Dynamic Range Processing and Digital Effects

Dynamic Range Compression

Compression is a reduction of the dynamic range of a signal, meaning that the ratio of the loudest to the softest levels of a signal is reduced. This is accomplished using an amplifier with variable gain that can be controlled by the amplitude of the input signal: when the input exceeds a preset threshold level, the gain is reduced. The figure below is a device input/output graph that shows the basic adjustable parameters found on compressors:

Compressor gain parameters

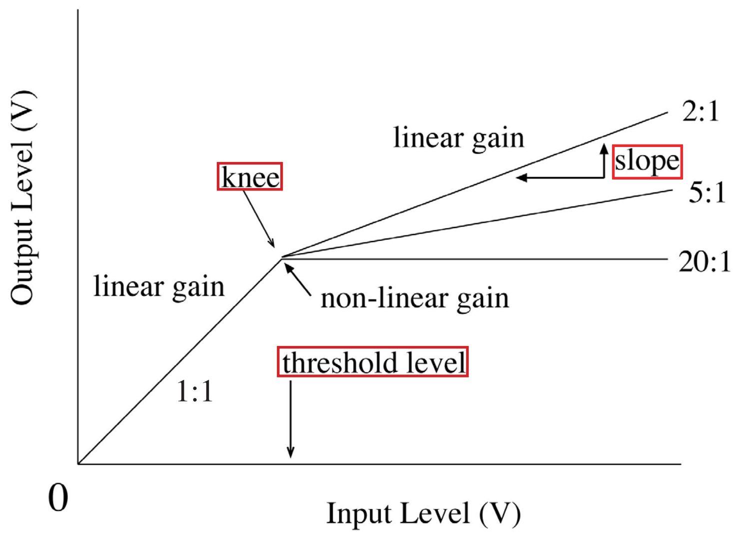

Gain is the ratio of the output to the input, the slope of the line. For inputs above the threshold level, the output level change for a given input level change is known as the compression ratio (the inverse of the above threshold slope.) The knee is the point at which the gain changes. It can be changed all at once ("hard knee" as in the figure above) or gradually with increasing input level ("soft knee"). Since this change is non-linear, it can result in distortion that sounds slightly different when changed smoothly or instantaneously. The time required for the gain to change after a signal crosses the threshold is referred to as the attack time and the time it takes to returns to the original gain when the signal drops back below the threshold is the release time. Generally, the attack is fast and the release is slow to prevent audible artifacts of the gain change, however too short an attack time reduces the transient onset of a sound and noticeably alters the character of the sound: sometimes this may be desirable, but often it is not. Overshoot is the result of the attack time setting, where the increased input amplitude is not yet reduced by the circuit (see figure below.) This often retains the character of the sound onset that identifies the type of instrument or source: a surprising amount of the distinctive character of an instrument's sound is contained in the initial transient, so allowing some time for the attack is often necessary. This can be tens of milliseconds. Interestingly, the human auditory reflex displays attack and release times of about 40 msec and 135 msec, respectively, and these values tend to sound natural when applied to electronic compression.

