



M2Tech S.r.l.

P.I.: IT01821100508

Via Carlo del Prete, 2 – I-56121 Pisa (PI) - Italy

High Resolution: Is It Really Necessary?

A white paper by Marco Manunta – M2Tech Srl, Italy

We're facing a seemingly unstoppable run towards very high resolution in digital audio: from CD to 96/24, to 192/24 up to 352.8/24 and even 352.8/32, not to mention DSD (2.8MHz/1bit or 5.6MHz/1bit). Considering how good some 44.1/16 systems (CD players or DAC) sound, the question is: do we really need very high resolution, and how high should "high resolution" be?

Bit depth

This parameter is maybe the easiest one to evaluate and justify. First, it's useful to remind readers that bit depth defines the signal-to-noise ratio of a digital playback or recording system, through a simple formula:

$$\text{SNR(dB)} = 1.76 + 6.02 * N,$$

Where N is the number of bits per sample or bit depth. In a CD-standard system (N=16), SNR is 98.08dB, a 24bit device will allow for 146.24dB. Literature often forgets the 1.76dB constant, so you may find 144.5dB instead of 146.24 and so on.

We can also define a "virtual" bit depth for analog devices and systems, through the ENOB parameter (Equivalent Number of Bits):

$$\text{ENOB (bits)} = (\text{SNR}_{\text{measured}} - 1.76) / 6.02$$

The ENOB of an amplifier with 105dB measured SNR is 17.15 bits (note that .15 bit has no real meaning).

It's easy to realize that a digital source delivering true 24 bits resolution connected to a system in which the overall ENOB is, say, 18 bits, is not guaranteed to fully exploit its performance. On the other hand, one may apply the law of balance and say that a digital source with 18 bits resolution (110dB) is sufficient to obtain the best overall performance from that system.

Things are not that easy. First of all, let's consider the nature of noise in digital systems. Noise is the sum of two contributors: quantization noise and thermal noise. The first one is related to bit depth and is very disturbing for humans ear-brain due to its correlation to the signal (in fact, it is sometimes called "quantization distortion", rather than quantization noise). The second one is related to the analog circuitry used in the analog-digital and digital-analog boundaries and is generally much more tolerated by human ear-brain due to its total non-correlation with the signal.

M2Tech S.r.l.

Sede legale: Via Carlo del Prete, 2 – I-56121 Pisa (PI) – Italy

Sede operativa: C/O Polo Tecnologico di Navacchio – Via Giuntini, 13/N – I-56021 Navacchio di Cascina (PI) – Italy

Tel: +39 050 754195 – www.m2tech.biz – info@m2tech.biz

Coord. Bancarie: Cassa di Risparmio di Parma e Piacenza, Lung. Gambacorti, 4, Pisa – IBAN: IT 41 U 06230 14000 000043151277



M2Tech S.r.l.

P.I.: IT01821100508

Via Carlo del Prete, 2 – I-56121 Pisa (PI) - Italy

Devices, and even systems, in which thermal noise buries quantization noise are generally well-sounding setups.

This explains why we cannot define a specific bit depth as a limit value to judge digital devices and systems: a 14-bit source will probably behave very well in a system in which the amplifier has only 70dB SNR, while it will sound awful in another system in which the amplifier has 100dB SNR. Also, this may explain why CD players with tube output have generally a pleasant sound: tubes are noisier than opamps and even than solid-state, discrete components buffers, thus their thermal noise buries the quantization noise from the conversion IC (I won't enter the mined field of harmonic composition of tubes' distortion or biasing issues as it's beyond our scope).

Conversely, an analog source with an outstanding 125dB SNR driving a 16bit ADC is a waste of money, as the quantization noise of the ADC will be easily heard over the source's thermal noise, giving a typical "digital" sound. On the contrary, an analog source with 100dB SNR driving a 24bit ADC will give very good results to listening, as the thermal noise of the source will act a dithering with regards to the quantization process, transforming the signal-correlated quantization distortion in a "quasi-thermal" noise which is much more tolerated by our brain.

Things are even more complicated when we have a more complex digital signal chain than a single-step unit. Let's consider, for example, a CD player with digital volume control. We all know that dithering and noise shaping are necessary to avoid distortion increase when we approach low levels (high attenuation factors). This is due to the fact that the processing engine which does the attenuation (multiply for a number less than one) has generally the same resolution as the incoming signal: to the usual errors due to the finite-resolution mathematic (errors = noise), the final truncation adds a lot of damage, leading to the raise of distortion products. If we measure the effective resolution of a signal passed through a digital volume control, it is generally lower than that of the incoming signal. It's easy that the quantization noise after the attenuation raises over the system thermal noise. What was a good sounding digital signal before attenuation has become a "digital" sounding one.

Things improve as resolution increases: computational and truncation noise can remain below thermal noise, so that no dithering nor noise shaping are necessary. This is already true with 24-bit systems, and even better with 32-bit systems: even if we use 4, 8 or even 16 bit only to sample noise, that noise helps us to keep the sound good while it travels through our system to the loudspeakers.

Summarizing, I dare to say that bit depth is very important for sound quality, even more than sampling frequency. To test the above "on the field", take a good 96/24 recording and get a 96/16 version and a 48/24 version using some editing software. Chances are that you will hear little differences between the 96/24 and the 48/24 versions, while you'll hear a bigger difference between the 96/24 and the 96/16 versions.

