Compression Exposed

A compressor is one of the most common outboard tools in a studio. All compressors perform the basic same function but, like microphones, various models perform it differently, giving each model a characteristic sound or personality. Here, we'll have a look at what a compressor does, what characteristics separate one model from another, and look at compressor use and abuse.

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 recording chain. The span between barely audible and physically painful is about 130 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 can we really use?

Few people listen in a totally soundproofed room. A well isolated 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, 100 dB or so is about all the dynamic range we can use. But consider: A very quiet living room has an ambient noise level about 25 dB higher than the threshold of hearing. The inside of an automobile is 60 dB higher. Since most consumer audio systems aren't capable of producing painful sound pressure levels (I'm rethinking that as a car drives by my house with the bass pumping loud enough to rattle my windows), a typical listening environment can only support a dynamic range of 65 to 75 dB.

Any 16 bit digital system worth its dither can provide a dynamic range of better than 90 dB. The theoretical limit is 96 dB (it's not really that simple but this is an accepted working figure) but necessities of life like mic preamps, mixers, and power amplifiers all add noise, eating into the low end of the theoretical dynamic range. So, allowing for the state of the art, we have to squeeze 115 dB of workable dynamic range into a 90 dB box (maybe around 110 dB for a 24-bit system). Practically, though, we have to squeeze harder so soft passages don't get lost when your neighbor starts up his lawnmower, or when playing the car radio over highway noise. So we can't record all the dynamic range that's available if we expect people to hear all the music we record.

Compression To The Rescue

A compressor reduces dynamic range. When used correctly, the dynamic range reduction of a good compressor is hard to detect, but a compressor has other applications where we may want to hear it working. For instance, it can become a useful sound shaping tool. A compressor may be inserted into a single channel of the recording chain when recording or mixing a track, or compression may be applied to an entire mix or sub-mix.

Let's look at a vocal for example. Hard consonants such as the letter 'T' create a high initial sound level before settling down, whereas most vowels tend to be more even in volume. The average volume level of a word may be fairly low, but because of an initial hard consonant, we can only raise the volume of that word 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 clipping the attack, a word (or a syllable) may be far enough below the level of the music to get lost or misunderstood.

If we reduce the gain momentarily during that loud attack, then bring it back up when we're safely past the peak level, the average level of the word may now be raised enough to be understood over the music. What we've done here is reduced the dynamic range of the word, the difference between the loudest and softest parts. Of course you can't adjust the compressor for every word in the song (well . . . you could on a digital workstation if you had the patience), but the combination of a properly set compressor and a singer with some control yields effective results.

Another use for compression, one that's prevalent today, is to make a recording sound louder. (To most listeners, louder equals better.) Often there's a single sound (a snare drum is a common example) that is somewhat 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 those loudest hits, the average level of the song can be raised.

Basic Theory and Buzzwords - Threshold

We need to be able to adjust the compressor so that it will reduce the level of signals above a certain volume level and not affect lower level signals. This level is called the Threshold, and nearly all compressors have a control for setting it. Those that don't, have a fixed internal threshold and you set the point where it starts compressing by adjusting the level of the signal sent to the compressor.

Except for those built into multi-function microphone processors or mixers, compressors are line-level input devices. The Threshold control is generally calibrated in dB relative to the nominal line level of the compressor (typically +4 dBu or -10 dBV), though it's rarely a precise calibration even at the 0 dB mark. But to keep the knob within a good working range, you should choose a compressor that's designed to operate at the nominal line level of your system.

Gain Reduction and Compression Ratio

Below threshold, a compressor has a linear gain characteristic, just like a good amplifier. Whatever goes in comes out unchanged except perhaps for a shift in level. When a below-threshold input signal increases by 6 dB, the output also increases by 6 dB. But a compressor's job is to reduce its gain whenever the input level goes above the threshold. If the compressor's output changes only 3 dB when an above-threshold input changes by 6 dB, we call this a Compression Ratio of 2:1. If we want to allow a peak coming in at 10 dB above threshold to come out 2 dB above threshold, we need a Compression Ratio of 5:1.

