to tell you that simple compression is not what you're looking for. Furthermore, if you've already tried a few different compressors and/or presets, then it's pretty unlikely that any of the other compression controls are actually going to salvage the situation. So the best thing is to step away from the compressor with your hands in the air, and look for alternative processing options instead.

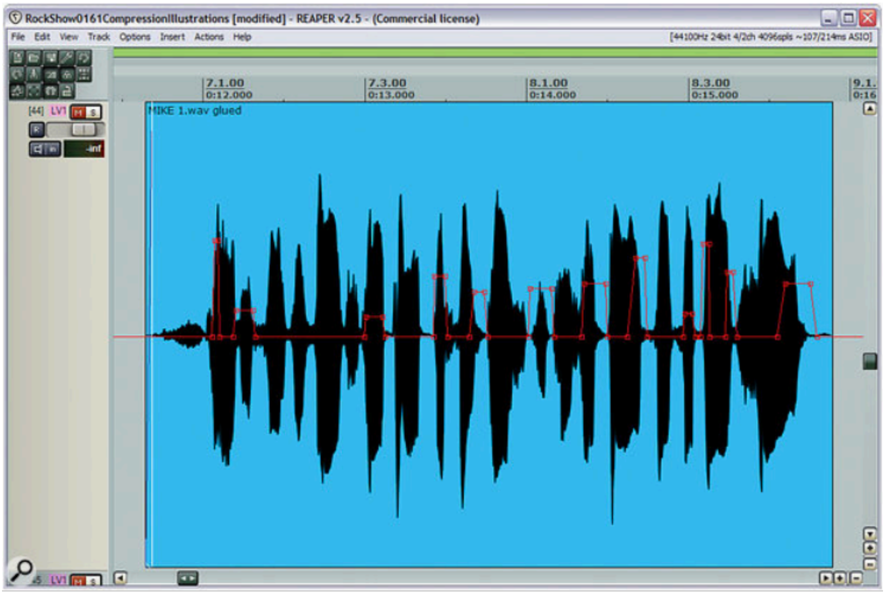

Compression is usually not enough to deliver the kind of up-front lead-vocal sound that many modern styles demand, so rather than trying to push your vocal compressor too hard, finesse any final balance tweaks using careful level automation. Here you can see some fairly typical level automation [overlayed lines with dots] for the lead-vocal phrase featured in the audio files that accompany this article on the SOS site. [Automation is a way for you to tell the software when to turn up or turn down the effect.]

An article about compression isn't the best place to go into all the other processes you might use at mixdown, but here's one example to demonstrate what I'm talking about. If you have a bass guitar recording with loads of very low frequencies, it'll be difficult in most mixes to find a fader level where the bass is audible enough in the mid-range without absolutely swamping everything else at the low end at the same time. No matter how much you compress that sound, you're unlikely to solve the problem because you won't be fundamentally changing the balance of the instrument's frequency content. It's much better to address this problem with EQ first. You'll be able to tell when you're doing the right things with the EQ when it starts getting easier to find a suitable fader level for the bass, and you might, again, discover that you don't need any compression at all.

Another very common occasion where compression can't provide a complete solution to mix balance issues is when dealing with very critical tracks, such as (typically) lead vocals. Commercial expectations for the audibility of lyrics are very high, and compression, no matter how expertly set up, is simply not an intelligent enough tool to keep a lead vocal exactly where you want it throughout most mixes. If you try to keep your vox parts up-front and audible in a mix entirely with compression, they'll usually sound over-processed, and it's a better tactic to keep the compression within musical-sounding limits before dealing with fine, moment-to-moment level tweaks manually, by moving the vocal fader during the mix. All the main sequencers now have good fader automation systems, allowing you to edit and refine fader moves until they sound exactly right, so if you're after the best vocal intelligibility possible, you should make a point of learning how these facilities work in your own software. Before I move on, let's quickly recap what we've covered so far.

All Those Other Controls

So why do we need all the other controls? If you've already taken the opportunity to try out a variety of compression presets on the same sound, you'll have noticed that some work more effectively in evening out the levels of the instrument in question than others, and that's because the deeper compression parameters in each preset tweak a variety of more subtle aspects of the compressor's gain-reduction action. If you can learn to adjust these parameters for yourself, you can match the compressor's action more closely to the specific dynamic-range characteristics of the input signal and more effectively achieve the static fader level you're looking for.

