



**Return to Articles Index**



**Back to Home Page**

---

## Contents

- [Introduction](#)
  - [Compression Versus Limiting](#)
  - [Additional Features](#)
  - [Why Use Compression](#)
  - [Contra Indications](#)
  - [Peak to Average Ratio & Dynamic Range](#)
  - [Typical Specifications](#)
  - [What To Listen For](#)
  - [What About Lossy Data Compression?](#)
  - [Even Worse Things That Can Happen](#)
  - [References](#)
- 

## Introduction

The term compression has several meanings in audio - there is lossy data compression (e.g. MP3), lossless data compression (a wave file compressed using a Zip program), and level compression and/or limiting. In this article, I concentrate on the last form, although a few words about MP3 are in order (later).

I recently read a wonderful article by a mastering engineer by the name of Bob Katz (see [references](#), below). Bob was adamant that many producers, engineers and musicians have joined a new race to see who's CD can be the "hottest" (i.e. loudest). There is a mistaken belief that this makes the sound more exciting - it doesn't - it makes it boring, and very tedious to listen to ...

Just like commercials on TV, which are compressed to within an inch of their lives. Does anyone find the sound satisfying? I'd be very surprised to hear someone (other than the producer) say yes - they are annoying, seem to be much louder than the programme you were enjoying, and cause a great many people to hit the mute button on the remote the instant they start (I'm one of them :-)

---

## Compression Versus Limiting

So what is the difference between these two effects? Depending on your outlook, either not much or a great deal. A limiter is usually set to a fixed threshold, and any signal that attempts to exceed the threshold is pulled back (attenuated) by exactly that amount needed to maintain the predetermined level. If the input gain is set way high, then all signals below the threshold (including noise) are boosted - in the extreme case so everything is the same volume (again including noise!). Limiters are "hard" compressors - the absolute level is fixed, and the compression ratio may be as high as 100:1 or more. This means that the input signal must increase by 100 "units" to make the output increase by one unit. Many limiters claim that the ultimate compression ratio is infinity, however this is probably an over estimation of the true figure.

A compressor uses much the same (or at least similar) circuitry, but will typically boost the level of signals below a preset threshold by a predetermined amount, and reduce the level of signals above the same threshold. Compression ratios of perhaps 2:1 are common, so the output will rise (or fall) by one unit for every two units of input change. A 50dB dynamic range is therefore reduced to 25dB (from softest to loudest signal). Unlike limiting, the compression threshold is typically set lower than the peak level - the actual threshold level could be anything from +8dB to -40dB, depending on the effect desired.

For example, a guitar (after suitable amplification) may produce transients of perhaps 5V peak, but yield an average level of only 500mV. This represents 10:1 or 10dB peak to average ratio. The peaks are produced as the pick strikes (or releases if you prefer) the strings, and the average is predominantly the normal decay of the note before the next is played. The VU meter (I refer here to real ones, not the stupid things you so often see that bear no resemblance to a proper VU meter) gives a good indication of the average level, and therefore the perceived loudness (VU = Volume Unit).

A PPM (Peak Programme Meter) shows the peak levels - no surprise there. Some meters, mainly electronic versions, provide both indications. A bar shows the average (VU) level, and a dot that "sticks" at some higher level shows the peak amplitude.

By using compression, the same guitar may have the maximum level reduced to perhaps 1V - the average level will now be higher as well (softer sounds are amplified, loud ones attenuated). Peak to average ratio may be reduced to 6dB or less, and the note will seem to just hang on forever ... well not quite, but you get the idea.

This sort of compression is common on percussion, strings, vocals - in fact almost anything. It is appropriate if (and only if) it provides the *sound* the artist wants - when compression is used just to make something sound louder, then it is better to just turn up the volume. This way, the original dynamics are preserved. Incorrectly

used, compressor/limiters will flatten the sound, and remove the life and soul of the music.

