

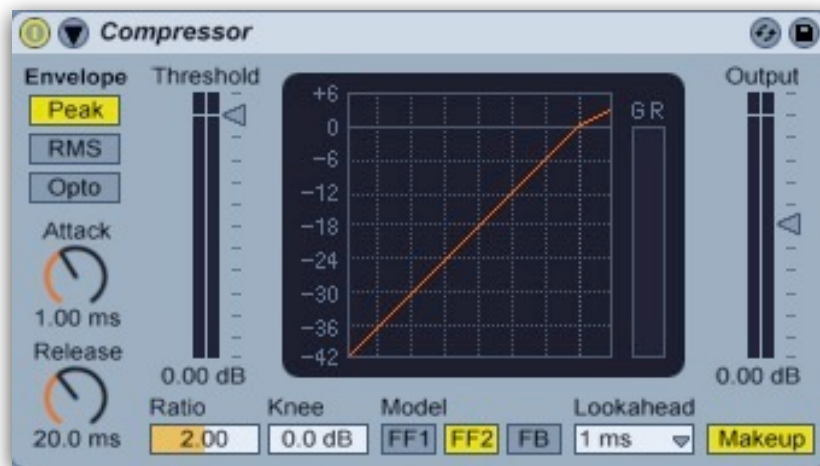
Dynamics Control

v1.1
July 28th, 2014

I've had the idea of writing a compression tutorial for awhile now, and it's definitely something many people have asked me to do. In a way I've resisted until now, because it's sort of like a vegan telling you how to grill a steak, it's just not something I use very much anymore. Or a painter trying to teach a student how to use the color red, there's a lot of different artistic ways to approach something like compression, none of them "right". However more and more I see people asking questions like "why does my song still not sound good, when I have compressors on every channel already?" Or, "what's the best preset for getting huge drum sounds in a club tune?" It's obvious there's a misunderstanding of what compression is, and a lot of misuse of it, so perhaps it's time I try and tackle the issue maybe clear up some of the mystique for those who just don't understand it very well.

As I said earlier, there's many different ways of using compression, but I'm going to focus on its traditional purpose, to help tame the dynamic range of your song, or more specifically, a part of the song. As with all my guides, I think it's important people realize I'm not saying "this is how it's supposed to be" or that there's no other way to use a compressor. Certainly they've been misused to great effect and fame. Think of the Benny Benassi basslines, or that Ministry would often use racks of compressors to achieve distortion for their guitars, and not dedicated distortion units. What I'm saying is, don't take my word for it, and don't just read this guide and think you have a handle on how to use compression. It's one of those things that's often very subtle to hear if done correctly, and it's something that takes a lot of hands on practice with to get a real understanding of. In this guide, I'll mainly be showing you how I learned to use compression, there's certainly many other ways to approach it though. That said, let's dive in.


For this guide I'm going to be using Ableton Live's Compressor plug in, though I will mainly be focusing on the controls that just about any compressor should have.



Those controls are Threshold, Ratio, Knee, Attack, Release, and Gain. Let's look at these one by one and get a rough idea of what they do.

Threshold - This is the THE most important control on a compressor, as it determines how much of your audio signal is actually getting acted on by the compressor. Think of your audio signal as a vertical graph, with the quietest parts at the bottom, and the loudest parts at the top. This is easy to visualize if you look at the screenshot above. With the threshold at 0dB, the compressor is not acting on your audio at all, it's passing through the compressor untouched by any dynamics control. As you slowly lower the threshold however, you're going to be compressing more and more of your audio, starting with the very loudest parts.