Although the technical raison d'être of compression is gain-reduction, compressors also change the tone of processed signals quite a lot, even when compressing comparatively little. So if you like the general tone of a compressor, but you can't find a suitable preset for the instrument you're processing, it's useful to be able to tweak the gain-reduction action manually to suit. And once you get some practice working with all the extra controls, it actually ends up being quicker and easier to set them up from scratch anyway.

So let's introduce some of the more advanced controls and look at how each can be used to adapt the compressor to specific tasks. As a first example, let's consider a slap-bass part. Now, as everybody knows, the best processing for slap bass is that button labeled 'Mute', but let's assume for the moment that this option has been ruled out… This particular slap-bass part is nice and dynamic and balances fine with the rest of the track, except that the odd slap note really takes off and leaps out of the track. You only want to turn down the sporadic signal peaks — but you want to turn them down pretty firmly in order to match the levels of the rest of the bass part.

Fig 1  Fig 2  Fig 3 

Fig 1: To see how the compression ratio control can work, let's take the example of a slap bass part, the waveform envelope of which might appear as shown here. The big spike is where a slap note has created a large level surge in an otherwise comparatively even line.

Fig 2: Setting a compression threshold above the majority of the note peaks allows you to compress just the rogue slap note, but if you used a normal moderate compression ratio (as in this waveform envelope) you wouldn't be able to contain the spike as well as you might like.

Fig 3: Increase the ratio higher, though, and the gain-reduction will stamp down much more firmly on the offending level spike, preventing it from leaping out unduly within the mix.

What compressors do is reduce the amount by which a signal level *exceeds* the compressor's threshold level, so in this case you want your compressor to put up a proper fight and all but stop the input signal from exceeding the threshold. That way you can set the threshold just above the level of the majority of the bass part, and it will then kick in at full force only when the over-zealous slap notes hit.

By contrast, imagine an electric guitar part where there are no dramatic level spikes, but where the overall dynamic range is still militating against finding a static fader level. You want your compressor to act more gently on signals overshooting the threshold level, so that you can set the threshold just above the level of the softest notes and then subtly squeeze the whole dynamic range down to a more manageable size.

Fig 4 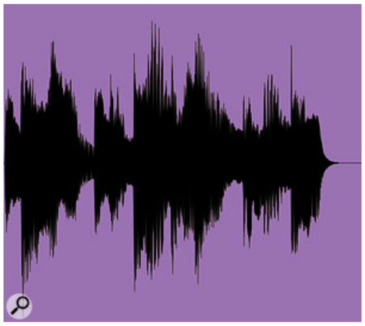 Fig 5 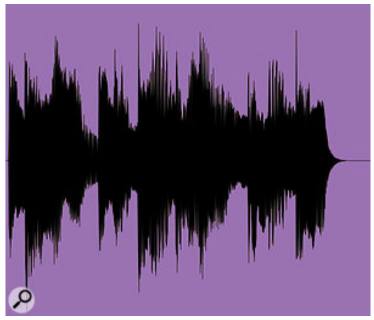 Fig 6 

Fig 4: In contrast to the slap-bass example, lower ratios tend to be better for instruments which have good musical dynamics, but simply have too wide a dynamic range.

Fig 5: Let's assume that the waveform envelope at the top represents such a part, then compressing with a low ratio can be used to gently squeeze the dynamic range such that it will maintain its position in the mix balance.

Fig 6: However, if the ratio is set too high, as in this bottom waveform envelope, the compression will iron out the part's internal performance dynamics and render it unmusical.

Ratio

It's a compressor's Ratio control (sometimes labeled Slope) that allows it to tackle these two contrasting problems, effectively setting how firmly the compressor reins in signals that overshoot the threshold level. At low ratio settings (something like 1.5:1) the overshoots are nudged politely back towards the threshold, whereas at higher settings (12:1, for instance), overshoots are beaten back by club-wielding thugs. At the highest settings — and some compressors offer infinity (or ∞):1 — overshoots are effectively stopped in their tracks, unable to cross the threshold at all. So for our slap bass, we'll be looking for high ratios, and for routine dynamic-range reduction tasks (like the electric guitar example) the lower ratios (up to about 3:1) will fix balance problems in a more natural-sounding way.