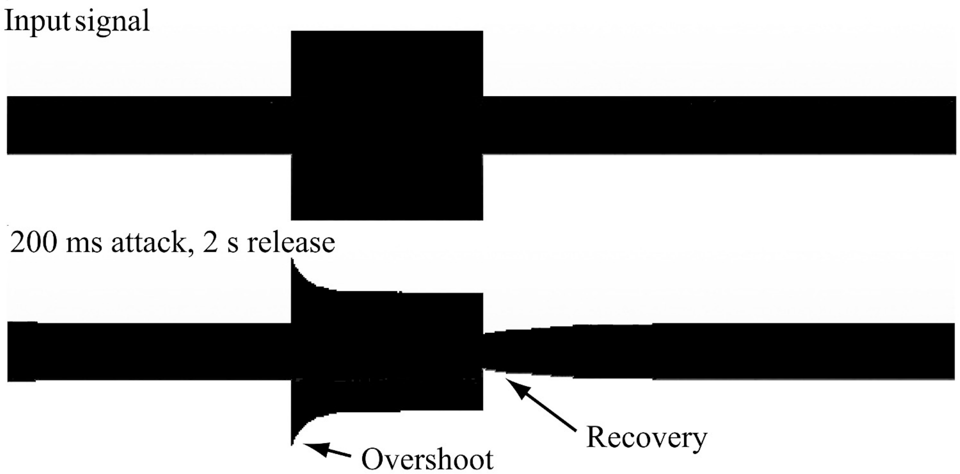

You might think of the use of compression in terms of viewing the skyline of a big city: the taller buildings hide the shorter ones behind them. This is analogous to how louder sounds mask softer ones in the same frequency range. The effect of compression would be to shrink the taller buildings relative to the shorter ones so that all the buildings attain more similar heights. You can imagine what this would look like: more of the shorter buildings would be visible although some of them would still be blocked by others. Similarly, compressing sounds makes them more equally prominent in the mix but there is still a limit to how much of the quiet sounds we will be able to perceive.

There are two basic circuit topologies used in compressors - feed-forward and feedback. As the names imply, the feed-forward devices use the input signal to generate the amplitude envelope signal while feedback circuits measure the already-compressed output to derive the control signal. Operation of the feedback device is simple: if the output is larger than the threshold it turns down the gain. That is all the circuit needs to do since compression is already being applied to the signal being measured. In the case of feed-forward devices, you need to know by how much the signal exceeds the threshold and turn down the gain by exactly the right amount. This requires calibrated measurement and gain circuits to work together to produce the proper gain reduction. The necessity of calibrating the measurement and gain circuits makes feed-forward devices inherently more complicated than the feedback type. The topology to some extent influences the sound of the device, but so do the type of gain control element used and the type of envelope tracking circuitry. In general, feedback designs reduce the dynamic range demands on the control circuitry as the dynamic range of their input is already compressed. The attack and release times of feedback compressors are affected by the compression ratio while feed-forward compressors time behavior is determined mainly by the speed of the level sensing circuitry. Some designs are program dependent, meaning their characteristics depend on the signal. Peak measuring circuits provide a fast measure of the instantaneous signal amplitude, making sure the signal maximum is always known. RMS and average circuits may more closely approximate how loud we would perceive a signal to be, but may not track short peaks that could overload the circuit. Examples of feedback compressors include LA-2A and 1176 models while dbx compressors are feed-forward types. Gain control elements include FETs (1176), optical (LA-2A, LA-3A), VCA (dbx 160, API 2500), variable-mu vacuum tube (Manley Variable Mu, Fairchild 670) and pulse-width-modulation [PWM] (Pye).

Compression may be employed at any stage in the recording process, although it can sound quite different depending on where in the chain it is used. It may be desirable to compress signals before they are recorded, especially if the recorder is noisy, so the reproduced signal will not also contain the compressed added noise that would appear if the signal were compressed after recording. Unfortunately, a poor adjustment of the compressor can cause undesirable changes to the signal that are then recorded, so compression must be used carefully, if applied before a signal is recorded, to avoid an overly “squashed” sound. This usually means using a threshold setting that results in only a few dB gain reduction on the signal peaks. It may also mean using a low ratio; something like 2:1 to 4:1, for example. If the ratio is set to 10:1 or greater, the effect is considered limiting because the output is effectively limited even if the input continues to increase. Limiting causes less alteration of the loudness relationship balance between elements of a sound than does compression because it leaves the signal unaltered until it hits a threshold higher than where a compressor's threshold would be set to achieve the same reduction of dynamic range. Limiters are often used on singers as they are being recorded to prevent recorder overloads from vocal plosives.