M2Tech S.r.l.

Sede legale: Via Carlo del Prete, 2 – I-56121 Pisa (PI) – Italy

Sede operativa: C/O Polo Tecnologico di Navacchio – Via Giuntini, 13/N – I-56021 Navacchio di Cascina (PI) – Italy

Tel: +39 050 754195 – www.m2tech.biz – info@m2tech.biz

Coord. Bancarie: Cassa di Risparmio di Parma e Piacenza, Lung. Gambacorti, 4, Pisa – IBAN: IT 41 U 06230 14000 000043151277



M2Tech S.r.l.

P.I.: IT01821100508

Via Carlo del Prete, 2 – I-56121 Pisa (PI) - Italy

Sampling frequency

It's widely known that the ear of a young boy from the countryside (grown far from discos) can easily hear 20kHz, while a mature music lover hardly catches 16kHz. Thus, a digital audio system with an upper frequency limit of 20kHz should be enough to enjoy the real high fidelity. As usual, things are more complicated.

Complex signals contain multiple frequencies which interact to produce intermodulation in all systems in which they travel. Our ear, together with our skull bones, is one of these systems. High frequencies intermodulate to the lower frequency range (for example, a 21kHz and a 22kHz tones can modulate down to 1kHz, well into our audible range). If we record and/or play a recording through a system with 20kHz limit, we miss those tones which should intermodulate into our ear and head, losing some of the original information content (that 1kHz tone which is part of the original signal, even if is produced into our body). Recording professionals may say that no microphone can capture frequencies higher than 40kHz, so this may state the real useful high frequency limit in the recording-playback chain. Even so, this means a minimum sampling frequency of 80kHz (according to Nyquist, the minimum sampling rate to accommodate a certain bandwidth is twice the bandwidth). Standards indicate 88.2 or 96kHz, with a usable bandwidth of 44.1 or 48kHz, respectively.

But there is something more. It's known by signal processing experts, and absolutely not popularized amongst music lovers, that converting an analog signal into a discrete-time one (as it happens when converting from analog to digital) destroys the phase information in the two top octaves of the resulting spectrum. In a CD-standard digital recording, all phase information are lost from 5.5kHz up to 22kHz, which is the highest frequency present in that recording. This affects mainly harmonics (very few fundamentals are available over 5kHz), disrupting the notes's envelope. This may explain why different instrument of the same kind recorded on CD's often sound very similar.

To raise the lower limit of the affected range to 20kHz, we need to record with a bandwidth at least 80kHz, so we need a sampling frequency at least 160kHz. Standards indicate 176.4 or 192kHz, for a usable band of 88.2kHz and 96kHz, respectively.

Then aliasing comes in. Aliasing is a phenomena due to sampling process which must be avoided in order to keep the original sound quality. The only way to do it is low-pass filtering the original analog signal before converting it into digital. This can be problematic when the signal's upper frequency limit is very close to half the sampling frequency. In this case, a very steep filter (commonly called "brickwall") is necessary, which is affected by a sever phase rotation down to audible frequencies. This is the case of the CD, in which the upper frequency limit (20kHz) is very close to half the sampling frequency (22.050kHz). In CD-like systems, at least 90dB attenuation must be obtained with a transition band of just 2kHz!

M2Tech S.r.l.

Sede legale: Via Carlo del Prete, 2 – I-56121 Pisa (PI) – Italy

Sede operativa: C/O Polo Tecnologico di Navacchio – Via Giuntini, 13/N – I-56021 Navacchio di Cascina (PI) – Italy

Tel: +39 050 754195 – www.m2tech.biz – info@m2tech.biz

Coord. Bancarie: Cassa di Risparmio di Parma e Piacenza, Lung. Gambacorti, 4, Pisa – IBAN: IT 41 U 06230 14000 000043151277



M2Tech S.r.l.

P.I.: IT01821100508

Via Carlo del Prete, 2 – I-56121 Pisa (PI) - Italy

Using higher sampling frequencies means having a larger transition band, thus less steep filters and more gentle phase behaviour, affecting higher frequencies. A 24-bit, 192kHz system handling an audio signal with 20kHz bandwidth can use a transition band of 76kHz to attenuate at least 120dB. This means that a simple 10-pole filter is sufficient. Much better and easier to implement than a 200-poles filter used in CD-like systems! Even better, a 32-bit, 384kHz system may use a 7-pole filter, something which is very similar to the natural band limiting in analog systems.

Conclusions

High resolution is not marketing hype. It can help to make a digital system or device sound more similar to an analog one, provided users and experts keep in mind the real meaning and usefulness of having large bit depths and high sampling frequencies.

M2Tech S.r.l.

Sede legale: Via Carlo del Prete, 2 – I-56121 Pisa (PI) – Italy

Sede operativa: C/O Polo Tecnologico di Navacchio – Via Giuntini, 13/N – I-56021 Navacchio di Cascina (PI) – Italy

Tel: +39 050 754195 – www.m2tech.biz – info@m2tech.biz

Coord. Bancarie: Cassa di Risparmio di Parma e Piacenza, Lung. Gambacorti, 4, Pisa – IBAN: IT 41 U 06230 14000 000043151277