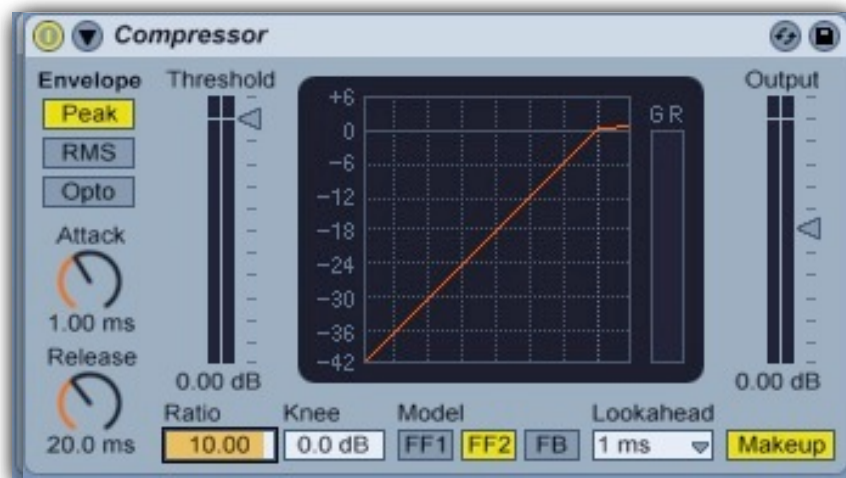
Typically, these parts are called transients, they are the sharp initial sounds at the beginning of any audio signal. Normally these are very short in nature, we're talking tens of milliseconds here, if not less. Think of the crack of a snare hit for instance, it's the loudest right after the drum stick hits it, and then it fades away very quickly.



The point is that these transients are often very, very short bits of audio, and yet they are taking up most of our dynamic range because they are so loud compared to the rest of the audio signal. So by lowering the threshold, you are USUALLY going to start compressing the transients first. This is important, because in effect it means we can look at compression in two ways. First, as a tool to help us tame transients, to reduce these millisecond burst allowing us to raise the level of the rest of the audio signal without clipping. Second, by allowing the transients to pass unaffected by the compressor, while using the compressor to raise the level rest of the audio signal. (I'll explain how this is possible shortly). Either way the end result is largely the same, we can now increase the level of that audio signal without clipping. All of this depends entirely on the setting of the threshold control, as every other control on a compressor depends on this one.

Ratio - Ratio is a fancy way of determining how much you are going to be affecting any signal that crosses the threshold. It's how much you are compressing the signal above the threshold. For example, with a 1:1 ratio, nothing will happen to the signal when it crosses the threshold, the signal will not be compressed at all. Your signal is exactly the same before and after compression. With a 2:1 ratio however, any signal over the threshold will be reduced in volume by half. For 3:1, the audio above the threshold will be 1/3rd as loud as it would be normally. The basic premise is that the higher the ratio, the more you will be compressing the audio above the threshold.

You can see the ratio represented graphically in the middle of Live's Compressor by the thin diagonal brown line.



At 1:1 ratio, it's just a straight line. At 2:1, you can see the threshold represented by an angle in the line, and above this angle the line now is half as loud as it was before. At 10:1 ratio, the line is bent even more. More than anything it's important to learn that the ratio just controls how much you're going to be compressing the signal. Very roughly, I'd say it's typical to use ratios from 2:1 to 10:1 for most uses. Any more than that, and we start getting into limiting, which I'll talk about later.

Knee - This control is directly related to ratio, and determines how immediate the compression begins on audio that crosses the threshold. For instance, at a 4:1 ratio and a knee of 0dB, as soon as the audio signal crosses the threshold, you're going to be compressing it with the 4:1 ratio. It's immediate, and often referred to as hard-knee compression, which you can see below as a very sharp angle on the graph. However if we raise the knee parameter to 4.5dB, you can see below that the sharp angle now becomes a curve. What we've done is slow down how quickly our audio above the threshold gets to the 4:1 ratio setting. Right at the threshold we no longer are immediately at a 4:1 ratio, but something smaller like 1.5:1 (just a guess), then 2:1 a couple dB above the threshold, and finally we hit our 4:1 ratio at about 4.5dB above the threshold. This is what we call soft-knee compression, it's a little more transparent sounding than hard-knee compression, not quite as noticeable. Some compressors don't display knee as dB, they might only have a knob labelled soft on one side, and hard on the other, but the effect is the same.

Attack & Release - These two parameters are time based controls, they affect how quickly the compression happens on any audio that crosses the threshold. Attack with a setting of 0 for instance, means that any time audio crosses the threshold, the compression begins immediately. As I mentioned earlier, since the transients of a sound are the first part of almost any audio signal, that means that it's the transients you're going to be affecting primarily. If you increase the attack setting to say 30ms, you're now going to have a 30ms window when the audio crosses the threshold when no compression will be taking place. In effect, you're letting the transients through untouched, and only going to be compressing the audio after the initial burst of the transient.

