

=====
ASPECTS OF SAMPLING, OVERSAMPLING, QUANTISATION, DITHER
AND NOISE-SHAPING, AS APPLIED TO DIGITAL AUDIO.
Copyright, Christopher Hicks, November 1994. V1.11
=====

0. COPYRIGHT

=====

This document is copyright 1994, Christopher Hicks. Permission is hereby granted to distribute or store it in electronic or magnetic form, and to generate paper copies for private study, educational use or other non-commercial purpose provided that both of the following conditions are met

- 1) the document must remain complete and unaltered,
- 2) no charge may be levied for copying it, or for access to it.

I reserve all other rights to this document. I disclaim responsibility for damage or loss that may arise from inaccuracies herein.

Christopher Hicks, November 1994.

1. INTRODUCTION

=====

The aim of this article is to dispel as many of the myths surrounding the conversion of audio signals to the digital domain, and back to the analogue domain, as possible, without the aid of mathematics and (much more difficult) without the aid of diagrams. All of the buzz-words in the title are directly related to these two processes and are, to a large extent, analogous in the two.

The conversion of an analogue signal, such as the signal from a microphone, to a form in which it may be digitally stored or manipulated requires two distinct processes - those of sampling and quantisation. Sampling and oversampling are concerned with the capture of an analogue quantity at a certain instant in time; quantisation, dither and noise-shaping are concerned with the representation of this quantity by a digital word of finite length.

2. SAMPLING

=====

Sampling can be (roughly) defined as the capture of a continuously varying quantity at a precisely defined instant in time. Most usually, signals are sampled at a set of sample-points spaced regularly in time. Note that this section says nothing about the digital word format used to represent this sample - that is considered later in the section on quantisation. The Nyquist theorem states that in order to faithfully capture all of the information in a signal of one-sided bandwidth B , it must be sampled at a rate greater than $2B$. A direct corollary of this is that if we wish to sample at a rate of $2B$ then we must pre-filter the signal to a one-sided bandwidth of B , otherwise it will not be possible to accurately reconstruct the original signal from the samples. The frequency $2B$ that is the minimum sample rate to retain all of the signal information is called the Nyquist frequency.

The spectrum of the sampled signal is the same as the spectrum of the continuous signal except that copies (known as aliases) of the original now appear centred on all integer multiples of the sample

rate. As an example, if a signal of 20 kHz bandwidth is sampled at 50 kHz then alias spectra appear from 30 - 70 kHz, 80 - 120 kHz, and so on. It is because the alias spectra must not overlap that a sample rate of greater than $2B$ is required. In digital audio we are concerned with the base-band - that is to say the signal components which extend from 0 to B . Therefore, to sample at the standard digital audio rate of 44.1 kHz requires the input signal to be band-limited to the range 0 Hz to 22.05 kHz. Strictly speaking the input signal must be band-limited to infinitesimally less than 22.05 kHz, but this is of no practical significance.

2.1 Nyquist Sampling

=====

The obvious, old and hard way of sampling an analogue voltage at 44.1 kHz is to do just that - feed the voltage into a conventional track and hold sampler running at a 44.1 kHz sample rate. As shown above, this requires that the input signal be band-limited to half the sample rate, in this case 22.05 kHz, else the aliased spectra will overlap and information will be lost.

For a practical implementation this may require an analogue filter of order 8 or 10 to be inserted upstream of the sampler to provide an audio bandwidth of 20 kHz and also the 80 dB or so of attenuation above about 24 kHz that is required for high-fidelity sound reproduction. It is possible to design such a filter, but it would require a number of closely-toleranced components and would suffer from all of the usual ailments associated with analogue electronics.

2.2 Oversampling Analogue to Digital Converters

=====

The less obvious, but easier and cheaper way (at least with the advent of cheap VLSI multipliers) is to sample the input at a higher frequency, thereby relaxing the constraints on the analogue input signal spectrum, and then low-pass filtering and decimating (reducing the sampling rate) in the digital domain. To do this the input is sampled at a higher frequency such as 4×44.1 kHz. As before, this requires the signal to contain no significant components above half the Nyquist frequency, but because of the increased sampling rate the Nyquist frequency is now 176.4 kHz. Since the analogue filter can still start to roll off at 20 kHz or so, but it does not need to be 80 dB down until the first alias spectrum starts at 156 kHz it can be of lower order.

