Letter to the editor

Marcel van de Gevel from Haarlem, The Netherlands, writes:

Dear Editor,

Mr van Maanen ("On the audibility of "high resolution" digital audio formats and how to test this"- Volume 5) is, of course, quite right when he claims that non-linear systems cannot be fully characterized by looking at the sine wave response and that the human auditory system is a non-linear system (at least that's how I've interpreted what he has written). However, I don't believe that the upper bandwidth limit of audio systems was chosen based on only sine wave tests, as van Maanen seems to imply. For example, a BBC report from 1957 by D.E.L. Shorter, W.I. Manson and E.R. Wigan, "Some experiments on the subjective effect of limiting the upper frequency range of programmes", report number L-034 (1957 / 10) discusses some tests done with actual music to find out what kind of audio bandwidth would be needed for audio lines for FM broadcasting. A more recent test with a different methodology and quite different results is T. Oohashi, E. Nishina, M. Honda, Y. Yonekura, Y. Fuwamoto, N. Kawai, T. Maekawa, S. Nakamura, H. Fukuyama, and H. Shibasaki, "Inaudible high-frequency sounds affect brain activity: Hypersonic effect", Journal of Neurophysiology, vol. 83, nr. 6, pp. 3548...3558, 2000.

Dr. van Maanen replies:

One of the major problems is that it is fundamentally impossible to determine the requirements for sound reproduction systems by sound reproduction systems: when something is "inaudible" is this because of the limitation of human hearing or because of the limitation of the sound reproduction system (including the microphone(s), sound recording and storage system)? By designing a sound reproduction system, you have to start somewhere and I have been told numerous times that the 20 kHz limit is based on the Fletcher-Munson curves. Apart from that, although I have deep respect for what people achieved 60 years ago, I seriously doubt that the equipment they had available in those days is superior to human hearing and any conclusions drawn from their work should be critically examined with our current knowledge, which, however, still leads to conflicting results. So far, I have never heard a sound reproduction system which comes even close to the live performance of a symphony orchestra. So there is still a lot of work to be done and we need deeper understanding of the workings of human hearing. In that perspective, I find the historic background of the 20 kHz limit less interesting; more interesting is the question whether we need an extended frequency response in order to bridge the gap with the symphony orchestra as this 20 kHz number has penetrated the whole audio business. Just look at the specifications of the different components from microphones to recording equipment to tweeters.

Mr van Maanen's definition of noise on page 65 is a very unusual one: the difference between the output of the anti-alias filter and output of the reconstruction filter. With imperfect filters, this means that even without any sort of quantization and without any added error signal, there will still be "noise", because the input and output signal of the reconstruction filter will not be quite the same. Surely with such a definition any analogue system with a bandwidth limit will also have a high "noise" level, even when it is completely free of hiss.

I have chosen the anti-aliasing and reconstruction filters such that with a perfect system the output of the reconstruction filter would be identical to the input to the digital system (= output of anti-aliasing filter). My aim was to find the contribution to the output signal due to the amplitude quantization and the interaction between the amplitude quantization and the reconstruction filtering. I found that the common way of defining the SNR of a digital system is incapable to describe these phenomena, so I used this definition to get a better description for the artifacts of digital systems without pretending this will be the final answer, but rather to show phenomena which had been overlooked, ignored or deliberately swept under the carpet. Being defined specifically for digital systems, it is incorrect to apply it to analog systems, but to me that is banging on an open door.

Regarding the figures in Mr Van Maanen's article, I notice that his sampled and quantized waveforms do not consist of discrete peaks but of staircase-like waveforms. That means that they have passed through a zeroeth-order hold filter. There is nothing wrong with that if this is only done to make the figures clearer, but if that zeroeth-order hold filter is accidentally put in the signal path, its roll-off might explain some of the error signals that Mr van Maanen finds. It would be interesting to see what happens to the error waveforms when the quantization is switched off, leaving only the filtering and the sampling. The errors should go to zero then, unless something is wrong with the filtering.

I think this question is a bit academic as all real life systems will use some form of reconstructed (continuous analog) signal as input for the reconstruction filter. So which conversion from the (theoretical) delta-peaks to a continuous analog signal should be chosen? As zero-order reconstruction is used by many systems, it is, in my view, a proper choice to illustrate the phenomenon of the interaction between the amplitude quantization and the reconstruction filtering (of which there are also numerous choices). Whether the phenomenon varies, and to which extent, by different choices of the numerous parameters of the reconstruction filtering is a subject in itself.