We can also say that this action represents 8 dB (10 minus 2) of gain reduction, and this is usually what's indicated on the compressor's Gain Reduction Meter. When expressing of the amount of compression in this way, we have to 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 that they applied light compression, where most of the peaks don't get more than 2-3 dB of gain reduction. This is typical of compression that would be applied to a singer with good dynamic control when tracking. It evens out sustained notes a bit and provides a small safety net against surprise overloads.

If we never want the output level to exceed the threshold, the compression ratio approaches infinity (10:1 is usually practically close), so that a large change in input level over threshold results in a very small change in output level. In this case, the compressor becomes a limiter, as the output level is limited to essentially the threshold level.

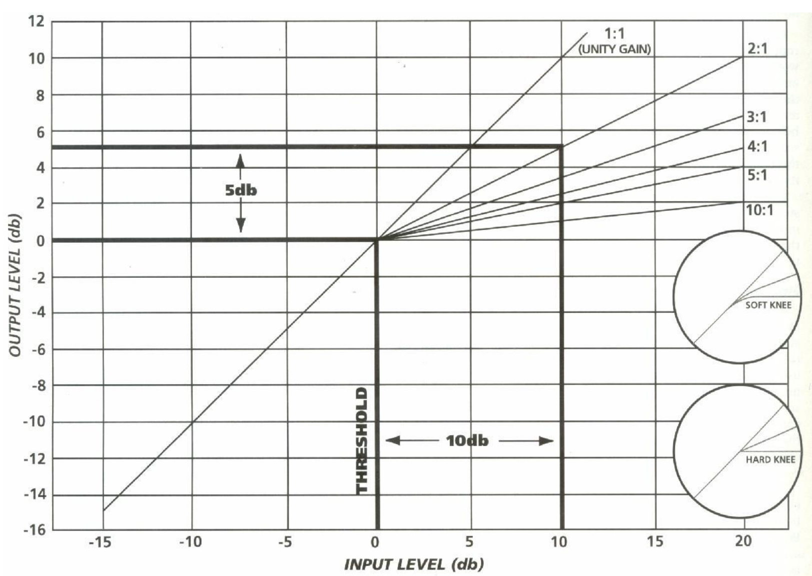

FIGURE 1 - COMPRESSOR ACTION

Figure 1 illustrates the basic action of a compressor graphically for several ratios. The slope of the line represents the gain. Notice that the line changes its slope at the threshold, in this case, 0 dB. This is what makes a compressor a compressor. Below the threshold level, the gain is unity, but above threshold, the gain is lower than 1, actually an attenuation. Look at the 2:1 line. Note that above threshold, for a 10 dB change in input, we get only a 5 dB change in output, a ratio of 2 to 1. Check out some points on the other compression ratio lines to convince yourself.

Knees

The point at which the slope of the 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. Some compresors offer you a choice of a hard or soft knee characteristic. (dbx holds the trademark on the term "over easy", their version of soft knee compression) The action of a soft knee compressor begins somewhere below threshold and the full amount of gain reduction isn't reached until the input level is somewhat above threshold. A hard knee tends to work better at catching transients while a soft knee characteristic tends to be less obtrusive on vocals. But these are only typical applications. Your voice or snare drum may vary.

Timing is Everything

A compressor doesn't know what's coming at it until it happens, so it needs some time to figure out how much gain reduction is needed. The compressor's response time is a function of the way the input level is detected. It's an integral part of the design and is one of the things that contributes to a compressor's personality. Another important time parameter 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 triggered by an over-threshold input, 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 signal. If the signal we're compressing has a loud initial attack (like almost any drum), we may want to allow the attack to get through unaffected even though it's louder than our desired average output level (though we may decide to limit it later). In this case, an attack time that's longer than the instrument's attack time is appropriate. On the other hand, if it's the transient that we want to sit on, we want 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.

