

LATEST SQUEEZE

The Real Benefits Of Parallel Compression

Parallel compression is a powerful mixing technique, but it's often misunderstood. Read on to find out what it really does — and how it can help you make better mixes.

HUGH ROBJOHN

When talking with sound-engineering students of all ages and experience, I often find that one area where most struggle is compression. Perhaps that's

not surprising, since the concepts are fairly challenging on their own — before even contemplating the subtleties of different kinds of physical implementation (or modelling). Compressors are available in myriad different forms, using feed-forward or feedback control paths, with RMS- or peak-weighted side-chains, and even more variations of audio-attenuation device, such as opto attenuators, diode rings, vari-mu valves, solid-state VCAs... and the list goes on!

Let's leave such complexities aside for the moment, though, because in this article I really want to try to clear up some common misunderstandings about the increasingly fashionable technique of parallel compression. Of course, when discussing an inherently complicated topic such as compression it's vital that we establish a common framework of terminology, so that we all understand what we're talking about — so that's where I'll begin.

The first thing to establish is what a compressor of any form actually

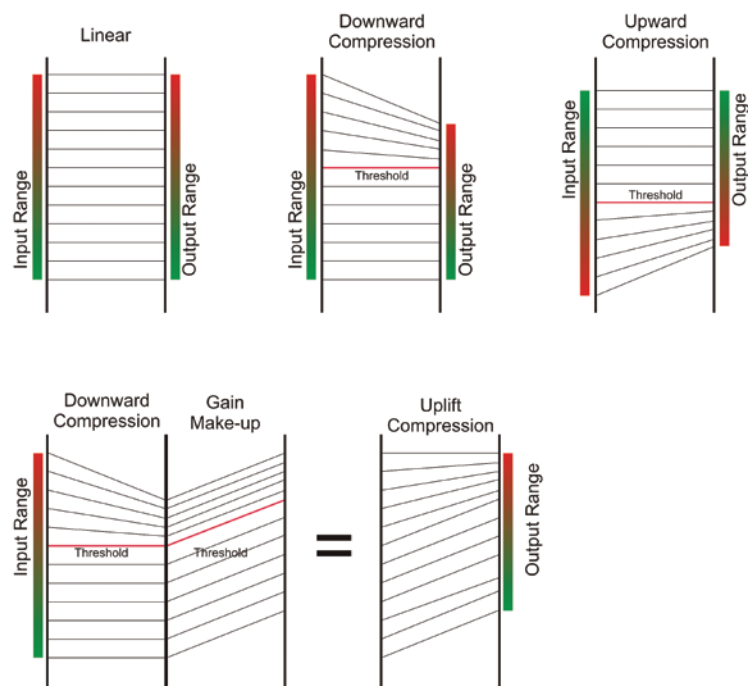


Figure 1: In these diagrams you can see the difference in effect on the dynamic-range reduction and output level of the various different basic types of compression discussed in this article. As you can see, 'upward compression' and 'uplift compression' do not achieve the same result, even if they do achieve a similar degree of dynamic-range reduction.




Figure 2: In this plot, you can see the effect of different ratios (the light blue is 2:1, the light green 40:1) of downward compression at the same threshold. The greater the ratio, the more attenuation is applied when the signal exceeds the threshold, and anything above 20:1 is usually considered to be 'limiting'. The red/green bars show just how compressed the dynamic range is.

does, and the answer is that it reduces the dynamic range of the input signal. Whether it's configured to make the loud bits quieter, or the quiet bits louder, fundamentally it exists to reduce the overall dynamic range from something large and unmanageable to something smaller and more appropriate for the intended application.

The term 'dynamic range' can be interpreted in a variety of ways, too, but for the purposes of this article I'm using it to refer to the range of average signal levels between the quietest musical element and the loudest — the difference between the *ppp* and *fff* sections of a performance, if you like.

Downward Compression

The vast majority of compressors apply 'downward compression' which means, in essence, that loud stuff is made quieter. More specifically, signals below the threshold level are left alone, while those above are 'squashed' by an amount determined by the ratio setting.