Compressor vs. Limiter

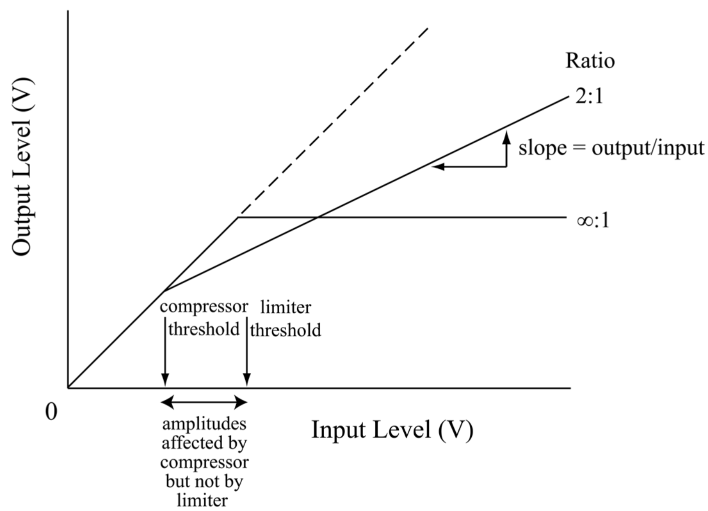

The difference between a limiter and a compressor, while one of degree, is clearly audible. Compression alters the loudness relationship between soft and loud elements of a signal more dramatically than does a limiter. By setting a low threshold, a compressor is active on more of the dynamic range of the signal: it reduces the gain over a relatively wide range of amplitudes. The limiter reduces the high amplitude elements of the signal, but leaves the soft and mid-amplitude elements alone. This makes the limited signal sound more like the original signal because only the peaks are reduced while the unaltered remainder of the signal is made linearly louder.

When a large amount of compression is desired, the limitations of the circuitry may be revealed. Because the compressor relies on the envelope of the signal to create a control signal, the time response of the compressor depends on the rate of change of the signal as well as on the time constant of the analysis circuitry. Many compressors allow adjustment of attack and release times. Other circuits (like the dbx 166) use the signal itself to determine the time response of the compression. In either case, a signal which has very large transients may cause overloading of the envelope measurement circuitry and create an audible pumping effect as the circuit recovers from the overload more slowly than the signal itself recovers. A way to avoid this is to use small amounts of compression at two or more stages in the recording process. This will make the signal very prominent without the overload problems because neither compressor is overdriven. Chaining dynamics processors is also a common technique. Sometimes, a limiter with a high threshold is used to reduce the big peaks in a signal followed by a compressor that reduces gain further without the overloading that the now-removed peaks might cause. Multiple compressors can also be used - each contributes a little compression. These techniques can produce a more natural compressed sound than a single processor alone.

Some compressors have a side-chain input. This allows a separate signal to control the gain of the compressed signal. For example, a vocal might be used to reduce the gain of a rhythm guitar so that the vocal always cuts through the mix. This technique is known as "ducking" and can be heard in commercials where the announcer's voice lowers the volume of the backing music. The effect allows for a sort of "automated mixing" in certain circumstances. It depends on finding optimal attack and release times as well as appropriate threshold and ratio to make the gain change sound natural, but sometimes this approach avoids the need for automation. Side chain inputs, commonly called "keying" inputs, can be used to allow selected signals to control the gain of several compressed signals. This technique is popular in electronic dance music (EDM), where the kick drum modulates much of the mix.

Many compressors allow filtering of the control signal to produce special frequency-related compression effects. The most popular of these effects is de-essing, which tends to reduce the sibilance from compressed speech or singing. By increasing the side-chain (control) signal gain in the frequency range of sibilance, between 2 and 8 kHz, for example, the s'es will be more compressed than the remainder of the vocal sound, resulting in a more natural sound. By decreasing the low frequency gain of the side-chain signal, the "pumping" effect of a bass drum on snare compression may be reduced.

It should be mentioned that compression may cause a signal to sound quite "weird" if heard in isolation, while in a mix it may still sound "correct". One must make final adjustments within the entire mix in order to get a complete picture of the overall sound. The masking phenomenon accounts for why this is so: parts of the compressed sounds interact with the other sounds in the mix, resulting in some masking of the compressed track itself. We do not hear this if we listen only to the compressed track in isolation. There is also the question of compressing individual elements of a mix or compressing the entire mix together. Both approached have their merits, but the end result will sound different and often both are applied. Compressing individual tracks makes them more "focused" in the mix, while compressing the whole mix tends to bring the elements together.

