5 Compression, processing and effects

> Compression plays a major part of my sound. I have them patched across every output of the desk.
>
> Armand Van Helden

Before looking at the principles behind designing sounds for different music genres, it's first necessary to examine the various processors and effects that are available and understand how they are accessed through the mixing desk. This is because effects are often used as part of the sound design process as well as during the mixing stage, so it pays to understand what they are, how they affect the audio and how a mixing desk can determine the outcome of the effect. This chapter concentrates on the behaviours of the different processors and effects that are widely used in the design and production of dance music, including reverb, chorus, phasers, flangers, delay, EQ, distortion, gates, limiters and compressors.

Of all the effects and processors available, a compressor is the most important tool in any dance musician's arsenal, so a thorough understanding of it is essential. Without compression, drums appear wimpy in comparison to the chest-thudding results heard in professionally produced music, mixes can appear lacking in depth, and basses, vocals and leads can lack any real presence. Despite its importance, the compressor remains the least understood processor of them all, largely because if properly applied the effect should be particularly subtle.

Compression theory

The fundamental application of a compressor is to reduce the dynamic range of a performance and this is vital when working with any form of music. As touched upon in the previous chapter, whenever you record any source you should aim to capture the hottest signal possible so that you can avoid artificially increasing the gain. This is because if you record a source that's too low in gain and then artificially increase it, not only will it increase the volume of the recorded source, it'll also increase the background noise. To prevent this, you need to record a signal as loud as possible, but vocals and 'real' instruments have a huge dynamic range, and the difference between clipping and a nominal level can be as much as 20 dB.

Consequently, it's impossible to set a good average recording level with so much dynamic movement, since if you set the recording level to capture the quiet sections, when it becomes louder the recording will clip. Conversely, setting the recorder so that the loud sections do not clip, any quieter sections will be exposed to more residual noise. Of course, you could sit by the recording fader and increase or decrease the recording levels depending on the section being recorded, but this would mean that you need lightning reflexes. Instead, it's much easier to employ a compressor to control the levels automatically. By routing the source sound through a compressor and then into the recorder, you can set a threshold on the compressor so that any sounds that exceed this are automatically pulled down in gain, thus allowing you to record at a more substantial volume overall.

A compressor can also be used to control the dynamics of a sound while mixing. For example, a dance track that uses a real bass guitar will have a fairly wide dynamic range, even if it was compressed during the recording stage. This will cause problems within a mix because if the gain is adjusted so that the loudest parts fit well within the mix, the quieter parts may disappear behind other instrumentation. Conversely, if the fader is set so that quieter sections can be heard over other instruments, the loud parts could be too prominent. Using compression more heavily on this sound during the mixing stage, the dynamic range can be restricted, further allowing the sound to sit better in a mix.

Although these are the key reasons why compressors were first introduced, a compressor's action also has more creative applications that are especially suited towards dance musicians. Since the signals that exceed the threshold are reduced in gain, the parts that do not exceed the threshold aren't touched, so they remain at the same volume as they were before compression. In other words, the difference in volume between the loudest and quietest parts of the recording are reduced, which means that any uncompressed signals will become louder relative to the compressed parts. This effectively boosts the average signal level, which in turn not only allows you to push the volume up further, but also makes it sound louder.

