

## Digital Audio: Some Myths and Realities

By Robert Orban

Chief Engineer

Orban Inc.

November 9, 1999, rev 1 11/30/99

I am going to talk today about some myths and realities regarding digital audio. I have been following a number of the USENET newsgroups devoted to professional and high-end audio, and it's clear that, even 20 years into the digital audio revolution, there are still a lot of myths and misconceptions out there.

---

The first myth is that there is **no information stored below the level of the least significant bit** in digital audio. This is only true if dither is not correctly used.

**Dither** is random noise that is added to the signal at approximately the level of the least significant bit. It should be added to the analog signal before the A/D converter, and to any digital signal before its word length is shortened. Its purpose is to linearize the digital system by changing what is, in essence, "crossover distortion" into audibly innocuous random noise. Without dither, any signal falling below the level of the least significant bit will disappear altogether. Dither will randomly move this signal through the threshold of the LSB, rendering it audible (though noisy).

Whenever any DSP operation is performed on the signal (particularly decreasing gain), the resulting signal must be re-dithered before the word length is truncated back to the length of the input words. Ordinarily, correct dither is added in the A/D stage of any competent commercial product performing the conversion. However, some products allow the user to turn the dither on or off when truncating the length of a word in the digital domain. If the user chooses to omit adding dither, this should be because the signal in question already contained enough dither noise to make it unnecessary to add more.

It is possible to apply so-called “**noise shaping**” to dither. In the absence of “noise shaping,” the spectrum of the usual “triangular-probability-function (TPF)” dither is white (that is, each arithmetic frequency increment contains the same energy). However, noise shaping can change this noise spectrum to concentrate most of the dither energy into the frequency range where the ear is least sensitive. In practice, this means reducing the energy around 4kHz and raising it above 9kHz. Doing this can increase the effective resolution of a 16-bit system to almost 19 bits in the crucial midrange area, and is very frequently used in CD mastering. There are many proprietary curves used by various manufacturers for noise shaping, and each has a slightly different sound.

Noise shaping was first popularized by Sony’s “Super Bit Mapping,” although the principle as applied to high-quality audio was published by Michael Gerzon and Peter Craven in the late ‘80s.

Aggressive noise shaping can improve the signal to noise ratio in the midrange by as much as 18dB. However, it is a myth that **noise shaping always helps audio quality**. The total noise energy in a noise-shaped dither is always larger than the total noise energy in garden-variety white, triangular-probability-function dither. In the case of aggressive noise shaping, it can be *much* larger—by perhaps 20dB.

It is very easy to destroy the noise shaping by downstream signal processing like re-equalization, which uses multiplication and increases the word length. A digital to analog converter that is non-monotonic will destroy it as well. What happens is that the spectral dip around 4kHz tends to get filled in, resulting in far higher noise than one would have gotten if one had used simple white dither in the first place. Aggressively noise-shaped dither should only be used at the final mastering stage when the final deliverable recording is being created.

In **production**, words with higher numbers of bits should be used for distribution throughout the plant, and these signals should be dithered with white TPF dither. 20 bit words (120dB dynamic range) are usually adequate to represent the signal accurately. 20 bits can retain the full quality of a 16-bit source even after as much as 24dB attenuation by a mixer. There are almost no A/D converters that can achieve more than 20 bits of real accuracy and many “24-bit” converters have accuracy *considerably* below the 20-bit level. “Marketing bits” in A/D converters are outrageously abused to deceive customers, and, if these A/D converters were

consumer products, the Federal Trade Commission would doubtless quickly forbid such bogus claims.

In **digital signal processing devices**, the lowest number of **bits per word** necessary to achieve professional quality is 24 bits. Since this represents 144dB dynamic range, one would think that this is overkill. However, there are a number of common DSP operations (like infinite-impulse-response filtering) that substantially increase the digital noise floor, and 24 bits allows enough headroom to accommodate this without audibly losing quality. This assumes that the designer is sophisticated enough to use appropriate measures to control noise when particularly difficult filters are used. The popular Motorola 56000-series DSPs have 24-bit signal paths and 56-bit accumulators, and this is one reason why they are very popular in pro audio.