The 'ladder diagram' (Figure 1) is a convenient way of describing the relationship between signal levels at the input of a compressor, and those at the output. The ladder at the top left shows the situation in bypass — the input signal appears at the output at exactly the same level, so all the 'rungs' are horizontal. The ladder to its right shows

the effect of downward compression, where signal levels below the threshold remain unchanged, while those above are reduced in level proportionally, according to the selected ratio (in this case 2:1).

The more conventional way of illustrating compression is with a 'transfer plot', which is a graph with the input level on the horizontal axis and the output level on the vertical axis. The graph in Figure 2 was obtained by measuring the amplitude response of a 'hard-knee' compressor plug-in in a DAW, using an Audio Precision test system. The dotted straight red line at 45 degrees shows the response with the compressor bypassed — clearly illustrating that what goes in comes out, unchanged in level!

The different coloured solid lines were obtained with the compressor switched in and the threshold set to -20dBFS. The light-blue line is the result of a 2:1 ratio, and it clearly shows that when the input is 10dB above the threshold (ie. a 'Generator Level' of -10dBFS on the horizontal axis), the output or 'Measured Level' on the vertical axis is at -15dBFS, which is 5dB above the threshold. In other words a rise of 10dB at the input results in a rise of only 5dB at the output — which is half as much, and hence a compression ratio of 2:1.

The other traces show progressively 'stiffer' ratios, of 3:1, 5:1, 10:1, and 40:1. Anything above 20:1 is generally referred to as 'limiting' because the output level

barely rises regardless of how much the input level exceeds the threshold.

Incidentally, a 'hard-knee' compressor, like the one used for these measurements, switches abruptly from doing nothing to squashing the audio, and that is revealed by the very distinct change of angle on the transfer curves. A 'soft-knee' compressor moves more gently from inaction to action, so its transfer plots would curve smoothly away from the 45 degree linear slope instead of diverting abruptly.

Dynamic Range Reduction

Now let's imagine a musical signal where the quietest element measures -35dBFS and the loudest is -5dBFS, so that we have a starting dynamic range of 30dB. If we were to pass that signal through a 2:1 compressor with a threshold at -20dBFS, the output signal will range between -35dBFS (this level is below the threshold and thus unchanged) and -12.5dBFS. The latter figure arises because the source peak level (at -5dBFS) is 15dB above the threshold, and thus will be reduced by half to 7.5dB above the -20dBFS threshold, which is -12.5dBFS).

Therefore the dynamic range has been reduced, in this case from 30dB to 22.5dB, and at the same time the peak level has been reduced in level by 7.5dB. To help illustrate those points, I've added coloured bars to the previous plot to illustrate how the dynamic range and peak levels have been reduced.

Useful though this form of compression is, often we want to reduce the dynamic range without reducing the peak level. In other words, we want to raise the level of quieter signal components rather than turn down the loud ones. The usual way to achieve this is to introduce 'make-up gain' at the output of the compressor. A good way of understanding the concept is to return to the ladder diagram.

I prefer to call this 'uplift compression' to avoid confusion with true 'upwards compression' (which I'll come back to in a moment). As the lowest ladder diagram in Figure 1 shows, the input signal is first downwardly compressed in the usual way, but then the output is raised in level by a fixed amount of 'make-up gain'. The overall result is that the dynamic range has been reduced, again, but this time the peak level is restored to the same as the input while lower-level signals have been raised — so the quiet bits have been made louder.

»

» However, this diagram reveals clearly that it is still the louder elements that have been 'squashed' — the quieter signals have simply been raised in level. That's the key difference between 'uplift compression' and 'upwards compression'. By the latter term, I mean processing that squashes the quieter elements while leaving the loud bits alone, as the final diagram in Figure 1 illustrates.