Compression is one of the more powerful techniques for manipulating sounds in a mixing environment and it is important to become familiar with its uses and limitations. Experimentation is the best way to appreciate what we can expect from dynamic range manipulation and what we cannot realistically achieve. In popular music, a lot of what separates professional-sounding mixes from more amateur-sounding ones is the effective use of compression and the use of recording techniques that are compatible with the compression that is to be used in mixing. There are a lot of compressors available, both in hardware and software, and each type has a distinctive sound. Solid-state devices like the Universal Audio 1176 compressor/limiter are good for aggressive drum compression while vacuum tube, optical gain cell devices like the LA-2A are favored for vocals. Although certain compressors are favored for a particular application, there are no absolutes as far a selecting an appropriate device.

Compressor plug-ins often try to emulate the behavior of the well-known analog compressors. As these approximations get better, they are growing in popularity, not unreasonably since you buy one plug-in and may then use several instances of it - something not possible with expensive hardware. Another advantage of digital compression programs is the ability to "look-ahead" at the signal to anticipate rapid changes in amplitude. Since there is inherent delay in all software effects, the mixing system must compensate for the delays so that all tracks in a mix occur at the same time at the output. While the entire output is slightly delayed, the individual elements remain in synchrony. This delay means that there is time to process a signal based on its future behavior, something not possible in the analog domain.

So how do we choose the right compressor when we have so many choices? When studios used hardware units, the decision was often made for us by the limited availability of devices. With plug-ins, there are so many available the choice can be baffling. The best long-term solution is to try the various options and decide which you like best for the range of signals you wish to mix. There is a lot of preference information on the internet, but there is no short cut from experimenting with the various possibilities to decide which implementations work best for the specific type of music you are working with at the moment. There are plenty of recording-related web sites that offer personal recommendations and these can be a clue about which software compressors might be worth consideration but they is no guarantee you will share the same opinions. As you try various plug-ins, you will soon find those that work best for you. You will also find that it is a moving target as new versions become available and offer different ways of treating the same tracks. The ability to adapt is an asset as technology moves rapidly and we do not wish to be trapped by options that disappear overtime as newer programs change what is available to us. It is the basic process we want to learn, not one particular method of accomplishing the task at hand.

Dynamic range processing gives the mix engineer great power to sculpt the individual sounds and the final combination as well. Learning how to maximize the effectiveness of these techniques requires experience and patience. The reward will be worth the effort.

Dynamic Range Expansion

Expansion is the complementary process to compression: it actually increases the dynamic range of a signal. When a signal falls below a threshold level, the gain is decreased. This causes quiet sounds to become even quieter, or inaudible, in the case of the noise gate. The most common form of expansion is the noise gate, which simply switches off signals that fall below the threshold. It is used to eliminate low-level noises from signals. The noise can be generated by the performance (as in breathing, rustling, etc. from human performers or noise from electronic sources like synthesizers or guitar amplifiers and pedals) or by imperfect recording techniques (tape noise or 60 Hz noise). Expanders also have attack and release controls, but here a fast attack opens up the gate immediately after a sound crosses threshold and release time reduces the gain of the sound some time after it drops below threshold. While a gate is simply on or off, expanders allow you to adjust the time course of the gain reduction and adjust how low the gain becomes while the signal is below threshold (it needn't be completely off).

While noise gates may be used at any stage of the recording chain, it can be risky to gate input signals. This is because misadjusted thresholds can cause the signal to be gated off when it should be allowed to pass through, clipping the beginning of phrases or sounds. Unlike compression, if expansion is applied after recording, it will eliminate some of the noise added in the recording process rather than increase it. Expansion can also be used to "tighten up" a signal by automatically shortening the decay time of a sound. Also, sends to reverb processors may be gated to create a special gated reverb effect or to eliminate low level parts of the signal from being processed. This can be helpful when applying heavy effects to a snare drum while there is significant bleed from the high-hat, for example. Expanders with selectable expansion ratios may be used to alter the dynamic range somewhat like a compressor, only in the opposite direction. You can set the gain for signals that drop below the threshold from none (as in a gate), to a ratio that only slightly reduces the gain of low-level sounds. When combined with a compressor or limiter, an expander gives an engineer the ability to finely sculpt the dynamics of sounds to better fit into a mix. Only through experimentation will you find the optimal combination of dynamics processing for a specific job, and this type of dynamics processing is not always necessary or appropriate.