---

## **Additional Features**

Compressor/limiters are usually fairly complex electronically. Since they already have a voltage controlled amplifier (VCA) circuit that must be of the highest quality to satisfy audio professionals, this can be used for other things as well, with little real increase in cost or complexity (they are already complex, so a little more won't hurt :-)

A common addition to these audio tools is a noise gate. This is provided with compressor/limiters, and is used to gate (or switch off) any signal below a preset minimum. Noise gates are used to remove unwanted low level signals, but are sometimes used to mess up the sound completely by removing the ambience. Better a little noise and a complete sound than a quietly decaying ambience that suddenly just stops. Used properly, a noise gate can seem to eliminate background hiss completely, while letting the signal through (the hiss is still there, but you can't hear it when the signal is present). Used improperly, the initial parts of sounds are cut off, and the natural decay is not present. This is (fortunately) rare in pro studios.

One final feature offered in many units is a "de-esser". The sibilants ("sssss" sounds) in vocals are often over emphasised by close microphone placement, mic characteristics, the vocalist, or equalisation - and often a combination of these. This can be very unpleasant, so the de-esser does exactly what its name implies - it reduces the sibilant sounds by an amount that the recording engineer can set according to need (or taste)

---

## **Why Use Compression?**

Compressors and limiters are used in music for a multitude of reasons. The first (and should be the only) reason is for the sound. Used properly, a compressor - or more correctly a limiter - will place an absolute cap on the maximum level that can be passed. This is invaluable for preventing a large PA system from distorting, or making certain that the ADC (Analogue to Digital Converter) does not clip (exceed the maximum conversion voltage). Digital distortion is extremely unpleasant, and is to be avoided, as with all forms of hard clipping.

There are many other reasons to use compression or limiting. Many instruments do not have the sustain that the musician desires, and this can be corrected by using a compressor to extend the note. As the signal fades, the compressor increases its gain, so the note lasts longer.

Another reason is to restrict the dynamic range. Movie soundtracks are a prime example. If the maximum level of a car bomb exploding or a shotgun fired at close range were to be reproduced, and all conversations were at the normal level, no-one in the theatre would hear anything that was said, or would be deafened instantly by the explosions. By reducing the dynamic range, both can be accommodated at levels that are appropriate, but limited to an acceptable maximum and minimum loudness.

By contrast, many trailers and theatre advertisements are heavily compressed - they have a consistent loudness that is greater than that of the main feature. This technique can work for a limited period (it gets your attention), but becomes very tiring very quickly.

The over use of compression results in a flat, lifeless reproduction. In his article, Bob Katz refers to "wimpy loud sound", and at some stage we've all heard it. You put on a CD, and it is LOUD, so you turn it down, so the loud parts don't leave you loudspeaker cones on the floor. You wait, you listen, you wait some more ... it never happens! There are no loud sections! There are no quiet sections. Everything is at the same volume from beginning to end, and the result is indeed wimpy. Certainly the CD is louder than others you own, but it is the same volume from start to finish and leaves you as flat as the sound.

How to make music bereft of life - compress it until it bleeds (to death). No hi-fi, regardless of cost or sophistication can make rubbish like that sound good, since there is virtually nothing that can be done. An expander (essentially the opposite of a compressor) may restore some vestige of what the artist intended, but compression is not easily undone unless you can obtain the exact reverse of the original settings - this is both difficult and time consuming to even attempt, assuming that you have the equipment in the first place (very few hi-fi systems incorporate an expander, so most of us are well and truly screwed).

---

## **Contra Indications**

Compression is commonly used in the final mix, and this is where things can go seriously wrong - everything is at the same volume, peak to average ratio is minimal, and the resulting sound is *almost always* worse than it was before the compression was applied. Used correctly, a small amount of compression may be useful with some musical styles, but it is completely unsuited to others. I have several CDs that sound "exciting" at first, but the sameness of having a constant barrage of sound at the same level becomes extremely fatiguing in only a short time. On some, I can hear the compressor/limiter acting ("breathing" or "pumping" are terms commonly used for this effect), which means that it has been over used, and the CD is then relegated to the "don't bother listening to this" pile. Most unfortunately, this pile is getting bigger, and many of the modern CDs are worse than older ones because of the stupid, unnecessary and pointless game of one-upmanship by the record companies, all trying to get the "hottest" CD on the block.