I disagree with van Maanen's claim in his footnote 7 that (non-subtractive triangular pdf) dither is only effective on infinitely long, static signals. This should not be the case as long as the dither is ergodic.

Suppose you take an ensemble of dithered quantizers, all properly dithered with 2 LSB peak-peak triangular pdf dither and all having statistically independent dither generators. If you apply a musical transient waveform to the inputs of all independently dithered quantizers in the ensemble and analyse the outputs, you will see that the average quantization error tends to zero when you make the ensemble big enough and that the standard deviation

(RMS value) of the quantization errors over the ensemble will tend to the signal-independent value that dither theory predicts. When each dithered quantizer produces an error with zero expected value and a signal-independent standard deviation, also on short transients, that means that each quantizer basically produces a small additive noise, not a distortion like an undithered quantizer would do.

This is precisely the point I am trying to make: on AVERAGE is works, but with transient signals you cannot (yet) average enough. I am not interested in the mathematical average of N dithers, I am interested in the sound I hear which comes from a single source! When does distortion become noise? Looking at short-term signals, the quantization steps will most likely happen at a different time, compared to a non-dithered version, but they are still there as can be seen in the **figure below**. The upper trace is without dithering, the lower one with. Only when the dithered signal is analyzed over an extended time lapse, the behavior of the quantization resembles that of noise.