This is why I mentioned above that there's two ways to achieve the same overall volume increase in a signal. You can use a short attack and compress the transient, thus allowing you to boost the entire signal later. Or, you can let the transient stay at the same level with a longer attack, and use the compressor to boost the audio after the transient. Both will let you use a louder signal, though the two methods can sound slightly different too, so it's worth playing around with both methods until you learn the differences.

Release controls how quickly the compressor stops working after the signal once again falls below the threshold. With a short decay and attack, your compressor will largely only be working on the transient of the signal, where as a longer Release means that more of the signal after the transient will be affected by the compression. A lot of the character of compression comes from how you set up the Release parameter. With a relatively shorter Release, the compression can be more or less transparent. With a longer Release, it's possible to get the signal being compressed to push and 'pump' in time with the song. Setting it too long however can just turn things into mush, or worse, completely negate the effects of the compressor in the first place. Conversely, setting too fast of an Attack and Release can often lead to distortion, as the compressor can't keep up and more or less starts tracking the actual waveform shape. After the Threshold, Release is probably the second most important control on the compressor for these reasons. In general, it's good to realise that the Release affects the character of the compression.

Gain - This determines the volume of your signal after the compression. If you used a short attack and compressed the transients, you can use the gain to now raise the level of the entire signal. Some compressors like Live's, even have an auto gain control (sometimes called Makeup) that raises the level of the signal by the same amount that you're compressing. For instance if the compressor is making the signal 3dB quieter, the Makeup gain will automatically raise the gain of the signal by 3dB to compensate.

As you can see, there's a lot of interaction going on in the compressor controls, and they all depend entirely on the signal you're feeding the compressor. This is one reason I believe there's no such things as presets for a compressor. If your compressor does have presets, use them only as a starting point, you really do need to adjust the controls still to get the compressor to do what it's supposed to be doing!

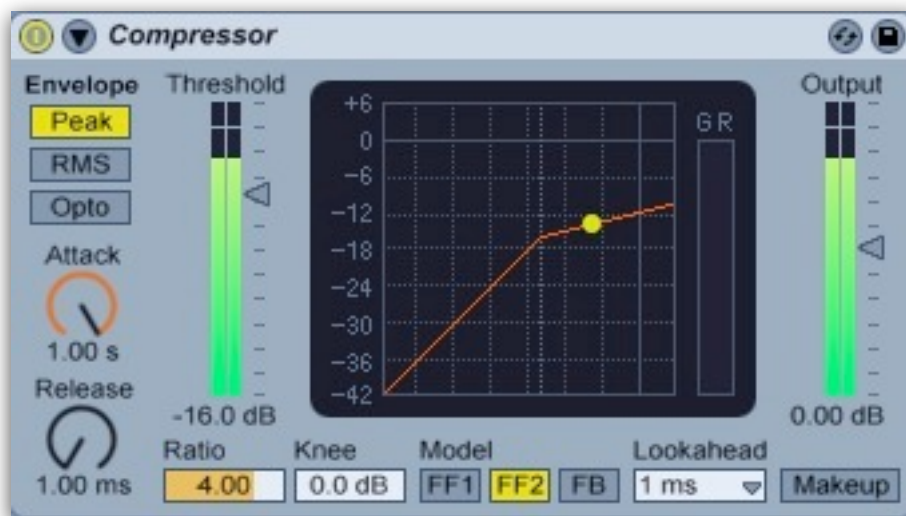
Ok, so let's put this into work and try a few examples to get a better feel for it. I think the best way to learn to use a compressor is by using it on an acoustic drum loop. Typically acoustic drums have very obvious transients, and they have very clear attacks and decays. I've gone ahead and made a quick loop in Live 7 using the acoustic drum samples that ship with the basic program. Nothing fancy, just something I cooked up real quick for demonstration purposes:

<http://tarekith.com/mp3s/compaudio/RawDrums.wav>

I highly recommend you download this audio file and follow along in your DAW of choice as I work through these examples. It's best to hear exactly how moving the controls affects the

sound, and not just the final examples I end up with. Also, you will likely need to turn up your monitoring more than usual, as most of the audio files are around -6dBFS.

