

Digital Level Control – done the right way

With sound examples!

In high-end HiFi circles a level control done in the digital domain is often viewed as being inferior to one operating in the analog domain. Let's look on how a digital level control works and why it can be an excellent solution if it is properly implemented.

A level control is a multiplication of the audio signal with a constant, the "gain factor". The gain factor usually is in the range of zero (signal fully off) to one (signal untouched). A factor of 0.5 then means that the audio signal is attenuated to half of its amplitude.

What exactly happens when we multiply two numbers? If we e.g. multiply a 2 digit and a 3 digit number, the resulting number can be up to 5 digits long (the sum 2 plus 3). As an example: 30 times 500 equals 15000. 2 digits times 3 digits yields a 5 digit result.

In digital audio, the numbers are represented in the binary system, not the decimal system. A decimal number consists of digits 0 through 9, a binary number of digits 0 and 1. So a binary number may look like this: 1011 0011 0101 1101. This is a 16 digit or 16 bit binary number, the grouping into 4 bit chunks is for better readability. The audio samples on a CD are represented with such a binary number system with each sample value represented with 16 bits.

Now let's assume we have a 8 bit gain factor for a level control. If we apply that to a signal coming off a CD we multiply a 8 bit gain factor with a 16 bit sample value. The result is up to 24 bits long (the sum of the wordlengths of the two factors). An example:

0100 1001 x 1001 0110 0111 1011 = 0010 1010 1110 1001 0001 0011

The question now is what do we do with the 24 bit long result? The digital to analog converter which converts the samples after the level control may only be capable to handle 16 bit wide samples. Thus what should we do with the excessive 8 bits? The simplest solution is to truncate the 24 sample to 16 bits, i.e. to cut off the 8 less significant bits. The truncated 24 bit result above then would look like this:

0010 1010 1110 1001 i.e. the first 16 bits of the 24 bit result above.

The remaining bits (0001 0011) are discarded. If these bits are discarded a error is introduced. This error is called a quantization error, because the 24 bit result is re-quantized to 16 bits.

Unfortunately the quantization error is part of the audio signal – and if we take that part away from the signal, the signal undergoes some distortion, the so called quantization distortion.

The sound example at the link below shows how such a distortion sounds. In this music example a 16 bit signal is truncated to 8 bits. 8 bits in order to clearly show the effect. Notice how the noise (distortion) is modulated by the music signal.

<http://www.weiss-highend.ch/computerplayback/nodither.mp3>

This is how a badly implemented digital level control works....

Fortunately there is a better way to handle the re-quantizing. One solution would be to use a D/A converter with a higher wordlength, e.g. a 24 bit converter, to accommodate for the 24 bit samples coming out of the level control. This of course would already help a lot, but there is another technique: dithering.

The idea about dithering is to de-correlate the quantization error from the audio signal. As we have seen in the example above, the quantization error depends on the audio signal (i.e. it is correlated with the audio signal). On the other hand, if dither noise is added to the 24 bit sample after the level control and before the re-quantization to 16 bits, the quantization error can be fully decorrelated from the signal. This means instead of distortion there is noise. The music is undistorted. The audio example at the link below is again a 16 bit signal quantized to 8 bits, but with dither noise added. A much more pleasant experience. Notice how the noise stays untouched by the music, i.e. there is no noise modulation. <http://www.weiss-highend.ch/computerplayback/flatdither.mp3>

Dithering does not stop here. More elaborate dithering schemes shape the noise such that it is mainly present at higher frequencies where the human ear is less sensitive. This means that the audible noise is much lower. The link below is again the 16 bit source quantized to 8 bits with noise-shaped dithering. Probably hard to believe that this is only an 8 bit system! Note that the music is not distorted at all, despite the 8 bit resolution. Remember, a 16 bit system has 65,536 quantization steps while a 8 bit system has only 256 quantization steps – a huge difference. And still, the properly dithered 8 bit system sounds great.

<http://www.weiss-highend.ch/computerplayback/shapeddither.mp3>


This is what a properly dithered level control is capable to do. You have heard the 8 bit version, imagine that with today's 24 bit converters – no question that a level control with a 24 bit wordlength easily rivals the best analog level controls. By the way, 24 bits means 16,777,216 quantization steps.

The last example below toggles the noiseshaping dither on and off to give a good contrast between dither / no dither versions.


<http://www.weiss-highend.ch/computerplayback/togglingdither.mp3>

Dithering is used in many disciplines. The pictures below show dithering applied when quantizing picture data.


Original:



Quantized w/o dithering:



Quantized with dithering:



© 2009, Daniel Weiss