Two Stage Compression

by Mike Caffrey

There is no such thing as too much compression. There's only Good Compression and Bad Compression.

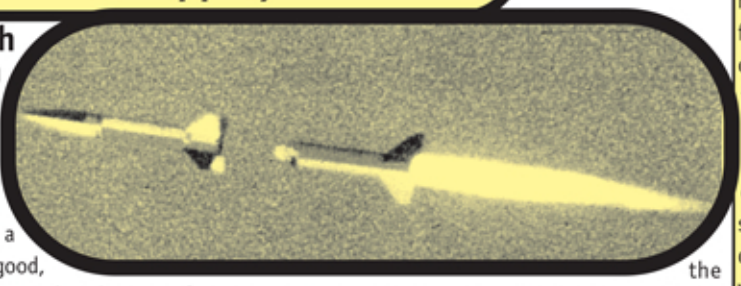

Two Stage Compression is a technique where you get good, transparent compression by compressing the sum of a compressed and uncompressed signal. You can do this through a bus routing method or with specific compressors, and it can be done in the analog and digital domains. The routing is a little complex in the abstract, so you may want to be in front of a console or DAW while reading this. Once the idea clicks, it's easy.

Hearing compression is a skill that one develops over time In issue #37 Michael Brauer explained that it took him three years to be able to hear the sound of compression, and he's known as one of the compression masters. The first step to understanding the relevance of the Two Stage Compression technique is recognizing Bad Compression, which is much easier if you know what to listen for. See the sidebar for some ear training tips on recognizing Bad Compression.

Two Stage Compression can be used any time you'd normally use a compressor and you can use any type of compressors. I'll use an drum subgroup for an example of how to set it up. Our hypothetical track will be a rock song with Nirvana-style dynamics - quiet verses and massively loud choruses.

Start by setting all of the drum channel faders so that they don't go to the stereo bus. Instead, assign them to two pairs of busses. We'll call the first stereo pair Bus A, and the second stereo pair Bus B. Bus A will send to the input of Compressor 1, can be single stereo compressor, or a pair of mono compressors used in stereo. The output of Compressor 1 should return to a pair of faders or stereo aux return. This return will not be sent to the stereo bus, but instead be sent to Bus B along with the uncompressed drum faders which we assigned initially. Bus B sends to Compressor 2 (again, one stereo or two monos) and its output returns to another pair of faders or stereo aux return. This return is the only part of the drum signal that gets sent to the stereo bus. None of the original channels, nor Bus A/Compressor 1 will be heard in the mix without being processed by Compressor 2 on Bus B. That's all there is to it. So, what is the point of this technique?

We've now set up Compressor 2 to compress the sum of our uncompressed drum tracks and a parallel compressed signal from Compressor 1. Understanding what parallel compression actually does is the key to understanding why there should be two stages of compression on the drums.

I usually consider parallel compression a "hack". Mainly it's in response to specific engineers I've seen use it mask a poorly set drum bus compressor. They smash the parallel signal and blend it the uncompressed sources. They like the sound of the fast release boosting the quieter parts of the signal, making the drum's decay and room ambience get louder very quickly and become closer in volume to the initial body of the sound. This can be exciting and explosive sounding, but it can be fatiguing and it often makes the cymbals flutter-pump. Their volume fluctuates very quickly with the compressors fast release making them sound extremely washy. In my experience, these same engineers like to use a fast attack, which removes all of the punch. The result is an unmusical, but loud, parallel compressed signal. They can get away with this by relying on the uncompressed signal to mask the Bad Compression and that's why I see this particular approach as a hack. All this does is make the soft parts, and average volume, louder without limiting the maximum volume of the uncompressed signal. This is the aspect of parallel compression that Two Stage Compression relies on. A little math explains how that works (see the sidebar).

The first step after setting up your routing is to set Compressor 1 on Bus A. Solo its return, and set the attack and release as if it's the only signal you're going to monitor the mix. Don't just smash the signal, set it so that it sounds great on the quiet parts of the song. Don't worry if this compressor compresses too much when your drums get louder when you hit the chorus. Once you're happy with the setting for the quiet sections, take that out of solo and blend it with the uncompressed drums. This way these two signals will take turns masking each other.