If floating point arithmetic is used, the lowest acceptable word length for professional quality is 32 bits. This word consists of a 24-bit mantissa and an 8-bit exponent, which is sometimes called “single-precision”.

A very pervasive myth is that **long reconstruction filters smear the transient response of digital audio**, and that there is therefore an advantage to using a reconstruction filter with a short impulse response, even if this means rolling off frequencies above 10kHz. Several commercial high-end D-to-A converters operate on exactly this mistaken assumption. This is one area of digital audio where intuition is particularly deceptive.

The sole purpose of a reconstruction filter is to fill in the missing pieces between the digital samples. These days, symmetrical finite-impulse-response filters are used for this task because they have no phase distortion. The output of such a filter is a weighted sum of the digital samples symmetrically surrounding the point being reconstructed. The more samples that are used, the better and more accurate the result, even if this means that the filter is very long.

It's easiest to justify this assertion in the frequency domain. Provided that the frequencies in the passband and the transition region of the original anti-aliasing filter are entirely within the passband of the reconstruction filter, then the reconstruction filter will act only as a delay line and will pass the audio without distortion. Of course, all practical reconstruction filters have slight frequency response ripples in their passbands, and these can affect the sound by making the amplitude response (but not the phase response) of the “delay line” slightly

imperfect. But typically, these ripples are in the order of a few thousandths of a dB in high-quality equipment and are very unlikely to be audible.

I have proved this experimentally by simulating such a system and subtracting the output of the reconstruction filter from its input to determine what errors the reconstruction filter introduces. Of course, you have to add a time delay to the input to compensate for the reconstruction filter's delay. The source signal was random noise, applied to a very sharp filter that band-limited the white noise so that its energy was entirely within the passband of the reconstruction filter. I used a very high-quality linear-phase FIR reconstruction filter and ran the simulation in double-precision floating-point arithmetic. The resulting error signal was a minimum of 125dB below full scale on a sample-by-sample basis, which was comparable to the stopband depth in the experimental reconstruction filter.

We therefore have the paradoxical result that, in a properly designed digital audio system, the frequency response of the system and its sound is determined by the anti-aliasing filter and not by the reconstruction filter. Provided that they are realized with high-precision arithmetic, longer reconstruction filters are always better.

This means that a rigorous way to test the assumption that high sample rates sound better than low sample rates is to set up a high-sample rate system. Then, without changing any other variable, introduce a filter in the digital domain with the same frequency response as the high-quality anti-aliasing filter that would be required for the lower sample rate. If you cannot detect the presence of this filter in a double-blind test, then you have just proved that the higher sample rate has no intrinsic audible advantage, because you can always make the reconstruction filter audibly transparent.

There is considerable disagreement about the audible benefits (if any) of **raising the sample rate** above 44.1kHz. **Stereophile Magazine** just reported a blind test of several different 20kHz lowpass filters applied to high sample-rate digital audio. Four experienced listeners first did blind A/B comparisons between full-bandwidth audio sampled at 96kHz, and filtered audio, still at 96kHz, using a digital audio workstation known to have very low jitter. None of them were able to identify the filtered audio; their results were equal to random guessing.

However, they then listened to a CD-R containing the same four selections, identified only as "1" through "4" with the order of the selections randomized.

Under the conditions where they always knew which cut they were hearing (but not the processing used, if any), they ranked their preferences for the sound of the four different cuts. It turned out that these preferences agreed exactly with the preferences they had earlier established in sighted tests, where they knew the processing applied to each cut. In the sighted tests, they preferred the unfiltered original.

An earlier test by well-known mastering engineer Bob Katz, using a somewhat higher-jitter workstation, resulted in Katz's being unable to hear any difference between the filtered and unfiltered signals. The four subjects of the current test reproduced this result; they reported that even moderate jitter completely masks the difference between the filtered and unfiltered signals.

This implies that 96kHz sampling may provide a subtle audible advantage. However, the fact that experienced listeners in the pro audio industry were unable to identify the filtered cuts in an A/B test means that the advantage is very subtle indeed, and is unlikely to be perceived by the average consumer. Moreover, four listeners and four cuts do not provide enough statistical data to rigorously prove anything, although the results are certainly suggestive.