I don't want it to be "hot", I want it to sound the way it should, with real dynamics, soft and loud passages, and things that make me jump! Fortunately, I am not alone, but unfortunately, record companies are still producing material for people with crap systems - "Make it sound good on a crap system - we'll sell more". This is rubbish - people with only a boom box don't care that much, otherwise they would strive for something better. Those among us with good or excellent systems should not have to listen to something that was mixed on a pair of near field monitors with the quality of a transistor radio, and compressed so heavily that it has lost all of the dynamics that make music what it is - or should be!

If the CD is a little quieter than expected, then I simply turn up the volume - even without a remote control, this is hardly an arduous task. Better that than have a whole pile of "hot" CDs that I can't bear to listen to because they have had all their life removed by an over zealous compressor-head.

---

### **Peak to Average Ratio & Dynamic Range**

All music has a peak to average ratio (and dynamic range - see below), since there are peaks and dips in the level (even when heavily compressed), and the average level must be lower than the peaks. The trick is to know what the peak to average ratio should be. It is commonly quoted as being between 10dB and 20dB (a power ratio of between 10:1 and 100:1). By this reasoning, music with a 10dB P-A ratio will need perhaps 50W to handle the peaks, but will provide an average power of only 5W. This is typical of a lot of music, and even some orchestral music will be at this ratio without any compression (relatively uncommon, but Bob Katz has experienced exactly this).

A ratio of 20dB is at the other extreme, so the same 50W amplifier will only produce an average power of 0.5W - this is where the use of high powered amplifiers for hi-fi is important - by the time the average power is high enough, the peak power is massive. A quick example ...

You want to listen to music at 90dB (SPL). Your speakers are rated at (say for convenience) 90dB/m/W, so with two of them, the effective sound pressure (at 1 metre) is 93dB SPL with 1W into each channel. You (the listener) are some distance away, so the level may be 3 to 12dB lower at the listening position. We shall assume 6dB as a reasonable guess for a typical listening room (although it may be considerably more than that, depending on room treatment, furnishings, etc.).

For 1W per channel, your SPL will be about 87dB SPL, so to get the extra 3dB, the power must be doubled to 2W per channel. If you have music with a peak to average ratio of 20dB, you will need 200W per channel to reproduce the music without distortion - assuming that you have that much power, the peak SPL will be in the order of 110dB ...

(90dB + 20dB P-A ratio = seriously loud).

Compression is your friend! Such a high P-A ratio will cause most high end systems (and their owner's ears) grief at high levels, so some degree of compression will make the reproduction less arduous on your system (and a lot less likely to frighten the cat to the point where it perches tightly on your head :-). The magic is to find the ratio that keeps the ratio to a reasonable figure (and there are no absolutes here!), while preserving the soul of the music. It can be done, and I have many CDs and vinyl albums that do it very well (Bob will also tell anyone who cares to listen how to do it well, too).

It can also be done very badly, and so many new releases do just that. Mind you, a lot of old releases were just as bad - this is not a new phenomenon, but has been happening ever since the compressor was first invented - or at least used in anger.

### **Dynamic Range**

Is there a difference between peak to average ratio and dynamic range? The answer depends entirely how long the averaging period is. Generally, the two are considered separate, but with enough compression they become equivalent. The dynamic range of a piece of music is the difference between the softest passage and the loudest - it takes little imagination to realise that if it is sufficiently heavily compressed there will be no difference at all.