Creative Uses of Dynamic Range Processing

In addition to allowing the engineer to make relatively natural-sounding changes to the dynamics of a program, these effects may also be used for special effects that are not necessarily natural. For instance, by feeding the control signal of a compressor from a different signal, one sound may be used to modulate the amplitude of another, a process known as "ducking". Since the compressor gain is controlled by the level of the second sound, the compressed signal's gain is now controlled by that other signal instead of its own level as in normal compression. While this effect is commonly used to allow an obnoxious announcer to hawk a product while the cheesy music's amplitude is automatically turned down as the words are spoken, it can be used to allow a vocalist to ride just above the backing instrumental tracks without having to carefully ride the gain of the music tracks. Using a similar technique with the gate allows one signal to gate another, an effect known as "keying": in fact, the control signal external input is often called the key input. This technique can tighten up a group of background singers, for instance.

Another interesting use of the control signal in a compressor is to insert an equalizer in series with the control signal input. This allows frequency-selective compression, since the equalizer can alter the frequency response of the control signal so as to allow different frequencies to cause more or less compression that they would otherwise. As previously mentioned, this can be used for de-essing and to prevent some frequencies from causing compression while other frequencies are compressed normally.

And inserting delay devices and distortion pedals into the control circuit can result in some "interesting" effects, although not generally as useful as the above techniques.

Digital Delay Effects

The advent of digital audio has, in addition to revolutionizing signal flow and recording processes, allowed for the development of sophisticated delay-based audio effects. These include reverberation, echo, chorus, and flange effects, all of which depend on signal delay. Before the availability of digital circuits to create these effects, they could only be synthesized using tape recorder delay and cumbersome mechanical devices like large tensioned steel plates with transducers to excite and capture the mechanical vibrations. Of course, the idea of using a real room as a reverberator, via speaker and microphone, is still of interest, but few recordists have that luxury. Even without a detailed understanding of digital signal processing theory, it is possible to use these devices, since their operation can be understood simply by examining the effect of combining direct and delayed signals.

Digital delay

The simplest of these devices is the digital delay or delay line. The time delay can be accurately determined and synchronized to the tempo of the music. The device samples the input signal and replays a delayed copy of the signal in combination with the direct (undelayed) signal. This provides a single discrete echo. If the delayed output is fed back to the input, the familiar series of decaying repeats is produced. The amplitude of the signal returned to the input determines how many repeats are heard. Digital delay is often used on vocals in a mix to give a larger vocal image. It is also commonly used on solo instruments like saxophone or guitar. The delay time can be adjusted to the tempo of the music, so that a discrete echo falls on an eighth note, for example. This results in an echo that is not obtrusive and produces a sense of space without being obvious. To determine the delay time (in milliseconds) for an eighth note, divide 30,000 by the tempo (in beats[quarter notes]/minute). Delay times of quarter notes, sixteenth notes, and eight-note triplets may sometimes be used, depending on the program material. Stereo delay can create a simulated stereo image by delaying the signal to the left and right channels by different times. Delay times in the range of 20-50 milliseconds work well for this on rhythm instruments, but longer times can be used as special effects on solos. It is recommended to check the effect in mono as it is possible to create some odd-sounding results when the multiple delays are mixed in the system rather than in the air as they are generally combined in the air between two loudspeakers.

Flanging

As we remember from multiple microphone technique, when time-delayed and direct signals are combined, a comb-filter effect is created. Generally, this is undesirable; however, the effect can be used to "spice-up" certain sounds. If the delay time is constantly altered slightly, a rich sweeping filter is created. This is known as flanging. The name derives from the original way of creating the effect: using a second tape recorder to delay the sound and slightly slowing the machines by placing hands on the reel flanges. Now, digital delay devices allow a slow oscillator to control the delay time. In addition, the depth of the effect can be controlled by changing the balance of the delayed and direct signals. For flanging, the delay time is in the range of .5 to 35 milliseconds and the modulation rate (which changes the delay time) is in the range of 1 to 10 Hz. A stereo effect can be created by placing the direct + delayed signal in one channel and direct - delayed signal in the other.