Uplift Compression

So, using the same conditions as the previous measured example, let's say that we want the same amount of dynamic-range reduction, but we also want the peak level to remain at -5dBFS. This is easily accomplished by keeping the same compressor ratio and threshold settings, but dialling in 7.5dB of make-up gain. The real-world result of doing that can be seen in the measured plot in Figure 3. The direct effect of adding 7.5dB of make-up gain is to raise the overall output level, which can be seen as the vertical shift of the 45-degree linear part of the transfer curve. In this case an input signal at -35dBFS emerges from the compressor at -27.5dBFS, and all of the quieter elements in our musical signal are raised by the same amount of 7.5dB.

Louder elements — those above the -20dBFS threshold — are compressed, as before, by a ratio of 2:1, but because the make-up gain is counteracting the compressor's peak level reduction, the highest input level appears at the output at the same -5dBFS.

Upwardly Mobile

Normal downward compression — whether it's used on its own or with make-up gain — inherently changes the character of loud signals to some extent by squashing them. Rule number one for any downward compressor is to squash anything loud! However, the action of turning the level down (and back up again afterwards) isn't instantaneous; it takes place over a timescale that is governed by the compressor's attack and release time constants. The inevitable result is that the sound and shape of complex but delicate and loud transient signals can be altered quite drastically. This is a significant part of the reason why different compressor designs can sound so different from each other, and why one compressor may be preferred in a given situation over another.

An alternative way of reducing the dynamic range, which could potentially




Figure 3: By adding make-up gain to downward compression, you shift the dynamic range, to give you greater level at the compressor's output, but the range itself is still compressed in the same way.

avoid this issue of damaged transients, is to compress the quieter signals and leave the loud ones alone. In other words, we could employ a device which turns the level up for signals below the threshold. This is true 'upwards compression' as described briefly above, and this kind of arrangement would — in theory, at least — preserve loud but delicate transient signals completely intact, while only squashing the quieter elements.

Unfortunately, though, it's extremely rare to find a physical device that can

actually do this, not least because it isn't particularly easy to achieve this goal with electronic hardware. (The Dbx Quantum, I believe, was one such device — but I'm not aware of any in current production.) Nevertheless, if I were to find and measure such a device, the transfer curves might look something like those I've mocked up in Figure 4, where loud signals above the threshold would remain completely unaffected and quiet signals would be turned up in level by an amount depending on the ratio setting (the




Figure 4: As you can see, true upward compression, shown on these plots, has a very different effect from downward compression, in that the loudest parts of the signal are left alone, and the quiet ones raised. There's still dynamic-range compression, but a different part of the signal is affected.

light-blue line being a 2:1 ratio, again). Once again, if the upwards compressor had a 'soft-knee' characteristic instead of the hard knee I've illustrated, the plots would curve gently upwards instead of changing angle abruptly. However, if you spend a few moments thinking about the idea of genuine upwards compression, you'll realise that it wouldn't actually be a lot of use in real-world applications, because it would probably raise the noise floor to wholly inappropriate levels. For example, take an upwards compressor with a gentle 2:1 ratio and a threshold at -20dBFS. If the ambient noise floor of the recording was at, say, -80dBFS, the act of compression would drag it up to an appalling -50dBFS, which probably isn't really what you'd want!

Parallel Compression

"Given the difficulty of finding a true upwards compressor," I hear you respond, "is there an alternative way of achieving something similar?" I'm glad you've asked that, because yes, there is. The idea of parallel compression was conceived to achieve much the same goal as upwards compression. In other words, its aim is to leave the delicate loud transients intact while raising the level of low-level signals — but to do so employing only the standard downward compressors found in conventional studios.

Given that description, you might logically assume that parallel compression somehow achieves a form of upwards compression — and a great many people do, indeed, often refer to it in that way (myself included on occasions). It's all a bit more complicated than that, though, because it really depends on how the parallel compressor is set up. In strict technical terms, a parallel compressor setup is actually a form of downward compressor, but it exhibits a unique and very useful characteristic that allows it to behave exactly like an upwards compressor over a *defined dynamic range*.

Before we examine real measurements of different parallel-compression configurations, it's probably worth revisiting quickly how parallel compression works. The basic concept is to split the input signal to feed two paths, one being a direct 'through' path, and the other feeding a normal downward compressor. The compressor's output is mixed at unity gain with the direct path to produce the 'parallel compressed' signal.

