Theory of Upsampled Digital Audio

Doug Rife, DRA Labs
Revised May 28, 2002

Introduction
This paper is a detailed exposition of the ideas on this subject first published as a brief Letter to the Editor in the December 2000 issue of *Stereophile Magazine*.

In the last few years, several DACs have appeared on the market and sold as "upsampling" DACs. Upsampling DAC manufacturers claim that their products improve the sound quality of standard CDs as compared to conventional DACs and most listeners agree. What can possibility account for the improved sound quality? After all, the digital data on the CD is the same no matter what DAC is used to convert it to analog form. The most popular theory to explain the superior sound quality of upsampling DACs is the time smearing theory. As will be explained shortly, all upsampling DACs employ slow roll-off digital reconstruction filters as opposed to conventional DACs, which employ sharp roll-off or brick wall reconstruction filters. Since slow roll-off filters show less time smear than brick wall filters when measured using artificial digital test signals, many have concluded that reduced time smearing is responsible for the subjective improvement in sound quality when playing CDs through upsampling DACs. Although this seems logical and is an appealing explanation, as will be shown, it falls apart upon closer examination. The main reason is that the digital audio data residing on a CD is already *irreparably time smeared*. No amount post-processing of the digital audio data by the playback system can possibility remove or reduce this time smearing. This paper concludes that the slow roll-off characteristic of upsampling digital filters is indeed responsible for the improved sound quality but that the mechanism has nothing to do with time smearing or the lack of it in the playback system.

In common usage, oversampling is the process by which the sampling rate of a digital audio pulse code modulated (PCM) data stream is increased to allow digital filtering to largely replace analog filtering in reconstructing the original analog audio signal. In the early days of CD playback, it was discovered that the analog reconstruction filters then used in CD players were expensive, cumbersome and prone to various forms of overload and distortion. It was soon recognized that these high-order analog reconstruction filters, also known as anti-imaging filters, could be replaced by a digital filter operating at a large multiple of the original sampling rate. Virtually all CD players today employ an 8-times oversampled digital filter driving the digital-to-analog converter (DAC). A non-critical low-order analog filter is usually added after the DAC to remove the very high frequency images created by the DAC's zero-order hold function.

For perfect waveform reproduction, sampling theory requires that the digital reconstruction filter remove all frequency components above the Nyquist frequency or ½ the sampling rate of the input digital data. Because CDs encode audio data sampled at 44.1 kHz, this requirement implies that the digital reconstruction filter remove all frequency components above 22.05 kHz. This filter therefore must be sharp, meaning that it has a steeply sloped transition band. The filter's frequency response must fall by over 90 dB in the narrow range from 20 kHz to 22.05 kHz. Let's examine this requirement for a sharp filter in a little more detail. Signal processing theory dictates the use of a sharp reconstruction filter only to accurately reproduce the exact shape of the original audio waveform. From a frequency domain perspective, such a brick wall reconstruction filter is not at all required to maintain perfect fidelity of the baseband audio frequency components lying below the Nyquist frequency of 22.05 kHz. Indeed, omitting the reconstruction filter altogether only has the effect of adding additional frequency components above 22.05 kHz with no effect whatsoever on the baseband audio spectrum below 22.05 kHz.

If the anti-imaging filter is omitted then ultrasonic images of the baseband audio frequencies will be present at the analog output. Analog audio components following the CD player, such as preamplifiers and power amplifiers, may not be able to handle the ultrasonic images and generate audible non-linear distortion. Because manufacturers of CD players cannot predict in advance what other components will be driven by their products they tend to choose the most conservative option which is to use a sharp digital anti-imaging filter followed by a slow roll-off analog filter having a cut-off frequency far above the audio band.

In light of the foregoing, it is now possible to define the term “upsampling”. An upsampling digital filter is simply a “poor” oversampled digital reconstruction filter having a slow roll-off rate. Some ultrasonic images of the baseband spectrum are allowed to appear at the output of the digital filter and this ultrasonic energy is input to the DAC. An upsampling digital filter still attenuates the ultrasonic images to a certain degree, the attenuation increasing with frequency. Upsampling is therefore a compromise between using a sharp anti-imaging filter and using none at all. An upsampling digital reconstruction filter purposely violates signal-processing theory. The reconstructed analog waveform is no longer the same as the original even though the baseband audio frequency components are not affected in any way. Waveform fidelity is compromised by the addition of ultrasonic images of the baseband audio signal. The first such image is a folded and frequency-shifted version of the baseband audio spectrum and lies between 20.05 kHz and 44.1 kHz. For example, a 20 Hz tone in the audio baseband gives rise to an image tone at 44.08 kHz. In contrast, a 20 kHz baseband tone gives rise to an image at 24.1 kHz which is much lower in frequency than the image of the 20 Hz tone at 44.08 kHz. This folding operation means that the baseband audio signal below 20 kHz and its first image lying between 20.05 kHz and 44.1 kHz are uncorrelated. Using a slow roll-off upsampling anti-imaging filter, as opposed to omitting the anti-imaging filter altogether, means that the first image dominates the ultrasonic content of the digital filter's output. All higher frequency images are necessarily more attenuated than the first image by the roll-off rate of the anti-imaging filter. The lack of correlation between the baseband audio signal and its ultrasonic images has important implications. These implications will be addressed after first considering some other features of slow roll-off anti-imaging filters.