Compressor 1 will have a narrower dynamic range and allowing you to increase the volume with make up gain. That will cause it to mask the uncompressed signal during the quiet sections and we now have great sound for our verses. I'm assuming we chose a slow attack for some punch and a release that's slow enough to keep the drums tight, but fast enough to open up for the next hit to be fully volume and punchy. Compressor 1 has raised the volume of the drums and altered their envelope through its attack and release. However, there is now a small problem.

When the big chorus hits, the compressed signal will hit its ceiling within the compressor. Now the balance will change in favor of the uncompressed signal which can get louder and will now mask our envelope created by Compressor 1. The big loud section will not be as punchy and tight. Compressor 2 on Bus B solves this problem. We're sending it the sum

The Sound of Bad Compression

First is what I call the "TV commercial snare sound," recognizable by it's pillowy attack from an excessively fast attack time. To hear this sound, put a compressor on the insert of a snare channel and set the attack to the fastest setting. Get about 6 dB of gain reduction and set the release as slow as you can, while still allowing the compressor to recover fully to zero dB of gain reduction before the next hit. You should have a sound that sounds a bit like hitting a pillow with a drumstick More of a fffffft! than a CRrack! That's fine if you want that sound as an effect, but that's Bad Compression because that's not what a snare sounds like. Try slowing down the attack and listening to how much fuller, and more powerful and punchy the sound gets. There will probably be less compression too, so you may want to readjust the threshold.

Another sound is too slow a release with a lot of gain reduction. My favorite demonstration of this is to have the same compressor on every channel of the mixer, set with the slowest attack and the slowest release and a very low threshold. Try this on a rock song where all the instruments hit together on the first downbeat. That first downbeat will be massively loud and then the compression will kick in and stay in making the entire mix really tiny sounding. That's more Bad Compression.

If you take this same setup and raise the threshold a little higher, so that there's no compression during a quiet verse but it kicks in the chorus in the same way as the previous bad example, you have what I call "hitting the ceiling". It's when there's supposed to be a dynamic change between sections, and the band actually played it that way, but the compression stops the dynamic change. It's like standing up and hitting your head unexpectedly against a low ceiling. This can happen with bad mastering too.

I mixed a track that was mostly a tiny sounding drum machine with vocals and a little bit of ambient guitar until the fourth chorus, where there was a massive drum fill and it turns into a big rock song for the last 30 seconds. When it came back from mastering, I was pretty convinced that the mastering engineer hadn't listened to the track all the way through. He had made the quiet part, which was the first 80 percent of the song, as loud as possible, so that the big drums couldn't be louder than the quiet sections of the songs. And since you could sense that it was a big section based on the performance and arrangement, the effect was to make the big section sound smaller that the section with the tiny drum machine.

Another example is when you hear clicks at the beginning of a word on a vocal track. That means the compressor was fully open at the beginning of the word, but then clamped down so far that the split second of uncompressed audio sounds like a click.

Two Stage Compression can prevent all of these types of Bad Compression.

(Continued on page 26)

of the uncompressed drum channels and the return from Compressor 1. We can now set Compressor 2 for sum of the two drum signals during the loud chorus we will have the proper envelope and volume, and any amount of dynamic contrast from the verse without Bad Compression artifacts.

This routing gives us two stages of compression, with Compressor 1 controlling the quiet verses and Compressor 2 controlling the loud choruses. For me, this type of control and resulting lack of compression artifacts is huge.

There's another way to look at Two Stage Compression. It can be used to reduce the dynamic range sent to a "primary" compressor. For instance, I really like the sound of my Pendulum 6386 compressor in its Fast mode on my guitar subgroup, and like to use it as my main guitar compressor. However, I don't like the results when the guitars push it to compress more than a dB or two.

If you want great sounding compression I think the ideal dynamic range to send to any compressor is around 3 dB. That's not too common in a musical performance. 6 dB is still manageable, but when you get beyond a 10 dB difference between the loudest and softest part of a signal, you're getting in to the range where you're going to start to hear compression artifacts no matter how great the compressor is.