The compressor is set up with a very high ratio — in fact, limiting is often used — and the threshold adjusted so that it provides a lot of gain reduction when the input signal is at its loudest. The more gain reduction the better, in fact: 20dB is a good start, but 40dB would be better! Of course, this requires a compressor that behaves nicely when applying a lot of gain reduction and doesn't generate any distortion products, which some do. The attack and release controls are set to suit the material and the effect required, in the usual way.

When configured for parallel compression in this way, the contribution from the compressor during the loudest signal peaks will be well over 20dB quieter than that of the direct path, because of the massive gain reduction it is being forced to apply. This means that at those points, its involvement in the mixed output signal is virtually insignificant; the output signal is completely dominated

»

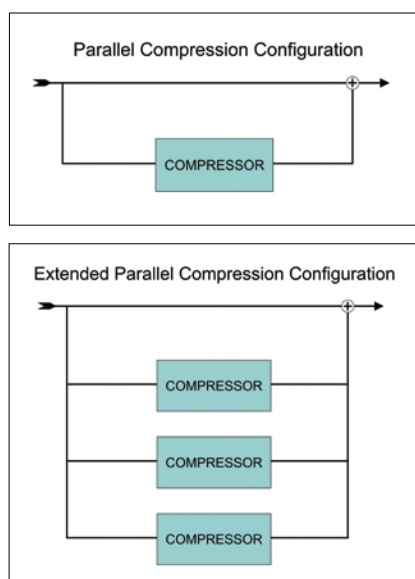


Figure 5: The way parallel compression is implemented is conceptually very simple — one merely splits a signal and passes one half through the compressor, and the other half remains uncompressed. The result is then blended together. However, it's also quite common, where resources are available, to run more than one compressor in parallel.

» by the original input signal coming via the direct path. As a result, those loud but delicate transients are left completely intact and unchanged — which is the primary aim of this technique. That's a tick on that one, then!

When it comes to quiet input signals — those which fall below the compressor's threshold — the compressor obviously won't be applying any gain reduction. This means that there will be identical signals via both the direct and compressor paths... and if two identical signals are mixed together, their combined level is 6dB greater than that of each individual signal. In other words, this simple parallel compression arrangement raises the level of quiet signals by 6dB. It's interesting to note that this raising of quiet signals is completely benign. There's no active gain manipulation going on, as there would be in a genuine upwards compressor: it's a straight make-up gain effect.

In terms of dynamic-range reduction, as this simple form of parallel compression leaves the loud bits unaffected and raises the quiet bits by 6dB, the total reduction in dynamic range is only 6dB. If more squashing than that is required, the solution is to stack up more parallel compressors. For example, adding another pair of compressors — making three in total plus a direct path — will provide four

identical signals for combination at the output, and hence an increase in low-level signals of around 12dB. Additional compression paths can be added for even more dynamic range reduction.

Although a little cumbersome and greedy of studio resources, this is actually quite a common technique, and was described by Michael Brauer in the 'Inside Track' feature of the November 2008 edition of SOS, as part of his vocal processing technique. Admittedly, part of his approach is about combining different compressors with different sound characters to achieve subtle tonal colour, but it's also about creating genuine upwards compression.

Although Brauer is probably the most frequently cited 'big-name' user of parallel compression, he's far from being the only one — our interview-based features are littered with similar mentions by successful producers and mix engineers, including Jaycen Joshua, who tells us that: "If you don't use parallel compression you risk things sounding thin. You begin with the original track that the producer loves, and if you manipulate that you may deviate too much from it, whereas if you blend a treatment in, it will be more pleasing to the ear."

Test Plots

I've just mentioned upwards compression again, so perhaps it's time to look at some real test measurements. Figures 6 and 7 were obtained using the

Audio Precision test set, measuring the audio signal going through a parallel compression configuration in a DAW. I used the same standard hard-knee compressor plug-in as in the previous examples, and adjusted it to have a threshold of -23dBFS and a ratio of 50:1, so that when the input signal neared 0dBFS the compressor was putting in just over 20dB of gain reduction.