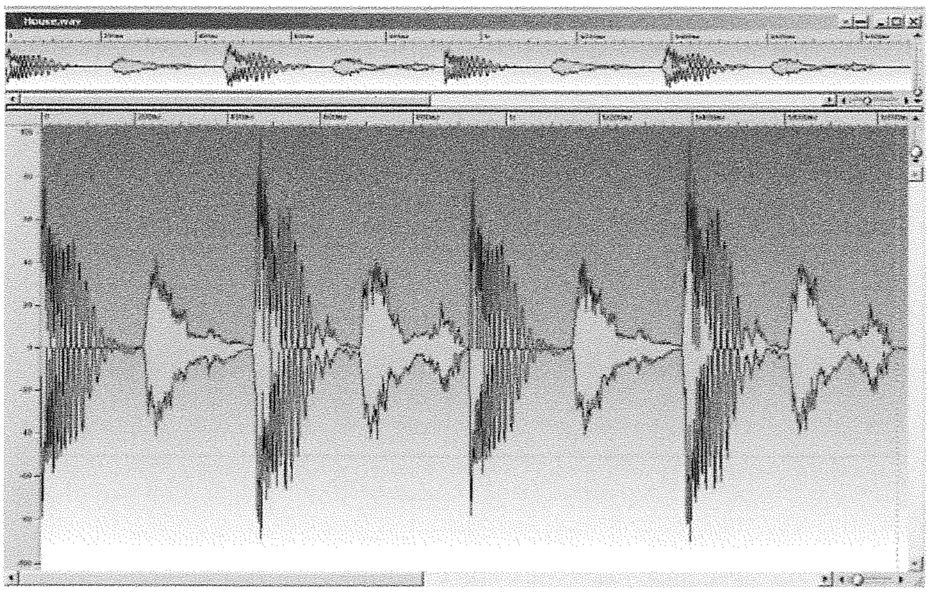

Figure 5.1 A drum loop waveform before compression.

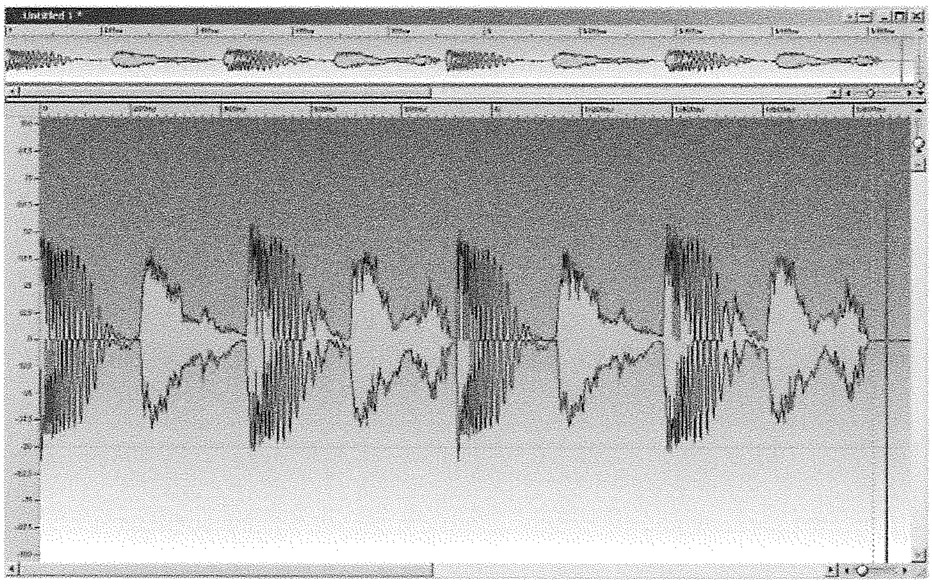

Figure 5.2 A drum loop waveform after compression (note how the volume difference (dynamics) between instruments has changed).

Note: After reducing the dynamic range of audio it may be perceived to be louder without actually increasing the gain. This is because we determine the overall volume of music from the average volume (measured in Root Mean Square), not from the transient peaks created by kick drums.

What's more, applying heavy compression to certain elements of a dance mix can change the overall character of the timbre, often resulting in a warmer, smoother and rounder tone – a sound typical of most dance tracks around today.

It is important to note, however, that although there are numerous publications stating the compression settings to use, there are no generic settings for any particular genre of music and its use depends entirely on the timbres that have been used. For instance, a kick drum sample from a record will require a different approach to a kick drum constructed in a synth, while if sampled from a CD, synthesizer or film, different approaches are required again. Therefore, rather than dictate a set list of compression settings, you can achieve much better results by knowing exactly what effect each control will have on a sound and how these are used to acquire the sounds typical of each genre.

Threshold

The first control on a compressor is the threshold, which as touched upon sets the signal level where the compressor will begin squashing the incoming signal. These are commonly calibrated in decibels and will work in direct relationship with a gain reduction meter to inform you of how much the compressor is affecting the incoming signal. In a typical recording situation this control is set so that the average signal level always lies just below the threshold, and if any exuberant parts exceed it, the compressor will jump into action and the gain will be reduced to prevent any clipping.