Next I'll add a Compressor plug in, set it to a fairly average ratio of 4:1, with Makeup gain off. One of the things I learned when I was first getting into compression, was to set the Attack control to its highest setting, and the Release to its lowest setting initially. This way we're making sure all of our transients are passing through still, and the compressor is working on each peak in our audio, and not still in the Release phase when the next drum hit comes. Next, we'll adjust the threshold to determine how much of the audio signal we're going to compress. In Live you can see the little yellow ball in the plug in graph that represents our actual audio signal entering the compressor. I'm going to lower the threshold until the signal is about 3-4dB above the threshold, with a Threshold setting of -16dB.



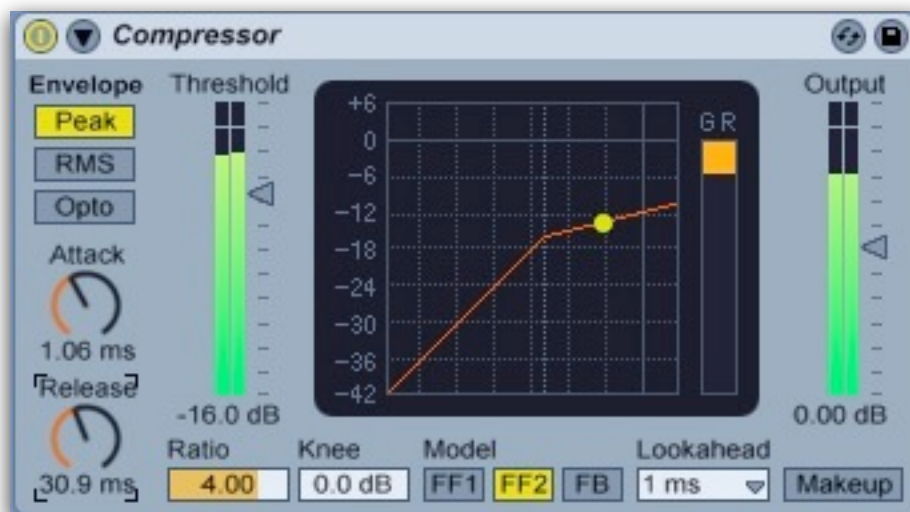
At this point you can see that the GR (Gain Reduction) Meter to the right of the graph is not showing any signal, and you really can't hear any difference in the audio either. With our Attack and Release settings being what they are, and given the level of the audio signal I'm using, the compressor is really not working yet, even though the threshold has been lowered. What we're going to do next, is slowly lower the Attack parameter, keeping an eye on the GR meter to see when we're actually affecting the audio above the threshold I set before. I think I'll stop right around 1.06ms, any more than that and I'm not really working the compressor too much, and less than that and I can hear it start to affect the attacks of my kick and snare slightly.

Now we're going to adjust the Release parameter. Before we go any further, quickly deactivate the compressor and listen to the raw audio again. Toggle the on off button on the compressor a few times to compare the audio before and after. With the Release set to its minimum currently, the drum sounds are kind of choppy, and they don't sound like they fade away very smoothly. The snare reverb is a good place to listen for this, in case you're having trouble hearing it.

While you have the compressor toggled off, take note of the peak level readings of the RawDrum loop on your audio track. It should be roughly around -6dB. As you play with the compressor settings, be sure to adjust your track fader to compensate for any changes you make. Try and keep the peak reading on the track meter around -6dB so you're comparing things more fairly. It's not an entirely accurate way of measuring loudness, but it serves our purpose for this tutorial. The Gain parameter of a Live Utility plug in can also be used, adjust the gain knob until you see -6dB on the track meter with the compressor on. I like to assign a single controller button or key command to the on off button of the Utility and Compressor plug ins, that way it's very easy with one button press to compare the sound of the compressed audio, with that of the original file, all at the same volume.

Anyway, back to the Release settings. Raise the Release until it's about 90ms, then toggle the compressor on and off. Can you hear how the reverb decay on the snare and hihats is more prominent now? I'm not saying that 90ms is the 'right' setting here, but what happens if you raise the Release even more? With a setting of 500ms, our drum loop sounds flat and boring, we've lost all the oomph and energy in the drums. Yes, we can actually raise the overall level of the drums to be much louder now, but we've lost all the impact as a result. What's actually happening is that the compressor is still pulling down the level from the previous drum hit when the next one comes, basically bypassing our Attack control and just compressing all of the audio signal. We lose the impact of transients, and the feel of each individual drum sound. We're left with a boring, dull sound now, albeit one that has been compressed. This is the most common mistake beginners make, they add so much compression that they suck the energy right out.