Now we have a digital data stream representing an analogue signal with components from dc to 88.2 kHz. This data is passed through a digital filter with a sharp cut-off at 22 kHz, which is relatively easy to implement, and can be made to have a precisely linear phase response. The filter output is a digital data stream representing the original analogue waveform, but with all the components above 22 kHz severely attenuated. Now we discard three out of every four of the samples to get our stream of samples at a rate of 44.1 kHz.

The scheme just outlined is conceptually fine. In a practical implementation the digital filter would be designed so that, rather than discarding the unwanted output samples, they do not have to be calculated in the first place. This represents a significant saving of computational effort.

2.3 Oversampling Digital to Analogue Converters

=====

Again the aim is to reduce the complexity of an analogue filter, this time the interpolation filter after the DAC, whose purpose is to remove the alias spectra from the converter output. For an audio signal of 20 kHz bandwidth the reconstruction filter has (ideally) to have a gain of one from dc to 20 kHz and a gain of zero from 24 kHz upwards. As in the case of the sampler this would require a complicated analogue filter.

If, however, we increase the sample rate by creating some new samples digitally then the first few alias spectra can be removed by a digital filter, relaxing the performance requirements of the analogue filter. For example a four times oversampling audio DAC runs at 176.4 kHz, so the first alias spectrum starts at around 156 kHz, and the analogue reconstruction filter can be of lower order since its transition-band is now 120 kHz wide.

The first step in this process is to insert three new samples in between each of the original ones. The value is unimportant but zero is often used as it enables an efficient hardware implementation. If zeros are used then the spectrum of the sampled signal at this point is unchanged, although the sample rate has been quadrupled.

Next, this faster sample stream is passed through a digital filter whose action is to make the new samples a smooth interpolation of the original data. The output of this filter is a sample stream at 176.4 kHz whose base-band spectrum (i.e. the music) is the same as the original, but the first three alias spectra of the original sampled

signal have now been removed.

Finally, the analogue signal is reconstructed with a DAC running at the four-times rate, and a low-order analogue filter, which removes the alias spectra centred around 176.4 kHz and multiples thereof.

2.4 Jitter

=====

Jitter is defined as the timing error at the transitions of a digital signal. Taking the ADC as an example, the effect of timing errors on the sample clock is to sample the input signal at slightly the wrong time instant, so although the average sample rate may be very accurately 44.1 kHz the samples are not necessarily taken exactly every 1/44100th of a second, but perhaps a little early or a little late. Since the input is constantly changing, a timing error on the sample clock translates to an erroneous sample level being captured.

The effects of sample clock jitter become more pronounced for high amplitude and high frequency input signals. The level and nature of jitter required for its effects to be audible is a current topic of research, debate and religious wars.

3. QUANTISATION

=====

At some point the sampled analogue quantity has to be converted to a finite-length digital word; this process is called quantisation. This will generally be done immediately after the sampler so that the subsequent data manipulation may be done digitally, but this is not

necessarily the case, for example if a switched-capacitor pre-filter is used. Common word-lengths used for digital audio are 16, 18, 20 and 24 bits. The best ADC and DAC chips around that are fast enough for digital audio have a resolution of 20 bits, though 16 and 18 bit parts are cheaper and far more common.

In a standard ADC (the "old, hard way" described above) the quantiser resolution and the output resolution are the same (since the digital output comes directly from the quantiser). Assuming a random input signal the errors associated with this quantisation process are white and uncorrelated, and yield a best-case signal to noise ratio of around $6n$ dB, where n is the number of bits in the output word.

3.1 Quantisation in Oversampling Converters

=====

Consider a 16-bit ADC clocked at 44.1 kHz. Its quantisation noise is approximately 96 dB below a full-scale sinusoid, and is spread evenly from dc to 22.05 kHz. If the ADC is clocked faster then the total quantisation noise power remains unaffected, but it is spread over a wider bandwidth. For example, if the converter speed is doubled, the quantisation noise power is spread from dc to 44.1 kHz. The desired signal is, of course, still in the band from dc to 22.05 kHz, and the quantisation noise power in this band is halved.

A digital low-pass filter with a cut-off of 22.05 kHz cuts out half of the quantisation noise, increasing the SNR by 3 dB, but leaving the audio-band signal unaffected. This filter is generally the same one as the digital anti-alias filter mentioned above. The process is extendible, and for each doubling of the sampling rate, the audio-band

quantisation noise is lowered by 3 dB. For example, the quantisation noise for a 4-times oversampling converter will be 6 dB lower (after filtering) than for the same converter operating without oversampling.