Ratio

The amount of gain reduction that takes place after a sound exceeds the threshold is set using a ratio control. Expressed in ratios, this control is used to set the dynamic range the compressor affects, indicating the difference between the signals entering the compressor that exceed the threshold to the levels that come out of the other end. For example, if the ratio is 4:1, every time the incoming signal exceeds the threshold by 4 dB, the compressor will squash the signal so that there is only a 1 dB increase at the output of the compressor. Similarly, if the ratio set were 6:1, an increase at the compressor's output of 1 dB will occur when the threshold is exceeded by 6 dB and likewise for ratios of 8:1, 10:1 and so on. Subsequently, the gain reduction ratio always remains constant no matter how much compression takes place. In most compressors these range from 1:1 up to 10:1 and may, in some cases, also offer infinity:1.

From this we can determine that if a sound exceeds a predefined threshold, the compressor will squash the signal by the amount set with the ratio control. The problem with this approach, however, is that we gain a significant amount of information about sounds from their initial attack stage and if the compressor jumps in instantaneously on an exceeded signal it will squash the transient, which reduces its high-frequency content. For instance, if you were to set up a compressor to squash a snare drum, the compressor will clamp down on the attack stage, which in effect diminishes the initial transient, reducing it to a 'thunk'. What's more, this instantaneous action will also appear when the sound drops below the threshold again as the compressor stops processing the audio. This can be especially evident when compressing low-frequency waveforms such as basses, since compressors can apply gain changes during a waveform's period. That is, if a low-frequency waveform, such as a sustained bass note, is squashed the compressor may treat the positive and negative states of the waveform as different signals and continually activate and deactivate. The result of this is an unpleasant distortion of the waveform. To prevent this from occurring, compressors will feature attack and release parameters.

Attack/release

Both these parameters behave in a similar manner to those on a synthesizer, but control how quickly the volume is pulled down and how long it takes to rise back to its nominal level after the signal has fallen below the threshold. In other words, the attack parameter defines how long the compressor takes to reach maximum gain reduction while the release parameter determines how long the compressor will wait *after* the signal has dropped below the threshold before processing stops.

This raises the obvious question that if the attack is set so that it doesn't clamp down on the initial attack of the source sound it could introduce distortion/clipping before the compressor activates. While this is true, in practice very short, sharp signals do not always overload an analogue recorder, since these usually have enough headroom to let small transients through without introducing any unwanted artefacts. This isn't the case with digital recorders, though, and any signals that are beyond the limit of digital can result in clipping, so it's quite usual to follow a compressor with a limiter or, if the compressor features a knee mode, set it to use a soft knee.

Soft/hard knee

All compressors will utilize either soft or hard knee compression, but some will offer the option to switch between the two modes. These are not controllable parameters but dictate the shape of the envelopes curve, and hence the characteristic of how the compressor behaves when a signal approaches the threshold. So far we've considered that when a signal exceeds the threshold the compressor will begin to squash the signal. This immediate action is referred to as hard knee compression. Soft knee, on the other hand, continually measures the incoming signal and when it approaches 3–14 dB (depending on the compressor) towards the current threshold, the compressor starts to apply the gain reduction gradually.

Generally, this will initially start with a ratio of 1:1 and as the signal grows ever closer to the threshold it's gradually increased until the threshold is exceeded, whereby full gain reduction is applied. This allows the compressor's action to be less evident and is particularly suitable for use on acoustic guitars and wind instruments, where you don't necessarily want the action to be evident.

Note: The action of the knee is entirely dependent on the compressor being used and some can be particularly long, starting 12 dB before the threshold, while others may start 3 dB before. As a matter of interest, 6–9 dB soft knees are considered to offer the most natural compression for instruments.

Peak/RMS

Not all compressors feature knees, so short transient peaks can sometimes catch the compressor unaware and 'sneak' past unaffected. This is obviously going to cause problems when recording digitally, so many compressors will implement a switch for peak or RMS modes. Compressors that do not feature these two modes will operate in RMS and this means that the compressor will detect and control signals that stay at an average level rather than the short sharp transient peaks. As a result, no matter how fast the attack may be set, there's a chance that the transients will overshoot the threshold and not be controlled. This is because by the time the compressor has figured out that the sound has exceeded the threshold it's too late – the peak's been and gone again. Therefore, to control short transient sound such as drum loops, it's often prudent to engage the peak mode. With this the compressor becomes sensitive to short sharp peaks and clamps down on them as soon as they come close to the threshold, rather than after they exceed it. By doing so, the peak can be controlled before it overshoots the threshold and creates a problem.