With a good mix, and a compressor/limiter that is correctly adjusted, the difference between the two will be less than in "real life", but great enough to create excitement - this is where experience and careful adjustment come in. The range of sounds we can hear (and will be assaulted by) is enormous, from the faint rustling of leaves on a very quiet night through to jack-hammers or jet planes (and certain motorcycles!). To expect to reproduce this range from a home hi-fi or theatre system is generally not possible and is undesirable anyway.

The difference between the two is blurred - there is an almost infinite grey-scale, rather than any black and white distinction between the two, so I shall leave it at that.

---

### **Typical Specifications**

A discussion of compressor/limiters would be incomplete without a brief explanation of the features and controls typically offered. A typical unit may offer the following (adapted from real specs for a typical unit) ...

Max. Input Level: 10V RMS (+22dBu)

Dynamic Range: 118dB

Signal to Noise Ratio: >100dB

Headroom: 18dB <0.05% at +4dBu, with 6dB compression  
Distortion: <0.05% at +4dBu, with 6dB compression  
Limiting Threshold: -40dBu to +20dBu  
Attack Time: 0.1ms - 200ms  
Release Time: 50ms - 3s  
Compression Ratio: 1:1 - 100:1 with selectable hard or soft compression knee

Some of these terms are self-explanatory, while others may need a little more information.

- **Maximum Input Level**  
This is the maximum signal (in Volts RMS and/or dBm or dBu \*\*). If this is exceeded, the input stage will distort due to clipping. This indicates that a unit with the above specification cannot be used for a direct speaker feed (few can, and a Direct Injection box is nearly always needed).
- **Dynamic Range**  
The range of the quietest sound (signal) to the loudest. The spec above is misleading, since it ranges from the noise floor to the maximum input level. The effective dynamic range is limited by the minimum signal to noise ratio the engineer will accept, so subtracting (say) 50dB from the above would be reasonable.
- **Signal to Noise Ratio**  
The level of noise referred to a specified output level. Again, the output level is not specified above, so we don't know if the actual noise floor is better than -100dBu or -78dBu, since the reference input or output level was not defined. Many manufacturers use A-Weighting for these measurements, which is also misleading for most applications. A-Weighting should be used only where the sound source is some distance from the listener (possibly hundreds of metres), or is expected never to exceed about 70dB SPL near field.
- **Headroom**  
This can be tricky, since there are many different ways to describe it. Most commonly, headroom is the difference (in dB) between the typical or specified output level, and the maximum the unit can provide without distortion.
- **Distortion**  
Another can of worms! Distortion measurements are not really meaningful if the level and frequency are not stated. It is even harder with anything that uses a VCA, since their distortion can vary considerably depending on the degree of amplification or attenuation. Distortion (probably) should be worst case - the frequency, amplitude and amount of gain that produces the most distortion. This would probably look bad in the specs, so is not used.
- **Limiting Threshold**  
This is the range of signal level where the unit can be set to function. Below the threshold, a limiter does nothing - it just passes the signal straight through (with maybe a little fixed gain or loss, as appropriate). Above the threshold, signals are attenuated by an amount determined by the compression ratio (see below).
- **Attack Time**  
A measure of how long it takes before the unit changes the gain when a signal is applied.

This may be very fast to prevent *anything* from exceeding the threshold, or deliberately slowed down to allow a percussive attack to the instrument. A relatively slow attack might be used with drums to increase the apparent dynamics, whilst actually reducing them.

- **Release Time**  
The release time determines how long the gain takes to return to "normal" after it has been reduced. A short release squashes the dynamics of the sound completely, and a very long release time ensures that all material remains at a constant fixed peak level, while still allowing normal variations. The overall dynamics (of a musical piece) are still compressed, since the soft passages will allow the unit to apply the maximum gain. This is all controllable.
- **Compression Ratio**  
As discussed above, this can be low (1:1 means no compression), almost infinite, or anywhere in between. An infinite (or high) ratio is a limiter, anything between 1:1 and about 10:1 is a compressor, although some would argue (correctly) that 10:1 is really limiting. Some units also offer a "soft" or "hard" knee (the threshold). A soft knee means that the onset of compression or limiting is gradual, ranging over a few dB, while a hard knee means that once the threshold is reached, the limiting action is immediate, with no gradual onset.