In Figure 6, the blue line shows the transfer curve of the direct path on its own, which, as you would expect, is a straight line at 45 degrees — confirming that the output level is the same as the input level via that route. The dotted purple line shows the output of the compression path on its own, where signal levels below the threshold follow the linear direct path, and above the threshold the output level is limited to give a virtually horizontal transfer curve. (Across the top 22dB of signal range the output from the compressor barely increases by 0.5dB.)

The red line is the result of parallel compression: mixing together the direct and compression paths at unity gain. As this plot clearly illustrates, low-level signals below the compressor's threshold exhibit a fixed level increase of 6dB — exactly as you would expect from combining two identical signals. For signals above the threshold, the transfer curve shows a gentle compression slope, with the overall gain gradually reverting back towards unity (although never quite reaching it).




Figure 6: The blue line shows the source signal, the dotted purple line is the compressed signal, and the red line is the blended — or parallel-compressed — signal. The shape of the red curve is beginning to display a similar characteristic to upwards compression.




Figure 7: The blue curve shows the effect of using three compressors in parallel, and the resemblance to upward compression is stronger still than in Figure 6. Note, however, that this is achieved without disproportionately boosting the very quietest part of the signal — ie. the noise!.

» That doesn't really look much like upwards compression, does it? It actually looks pretty much the same as conventional downwards compression... so I tried to match the dynamic-range reduction using the same compressor applied directly on the input signal with make-up gain in the usual way, to provide 'uplift compression'. However, the closest I could get to the same compression characteristic was with a ratio of 1.3:1, a threshold of -22dBFS, and 6dB of make-up gain. The result is shown by the green line. Although the 'knee' into compression occurs slightly lower in level, the peak level is very close and the overall gain reduction is very similar.

Something interesting becomes apparent, though, when you compare the compression slopes of the parallel (red) and uplift (green) arrangements. Look closely and you'll see that the uplift compressor's trace is a perfectly straight line, but the parallel compression slope actually curves or 'bows' downwards. In fact, from 0dBFS down to the parallel compressor's threshold level, it actually looks just like the gently curving slope of an upwards compressor.

To try to make this 'bowed' compression slope more obvious, I increased the amount of dynamic range reduction by adding another pair of parallel compressors, totalling three in all, and the result is shown in the Figure 7. The blue line shows the transfer curve with all three compressors set with their

thresholds at -20dBFS. As you can see, quiet signals are raised by 12dB now, and immediately above the threshold, the compression curve is quite distinctly bowed before trending towards the 45-degree angle of linearity as it nears 0dBFS. Above the threshold, then, the curve follows exactly the same general shape as true upwards compression. There is one very obvious difference, though, which is that the parallel compression slope reverts to a linear gain shift (similar to make-up gain) below the threshold.

A true upwards compressor's slope would continue more horizontally.

In the same plot, the green line is the result of setting the compressor's threshold to -40dBFS, and this makes the upwards compression characteristics much more obvious. To help reinforce the point further, Figure 8 illustrates exactly the same parallel compression setup, still with the threshold set to -40dBFS, but with yet more parallel compression paths.

The red dotted line is the transfer curve for the direct path on its own again, and the purple line is the output from a single parallel compressor configuration. That looks like an upwards compression curve, doesn't it? Well, in fact, it is, and it becomes even more obvious when the ratio is increased by adding extra parallel compressors. The dark blue line is from a parallel compressor with three compressor paths plus a direct path, the green line is with five compressors, and the orange line with seven compressors (all set with thresholds of -40dBFS).

Quite clearly, input signals above -40dBFS are being upwardly compressed by the equivalent of a gentle soft-knee compressor, with the ratio being determined by the number of parallel compressors in the configuration! However, it is important to remember that signals below the threshold are also raised in level, but by a fixed gain shift dependent on how many parallel paths are being summed — and that behaviour is quite different to that of a true upwards compressor.




Fig 8: The red dotted line is the uncompressed signal. The coloured lines show the characteristics of adding a single parallel compressor (purple) and more compressors: three (blue), five (green) and seven (orange), all set with thresholds of -40dBFS.