Once the signal drops back below the threshold level, the compressor starts working like a piece of wire again, but this doesn't happen instantly. Instead, the compressor's gain rises gracefully (we hope) back to unity over some period of time. This period between the signal returning to threshold and when this gradual gain increase begins is called Release Time. We want a finite release time since music is dynamic and transient in nature - we don't want the gain jumping up only to have to jump back down again to sit on an immediate following transient. The compression action is far less detectable if the gain isn't jumping all over the place.

Some compressors have no adjustments for attack and 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 two.

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

Metering

Compressors usually have a meter which looks like a VU meter only it works backwards, indicating the amount of gain reduction rather than the signal level. A typical compressor meter reads 0 dB when the input is below threshold and moves down scale as the input level goes above threshold and the compressor does its thing. Often there's a separate meter or a switch to allow the meter to read input level as a guide to setting the threshold, and some also allow you to read output level.

Stereo or "Program" Compressors

A compressor is basically a single channel device, but we also have stereo compressors. They're often called "program compressors" when used for compressing a total stereo mix. What makes a stereo compressor different from just patching one compressor into each channel is that the signal that controls the amount of gain reduction is shared by both channels.

Turning down the gain of one channel of a stereo mix causes the balance to shift to the louder side. We don't want the image to flop around due to the independent action of two compressors, so we connect them together such that when either channel requires gain reduction, that same amount of gain reduction gets applied to both channels. When so connected, one set of controls becomes the master and works for both channels.

Nearly all stereo compressors have a switch to turn the stereo link off so you can use it as two compressors. Quite a few single channel compressors have a "link" connector for stereo operation, but be warned - there's no interface standard for compressor links. You can link two of the same model compressor if so equipped, but if you have two different compressors, don't expect them to be friends.

The Guts

On the surface, a compressor is a fairly simple device. All the action takes place in the gain control element with an amplifier on the front and back end to match up signal levels to the outside world. There are several different devices that can be used 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".

There are two signal paths in a compressor, the main audio path and the sidechain. The audio path is what you put in and what you expect to get out. But in order to derive the voltage used to adjust gain, the input signal must be split off and detected. This path for this control signal is called the sidechain.

We aren't normally interested in the instantaneous level cycle by cycle, but rather, want the control voltage to follow the average level, or envelope, of the input waveform. The method for deriving the sidechain control voltage is another personality builder. Different designs have employed simple averaging detectors, peak level detectors, and true RMS averaging detectors. Each yields a control voltage that follows the input signal a little differently, so each type of sidechain detector imparts a different control action on the main audio signal. In addition, some designers have applied their own "corrections" to the sidechain signal to compensate for non-linearity of the gain control element.

A couple of terms associated with compressor design are "feed-forward" and "feed-back". In a feed-forward design, the input signal branches off to the sidechain detector. In a feed-back design, the output after the gain control element is fed back through the detector to control the gain.

One popular compressor design uses a light dependent resistor (LDR, or photocell) as one component of a voltage divider. An incandescent light bulb, LED, or electroluminsecent (EL) panel attached to the LDR is driven by the sidechain control voltage (not the input signal itself). The actual audio signal passes through the LDR. As the signal gets louder, the light gets brighter causing the LDR to change resistance. This changes the voltage divider ratio, changing the gain. A Vactrol(R) is a sealed module consisting of a light source and LDR that was used in several "classic" compressors. Today the term "Vactrol" is associated with a certain flavor of compression whether there's a genuine one in there or not. This is often called an "optical" type compressor, examples being the Teletronix/UREI LA series, some early Tube-Techs, ADL, and the Manley Electro-Optical.

Another design uses a vacuum tube as the gain control element, lending the name "Variable-mu" (mu is the abbreviation for a tube's gain) to another type of compressor. The classic variable-mu compressor is the Fairchild 670, now selling on the vintage gerar market for over $20,000! Altec also made one, and Manley Labs currently builds one of their own design using the same gain control principle, and using a full differential signal path which cancels the second harmonic distortion introduced by the gain control tube's action.