While this can be particularly useful when working with drum and percussion sounds, it can create havoc with most other timbres. Keep in mind that many instruments can exhibit a particularly short, sharp initial attack stage and if the compressor perceives these as possible problems it'll jump down on them before they overshoot. In doing so, the high-frequency elements of the attack will be dulled, which can make the instrument appear less defined, muddled or lost within the mix. Therefore, for all instruments bar drums and percussion, it's advisable to stick with the RMS mode.

Make-up gain

The final control on a compressor is the make-up gain. If you've set the threshold, ratio, attack and release correctly, the compressor should compress effectively and reduce the dynamics in a sound, but this compression will also reduce the overall gain by the amount set by the ratio control. Therefore, whenever compression takes place you can use the make-up gain to bring the signal back up to it's pre-compressed volume level.

Side-chaining

Alongside the physical input and output connections on hardware compressors, many also feature an additional pair of inputs known as side-chains. By inputting an audio signal into these, a sound's envelope can be used to control the action that the compressor has on the signal entering the normal inputs. A good example of this is when a radio DJ begins talking over a record and the volume of the record lowers so their voice becomes audible, then when they stop speaking the record returns to its original volume. This is accomplished by feeding the music through the compressor as normal, but with the microphone connected into the side-chain. This supersedes the compressor's normal operation and uses the signal from the side-chain rather than the threshold as the trigger. Thus, the compressor is triggered when the microphone is spoken into, compressing (in effect lowering the volume of the music) by the amount set with the ratio control. This technique should only be viewed as an example to explain the process, though, and more commonly side-chaining is usually used to make space in a mix for the vocals.

In a typical mix, the lead sound will occupy the same frequencies as the human voice, resulting in a cluttered mid range if the two are to play together. This can be avoided if the lead mix is fed into the main inputs of the compressor while the vocal track is routed into the side-chain. With the ratio set at an appropriate level (depending on the tonal characteristics of the lead and vocals) the lead track will dip when the vocals are present, allowing them to pull through the mix.

Hold control

Most compressors that feature a side-chain are likely to also have an associated 'hold' control on the fascia or employ an automated hold function. This is employed because a side-chain measures the envelope of the incoming signal and if both the release and attack are too fast, the compressor may respond to the cycles of a low-frequency waveform rather than the actual envelope. As touched upon previously, this can result in the peaks and dips of the waveform activating and deactivating the compressor, resulting in distortion. By using a hold the compressor is forced to wait a finite amount of time (usually 40–60 ms on automated hold) before beginning the release phase, which is longer than the period of a low-frequency waveform.

Practical compression

Despite the amount of control offered by the average compressor, they are relatively simple to set up for recording audio. As a generalized starting point, it's advisable to set the ratio at 4:1 and lower the threshold so that the gain reduction meter reads between −8 and −10 dB on the loudest parts of the signal. After this, the attack parameter should be set to the fastest speed possible and the release set to approximately 500 ms. Using these as preliminary settings, they can then be adjusted further to suit any particular sound.

Note: It's advisable that compression is applied sparingly during the recording stage, because once applied it cannot be removed. Any exuberant parts of the performance should be prevented from forcing the recorder's meters into the red while also ensuring that the compressor is as transparent as possible. Solid-state compressors are more transparent than their valve counterparts and so are better suited for this purpose.

As a general rule of thumb, the higher the dynamic range of the instrument being recorded, the higher the ratio and the lower the threshold settings need to be. These settings help to keep the varying dynamics under tighter control and prevent too much fluctuation throughout the performance. Additionally, if the choice between hard or soft knee is available, the structure of the timbre should be taken into account. To retain a sharp, bright attack stage, hard knee compression with an attack setting that allows the initial transient to sneak through unmolested should be used, provided of course that the transient is unlikely to bypass the compression. In these instances, and to capture a more natural sound, soft knee compression should be used.

Finally, the release period should be set as short as possible, but not so short that the effect is noticeable when the compressor stops processing. After setting the release at 500 ms, the time should be continually reduced until the processing is noticeable, then increased slowly until it isn't.

Note: Some compressors feature an auto-mode for the release that uses a fast release on transient hits and a slower time for smaller peaks, making this task easier.