What about a setting of less than 90ms for the Release now? We lose the ability to hear the snare tail as clearly, but the loop tightens up more, and we can still raise the overall level of the drum loop some. Personally, I'd go for a setting of about 30ms, and then I can raise the channel fader by 3dB to get my drums back to that -6dB level. Here's the final compressor settings for the drum loop:



If you compare the RawLoop with the new compressed audio file, you can clearly hear how much punchier the drums are. They both have the same peak level, but the compressed drums feel more weighty.

<https://tarekith.com/mp3s/compaudio/RawDrums.wav>

<https://tarekith.com/mp3s/compaudio/CompDrums.wav>

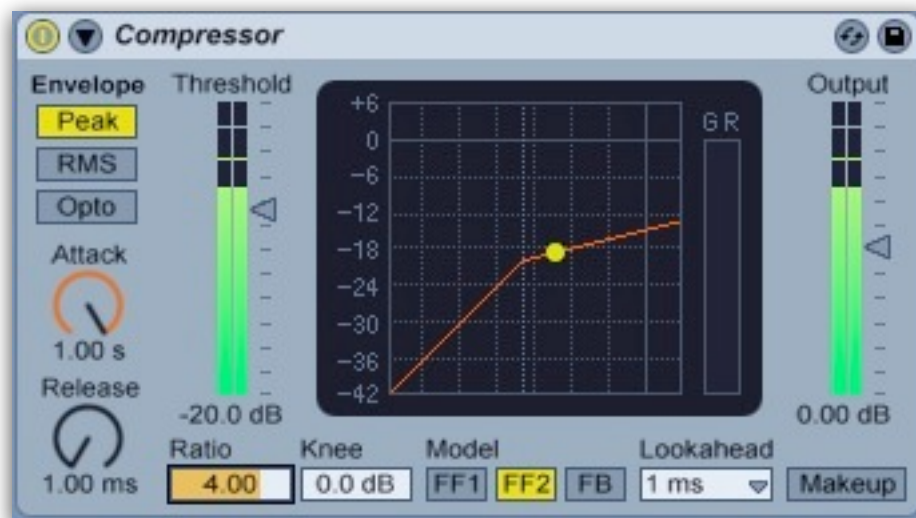
It's not a huge difference, and it shouldn't be. We added some meat to things, without losing any punch or our initial transients, and we didn't squash the whole thing into oblivion. Feel free to play around with all the different settings we've talked about so far though, they really do all interact with each other. Only by experimenting will you understand in what ways they interact.

Try playing with the Knee parameter too, and listen to how that affects the sound of the drums while not changing the overall level of the audio post-compression.

Let's try another example with a bassline that was played with a lot of volume discrepancies, and we want to use compression to even things out. Again, I've just whipped up a quick bassline to use for this example, nothing fancy:

<https://tarekith.com/mp3s/compaudio/RawBass.wav>

This audio file is also peaking around -6dB, again be sure to compensate for any volume changes the compression is adding when comparing the compression effects with the raw audio file. Once more, add a compressor, set the ratio to 4:1, the attack to full, and the release to its minimum. Lower the threshold until we're getting about 5-6dB or so of signal above the threshold point, in this case -20dB.