By the same principle, if we ran a 15 bit ADC at 4 x 44.1 kHz we would get the same audio-band performance as with a 16-bit device sampling at 44.1 kHz, since the increase of quantisation noise due to the poorer resolution is balanced by the SNR improvement brought about by oversampling. Again the process is extendible, and for each factor of four by which we increase the sample rate we can drop one bit of resolution off the converter.

So if we oversample by a factor of 4^{15} then theoretically we can drop 15 of the 16 bits, and use a one-bit converter. Alas, this implies a sample rate of about 50×10^{12} Hz, which is well into the infra-red. However, noise-shaping can be used to reduce the sample rate required to a such a level that the use of very low resolution converters is practical.

Oversampling gives similar benefits in the digital to analogue conversion process. For each factor of four by which the sample stream is oversampled, one bit may be dropped from each data word without significantly degrading the audio-band performance. Dropping bits introduces noise into the signal, but if the signal has been sufficiently oversampled then the power of this new noise in the audio band is lower than the noise in the original recording. Hence the original noise dominates and the wordlength truncation introduces insignificant amounts of extra noise.

Again, to reduce the word-length to one bit in this simplistic manner requires the same impractical sample rate as in the A to D case, but once again high-quality audio performance is achieved at practical sample rates with noise-shaping.

4. DITHER

=====

Much of the preceding text on quantisation and requantisation noise assumes the signal to be random. For many signals this is not the case, and the result is that the quantisation noise, rather than being white, is found to be highly correlated with the signal. This manifests itself as very nasty-sounding level-dependent distortions which become more prominent as the signal is decreased in amplitude.

To avoid this problem a small amount of noise is added to the signal before quantisation, this process being known as dithering. The dither values are often drawn from a triangular distribution, and it is desirable that they be uncorrelated. The dither power is chosen such that the quantiser transfer function is just linearised, without adding excess noise to the quantised signal. This has the effect of decorrelating the quantisation noise from the signal, and avoiding quantisation-related distortion.

5. NOISE-SHAPING

=====

Quantisation (after the addition of white-noise dither) ideally results in a white power spectrum - that is, the noise floor has

constant noise power spectral density (NPSD). This is a direct result of rounding to the nearest value when performing the quantisation, with a random input signal.

However, if we base our decision of whether to round up or down not upon which is the nearer value but upon some other criterion, then we can make the output quantisation noise spectrum have almost any form we desire, but still have (roughly) the same total power. We are not going to hear noise above about 20 kHz, so we force as much of the quantisation noise as possible into the band above 20 kHz.

The Nyquist frequency for a 256 times oversampling converter is about 11.2 MHz, (compared with 50 THz calculated for the hypothetical 1-bit converter above). If we oversample by 256 times then we can put as much noise as we want into the band from 20 kHz to 5.6 MHz, where it will not interfere with the audio. By keeping the quantisation NPSD low enough in the audio band we are able to achieve an audio-band SNR of 90 dB or more; the NPSD above 20 kHz will be very much higher, but since there is no desired signal at those frequencies it really does not matter.

This is accomplished in practice by placing the quantiser in a feedback loop with a digital filter, such that the filtered quantisation error is subtracted from a subsequent input sample. A detailed explanation of this is unfortunately beyond the scope of this article, being a virtual impossibility without pictures and a few equations.

5.1 Noise Shaping in Mastering

=====

Noise shaping is used in a very similar manner in mastering processes such as Deutsche Gramophon's ABI and Sony's Super Bit Mapping. Using a 20-bit ADC to record the master tape and then requantising to 16-bits with a suitable (proprietary) noise-shaping algorithm then, rather than have a white noise floor, they achieve low NPSD where the ear is most sensitive at around 3 kHz, and a much higher NPSD at high frequencies. The shaped noise floor is found to be subjectively quieter, though the total quantisation noise power is actually slightly higher.

6. CONCLUSION

=====

Various.htmllects of analogue to digital, and digital to analogue conversion have been discussed, with particular reference to digital audio applications. It was seen that A to D conversion is performed as two distinct processes - sampling and quantisation.

Criteria were stated for adequate performance of both processes, and various schemes to circumvent associated problems were put forward. Emphasis was placed on the technique of oversampling and the application of dither and noise-shaping was also mentioned.