Table 5.1 Compression settings

Compression settings	Ratio	Attack parameter (ms)	Release parameter (ms)	Gain reduction (dB)	Knee
Starting settings	5:1 to 10:1	1–10	40–100	−5 to −15	Hard
Drum loop	5:1 to 10:1	1–10	40–100	−5 to −15	Hard
Bass	4:1 to 12:1	1–10	20 or auto	−6 to −13	Hard
Leads	2:1 to 8:1	3–10	40 or auto	−8 to −10	Hard
Vocals	2:1 to 7:1	1–7	50 or auto	−3 to −10	Soft
Brass instruments	4:1 to 10:1	1–7	30 or auto	−8 to −13	Hard
Electric guitars	8:1 to 10:1	2–7	50 or auto	−5 to −12	Hard
Acoustic guitars	5:1 to 9:1	5–20	40 or auto	−5 to −12	Hard

The settings shown in Table 5.1 are naturally only starting points and too much compression should be avoided during the recording stage, something that can only be accomplished by setting both the ratio and threshold controls carefully. This involves setting the compressor to squash audio but ensuring that it stops processing and that the gain reduction meter drops to 0 dB (i.e. no signal is being compressed) during any silent passages.

As a more practical example, with a simple four-to-the-floor kick running though the compressor and the ratio and threshold controls set so that the gain reduction reads −8 dB on each kick, it's necessary to ensure that the gain reduction meter returns to 0 dB during any silent periods. If it doesn't, then the loop is being over-compressed. If the gain reduction only drops to −2 dB during the silence between kicks, then it makes sense that only 6 dB of gain reduction is actually being applied. This means that every time the compressor activates it has to jump from 0 to 8 dB, when in reality it only needs to jump in by 6 dB. This additional 2 dB of gain will distort the transient that follows the silence, making it necessary for the gain reduction to be adjusted accordingly.

Creative compression

While it is generally worth avoiding any evident compression during the recording stage, deliberately making the compressor's action evident forms a fundamental part of creating the typical sound of dance music. To better describe this, we'll use a drum submix to experiment upon. This is a typical dance drum loop consisting of a kick, snare, closed and open hi-hats.

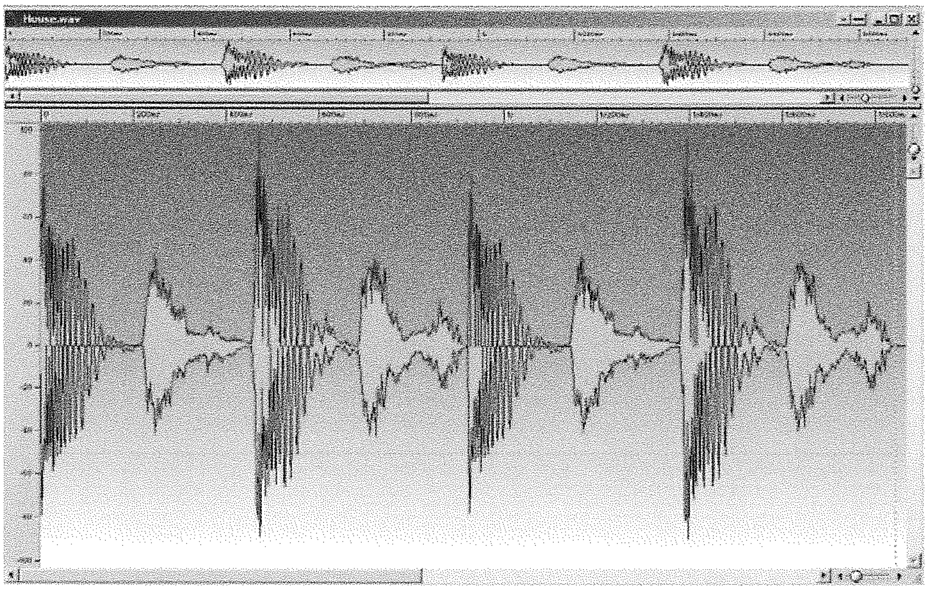

Figure 5.3 The waveform of the drum pattern.

Note: On the audio CD you can hear the basic drum pattern.

