Using Compressors

Compression has become such an integral part of mixing that we thought that this chapter might be helpful.

Dynamic Range

To use a compressor effectively, it's important to understand the concept of dynamic range. Dynamic range is the difference between the loudest and quietest signal levels passing through the system chain. The span between barely audible and physical discomfort is about 120 dB, so we consider that to be the dynamic range of human hearing. Anything below the threshold of hearing will be lost, as will anything above the threshold of pain. But how much dynamic range do we need in our recordings, or, more importantly, how much can we really use?

Few people listen in a totally soundproofed room. A very well-designed studio control room has an ambient noise level 10 to 15 dB above the threshold of hearing. Since we want to keep ourselves safe from hearing damage, 95 dB or so is about all the dynamic range we can use. But consider this: A very quiet living room has an ambient noise level 25-30 dB above the threshold of hearing. The inside of an automobile at reasonable driving speed is upwards of 60 dB above the threshold of hearing, and with a window open, all bets are off.

Since most consumer audio systems aren't capable of producing painful sound pressure levels (with the exception of those cars that cruise by with the bass pumping loud enough to rattle the windows in your house), a typical listening environment can only support a dynamic range of 65 to 75 dB.

The electronics in our recording chain are much better than that, with digital recorders and analog or digital consoles capable of 105 to 115 dB of dynamic range. This means that the listening environment, not the gear, sets the practical limit on dynamic range. If you see a number greater than that on a spec sheet, it's purely theoretical – a 24-bit system is theoretically capable of 144 dB of dynamic range but practical things like electronic components stand in the way of actually achieving it in practice, at least today.

Even in the best of control rooms, we have to squeeze as much as 115 dB of obtainable dynamic range into a 95 dB box. For the rest of the world, though, we have to squeeze harder so soft passages don't get lost when your neighbor starts up his lawnmower, or when listening to the car radio at 65 miles per hour with the top down. Therefore we can't record all the dynamic range that's available if we expect people to hear all the music we record.

What a Compressor does

A compressor reduces dynamic range. When used properly, it's hard to detect the action of a good compressor, but a compressor can also be a used creatively as a sound shaping tool. A compressor can be inserted into a single channel in the recording chain when recording or mixing a track, or compression can be applied to an entire mix or sub-mix. The hookup diagrams illustrate both applications.

Let's look at applying compression to a vocal track. Hard consonants such as the letter 'T' create a high initial sound level before settling down, whereas most vowels tend to be more even. The average volume level of a word may be fairly low, but because of an initial loud consonant, we can only raise the word's volume so far before running out of headroom. If there's music playing under the voice, even when boosting the vocal level as high as possible without distorting the attack, a word (or a syllable) may be far enough below the level of the music to become lost or misunderstood.

If we reduce the gain momentarily during the loud attack, then bring it back up when the distortion-risking blast is over, we can now boost the average level of the word so that it can be understood clearly over the music. What we're doing here is reducing the dynamic range of the word, the difference between the loudest and softest parts. By doing that, plus boosting the level, we've made the vocal track louder in the mix without reaching a volume at any point that will cause distortion. To do this requires a very fast compressor or limiter, but it's a good example of the principle.

Since to most listeners, louder equals better, another use for compression is to make a mix sound louder. Often there's a single sound - a snare drum for example – that's noticeably louder than anything else in the mix. A drummer hits the snare louder on some beats, and the loudest hit determines the maximum level that can be recorded. By compressing the overall mix and sitting on the loudest hits, the average level of the song can be raised.

Basic Compressor Theory and Buzzwords

Threshold

We need to be able to set the input gain of the compressor so that it will reduce the level of signals above a certain volume level and not affect lower-level signals. The level at which the compressor starts working is called the threshold, and nearly all compressors have a control for it. With a compressor that has a fixed threshold (no adjustment), you set the level where compression begins by adjusting the level of the signal going to the compressor.

The Threshold control is generally calibrated in dB relative to the nominal operating level of the unit (+4 dBu or -10 dBV), though it's rarely a precise calibration. If you wanted the compressor to work on everything that's 6 dB below normal operating level, you'd set the Threshold control to –6 dB.

Compression Ratio

Below threshold, a compressor has a linear gain characteristic, just like a good amplifier or a piece of wire. Whatever goes in comes out unchanged. When the signal going into the compressor increases by 6 dB, as long as that increase doesn't put it over the threshold, the output also increases by 6 dB.

Once the signal goes over the threshold however, the compressor goes to work and starts reducing its gain, or more accurately, adds attenuation. If the compressor's output changes by only 3 dB when an above-threshold input changes by 6 dB, we call this a 2:1 Compression Ratio. A compression ratio of 5:1 means that a signal that's 10 dB above threshold will come out only 2 dB above threshold.

Gain Reduction

We can also say that this action represents 8 dB (10 minus 2) of gain reduction. When expressing the amount of compression in this way, we take an "eyeball average" since the actual amount of gain reduction at any instant depends on the input level at that instant. When someone says "I compressed vocals 2 to 3 dB," they mean they applied enough compression so that most of the over-threshold peaks were attenuated by 2 or 3 dB.