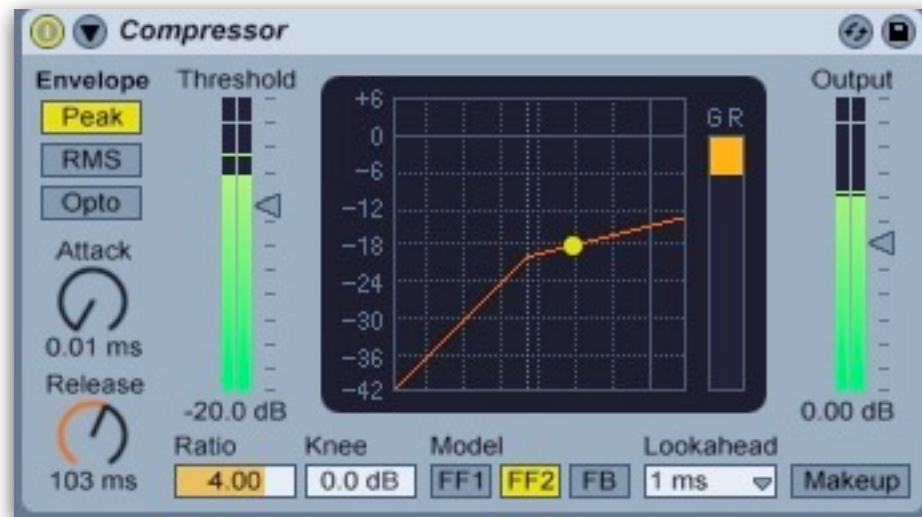


In the uncompressed file, can you hear how the initial pluck of the bass string seems to be almost missing in the quieter notes? Try lowering the Attack of the compressor to its lowest value. While we get back a lot of the pluck sound on the quieter notes, you can also hear that

distortion I was talking about earlier, since the release is still so short too. (its easier to hear if you make sure that the compressed signal is at -6dB on your track meters.)

<https://tarekith.com/mp3s/compaudio/DistBass.wav>

Let's raise the Release up to about 100ms now, which not only removes the distortion, but also lets us raise the overall level of the compressed bassline by 10dB without affecting our peak level



Compare the two volume-matched audio files again:

<https://tarekith.com/mp3s/compaudio/RawBass.wav>

<https://tarekith.com/mp3s/compaudio/CompBass.wav>

It's admittedly an extreme example, but it shows how you can not only use compression to fix playing mistakes to some extent, but also how much weight you can really add back to an audio file. Again, this is probably more compression than I'd use for real, but I wanted to illustrate the point and make the effect easier to hear.

What about limiting, which I mentioned earlier? Basically, a limiter is nothing more than a compressor with an infinite ratio. Which is a fancy way of saying that no signal will ever go above the Threshold setting. Typically limiters have a very fast attack time as well, as you're mainly going to be using the limiter on the transients of an audio signal, and thus you want the limiter to work right away when a signal reaches the threshold. Some mastering limiters also have an automatic release, which tracks the audio and dynamically changes the release setting based on the signal level.

This makes limiting a very transparent form of compression, albeit one that you can only realistically use to decrease the dynamic range of a signal by 3-6dB at most. It also only really works on material with very short, clear transients. Any more than that and the limiter starts to degrade the sound, and you're likely better off going with a compressor where you have more control over the way it reacts to the audio. For small amounts of dynamic reduction though, limiting can often be the most transparent method. Here's a comparison of the RawDrum loop we used before, along with the compressed version, and a limited version. I used the new Live 8 Limiter for this example, the Utility plug in after is only to get the signal back down to -6dB for audible comparison:



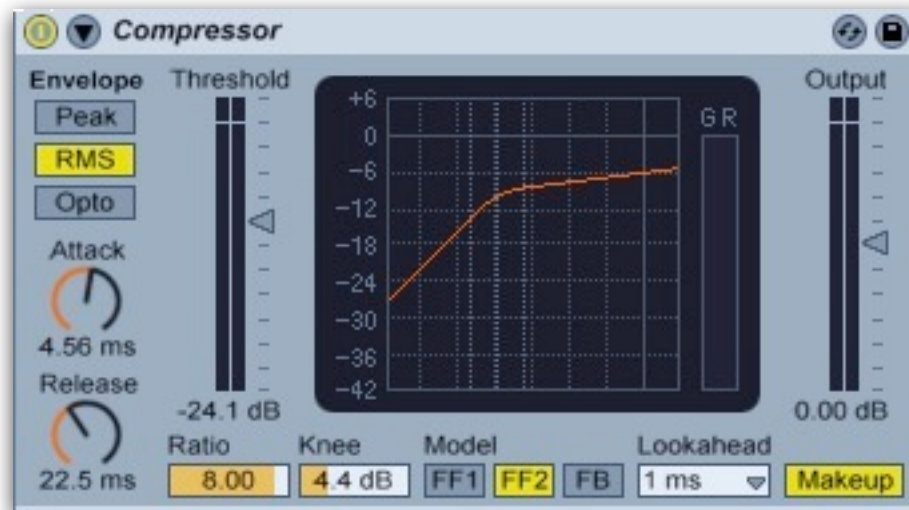
<https://tarekith.com/mp3s/compaudio/RawDrums.wav>