It's clear from listening to the drum loop (or simply looking at its waveform) that the greatest energy – that is, the loudest part of the loop – is derived from the kick drum. With this in mind, if a compressor is inserted across this particular drum loop and the threshold is set just below the peak level of the loudest part, each consecutive kick will activate the compressor.

If the release time is set short, the compressor will activate on the kick then quickly release when the kick drops below the threshold. This results in rapid changes in volume, producing a pumping effect as the compressor activates and deactivates on each kick. This is called 'gain pumping' which, although frowned upon in some areas of music, is deliberately used in dance and popular music to give a track a more dynamic feel. The exact timing of the release will depend entirely on the tempo of the drum loop and must be short enough for the compressor to recover before the next kick. Similarly, the release must be long enough for this effect to sound natural, so it's best to keep the loop repeated over four bars and experiment with the attack and release parameters until the required sound is achieved.

Note: On the audio CD you can hear the effect the release has on a sound.

If we expand on this principle further and add a bass line, pad and chords to the previously compressed loop and compress it again in the same way (i.e. with a short release), the entire mix will pump energetically. As the kick is still controlling the compressor (and provided that the release isn't too short or long), every time a kick occurs the rest of the instruments will drop in volume, which accentuates the overall rhythm of the piece.

Note: On the audio CD you can hear the effect of compression on the bass, pad and chord structure.

Gain pumping is also useful when applied across the whole mix, even though each element in the mix may already have been compressed in isolation. Gain pumping across the whole mix is used to balance areas in the track where instruments are dropped in and out. When fewer instruments play, the gain of the mix will be perceived as lower than when all the instruments play simultaneously. The overall level can be controlled by strapping a compressor across the mixing desk's main stereo buss (more on this in later chapters), to make the mix pump with energy.

Gain pumping across the entire mix ('mix pumping') should be applied with caution, because if the mix is pumped too heavily it will sound strange. Setting a 20–30 ms attack with a 250 ms release and a low threshold and ratio to reduce the range by 2 dB or so should be sufficient to produce a mix that has the 'right' feel. That said, there are many instances, for example on *MTV Dance* or the more recent *Rapture* channel, where gain pumping is so forceful that it's blindingly obvious, so the degree to which you employ these techniques in your own mix is entirely up to you.

Compression can also be used on individual sounds to change the tonal content of a sound. For example, by using heavy compression (a low threshold and high ratio) on an isolated snare drum, the snare's attack avoids compression but the decay is squashed, which brings it up to the attack's gain level. This technique is often employed to create the snare '*thwack*' typical of trance, techno and house music styles. Similarly, if a deeper, duller, speaker-mashing, kick drum 'thud' is required, the compressor's attack should be set as short as possible so that it clamps down hard on the initial attack of the kick. This eradicates much of the initial high-frequency content and as the volume is increased with the compressor's make-up control a deeper and much more substantial 'thud' is produced.

Note: Because hip-hop, rap, house and big beat often have particularly loud bass elements that play consecutively with the kick, the different sounds can conflict, muddying the bottom end of the mix. This can be prevented by feeding the kick drum separately into the side-chain inputs of the compressor, with the ratio set to 3:1, with a fast attack and medium release (depending on the sound). If the bass is fed into the compressor's main inputs, every time the kick occurs the bass will drop in volume, making space for the kick and thereby preventing any conflicts.

It is, however, important to note that the overall quality of compression depends entirely on the type of compressor being used. Compressors are one of two types: solid-state or valve. Solid-state compressors use digital circuitry throughout and will not tend to pump as heavily or sound as good as those that utilize valve technology. Some valve compressors will be solid-state in the most part, using a valve only at the compressor's make-up gain. Solid-state compressors are usually more transparent than their valve-based counterparts and are usually used during the recording stages. Valve compressors are typically used after recording to add warmth to drums, vocals and basses, an effect caused by small amounts of second-order harmonic distortion that are introduced into the final gain circuitry[1] of the compressor. This distortion is a result of the free movement of electrons within the valves and occurs at exactly twice the frequency of the amplified signal. Despite the fact that this distortion only contributes 0.2% to the amplified signal, the human ear (subjectively!) finds it appealing.

More importantly, the characteristic warmth of a valve compressor differs according to the model of valve compressor that is used, as each will exhibit different characteristics. These differences and the variations from compressor to compressor are the reasons why many dance producers will spend a fortune on the right compressor and also why it isn't uncommon for producers to own a number of both valve and solid-state types.