When I'm talking about a ratio of 3:1, for example, you might wonder what that figure actually means. Put simply, for every 3dB by which the input signal exceeds the threshold, only 1dB will be allowed to pass by the compressor. I could give you some lovely graphs, but I don't think it'd be a lot of practical help, because some compressors don't label their Ratio controls and different compressors can react quite differently for the same Ratio setting. A much more practical and intuitive approach is simply to use a compressor with a gain-reduction meter so that you can see when and how much the compressor is working as you juggle the Threshold and Ratio controls.

In the case of the slap bass, you'd set the ratio up fairly high to start with, and then find a threshold setting that caused the gain reduction to kick in only on the slap peaks. Once you'd done this, you'd listen to ascertain whether you'd solved the balance problem, and then adjust the Ratio control accordingly. Still too much slap? Increase the ratio to stamp on the peaks more firmly.

With the electric guitar example, you might start off with a fairly low ratio (maybe 2:1) and then set the threshold so that gain-reduction happens for all but the quietest notes. With the threshold in roughly the right place, you could then turn back to the ratio control and tweak it one way or the other to achieve your static fader level. If some quieter notes are still too indistinct, increase the ratio to reduce the dynamic range further. Why not just max out the Ratio control? The danger is that if you turn it up too high, you'll iron out the important performance dynamics that make the part sound musical, leaving it a bit flat and lifeless — so try to turn up the Ratio control only as much as is required to get the balancing job done.

At this point you might very well ask: what would I do if that slap-bass part needed not only high-ratio control of the slapped notes, but also more general low-ratio dynamic-range reduction? The answer is that you could deal with the problem by chaining more than one compressor in series. This is quite common in commercial practice, and lets you dedicate each specific compressor to a different task. If you're wondering what order to put the different processors in, though, the answer isn't quite as clear. The best solution is to try both ways and choose the one that best resolves the balance.

Why Attack & Release Matter

Up to this point, we've been dealing with the compression controls that are the easiest to get a handle on, but when it comes to a compressor's Attack Time and Release Time parameters, a lot of newcomers quickly become confused. So let's once again examine some examples of real-world balance problems to illustrate the practical purpose of these controls.

Let's say that we're mixing a song where a strummed acoustic guitar has a nice, natural sustain that works really well when it's at the right level in the mix, but you find that you have to turn the fader down whenever the player digs in more during your song's choruses. "Sounds like a job for Captain Compressor!", you cry, but when you actually start dialing in the processing you find that, rather than just reducing the level differences between song sections, the compressor seems intent on evening out the much shorter-term level differences between the attack-transient and sustain portions of each strum. Although you can sort out your overall balance problem, you're having to pay an unacceptable price: the impact of each strum is softened, or the instrument's sustain is over-emphasized.

The Attack and Release controls provide a remedy to this ailment, because they determine how quickly the compressor's gain-reduction reacts to changes in the input signal level: the former specifies how fast the compressor reacts in reducing gain, while the latter specifies how fast the gain reduction resets. The reason why the compressor in our example isn't doing the job it's being asked to do is that it's reacting too fast to changes in the signal level. In other words, its attack and release times are too short. Increase these and the compressor will react more slowly, which means that it's likely to deal with this particular balance problem more effectively, because it'll track longer-term level variations (such as those between our verse and chorus) rather than short-term ones (such as those between the individual strum transients and the ringing of the guitar strings between them).

If you look at these controls' legending [markings], you'll notice that the times are usually expressed in milliseconds, although you do occasionally find microseconds and whole seconds. However, as with the Ratio control, I wouldn't recommend getting too hung up on exact numbers, because they're only ever a rough guide to how a specific compressor responds in practice. A much better tactic is to focus on finding the best balance with the fewest unmusical side-effects, adjusting the Attack and Release controls by ear. A compressor's gain-reduction meter can be a very good visual guide here, as it'll show you not only how much compression is being applied, but also how fast it's changing. Because the effects of gentle compression can be subtle, the visual feedback from the meter can be a great aid in setting up the controls.

Snare Compression: Three Different Settings

The ability to adjust Attack and Release controls independently significantly increases the range of balance problems than can usefully be tackled, so let's look at another common example: a snare-drum backbeat. Set the attack time too fast and the compressor will respond quickly to the fleeting initial drum transient, reducing the gain swiftly. If you then set the release time very fast, the gain reduction will also reset very rapidly — well before the drum sound has finished, such that the lower-level tail of the drum hit won't be compressed as much. The drum transient will be de-emphasised relative to the overall snare sound.