\*\* dBm is a reference level based on 1mW into 600 ohms. This represents a voltage of about 775mV. dBu is based in a reference level of 1V RMS (also sometimes seen as dBV).

---

## What To Listen For

An excellent way to hear compressor/limiters in action is an outside broadcast on TV. While the presenter is speaking, the level is constant, with very little variation - even when they are off-axis from the microphone. When there is a break in the commentary, the background noise can be heard to increase at a (relatively) fixed rate, until it is as loud as the presenter's voice, or someone starts speaking again. It is then instantly reduced to where it was before.

As you get used to what to listen for, you will hear many CDs where the level of the backing track falls when the singer (or someone else) starts making their noises - this (to many audio professionals) is quite normal, but it is not - it is a typical case of over application of compression in the final mix. When an additional sound is added to those already present, it is *supposed* to get louder - this is called dynamics (or even "micro-dynamics" - a reduced scale version of the real thing).

These effects are especially noticable with commercials on radio or TV - listen for them so the sound can be identified. Should you purchase a CD that does the same, complain to the record company - they have ruined your music!

---

## What About Lossy Data Compression?

MP3 - love it or hate it, it is here (probably) to stay. As can be determined from a multitude of sources, MPEG Layer 3 (or MP3 for short) discards information that theory (and a lot of experimental testing) indicated would be inaudible. It uses a

well known characteristic of our hearing called "masking", where it is known (and can be proven) that certain frequencies and levels are completely inaudible when accompanied by another signal at a higher level. The points where masking take effect are beyond the scope of this article, but differ according to relative levels and frequency, and the frequency band itself. An MP3 encoder breaks the signals into sub-bands using filters, and each is treated differently according to a set of rules built into the encoder signal processor.

While it is generally considered that a high bit rate (128kb/s or above) MP3 track is of "near CD quality", many people will dispute this vehemently. My own experiments and listening tests indicate that imaging is poor, and the precise placement of instruments and vocals is missing. Some instruments - especially the harpsichord (so I hear) sound completely different when encoded, almost regardless of the bit rate. A good test is to "rip" some pink noise (preferably generated by an analogue source), and compare the difference.

There should be no difference - or at least it should be inaudible, but this is not the case! At 320kb/s the difference is barely audible - one has to listen carefully to hear it, or figure out just *what* to listen for ... but there is a difference, and it also shows up very clearly on the analyser of Winamp. The peaks are flattened, so the dynamic range (or peak to average ratio) is degraded, and the sound of the noise lacks "life" compared to the original recording.

If we can hear a difference with noise, why would music be any different? It isn't! Ok, noise has a relatively constant bandwidth (DC to daylight in the extreme), and excites all frequencies more or less simultaneously. Well, so does a lot of music, albeit for short periods at a time.

Will my comments here make MP3 go away? Of course not, and nor should it go away, because it is a useful way to archive recordings, or provide people (who insist on *not* hearing approaching traffic while they run or cycle) a convenient medium for portable sound.

---

## **Even Worse Things That Can Happen**

On a final note (pardon the pun :-)) a reader recently sent me an out-take of a CD. It was clipped! Not just compressed and limited to the maximum (that too), but with actual clipping - flat tops on some peaks. I asked him to check his setup very carefully to ensure that the record level was not set too high, and he assured me that he had at least 3dB of headroom above the peak CD level.

At some stage, I shall check some of the CDs I have that annoy me because of the constant loudness to see if they have the same problem. Probably not, since the clipped CD was from an "indie" (independent producer) so would not have had the controls in place one would expect from an established mastering house.

Be that as it may, there is not really much point in setting up the "ultimate" hi-fi system, with headroom to spare and almost zero distortions of any kind, only to have the music CD pre-distorted, compressed, limited and bent so far out of shape that it is no longer useful for anything other than a coaster.