That's a typical amount of compression applied when tracking a singer with good dynamic control. It evens out sustained notes and provides a small safety net against surprise overloads. A singer with less control may require 10-12 dB of gain reduction to level out his track.

Limiting

If we never want the output to exceed a preset level, the compression ratio must approach infinity (10:1 is usually close enough for rock and roll), so that a large over-threshold signal results in a very small change in output level. A compressor set to a very high ratio becomes similar in action to a limiter, as its output level is limited to just a little above the threshold level. Limiting the output of a PA system is a good way to protect your speakers and your audience's ears.

The graph shown below illustrates the action of a compressor for several ratios. The slope of the line represents the gain. Notice that for any compression ratio other than 1:1 (no compression) the line changes its slope at the threshold, which we've chosen to be 0 dB.

This change in slope is what makes a compressor compress. Below threshold, the gain is unity – the output change is equal to the input change. But above threshold, the gain is less than 1, which is actually attenuation. Look at the 2:1 line. Note that above threshold, we get only a 5 dB change in output for a 10 dB change in input – a ratio of 2 to 1.

The Compressor Knee

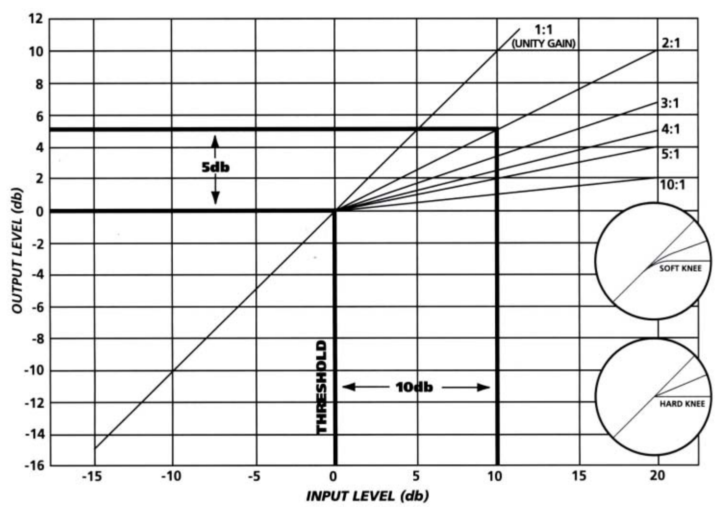

COMPRESSOR ACTION

The point at which the slope of the gain line changes is called the knee. A compressor is said to have a "hard knee" characteristic when the slope changes abruptly at threshold. A "soft knee" characteristic is one in which the gain change is gradual over some range of input level. With a soft knee compressor, gain reduction actually begins a little below threshold and the full amount of gain reduction isn't achieved until the input level is somewhere above threshold. A hard knee compressor tends to work better at catching transients while a soft knee tends to be less obtrusive on vocals. But these are only typical applications. Your voice or snare drum may vary.

Controlling the Gain

On the surface, a compressor is a fairly simple device. All the action takes place in the compressor's gain control element with an amplifier on the front and back end to match up signal levels to the outside world.

Several different devices can function as the variable gain element, and to a large extent it's the characteristics of these different devices that give each different compressor its "personality."

The signal level must be detected so the gain reduction element knows how much gain to drop. Some compressors detect the RMS (average) level of the signal, others detect the peak, and still others detect both. In some compressors, the level at the input is detected and the gain control element decides what to do based on that. Another design measures the signal level at the compressor's output and feeds that information back to the gain control element. All of these things affect the sound of a compressor.

Today you can buy compressors with multiple personalities, where the gain element's response curve shape with respect to time is selectable, thus emulating the action of different compressor designs.

Attack Time

An important compressor characteristic is Attack Time. Its definition is a bit loose. To some, it's the amount of time it takes for the compressor to reach full gain reduction when hit with an over-threshold signal. Others define it as the time required to get most of the way (typically 67%) there.

Attack time makes a big difference in how a compressor affects the sound of the signal you're compressing. If the signal has a loud initial attack (like a drum), we may choose to allow the attack to get through unaffected even though it's louder than our desired average output level. To do this, adjust the compressor's attack time to be longer than the instrument's attack time. On the other hand, if it's the transient that we want to squash, we need a fast attack (short time) so that gain reduction will begin as soon as possible after the input crosses the threshold.

Attack time is not the same as the shape of the compressor's knee, but their effects are related.

Release Time

Once the signal drops back below the threshold level, the compressor needs to return to unity gain, but this doesn't happen instantly. Instead, the compressor's gain rises gracefully (we hope) over some period of time. This time period is called Release Time.

Some compressors have no adjustments for attack or release times. Either they're fixed by design or they're "program dependent" (automatic), which means the compressor decides how fast it should respond based on the envelope of the input signal. Each design has its place in the universe and you're not necessarily being cheated if your compressor is missing a knob or four.