Fig 7 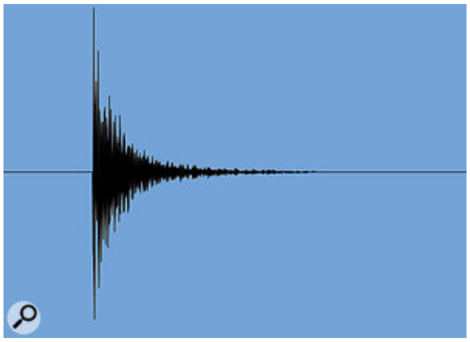

Fig 8 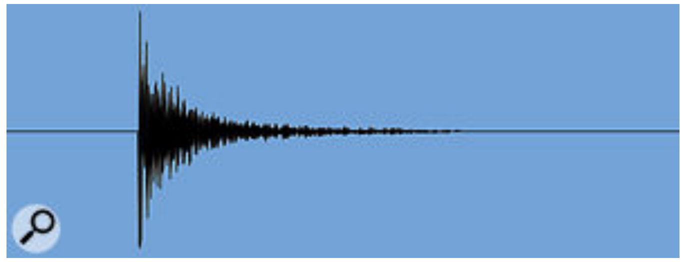

Fig 9 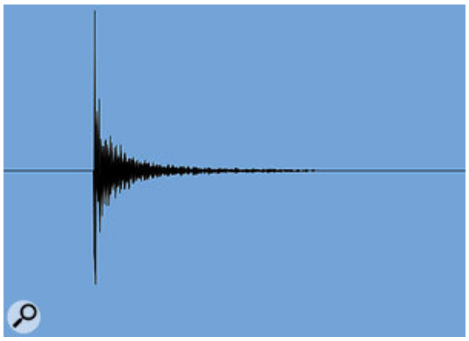

Fig 10 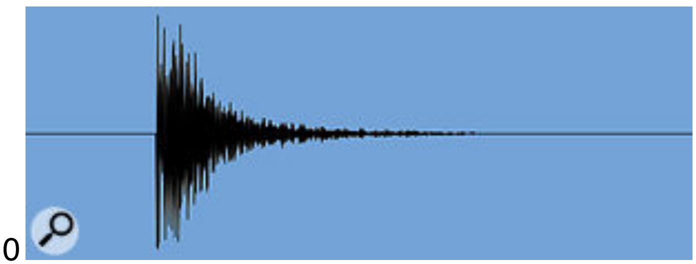

Fig 7: A compressor's attack and release times can have very different effects on the waveform envelope of a snare-drum hit. Here's the unprocessed snare hit.

Fig 8: Here you can see the transient-suppressing effect of very short attack and release times.

Fig 9: The third waveform shows how combining a fast attack with a slower release gives you pretty much just an overall level change, with little change in the nature of the snare sound.

Fig 10: In the last waveform, increasing the attack time a little has boosted the level of the initial percussive transient in relation to the sustain.

On the other hand, if you partner your fast attack with a slower release, the gain-reduction will reset very little during the drum hit itself, instead resetting itself mostly between the hits, so the balance between the transient and sustain phases of the drum will remain pretty much unchanged. The compressor in this case is simply making the level of each drum hit appear more consistent. However, if you then increase the attack time, you'll find that some of the drum transient begins to sneak past the compressor before its gain reduction clamps down, effectively increasing the level difference between the transient and the rest of the snare sound.

So the attack and release controls have made possible three different balance results — less transient level; more consistent hit level; and more transient level — all from the same compressor. This ability to achieve very different effects is partly what confuses some newcomers to compression, and it's also one of the reasons why promisingly-named compressor presets often don't do the trick: if your 'Snare' preset has been set up to reduce the drum transient, it won't help if you actually need more transient in your mix!

Side-effects

Although thinking in terms of balance answers most questions about attack and release times, in certain circumstances you may find that some settings produce unwanted side-effects. The first problem occurs when you set a fast enough attack and release that the compressor begins to react to individual waveform cycles, rather than the overall signal-level contour. The gain reduction then effectively changes the waveform shape, producing distortion — the nature of which will depend on the sound being processed and the compressor you're abusing. Bass sounds, with their slow-moving waveforms, are particularly prone to this, but delicate acoustic instruments can also present difficulties because they'll ruthlessly expose the smallest of distortion artifacts.

