

24bit vs 16bit, the myth exploded!

It seems to me that there is a lot of misunderstanding regarding what bit depth is and how it works in digital audio. This misunderstanding exists not only in the consumer and audiophile worlds but also in some education establishments and even some professionals. This misunderstanding comes from supposition of how digital audio works rather than how it actually works. It's easy to see in a photograph the difference between a low bit depth image and one with a higher bit depth, so it's logical to suppose that higher bit depths in audio also means better quality. This supposition is further enforced by the fact that the term 'resolution' is often applied to bit depth and obviously more resolution means higher quality. So 24bit is Hi-Res audio and 24bit contains more data, therefore higher resolution and better quality. All completely logical supposition but I'm afraid this supposition is not entirely in line with the actual facts of how digital audio works. I'll try to explain:

When recording, an Analogue to Digital Converter (ADC) reads the incoming analogue waveform and measures it so many times a second (1*). In the case of CD there are 44,100 measurements made per second (the sampling frequency). These measurements are stored in the digital domain in the form of computer bits. The more bits we use, the more accurately we can measure the analogue waveform. This is because each bit can only store two values (0 or 1), to get more values we do the same with bits as we do in normal counting. IE. Once we get to 9, we have to add another column (the tens column) and we can keep adding columns add infinitum for 100s, 1000s, 10000s, etc. The exact same is true for bits but because we only have two values per bit (rather than 10) we need more columns, each column (or additional bit) doubles the number of vaules we have available. IE. 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024 If these numbers appear a little familiar it is because all computer technology is based on bits so these numbers crop up all over the place. In the case of 16bit we have roughly 65,000 different values available. The problem is that an analogue waveform is constantly varying. No matter how many times a second we measure the waveform or how many bits we use to store the measurement, there are always going to be errors. These errors in quantifying the value of a constantly changing waveform are called quantisation errors. Quantisation errors are bad, they cause distortion in the waveform when we convert back to analogue and listen to it.

So far so good, what I've said until now would agree with the supposition of how digital audio works. I seem to have agreed that more bits = higher resolution. True, however, where the facts start to diverge from the supposition is in understanding the result of this higher resolution. Going back to what I said above, each time we increase the bit depth by one bit, we double the number of values we have available (EG. 4bit = 16 values, 5bit = 32 values). If we double the number of values, we halve the amount of quantisation errors. Still with me? Because now we come to the whole nub of the matter. There is in fact a perfect solution to quantisation errors which completely (100%) eliminates quantisation distortion, the process is called 'Dither' and is built into every ADC on the market.

Dither: Essentially during the conversion process a very small amount of white noise is added to the signal, this has the effect of completely randomising the quantisation errors. Randomisation in digital audio, once converted back to analogue is heard as pure white (un-correlated) noise. The result is that we have an absolutely perfect measurement of the waveform (2*) plus some noise. In other words, by dithering, all the measurement errors have been converted to noise. (3*).

Hopefully you're still with me, because we can now go on to precisely what happens with bit depth. Going back to the above, when we add a 'bit' of data we double the number of values available and therefore halve the number of quantisation errors. If we halve the number of quantisation errors, the result (after dithering) is a perfect waveform with halve the amount of noise. To phrase this using audio terminology, each extra bit of data moves the noise floor down by 6dB (half). We can turn this around and say that each bit of data provides 6dB of dynamic range (*4). Therefore 16bit x 6db = 96dB. This 96dB figure defines the dynamic range of CD. (24bit x 6dB = 144dB).

So, 24bit does add more 'resolution' compared to 16bit but this added resolution doesn't mean higher quality, it just means we can encode a larger dynamic range. This is the misunderstanding made by many. There are no extra magical properties, nothing which the science does not understand or cannot measure. The only difference between 16bit and 24bit is 48dB of dynamic range (8bits x 6dB = 48dB) and nothing else. This is not a question for interpretation or opinion, it is the provable, undisputed logical mathematics which underpins the very existence of digital audio.

So, can you actually hear any benefits of the larger (48dB) dynamic range offered by 24bit? Unfortunately, no you can't. The entire dynamic range of some types of music is sometimes less than 12dB. The recordings with the largest dynamic range tend to be symphony orchestra recordings but even these virtually never have a dynamic range greater than about 60dB. All of these are well inside the 96dB range of the humble CD. What is more, modern dithering techniques (see 3 below), perceptually enhance the dynamic range of CD by moving the quantisation noise out of the frequency band where our hearing is most sensitive. This gives a perceivable dynamic range for CD up to 120dB (150dB in certain frequency bands).

You have to realise that when playing back a CD, the amplifier is usually set so that the quietest sounds on the CD can just be heard above the noise floor of the listening environment (sitting room or cans). So if the average noise floor for a sitting room is say 50dB (or 30dB for cans) then the dynamic range of the CD starts at this point and is capable of 96dB (at least) above the room noise floor. If the full dynamic range of a CD was actually used (on top of the noise floor), the home listener (if they had the equipment) would almost certainly cause themselves severe pain and permanent hearing damage. If this is the case with CD, what about 24bit Hi-Rez. If we were to use the full dynamic range of 24bit and a listener had the equipment to reproduce it all, there is a fair chance, depending on age and general health, that the listener would die instantly. The most fit would probably just go into coma for a few weeks and wake up totally deaf. I'm not joking or exaggerating here, think about it, 144dB + say 50dB for the room's noise floor. But 180dB is the figure often quoted for sound pressure levels powerful enough to kill and some people have been killed by 160dB. However, this is unlikely to happen in the real world as no DACs on the market can output the 144dB dynamic range of 24bit (so they are not true 24bit converters), almost no one has a speaker system capable of 144dB dynamic range and as said before, around 60dB is the

most dynamic range you will find on a commercial recording.