Chorus

The chorus effect is qualitatively similar to the flange effect; it attempts to simulate the effect of several distinct sound sources producing nearly the same sound, like a choir does with multiple singers in unison. Unlike flanging, chorusing often employs amplitude modulation as well to simulate the way singers' volumes vary in time. Electronically, it is achieved using small random variations of the time delays and amplitudes and uses several separate such channels which are recombined in stereo to produce a very rich sound. Modulation rates are longer than for flanging, typically 0.1-0.5 Hz, with similar delay times of 1 to 50 milliseconds.

Reverberation

By far the most complicated of the digital effects is simulated reverberation, which attempts to create artificially the complicated sound field created by sounds reflecting off the walls, ceiling, and floor of a room. Sound, traveling at about 1100 feet/second, bounces off surfaces and slowly decays, producing what we recognize as the sound of a room. Each bounce alters the spectral content of the sound, as frequencies are absorbed or reflected depending on the physical nature of the reflective surfaces. Just after the onset of a sound, discrete echoes, known as early reflections, are audible. Soon, however, these reflections build into a less discrete, but denser sound. The complexity of this process has made reverberation programs extremely complicated, and early attempts at synthetic reverberation were obviously poor imitations of the real sound. Newer devices have improved dramatically, to the point where it is difficult to tell whether real or synthetic reverb has been employed in a recording. Digital reverbs allow the user to choose between different programs, each of which seeks to duplicate the behavior of one type of sonic environment, such as room, hall, auditorium, etc. Within a program, adjustment of decay time, early reflections, pre-echo time, reverberation equalization and many other parameters may be used to tailor the sound, however the basic program still simulates the same physical space. While the spatial size may be adjustable, the basic character of the sound remains similar within a program. The time required for the reverberations to fall to -60 dB is known as the RT60 and is often referred to as the reverberation time.

An alternative to simulating the physics of a reverberant space is sampling the impulse response of real spaces and using a digital process called convolution to simulate the sound. Convolution mathematically combines the input signal with the sound of the space captured by the impulse response to produce the reverberation that would have occurred in the actual space. Some of the character can often be adjusted in the process.

Convolution reverbs are increasingly popular.

In addition to digital reverberators, there exist mechanical reverberation systems, using springs, plates, and chambers to create reverb effects. These systems were popular before digital systems became available and are frequently simulated using digital techniques. The simplest systems use springs connecting transducers to create the boingy reverb still used in many guitar amplifiers. More sophisticated plate reverbs use tensioned steel plates (several feet in length) with transducers carefully placed to drive the plate and convert the vibrations back into electronic signals. They are very heavy, bulky, and cannot be located where there are physical vibrations or extraneous sounds. They often sound good on vocals and drums, so they are also simulated on many digital reverberation systems. Chamber reverbs used actual chambers to create reverberations, via speaker and microphone. These have a characteristic sound that is sometimes desirable and can be selected on many digital reverbs.

Precisely which controls are present on reverb processors vary according to the complexity of the device, often in proportion to the price. Many low priced machines allow for simple programming of reverb time, equalization, and pre-delay only, while more expensive machines allow many, more complex parameters to be adjusted. It is important to understand thoroughly each specific device in order to maximize the performance of the reverberation algorithms employed. Reverberation can be used to simulate realistic sounds, as in the case of enhancing live stereo recordings, or it can be used as a special effect in multi-track recordings and on synthesized sounds to impart a sense of realism. Each of these applications requires a bit of experimentation to obtain the best possible sound.

Making choices with effects

There is a tendency to think that if an effect sounds good it should be used everywhere. The truth is that sparing use of effects is often more noticeable and preferable to their overuse. Dynamics processing is widely used but the intent is often to make things sound natural rather than obviously enhanced. Particularly with the wide availability of plug-in effects, it's tempting to over apply them. How we listen to our mixes also affects how much of an effect we choose to use - headphones enhance our perception of reverberation, for example, leading to less reverb in the mix than we would choose if listening on loudspeakers. Deciding on which effects to use and how much is an artistic decision that often changes as we become more adept at mixing and learning from our past choices (gated snare reverb, anyone?) It is an art we can continue to refine though out our lives.