Another common problem is with percussive bass sounds, such as kick drums, which can appear to lose bass content if you compress them with attack times under about 50ms. This is because the compressor begins clamping down during the first couple of waveform excursions, which seems to affect lower frequencies more than higher ones, shifting the tonal balance of the sound. Once you know this is a danger, it's not that tricky to avoid, but if you're not listening for it, it's easy to miss while you're concentrating on balance issues.

One final thing to say is that changing the attack and release times will affect the amount of gain reduction that's occurring for a given combination of threshold and ratio settings. For example, a side-stick sound (which comprises a short transient and very little sustain) might completely bypass a compressor that has a long attack, even if its level shoots way over the compressor's threshold. So it's not uncommon to keep adjusting Threshold and Ratio controls alongside your attack and release, to take account of these kinds of changes.

Summing Up

There's more to using compressors than just solving mix-balance problems, but until you're confident with their fundamental gain-reduction properties, their more advanced and creative applications will be a bit baffling. I hope this article has clarified these basic functions in such a way that you can start putting them to use sensibly straight away, while avoiding the most common processing mistakes. Once you're confident of which controls you need to reach for in any given case, you'll find that the more subtle differences between compressors begin to become more relevant, and that the purpose of more advanced multi-band and parallel-processing techniques becomes more logical. All of which can be exciting and interesting stuff — but that's a very different article!

Former SOS Reviews Editor and regular contributor Mike Senior has worked professionally with artists such as The Charlatans, Nigel Kennedy, Therapy, and Wet Wet Wet. He now runs Cambridge Music Technology, delivering modular training courses based on the studio techniques of the world's most famous producers.
www.cambridge-mt.com

Compressor Or Limiter?

Compressors that are specifically designed to offer very high-ratio compression are often called limiters, so if you find that your compressor simply can't muster a high enough ratio to do a particular job, don't be afraid to try a limiter instead. If you do, though, you'll probably find that it uses the 1176LN-style input-gain control setup, and in some cases the threshold may be set to the digital clipping point for mastering purposes, without any post-compression gain control. This means that you can end up sending your overall signal level into orbit before you've brought about the gain-reduction you require. It's pretty easy, though, to add another little utility gain plug-in (GVST's free GFader, for example) after a limiter, to bring the overall level back down to earth.

Which Parts Do I Need To Compress?

Each mix will be different, but some instruments are more likely to need dynamic-range control than others. Top of the list are vocals, because although they naturally have a very wide dynamic range, they're the main carrier of the vital melody and lyrics in most mixes and thus actually need to maintain a very small dynamic range. Even in natural-sounding acoustic mixes, some control of vocal levels will usually be required, and although it can be achieved entirely though fader automation, compression typically plays some role.

Bass parts are also usually compressed. Bass guitars can have quite a wide natural dynamic range, but even where the dynamics are already quite restricted compression is quite commonplace because of the importance of controlling the levels of low mix frequencies. Pianos often present problems, not just because of their wide dynamic range, but because the complexity and purity of their sound tends to expose compression side-effects. I've recently discovered that scratch DJs are also tricky customers, because the details within scratching parts tend to be as important as the higher-level signal peaks.

Some parts often need no compression at all. Anything heavily distorted will already have been leveled out by nature of the distortion process. Electric guitar parts can often be left uncompressed. In fact, compressing them can sometimes remove the last remaining vestiges of musical dynamics. Synths can frequently be left to their own devices, too, particularly the more static, pad-like sounds.

EQ & Compressor: Which First?

If you use a chain of multiple processes on an instrument, you might wonder where you should put the compressor. Equalization is primarily about changing signal levels, albeit in carefully specified frequency regions, so pre-compression EQ [EQ before the compressor] can alter the gain-reduction action of the compressor, but post-compression EQ won't.

If you're happy with the way your compressor is working, just put any equalization after it in the processing chain, but if you find that frequency-based problems make it difficult to achieve the compression you want, dealing with this problem via pre-compression EQ makes sense. For example, extreme low-frequency thumps from a vocalist tapping their foot on the mic stand can play havoc with attempts to compress the vocal itself. Filtering out these low-frequency level peaks with EQ, pre-compression, can immediately make the compression sound much more predictable. Sporadic low-frequency resonances from acoustic guitars or guitar/bass cabs can also be tackled in this way.

Audio Examples On-line

It can be difficult reading about compression without hearing what it does. We've placed a number of audio files to accompany this article on the SOS web site, with a detailed explanation of what they demonstrate.

Published September 2009