

Dynamic Range Processing and Digital Effects

Dynamic Range Compression

Compression is the automatic reduction of gain as the signal level increases beyond a preset level called the threshold. It is an automatic version of “gain riding”, where one might manually lower a fader as the signal grows larger to maintain a more constant output level. For signals above the threshold, the increase in output amplitude for a given increase in input amplitude is known as the compression ratio. Below threshold, the ratio is 1:1 while above threshold the ratio may be set anywhere up to ∞ :1, which is considered limiting since in that case the output signal does not increase at all as the input signal increases. The time required for the gain to reduce when a signal crosses the threshold is referred to as the attack time and the time it takes to return 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. The effect of compression is to reduce the overall dynamic range of a signal. This can be for musical effect, to make a sound stand out more in a complex mix, or for the purposes of noise reduction in noisy recorders or transmission systems.

Compressors are variable-gain amplifiers that use a control signal derived from the amplitude envelope of the signal being compressed. The control signal is generated either by a rectifier and filter or through more elaborate circuitry that is able to calculate the R.M.S. value of signal amplitude. The accuracy of the amplitude envelope so generated affects the sound of the compressor. There are two topologies for compressor circuitry, one in which the input signal is fed to the detector circuitry (feed-forward) and one in which the output signal is sent to the detector (feed-back). Each type has its characteristic sound and they behave somewhat differently although they both result in a compressed output. Many compressors allow external access to the detector circuitry, allowing creative uses that will be discussed later.

You can think of the use of compression in terms of viewing the skyline of a big city, where 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 and increase the height of shorter ones so that all the buildings attain more similar heights. You can imagine what this would look like: more of the buildings would be visible although some of them would still block others. Similarly, compressing sounds makes them more equally audible in the mix but there is still a limit to how much of the quietest sounds we will be able to perceive.

Compression may be employed at any stage in the recording process, although it can sound 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 later contain the compressed added noise that would appear if the signal were to be compressed after recording. Unfortunately, a poor adjustment of the compressor can cause undesirable changes to the signal that are then committed to storage, so compression should be used carefully, if applied before a signal is recorded, to avoid an overly "squashed" sound. This 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.

The difference between a limiter and a compressor, while one of degree, is 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. A 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 simply made louder. Limiters are useful on singers as they are being recorded to prevent recorder overloads from vocal plosives. They are also increasingly used for making mixes louder, a practice that has definite potential for abuse.

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 that has very large transients may cause overloading of the 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 smaller amounts of compression at two or more stages in the recording and mixing process. This can make the signal sound very prominent without the overload problems because neither compressor is overdriven.

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 "normal" sounding compression. 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 and/or limiting might cause a signal to sound quite "weird" if heard alone, while in a mix it sounds "correct". One must make final adjustments within the entire mix in order to get a complete picture of the overall sound. There is also an issue of compressing individual elements of a mix and compressing the entire mix together. Both approaches 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 as selecting an appropriate device.

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. 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 guitars) or by imperfect recording techniques (tape 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 is especially important 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.

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 than 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. 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. 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 set 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 creates 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.

Flange

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 an oscillator to control the delay time. In addition, the depth of the effect can be controlled by changing the balance between 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. Electronically, it is achieved using small random variations of the time delay and uses several separate delay 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. Unlike flanging, chorusing often employs amplitude modulation as well to simulate the way singers' volumes vary in time.

Reverberation

By far the most complicated of the digital effects is simulated reverberation, which attempts to create electronically the 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.

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.