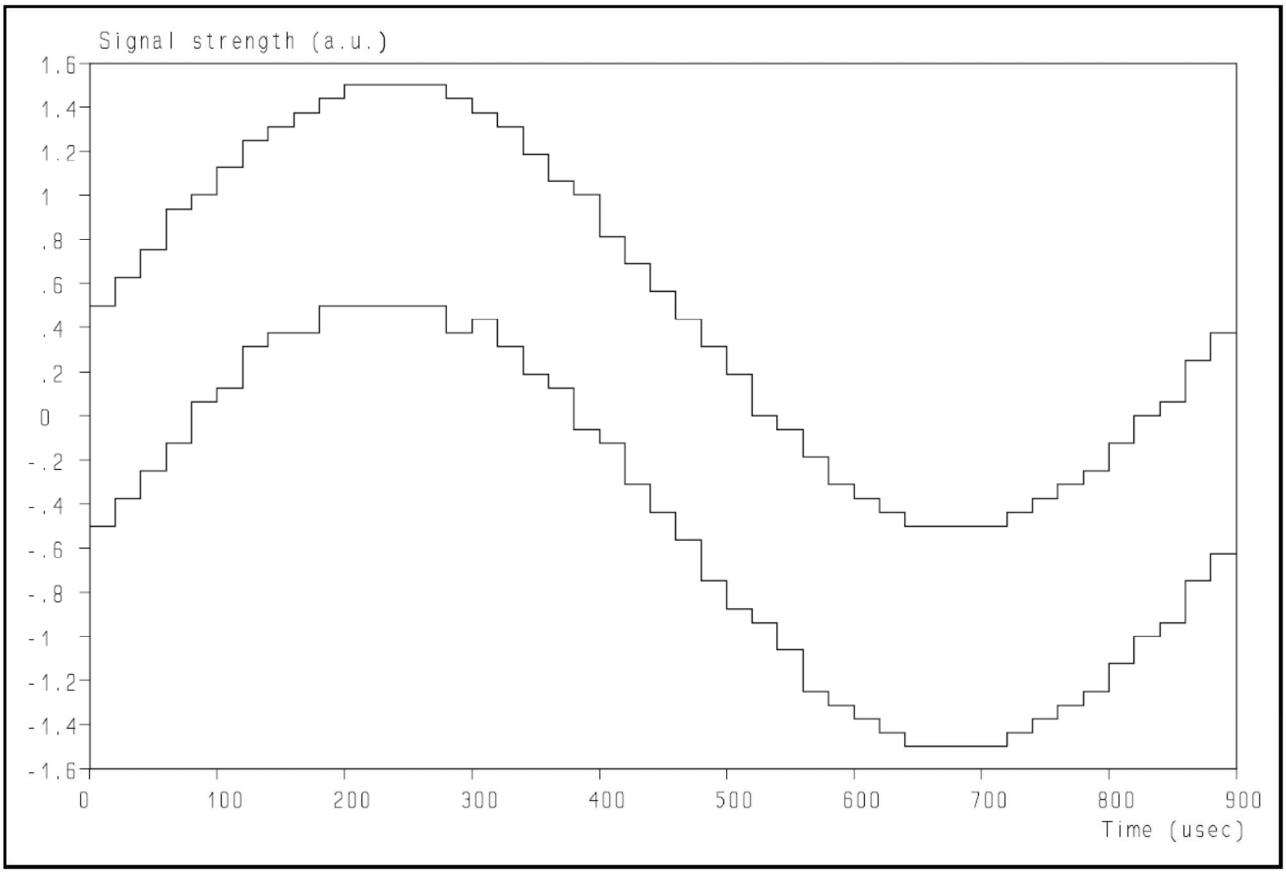

The noise of a dithered quantizer is not quite the same as additive analogue noise, though. For example, when the momentary value of the short musical transient is 12.2 LSB, each dithered quantizer in the ensemble will produce an error of -0.2+n LSB, where n is an integer that will be different from quantizer to quantizer, depending on the signal coming out of its dither generator. When the momentary value of the musical transient is 9.47 LSB, each dithered quantizer will produce an error of -0.47+n LSB. There always remains a statistical dependence between the input signal and the round-off error, even if the average and standard deviation are properly made signal-independent by dither. This again holds for musical transients as well as infinitely long signals.

I had a discussion about the same issue with Prof. Vanderkooij at one of the AES conventions. He argued similar, but I showed him that with short signals, the averaging is insufficient and he admitted that he had overlooked that point. So there is a difference between "continuous" and "transient" signals.

I don't understand why Mr van Maanen's figure 11 doesn't show any quantization error. If this figure is supposed to represent SACD, it should have a very coarse quantization into only two levels. Two levels is actually too few to properly dither it, which is one of the technical disadvantages of SACD compared to PCM. The other disadvantage is the relatively high level of ultrasonic noise. 24-bit 192 kHz PCM can cleanly reproduce the audio spectrum up to just below fs/2. The sigma-delta modulator on an SACD can't, because it has a quantization noise spectrum that rises enormously above 20 kHz.

Figure 11 shows the effect of a PCM system with an 1800 kHz sampling frequency (as described in the text and in the caption of the figure). As can be calculated, the single cycle of the 18 kHz signal includes 100 samples. At the resolution of this figure, these individual quantized steps are no longer distinguishable. In my **original figure 11, (see below and the enlargement added)** the quantization is still visible, albeit barely. So it is there, but because of the higher sampling frequency and cut-off frequencies of the filters, the signal is far better reproduced with a much smaller contribution of artifacts, which is the point I was trying to illustrate.