Note: Most dance producers agree that solid-state circuitry tends to react faster, producing a more defined, less forgiving sound, while valve compressors add warmth that improve the overall timbre.

While it isn't essential to know why these differences exist from model to model, it's worth knowing which compressor is most suited to a particular style of work. Failing that, it also makes for an excellent conversation stopper, so what follows is a quick rundown of the five most popular methods of compression:

- Variable MU.
- Field effect transistor (FET).
- Optical.
- VCA.
- Computer-based digital.

Variable MU

The first compressors ever to appear on the market were called variable MU units. This type of compressor uses valves for the gain control circuitry and does not have an adjustable ratio control. Instead of an adjustable control, the ratio is increased in proportion to the amount of the incoming signal that exceeds the threshold. In other words, the more the level overshoots the threshold, the more the ratio increases. While these compressors do offer attack and release stages, they're not particularly suited towards material with fast transients, even with their fastest attack settings. Due to the valve design, the valves run out of dynamic range relatively quickly, so it's unusual to acquire more than 15–20 dB of gain reduction before the compressor runs out of energy. Nonetheless, variable MU compressors are renowned for their distinctive, phat, lush character, and can work magic on basses and pumping dance mixes. The most notorious variable MU compressors are made by Manley and can cost in excess of £3500.

---

[1] John Ambrose Fleming originally developed the valve in 1904, but it was 2 years later before Lee De Forest constructed the first Triode configuration. Edwin Howard Armstrong then used this to create the first ever valve amplifier in 1912.

Field effect transistor (FET)

FET compressors are often recognized as simulating the action of a typical valve compressor because they use a field effect transitor to vary the gain. They provide incredibly fast attack and release stages, making them an excellent choice for beefing up kick and snare drums, electric guitars, vocals and synthesizer leads. While they suffer from a limited dynamic range, if they're pushed hard they can pump very musically and are perfectly suited for gain pumping a mix. The only major problem is getting your hands on one. Original FETs are as rare as rocking horse manure and consequently second-hand models are incredibly expensive. Reproduction versions of the early FETs, such as the UREI 1176LN Peak Limiter (approx. £1800) and the LA Audio Classic II (approx. £2000), are a worthwhile alternative.

Optical

Optical (or 'opto') compressors rely on a light bulb and phototransistor assembly to automatically adjust the gain of the compressor. The phototransistor monitors the current illumination level of the light bulb and subsequently adjusts the gain to suit. Because the phototransistor must monitor the light-bulb before it takes any action, some latency is created in the compressor's response, so the more heavily the compression is applied, the longer the envelope times tend to be. Consequently, most optical compressors utilize soft knee compression. This creates a more natural attack and release, but also means that the compressor is not quick enough to catch many transients. Despite this, optical compressors are great for compressing vocals, basses, electric guitars and drum loops, providing that a limiter follows the compression (limiters will be explained later).

There are plenty of opto compressors to choose from, including the ADL 1500 (approx. £2500), the UREI LA3 and UREI Teletronix LA-2A (approx. £2900 each), the Joe Meek C2 (approx. £250) and the Joe Meek SC2.2 (approx. £500). Both Joe Meek units sound particularly smooth and warm considering their relatively low prices, and for the typical gain pump synonymous with dance then you could do worse than to pick up the SC2.2.

> Note: Incidentally, all Joe Meek's compressors are green because, after designing his first unit, he decided to spruce it up by colouring it with car aerosol paint and bright green was the only colour he could find in the garage at the time.

VCA

VCA compressors offer the fastest envelope times and highest gain reduction levels of any of the compressors covered so far. These are the compressors most likely to be found in a typical home studio. As with most things, the quality of a VCA compressor varies wildly in relation to its price tag. Many of the models aimed at the more budget conscientious musician reduce the high frequencies when a high gain reduction is used, regardless of whether you're clamping down on the transients or not. When used on a full mix these also rarely produce the pumping energy that is typical of models that are more expensive. Nonetheless, these types of compressor are suitable for use on any sound. The most widely celebrated VCA compressor is the Empirical Labs Stereo Distressor (approx. £2500), which is a digitally controlled analogue compressor with VCA,