
# Compaction and compression

**A couple of issues ago, the editor drew attention in his leader to a problem that has been occupying Master Audio's MARTIN SPENCER's working days (and nights) for well over a year. Why does compressed audio sound so downright annoying, and is there anything that can be done to make it sound better?**

IT'S MY CONTENTION that more than one type of compression is responsible for the awful sound that we all have come to know and detest; further, that the combination of two different types of 'compression' is particularly deadly.

When we refer to compressed audio, the delivery media that spring to mind will probably be MP3 players, audio streaming via the Internet, and DAB digital radio. But rather tellingly, Zen also includes 'CD re-masters and a whole catalogue of modern releases' in his tirade. Last time I looked, CDs still delivered an audio bit rate of some 1411kbps, the same as they did on their introduction to the market back in the early 1980s.

So what precisely do we mean by 'compressed audio' anyway? Let's start by eliminating some slippery language (© Bob Katz). What we have here is one word, 'compression', used for two entirely different processes. On the one hand, we have good old dynamic-range-reducing compression. The word has been used this way since the 1950s or probably earlier, and there seems no good reason to break with tradition.

On the other hand is the modern connotation of stripping allegedly redundant information out of digitally coded audio or video to reduce the number of bits required to transmit it. I'd like to make a proposal. Could we call the bit-rate-reduction process 'compaction' instead, so we have a distinction between the two processes? Thus, an MP2 or MP3 encoder becomes known as a compactor.

There's more slippery language to deal with. Bear with me just a little longer — getting our terms right is going to be worth it. 'Compression', as a catch-all term for manipulation of dynamic range, is still so vague as to discourage proper understanding. Compression, as we all do know really, is when we squash the dynamics but leave a fraction of the original change in place; done by fiddling with the 'slope' knob of course. Contrast this with a limiter, which once above its threshold would erase all trace of the original dynamic 'eventually', because the time constants that both compressors and limiters use will ensure any changes in gain are applied gradually. Another term commonly encountered in discussion of broadcast audio processing is AGC or automatic gain control — this is a compressor or limiter with very long time constants indeed.

Moving to the weapons-of-mass-destruction department in the armoury of brutal audio processing tools, we have 'clipping'. This is a compressor or a limiter, but without any time constants to soften its impact. A compressor without time constants is called a soft clipper. A limiter without time constants is called a hard clipper. A relationship without time, Constance, is called a one night stand.

Before we look at exactly what can be done to improve the quality of compressed and compacted audio, it is worth having a look at the broadcast processors of the past, tracing the evolution of techniques for radio processing that have now found their way into the mastering house.

In the beginning, there were just two processors: a gain rider and a peak limiter. Gain riding was not carried out in the racks room, it was a job for a professional in the control room, who might even be armed with scripts and scores, and who certainly made a judgement of level based on his ears, with metering mainly for confidence. Final peaks had to be controlled to prevent over-modulation, and this was done by a 19-inch rack box, used to control brief peaks that were too fast for the manual operator to handle, and to minimise the effect of human error. This system gives good audio quality, and the dynamic range can be artistically tailored to fit pre-determined limits. Although probably preferable to any automatic system, it's not really a viable option in the multichannel age. The thought of serried ranks of Levellers sitting at Master Control all with their headphones on is a nice one though.

In the 1950s and 1960s along came increasingly crowded AM bands and FM with pre-emphasis. This is where the trouble started for the audio. Now don't get me wrong, FM was a wonderful invention, and we should all be grateful to Edwin Armstrong for pioneering it. It was just an unfortunate side-effect that narrowing AM receiver bandwidths to cope with interference, and introducing de-emphasis on FM to cope with its noise spectrum, both contributed to an eventual need some 30 years later to make the broadcast programme signal brighter. Of course, the root of this was back in the days when popular music meant big-band, cabaret, musicals, rock and roll, and then the Folk and Blues revivals. In the days when the system standards were developed, high-frequency components of the programme did not often tax the transmission system in the way they later would but the seeds of an unfortunate future problem had been planted.

A succession of broadcast audio processing techniques were developed from the late 1950s; we'll review these very briefly here, since a full history of broadcast audio processing would be an article (heck, a book!) in its own right. Many of us in Europe will remember Radio Caroline first running an Optimod on AM, and the UK landmark came when Capital Radio first installed its Inovonics 250 for FM; exciting times, but outside the scope of this article.

Automatic processing for broadcast first started with wide-band AGCs with sufficiently adaptive response to handle a really wide dynamic range at the input; still essentially an 'engineering' tool rather than something that radically changed the sound, but very useful nonetheless. Devices of this type, such as the CBS Audimax, were commonplace in the USA during the early years. In the UK the BBC and even the IBA continued to frown on audio processing, insisting that all stations stuck to what Aunty thought best: conscientious manual gain riding at the studio and automatic safety limiters at the transmitter sites.

In the 1970s two novel techniques were pioneered in the USA: multiband audio processing (Mike Dorrough) and overshoot compensation (Bob Orban). Both techniques were pivotal in the advance of FM radio.

Multiband audio processing was originally conceived to get around the problems of wide-band limiting — pumping, and bass peaks causing hole-punching and ducking. But multiband processing soon became, perhaps more importantly, a means to add boom and sizzle for those (presumed the majority) listening on poor speakers. Further, it was possible automatically to tailor the amount of EQ to the source material, previously impossible with conventional 'static' EQs. This amounted to a form of remastering tracks that were either 'EQ thin' or 'EQ rich' towards a common standard, namely the signature sound of the radio station in question.