<https://tarekith.com/mp3s/compaudio/CompDrums.wav>

<https://tarekith.com/mp3s/compaudio/LimitedDrums.wav>

There is one more technique I'd like to talk about for dynamics control, and that's Parallel Compression. Sometimes called the "New York Drum Trick", this is a great method having your cake and eating it too when it comes to compression (within reason). The basic premise is that you slightly mix a heavily compressed version of an audio signal with a completely dry version of it. This way you don't affect the transients at all, and compressed signal adds the weight you typically want when compressing things. Most often used on drums or percussion parts, but it's also used in mastering to help increase the level of a song without overly compressing all the life out of it. In this example I'll use it on a combination of the drum and bass examples I've created so far.

The basic idea here is to buss all your drum sounds (or whatever you want to use this technique on) to a compressor, and then set up the compressor so that you are heavily compressing the audio. I mean really squash it, use a lower threshold than you otherwise might, set the compressor to detect RMS levels and not Peak if that's an option, and use a



higher ratio too. Here's the settings I used in this example:

Then you blend in this compressed version of the audio with the dry, uncompressed drum tracks. You only need to use a little bit of the compressed version, you don't want it split evenly between the two, that's way too much. Usually around 20-30% of the compressed signal is enough. Here's the audio examples showing just the raw drum and bass sound I used, and then the Parallel compressed version:

<https://tarekith.com/mp3s/compaudio/ParallelDry.wav>

<https://tarekith.com/mp3s/compaudio/ParallelComp.wav>

You can hear how the transients are all still there, but there's definitely an increase in volume (but not peak level) and weight in the low end of the audio now. It's a great trick, but like all things relating to compression, it's easy to overdo it, so use it sparingly. If you're an Ableton Live user, I created an Audio Rack for this example, which you can download here:

<https://tarekith.com/assets/ParallelCompression.adg>

You will still need to adjust the compressor Threshold, Attack and Release times to fit your own music though. I find it's best to do that while listening only to the compressed channel, so I've included a control to turn the volume off on the dry channel while you do this. The first knob controls the mix of the compressed signal with the dry signal. Fully left and all you hear is the

dry signal, while raising the knob blends in more of the compressed signal. This rack is likely better suited for drum sounds than mastering an entire mix, but try running an entire track through it too, just to hear what might be possible with further tweaking of the compressor.

As I mentioned at the start of this guide, this is only a brief introduction into compression, and a few practical examples to help people get started. There's so many ways to use a tool as versatile as a compressor, it would literally take me a book to even attempt to cover all the different uses (and it feels as if I've written one already!). And I also want to reiterate that the examples I've used here are in no way what I would consider to always be the 'right' way to use compression by any means. As with anything related to music production, it all depends on your own song and goals in terms of how you use the tools at your disposal. The best way to get to grips with compression is to use it on some really raw and dynamic audio files, to play around with all the controls and understand that the settings you use will always change and depend entirely on the audio you're treating. Even then, two people will likely approach compression two different ways for the same audio file, because it can be used for so many different purposes.

And before you start throwing compressors on everything, really take a second to stop and think about what your purpose is before you start. Using compression too much in a mixdown, even if used properly, can be the quickest way to a mushy and flat sounding track. Music NEEDS dynamics, it's what gives it energy, and power, and a sense of space and depth. I mentioned before that I only use maybe 1 or 2 compressors in an entire tune, and only if I can't solve an issue another way. So before you go crazy applying what you may have learned here, recognize that it's not always the answer, and sometimes the best compression is no compression at all. Learn to recognize when your audio is already compressed to begin with. If you sample a lot from vinyl or use prepared loops, likely these have already been compressed, and adding more compression will rarely make them any better. It's an incredibly powerful tool, but the hardest part is knowing when to use it, and when to leave well enough alone.

Erik M.

tarekith@icloud.com

Feel free to pass this document on as you see fit, though I ask that you do not modify it from it's current form, and give proper credit. If you see any errors, please let me know so I can correct them asap.