Today's garden variety (and some not-so-garden variety) compressors use a voltage controlled amplifier or attenuator (VCA) for gain control. VCA compressors are common due to the availability of inexpensive and reasonably good VCAs in IC form. Since this style of VCA is built on a single chip, the designer's challenge is to keep the gain control signal out of the audio signal path. The VCA is the fastest responding and most linear (or rather, most predictable) of the various gain control elements, so it lends itself well to the design of a new breed of digitally controlled analog compressors. Here, the sidechain control voltage can be shaped digitally to produce any imaginable response curve, allowing the compressor's gain control element to emulate the sound of any "classic" which can be measured.

Virtually everything inside the box affects the sound of a compressor - the gain control element, the way the sidechain signal is derived and processed, the sound of the input and output amplifier stages and (if included) transformers, even the power supply. Today's much lusted-after "tube compressor sound" is really a new development. The often emulated Teletronix LA-2 was the only one of its lineage with tube amplifiers. The later Teletronix/UREI compressors (LA-3, LA-4) were all solid state, though using the same basic (with seasonal variations) gain control element. Much of what we think of as the "warm tube sound" of a compressor is a result of a tube input and/or output stage, not the compressing action itself.

Compressor Warts

Two terms often used to describe a compressor's action, unfortunately uncomplimentary ones, are pumping and breathing. Breathing is most noticeable on a solo voice and is often, in fact, the sound of the vocalist breathing. If release time is very short, the gain will rise quickly during pauses between words, just when the singer breathes. The increased gain makes the breath more audible. Hearing a singer take 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 (which may also include leakage from the singer's headphones) will be boosted by the gain increase creating undesirable noises where there should be quiet. All compressors will exhibit some breathing, but careful adjustment (which includes controlling room noise by careful mic positioning) can minimize it.

Pumping is another compressor artifact. It's most apparent when compressing an overall mix rather than a single track. If one instrument in the mix is louder than the others, this is what 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 level of the mix up and down. Compressors that work best on full program material generally have very smooth attack and release curves and slow release times 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. By lowering the threshold while keeping the compression ratio fixed, you can reduce the maximum output level by compressing over a larger portion of the input signal's range. By keeping the threshold fixed but increasing the compression ratio, you'll reduce the output level by only affecting the loudest signals. There are no rules for this, let your ears be your guide, with the meters as a sanity check.

Adjusting attack and release times can change the timbre of a compressed instrument by rounding off the attack or stretching out the sustain portion of the note's envelope. A drum hit can be "stretched out" be using a long release time, a fairly high compression ratio, and a healthy gain boost. If an instrument or singer produces a soft note following a loud note, release time should be short to let the gain come back up and let that soft note through.

Material with significant low frequency content requires special care when compressing. The attack and decay portions of a kick drum run 60 to 80 milliseconds, but a low pitched kick can have a fundamental frequency of about 40 Hz. This means that only three of four cycles of the kick's fundamental are heard on each hit, much of that being in the decay portion of the envelope. Stretching this out with a fast attack and high compression ratio can make more cycles of the fundamental audible. The beater attack is a higher frequency (1 to 3 kHz) so a moderately fast attack will let a few cycles of beater through while reducing the low frequency "whump". Slowing down the attack allows more of the beater sound to get through, often allowing you to bring down the overall level of the kick in the mix. But this won't give you the chest slamming kick sound that's an integral part of some forms of music.

Very fast attack times often work well on vocals but don't work well on kick drum or 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.

Squeezing Out

Compression isn't a by-formula thing. No article can honestly tell you how to set a compressor for a particular instrument, because there are so many things that can be different that rules don't work. But applying your knowledge of how a compressor works and what the knobs do will help you understand what you're hearing and twiddle the knobs more effectively.