Regardless of whether further, more rigorous testing eventually proves that 96kHz sampling is audibly beneficial, it has no benefit in BTSC stereo because the sampling rate of BTSC stereo is 31.47kHz, so the signal must eventually be lowpass-filtered to 15.734kHz or less to prevent aliasing.

Sample rates of 48kHz are beneficial in DTV, which uses this sample rate internally, but higher rates provide absolutely no further benefit.

Let's briefly discuss **jitter**, which has been on many people's minds lately.

One of the great benefits of the **digitization of the signal path** in broadcasting is this: Once in digital form, the signal is far less subject to subtle degradation than it would be if it were in analog form. Short of becoming entirely un-decodable, the worst that can happen to the signal is deterioration of noise-shaped dither, and/or added jitter.

Jitter is a time-base error. The only jitter that cannot be removed from the signal is jitter that was added in the original analog-to-digital conversion process so that the original samples were not quite uniformly sampled in time. *All* jitter added downstream from the original conversion can be completely removed in a sort of

“time-base correction” operation, accurately recovering the original signal. The only limitation is the performance of the “time-base correction” circuitry, which requires sophisticated design to reduce added jitter below audibility. This “time-base correction” usually occurs in the digital input receiver, although further stages can be used downstream. It is hard to build digital hardware that’s perfectly jitter-free, although the state of the art constantly advances. But always remember that the only place where jitter counts is right at the sample clocks of the A-to-D and D-to-A converters. Provided that the digital words themselves can be recovered, an arbitrary amount of jitter can be introduced elsewhere in the digital signal path, and it can be completely removed before D-to-A conversion, provided that your hardware is well enough designed.

Finally, let’s consider the myth **that digital audio cannot resolve time differences smaller than one sample period, and therefore damages the stereo image.**

People who believe this like to imagine a step function moving between two sample points. They argue that there will be no change until the step crosses one sample point.

The problem with this argument is that there is no such thing as an infinite-risetime step function in the digital domain. To be properly represented, such a function has to first be applied to an anti-aliasing filter. This filter turns the step into an exponential ramp, typically having equal pre- and post-ringing. This ramp can be moved far less than one sample period in time and still cause the sample points to change value.

In fact, assuming no jitter and correct dithering, the time resolution of a digital system is the same as an analog system having the same bandwidth and noise floor. Ultimately, the time resolution is determined by the sampling frequency and by the noise floor of the system. As you try to get finer and finer resolution, the measurements will become more and more uncertain due to dither noise. Finally, you will get to the point where noise obscures the signal and your measurement cannot get any finer. But this point is orders of magnitude smaller in time than one sample period.

---

---

So let's review the myths I discussed today.

- First is the myth that there's no information below the least significant bit in digital audio. With proper dither this is completely untrue.
- Second is the myth that noise-shaped dither gives you a free lunch. In fact, noise shaping is easy to destroy by downstream signal processing or imperfect conversion. So it should be used with considerable discretion.
- Third is the myth that long reconstruction filters cause smearing of transient information, and that short reconstruction filters therefore sound better. I have shown that this is completely incorrect, provided that all of the energy passed by the anti-aliasing filter falls in the passband of the reconstruction filter.
- Fourth is the myth that jitter matters anywhere in a digital audio system. In fact, the only places it matters are at the input and output converters. If it matters anywhere else, it means that your hardware is inadequate and has not completely removed the time base error.
- The last myth is that the time resolution of the digital system is limited to one sample period. This ignores the fact that all data in a digital system have been bandlimited by the anti-aliasing filter, so no sharp transitions occur between samples. The time resolution of a digital system is instead limited by the sample period and by the noise floor of the system, and can easily be nanoseconds, not microseconds.
- And finally, the jury is still out on the issue of sampling rates higher than 48kHz. One small study suggests that 96kHz provides slight audible benefits to expert listeners using the finest equipment. But no one claims that the advantages are large, or even moderate.