Overshoot compensation addressed the problems of loudness-robbing overshoots, combining hard clippers with low-pass filters in a way that optimised time-domain (overshoot and thus loudness) and frequency-domain (bandwidth) performance simultaneously. This was the first really powerful loudness boosting technique that did not rely on extending the source bandwidth — otherwise known as splattering the band. Of course, it came with a fairly hefty distortion penalty when pushed hard.

In the 1990s, the availability of suitable DSP allowed the introduction of even more powerful techniques for greater loudness (RMS level) with reduced distortion. These techniques mostly revolve around two novel abilities: look-ahead and complex adaptive systems. Both were theoretically possible in analogue, but only really became practical with digital.

So what are the problems inherent in 'traditional' audio processors? I'm going to simplify here in the interests of brevity but first is the widespread use of limiting (as opposed to compressing) even in slow automatic gain control and compressing functions. This tends to erase programme and musical dynamics; not a good idea. This is presumably because of a perception that loudness is of such dominating importance that we must throw subtlety out with the bathwater. Perhaps, in this age of HD video and audio, it's an idea whose time has come and now gone.

Second is the unholy alliance of the multiband compressor with the clipper. Any consideration of multiband processing must focus on how the signal is treated after the bands are recombined. Unfortunately, it turns out that multiband compression does not generate anywhere near competitive loudness on its own, so it's invariably followed by a wide-band dynamic limiter and/or a clipper. Sonically, this is a very interesting combination. Multiband and dynamic processes tend to make the sound mushier, reducing impact and transient clarity. However, passing the recombined multiband output through a clipper, a device that can boost the fundamental output levels beyond 100% while also adding harmonics, can go loud while restoring some of the 'slam', perceived transients and brightness to the signal, at the expense of gross intermodulation.

This is, in essence, what most processors do (even your humble single-band outboard compressor-limiter typically includes a clipper, if you engage the limit function). It's one reason you can still hear drums reasonably clearly despite all the dynamic squashing going on, but it's also the reason for all the distortion. Backing off the processing to get rid of distortion is often unsatisfactory — the loudness benefit seems to disappear almost as fast as the distortion!

Third, it is common knowledge in radio engineering circles that clipping should not be used in an audio processor that is feeding a compactor. The reason usually given is that encoding the clipping distortions (harmonics and intermodulation) will waste some of the bit-budget of the encoding engine in the compactor. The comment was, I believe, originally voiced by Bob Orban many years ago in a technical paper, but has now been used by more than one company in the advertising of their digital radio processors. These digital radio processors use various forms of look-ahead and other 'distortion cancelled' techniques allegedly to mitigate the encoder stress effect. However, in many cases the waveform is changed sufficiently rapidly that a following compactor can't cope without introducing considerable additional audible distortions — the 'burbling', washy, smeared sound. This is particularly a problem for the UK DAB system, encumbered as it is with ancient Layer II at, typically, only 128kbps.

There's no doubt that avoiding clipping is a good idea but there is a lot more to the story than this. Taking what is really a 95% traditional radio processor and bolting it into the digital transmission path is the single avoidable act most responsible for the awful sound we are becoming disgusted with. And unfortunately, because 'loudness wars' are now being fought in the trenches of mastering as well as broadcast, the poison gas of over-processing is tending to affect every single delivery channel.

But let's be positive. The first part of the solution is to eliminate all unnecessary processes that are exacerbating the problems of compaction. This is a system-wide problem that extends to all links in the chain, and particular attention must be paid to maintain signal quality in contribution feeds, storage, mixing, distribution and transmission — a lot of places where the audio quality can be impaired. Here we will concentrate on the audio processor, which is one of the more critical links in the chain. Out go multiband compressors and out go 'main' (combining) clippers.

What is then needed are specific algorithms to deal with the following problems: normalising average level, powerfully if called for; retaining some musical dynamics and natural flow; avoiding hole-punching and pumping; and peak control without clipping, aliasing or inter-sample overload.

If we pay heed of all the above points, we can have an audio processing system that is capable of smoothing out levels without sounding awful, and that therefore makes the most of whatever delivery channel is available. Of course, the trade-off between the first two of those requirements must be set by the user, and in doing that, best results will be obtained if the operator is asked to be aware of his levels and machine intervention reduced (but not eliminated). The fader position will still have some influence on the final loudness, albeit with extremes of loud and quiet ameliorated, and this can be used to good effect in the right hands.

Returning in closing to the more general question of compressed audio, the current situation is indeed awful, but there are also one or two possible rays of hope. One is AAC plus compaction, which does seem to offer a substantially better listening experience than the others, especially at low bit-rates like 128kbps. Even if digital radio broadcasters insist on subdividing their digital delivery channels until they sound dire (the UK is already there, and it's looking pretty dicey for the US 'HD Radio' system with maybe 48kbps of AAC+, due to the likely adoption of multicasting), the competition may not follow suit. There is always the possibility of something better supplied by alternative delivery channels: iPods, digital TV in various forms, satellite radio, and Internet streaming are all capable of supporting at least 128kbps per stereo audio channel. That's not to mention DMB and other pioneering forms of 3G delivery.

As new means of delivery become increasingly accessible, competition may act as a self-regulatory mechanism. The new fear becomes the competitor who offers similar content, with the same accessibility, and better quality. Remember, it first happened that way when microgroove records ditched the 78, and it certainly happened that way again when FM radio was the new kid in town. Most recently, the iPod has made much greater impact than cheap MP3 players. Perhaps digital delivery does have to maintain decent sound quality after all. ■

**Contact**

**MASTER AUDIO LTD, UK:**
**Email:** martin@master-audio.com