Make-up Gain

Since a compressor reduces gain, it is usually necessary to amplify the signal following the gain reduction circuit in order to get the output back to nominal operating level. Most compressors have an Output or Make-up Gain control which allows you to adjust the signal output level to match up with the next device in the signal chain.

The Gain Reduction Meter

Compressors usually have a meter that looks like a VU meter only it works backwards, indicating the amount of gain reduction rather than the signal level. A typical compressor's meter reads 0 dB whenever the input is below threshold and moves down scale as the input level goes above threshold and gain reduction starts. The meter on a compressor may be switchable to indicate either compression or output level.

Stereo or "Program" Compressors

A compressor is basically a single channel device, but there are stereo compressors, often called "program compressors," most often used to compress a full stereo mix.

The thing that differentiates a stereo compressor from simply patching one compressor into each channel is that the signal that controls the amount of gain reduction is shared by both channels.

Reducing the level of one channel of a stereo pair will cause the balance to shift to the louder side. We don't want the stereo image to wander around when one channel goes over threshold and the other one doesn't, so we connect the level detectors of the two compressors together. Now when either channel requires some gain reduction, that same amount of gain reduction is applied to both channels.

Compressor Artifacts

Two uncomplimentary terms often used to describe the sound of a compressor are breathing and pumping.

Breathing

Breathing is most noticeable on a solo voice and is often, in fact, the sound of the vocalist breathing. If release time is short, the gain rises quickly during pauses between words, just as the singer takes a breath. The increased gain makes the breath more audible.

Hearing a singer taking a breath may not always be desirable or dignified, but at least it's organic. Few recordings are made in an absolutely silent environment, however. Any ambient noise in the room will be boosted when the gain rises, creating an artificial "breathing" sound, perhaps even bringing leakage from the singer's headphones along with it.

All compressors exhibit some breathing, but careful adjustment (which includes controlling room acoustics and mic positioning) can minimize it.

Pumping

Pumping is another compressor artifact. It's more apparent when compressing an overall mix than a single track.

One instrument in the mix that's louder than the others will trigger the compressor into action. If that instrument stops playing, even for an instant, the level of the mix will increase noticeably. Each time the dominant instrument starts or stops, it "pumps" the average level of the mix up and down.

Compressors that work best on full program material generally have very smooth attack and release curves and a slow release time to minimize the pumping effect.

Working the Knobs

If your signal has peaks up to +15 dB and you want to reduce those peaks to a more manageable +5 dB, you might set the threshold at -5 dB and compress using a gentle 2:1 ratio. Or if you want to use a stiffer ratio, say 6:1, you'd set the threshold at +3 dB. As an exercise, try plotting out a few combinations yourself. Then set up your compressor and listen to the differences.

Lowering the threshold while keeping the compression ratio fixed reduces the maximum output level, since you're compressing over a larger portion of the dynamic range of the input signal. By keeping the threshold fixed but increasing the compression ratio, you'll reduce the output level by compressing only the loudest signals. There are no rules for this. Let your ears be your guide, with the meters as a sanity check.

The Compressor as a Tone Modifier

Adjusting the attack and release times of a compressor applied to an instrument can change its timbre by rounding off a sharp attack or stretching out the sustain portion of the note's envelope. It's sort of like having an equalizer with a different parameter.

A drum hit can be "stretched out" by applying a long release time, a healthy gain boost, and fairly high compression ratio.

Compressing low frequency program material requires special care. The attack and decay portions of a kick drum are 60 to 80 milliseconds long, but a low-pitched kick has a fundamental frequency of about 40 Hz. This means that only three of four cycles of the kick's fundamental frequency are heard on each hit, much of that being in the decay portion of the envelope.

Compressing the drum with a fast attack time and high compression ratio will make more cycles of the attack portion of the drum audible, making its "thump" sound much more pronounced without having to boost low frequency gain with an equalizer. Try this next time you want a kick in the chest.

The beater attack is at a higher frequency (1 to 3 kHz) so a moderately fast attack will let a few cycles of beater through while working on the low "whump." Slowing down the attack lets more of the beater sound pass before being compressed, often allowing you to increase the impact of the kick drum without increasing the level of the track in the mix.

Compressors with very fast attack time that often work well on vocals, work poorly on a bass because the compressor actually tries to follow the individual cycles of the waveform rather than the envelope of the note. This characteristic can be used as a special effect, but usually it just takes all the life out of a bassy instrument.

Some people put a compressor in-line with a signal with the controls set for no compression action at all. One reason is that they like what the electronics of the input and output amplifiers do to the sound. This is most noticeable with “vintage” units that use tube electronics – a mighty expensive tube-in-a-box, but if it’s there and it works, there’s no reason not to use it.

The Bottom Line

A compressor can never be used by the book – not even this one. You need to listen carefully when you make adjustments. By understanding the effect of each of the adjustable parameters, you’ll be better able to reach your goal quickly.

Notes