So, if you accept the facts, why does 24bit audio even exist, what's the point of it? There are some useful application for 24bit when recording and mixing music. In fact, when mixing it's pretty much the norm now to use 48bit resolution. The reason it's useful is due to summing artefacts, multiple processing in series and mainly headroom. In other words, 24bit is very useful when recording and mixing but pointless for playback. Remember, even a recording with 60dB dynamic range is only using 10bits of data, the other 6bits on a CD are just noise. So, the difference in the real world between 16bit and 24bit is an extra 8bits of noise.

I know that some people are going to say this is all rubbish, and that "I can easily hear the difference between a 16bit commercial recording and a 24bit Hi-Rez version". Unfortunately, you can't, it's not that you don't have the equipment or the ears, it is not humanly possible in theory or in practice under any conditions!! Not unless you can tell the difference between white noise and white noise that is well below the noise floor of your listening environment!! If you play a 24bit recording and then the same recording in 16bit and notice a difference, it is either because something has been 'done' to the 16bit recording, some inappropriate processing used or you are hearing a difference because you expect a difference.

G

1 = Actually these days the process of AD conversion is a little more complex, using oversampling (very high sampling frequencies) and only a handful of bits. Later in the conversion process this initial sampling is 'decimated' back to the required bit depth and sample rate.

2 = The concept of the perfect measurement or of recreating a waveform perfectly may seem like marketing hype. However, in this case it is not. It is in fact the fundamental tenet of the Nyquist-Shannon Sampling Theorem on which the very existence and invention of digital audio is based. From WIKI: "In essence the theorem shows that an analog signal that has been sampled can be perfectly reconstructed from the samples". I know there will be some who will disagree with this idea, unfortunately, disagreement is NOT an option. This theorem hasn't been invented to explain how digital audio works, it's the other way around. Digital Audio was invented from the theorem, if you don't believe the theorem then you can't believe in digital audio either!!

3 = In actual fact these days there are a number of different types of dither used during the creation of a music product. Most are still based on the original TPDFs (triangular probability density function) but some are a little more 'intelligent' and re-distribute the resulting noise to less noticeable areas of the hearing spectrum. This is called noise-shaped dither.

4 = Dynamic range, is the range of volume between the noise floor and the maximum volume.

Source: head-fi.org

Comments:

1.) The Nyquist theorem does not include amplitude quantization meaning infinite resolution, so it doesn't discuss quantization effects at all.

2.) There is no audio system in the world giving more than 20 clear bits of signal due to resistance and semiconductor noise characteristics.
3.) You cannot "recreate" a single thing by dithering, just make it sounding more natural to the ears, especially when using noise shaping filters for the dither signal.
4.) It's not true ADC's do any dithering. Some of them do some lowpass filtering with noise shaping involved when delta-sigma type which happens not for the dithering purposes but for the signal itself.
5.) You cannot increase dynamic range by dithering.
6.) It's not dynamics killing people and affecting hearing but sound pressure with the given numbers of 140dB = pain, 160~180dB = death, respectively. You can listen to the signal with the 144dB dynamics not exceeding safe sound pressure limits, just set the volume appropriately. Sure, you won't hear the bottom of your dynamic range then.

Answers:

1. True.
2. True. Most people believe that their 24bit DAC is actually a 24bit DAC, just marketing I'm afraid.
3. True. Dithering is just a process which should be used whenever a quantisation or re-quantisation is performed, to convert quantisation errors into un-correlated noise.
4. This one is not true. All ADCs use dither. Some 24bit ADCs use self-dither, in other words because the digital noise floor is so low (-144dB) the noise generated by their own internal components is enough to dither, but one way or another, they all dither. Also, all ADCs use a low-pass brick wall filter (anti-alias filter). Noise-shaped dither is not and should never be used in an ADC or when mixing. As the recorded channels are mixed the re-distributed noise is summed and can cause problems. The only time noise-shaped dither should be applied is during the last quantisation process. This usually means when converting the 24bit master from the recording studio into 16bit for CD release.
5. Sort of true. In an absolute sense CD has 96dB dynamic range, however if we move the noise that is down at the -96dB level to areas of the hearing spectrum where we are less sensitive (for example below 60Hz or above 12kHz). This gives a perceived improvement of dynamic range for 16bit. Bob Katz, the leading expert, reckons that about 120dB is the perceived dynamic range achievable with today's dithering technology.
6. True. Though of course by turning down your amp and not hearing the quietest sounds, then you are not hearing all the detail or the whole dynamic range, so it rather defeats the whole purpose of more dynamic range (more bits) in the first place.

Revision #3

Created 2026-04-01 17:14:48 CEST by Philip

Updated 2026-04-13 19:23:43 CEST by Philip