I tried to get a couple of quotes from compressor designers as to what they think the maximum dynamic range is that you can send to a compressor and still expect it to sound good. Several pointed out that you can have far more than 10 dB of gain reduction and have a great sound. That's true, however gain reduction is not always a reflection of the dynamic range being sent to the compressor's input. You could send a steady state signal with a dynamic range of zero to a compressor and set it for 50 dB of gain reduction and not get compression artifacts. It's very easy to send a compressor a wider dynamic range than it can handle well.

Two Stage Compression allows me to use the 6386 as the primary compressor on the guitar bus without having more than a dB or two of gain reduction. Just add in a parallel compression stage before it and that will narrow the dynamic range hitting the input of the 6386. Since parallel compression brings up the low and average levels, you can send a dynamic range to the second compressor that's easier for it to handle because it's been narrowed.

Two Stage Compression is great for vocal tracking too. The parallel compression does what good mic technique is meant to do. Now your regular vocal compressor sees an appropriately narrow dynamic range. Great mic technique alone is not always enough for singers with very powerful voices. Some would have to move so far off the mic to control their dynamic range that it would ruin the vocal sound through loss of proximity effect and addition of ambient reflections from the room. So in some cases this will be the only way to get the compression right. For engineers who struggle to get good vocal compression, their struggle may actually be against the limitations of using a single stage of compression not the limitations of their gear.

Daisy chaining two compressors in series is far from a new technique. That's not the same thing as Two Stage Compression. The difference is, when daisy chaining, you send a compressed signal to the second compressor. With Two Stage Compression, you send both a compressed and an uncompressed signal to the second compressor. That distinction is very important to get the benefit of this technique.

Until recently the complexity of the routing has been a drawback to this technique. When I demoed this technique at TapeOpCon, there was a pair of Chandler Germanium Compressors in the Potluck Studio. (See my review this issue) This compressor is unique in that it has a Wet/Dry Mix knob that allows you to blend the uncompressed signal with the compressed signal, within the compressor. This eliminates the need for bussing to do Two Stage Compression. Instead, simply daisy chain the Germanium Compressor in front of another compressor and you have the parallel compression in front of a regular compressor giving you the Two Stage Compression control without all of the bussing.

This is done identically when working in the digital domain. The only problem is if you don't have automatic delay compensation, and that's easily solved by adding a third "dry" bus. Set up Busses A and B the same way and add Bus C. This will be set up with the same plug-ins as Bus A, but bypassed. Technically the plug-ins will be on aux returns, but regardless, the signal running through the bypassed plug-ins will have the same latency, thus avoiding the comb filtering that would otherwise happen. Bus C is the uncompressed drums, so the returns for Bus A and C are both sent to Bus B, and the Two Stage Compression routing is complete.

Currently I mix with at least three compression subgroups - drums/bass, guitars/keys and vocals. I always make them all Two Stage subgroups. The result is a very natural sounding mix with as big or small a dynamic range as I choose. Geoff Daking calls dynamic range "the enemy" and he explanation for that is right. We need a fairly narrow dynamic range for the context we most often listen to music in, or we'll constantly be reaching for the volume knob. Controlling dynamic range through using a lot of compression or peak limiting is easy. Doing it and making it sound good ranges from hard to impossible. Two Stage Compression is the best technique I've found for controlling dynamics, and keeping them musical and artifact free. ☮

www.monsterislandrecording.com
Watch Mike's video of this technique at MonsterIslandTV.com

The Math Behind Parallel Compression

Suppose we have a volume scale of 1-10 with 10 being the loudest. In the verses the drums are quiet, around 2. In the chorus they're loud, around 8. Our average volume is 5 [(2+8)=10/2=5]. Now, add heavily compressed parallel signal, with a lot of make up gain. It's level is 8 and never fluctuates, so 8 is also the average. Figure out the new average [(8+5=13) and divide by 2, because it's two sources and you get an average level of 6.5. Blending the compressed signal in parallel brings up both the average and the minimum levels.