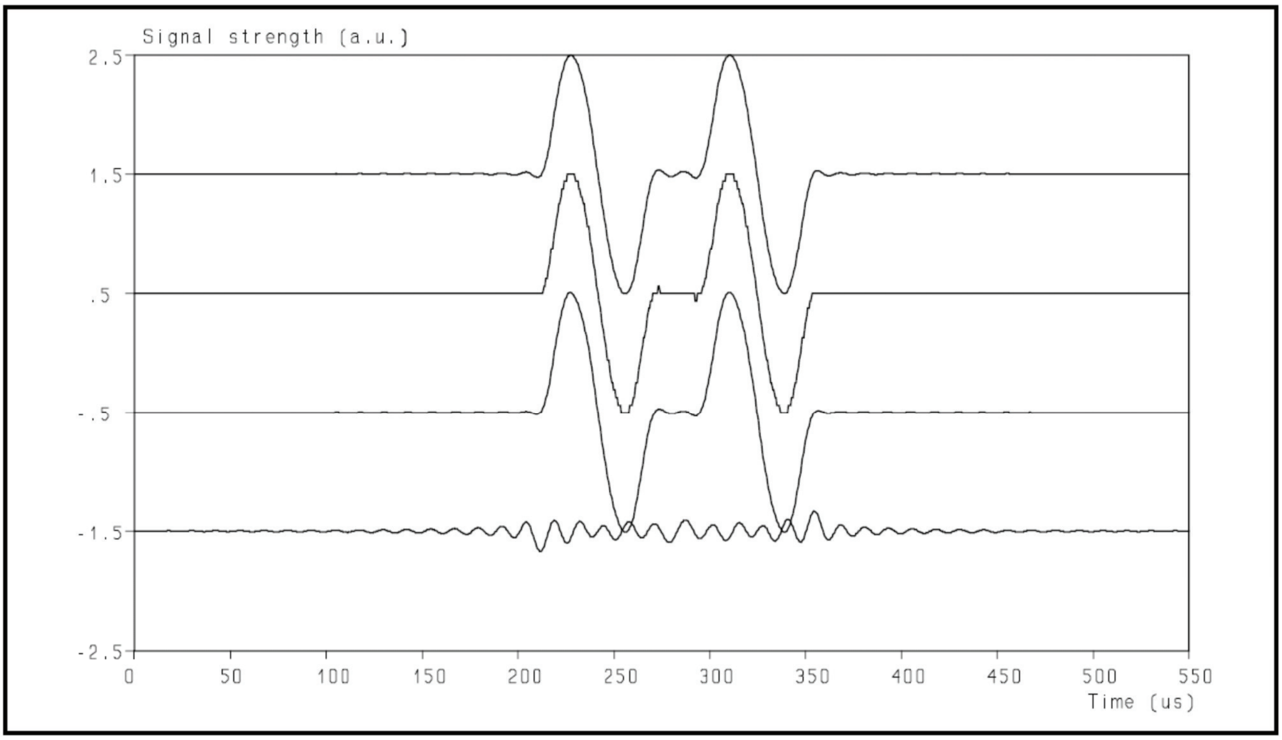

[Original Figure 11 from the article]

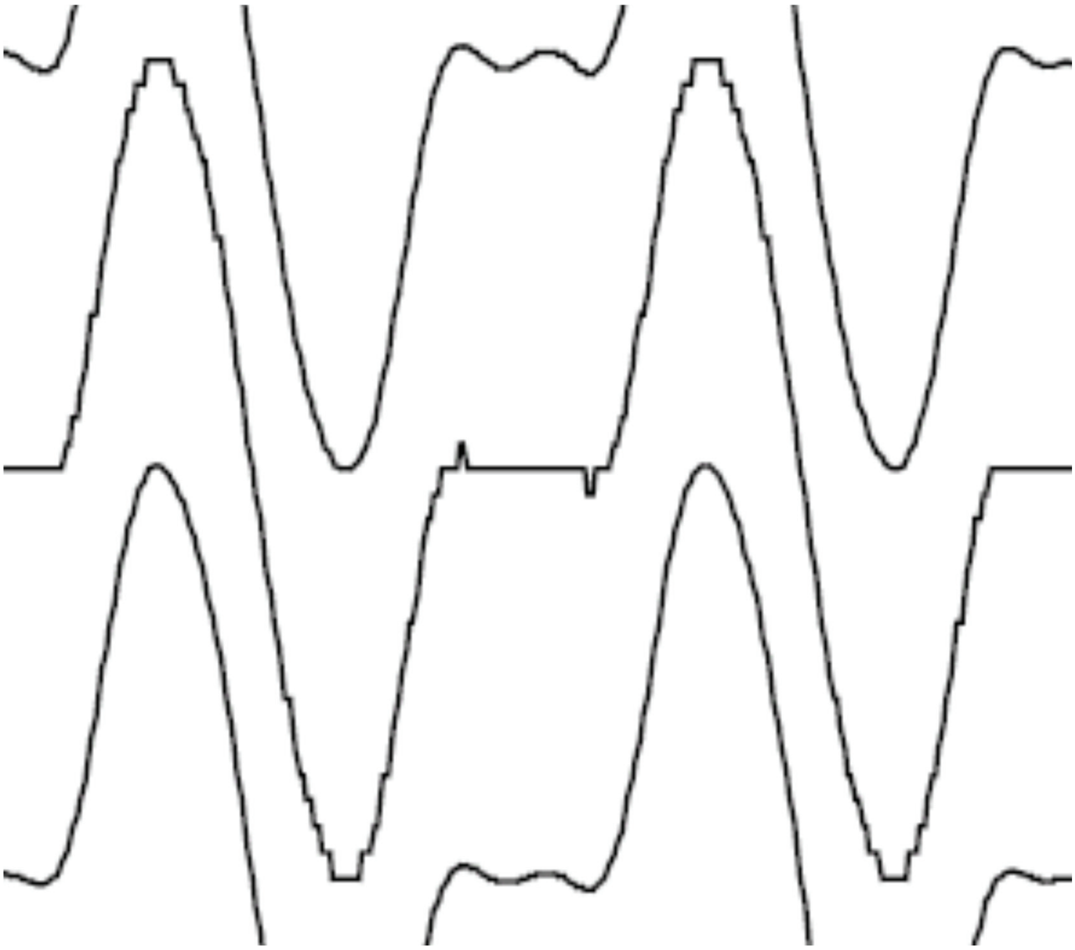

[Enlarged portion of the original Figure 11 from the article]