I pointed out earlier how a true upwards compressor with a gentle 2:1 ratio and a threshold at -20dBFS would raise the ambient noise floor from an acceptable -80dBFS to an appalling -50dBFS. Presumably this is a reason there are so few plug-in upward compressors, despite the fact that DSP designers are free from the constraints of electronics. Using parallel compression with a single compressor would raise the noise floor only to -74dBFS, and even using three parallel compressors for more squash would only raise it to -68dBFS — which is far more acceptable!

Conclusions

Exploring the real-world measurements obtained from compressors in different configurations has revealed that, taken as a complete system and considering all signal levels from the noise floor to the clipping point, parallel compression is actually a variation on conventional downwards compression. In the strict technical sense, parallel compression is a variation of *uplift*

compression — combining downwards compression with make-up gain, albeit in a non-standard way.


However, the benefit of this approach is that the compressed region above the threshold doesn't exhibit the normal straight-line compression-slope characteristic that would be expected of a conventional compressor. It is 'bowed' downwards in a way that resembles the typical characteristic slope of a true upwards compressor — at least over a restricted dynamic range.

Consequently, any signals in the region above the compressor threshold are being compressed in a manner which is exactly the same as that of a true upwards compressor: lifting the level of quieter signals just above the threshold, while leaving the louder elements alone in a very gentle and progressive manner. To all intents and purposes, therefore, parallel compression does bestow a very similar character to upwards compression, within a specific range of signal levels (ie. those above the threshold).

From a practical perspective, then, the

key to effective parallel compression is to ensure that the compressor threshold is set comfortably below the lowest signal level to be processed. In that way, the source's dynamic range is fully contained within the true upwards compression region of the transfer characteristic.

Furthermore, adding additional parallel compressors increases the effective compression ratio and makes the processing more obvious and effective, yet without introducing disastrous hikes to the noise floor. While configuring multi-path parallel compression systems might be impractical in hardware-based studios, it is trivially simple to achieve with modern DAW setups! All you have to do is dial in the appropriate settings for one compressor, and then copy it to as many additional parallel channels as you require.

So now you know: if you are ever asked whether parallel compression is 'upwards' or 'downwards', the true answer is that it is both, depending on how wide your dynamic range window is. And hopefully you'll now have a better idea of how to make the most of it in your mixes! 

Mix with the best!



"Besides the excellent interviews and fascinating, in-depth recording and mixing articles, I can always depend on Sound On Sound for complete, unbiased reviews of the latest pro-audio gear."

Bob Clearmountain, engineer, producer and mixer, Grammy Award winner (Bruce Springsteen, The Rolling Stones, Paul McCartney, INXS)



"As a professional I admire Sound On Sound as one of the most trusted and credible sources of inspiration and information."

Jack Joseph Puig, mixer, producer, Grammy Award winner (Rolling Stones, U2, Mary J Blige, Black Eyed Peas)

SOUND ON SOUND

The World's Best Recording Technology Magazine



This article was originally published in Sound On Sound magazine, February 2013 edition



Available on the App Store



follow us on Twitter



find us on Facebook



go to the SOS YouTube channel



visit the SOS forum

Subscribe and Save Money!

Visit our subscriptions page at www.soundonsound.com/subscribe for more information on the Sound On Sound App go to: www.soundonsound.com/app

Sound On Sound, Media House, Trafalgar Way, Bar Hill, Cambridge, CB23 8SQ, United Kingdom
Email: subscribe@soundonsound.com Tel: +44 (0) 1954 789888 Fax: +44 (0) 1954 789895

All contents copyright © SOS Publications Group and/or its licensors, 1985-2013. All rights reserved.

The contents of this article are subject to worldwide copyright protection and reproduction in whole or part, whether mechanical or electronic, is expressly forbidden without the prior written consent of the Publishers. Great care has been taken to ensure accuracy in the preparation of this article but neither Sound On Sound Limited nor the publishers can be held responsible for its contents. The views expressed are those of the contributors and not necessarily those of the publishers.