The Time Smearing Theory

A slow roll-off anti-imaging filter also has another consequence. Because of its slow roll-off rate, its impulse response exhibits less ringing and time smearing as compared to a conventional sharp oversampled anti-imaging filter. Fourier analysis shows that the time resolution and bandwidth of all filters are inversely related. A slow roll-off filter allows more high frequency energy through compared to a fast roll-off filter. Therefore, the slow roll-off filter has better time resolution, or equivalently, less time smearing.

The conventional wisdom as to why gentle oversampling filters sound better than sharp oversampling filters focuses on these differences in the time domain response. The idea is that the since slow roll-off filters show less time smearing this *must* be the explanation for the subjective improvement in the sound quality of upsampled digital audio playback. There are two serious problems with this explanation. First, we must consider how a CD is made. The counterparts of reconstruction filters used in CD playback, are anti-aliasing filters used in recording and CD production. Unfortunately, the roll-off rate of anti-aliasing filters for 44.1 kHz sampling cannot be made gentle without severely compromising the baseband audio signal. If a sharp anti-aliasing filter is not used during recording at 44.1 kHz or when down-converting a 96 kHz master to 44.1 KHz, folded images of the baseband frequencies will fall right back into the audio baseband instead of being frequency-shifted to ultrasonic frequencies as occurs in playback. These in-band images or aliases are clearly audible even at extremely low levels. To prevent audible in-band images from appearing, the decimating anti-aliasing filter used to down-convert 96 kHz to 44.1 kHz, must be very sharp indeed and must also severely attenuate all frequencies above 22.05 kHz. Thus, anti-aliasing filters as well as decimation filters must introduce time smear at least as severe as that introduced by sharp anti-imaging filters used in playback. What's more, the time smearing due to the anti-aliasing filters cannot be removed. This is so because the time smearing on CDs is a consequence of the fact that the signal has been strictly bandlimited to 22.05 kHz. Any frequency components above that frequency which may have been recorded on the original the 96 kHz master are gone forever.

In addition to anti-aliasing and decimation filters, the time smearing theory also disregards even more severe time smearing that is routinely introduced by loudspeakers. The sonic improvements of slow roll-off anti-imaging filters are audible through time-coherent speakers but just as readily audible when listening through the vast majority of loudspeakers many of which show relatively poor time domain response. Whatever role time smearing plays in audio quality it is not sufficient to explain the increase in sound quality due to upsampled digital audio. The main reasons are: a) the time smearing of the kind produced by a sharp anti-imaging filter is already present in the digital data on every CD ever made and; b) the time smearing added by the vast majority of loudspeakers is much larger than the time smearing caused by the combined time smearing of both the anti-aliasing and anti-imaging filters.

Differential Non-linearity in DACs

What all successful upsampling DACs have in common is that they allow a small but significant amount of ultrasonic energy through their digital signal chains. This fact is key to understanding why sound quality is improved by upsampling conversion products. The improvement in sound quality is due to the presence of these ultrasonic images and not due to any minor changes in time smearing. In other words, the reduced time smearing exhibited by slow roll-off digital filters is a just a big red herring that has mislead a great many people.

If the improvement in sound quality is due to the ultrasonic images then what possible mechanism can explain it? DACs are known to exhibit at least two forms of non-linear distortion. The first, integral non-linearity is similar to the non-linearities typically found in loudspeakers, phono cartridges and tube amplifiers. Integral non-linearity is a large-scale and smooth deviation of the DAC's transfer function from a perfectly straight line. The second type of non-linearity found in DACs is differential non-linearity. This form of non-linear behavior has a finely detailed grainy structure. In an ideal DAC, each transition from one digital input code to the next higher code results in exactly the same increase in the analog output current or voltage, as the case may be. For example, if an ideal DAC's digital input is changed from 1000 to 1001, a DVM monitoring its output may register a change of 100 microvolts, just to pick a round figure. In that case, one would expect that changing the digital input code from 1001 to 1002 would again increase the output voltage by exactly another 100 microvolts. However, real world DACs do not show the same change in the analog output for every equal change in the digital input code. In fact, absolutely no change in the output voltage is possible for a one LSB (least significant bit) increase in the digital input code with most commercial DACs and this would still be within their published specifications.

What's important here is that while differential non-linearity is found in all DACs, it has no counterpart in the analog domain. Differential non-linearity, like jitter, is a uniquely digital form of non-linear distortion. Just as very small amounts of jitter can degrade digital sound quality so too can very small amounts of differential non-linearity.

In mutibit DACs, distortion due to differential non-linearity tends to increase with signal level. This occurs because the more significant bit transitions show more differential non-linearity error than the less significant bit transitions. The reason has to do with the architecture of most mutibit DACs and is beyond the scope of this paper. What's important here is that the most significant bits are in use only during loud musical passages. Therefore, distortion due to differential non-linearity in multibit DACs tends to increase during loud musical passages, which can obscure low-level detail and result in a subjectively grainy, harsh and sterile sound.

Upsampling Ameliorates Differential Non-linearity

One way to reduce differential non-linearity is to average the outputs of two or more DACs connected in parallel. Indeed, several CD players use this technique to improve sound quality. The idea is that the differential non-linearity errors in each individual DAC will tend to be random and thus tend to average out when several DACs are connected in parallel. Upsampling provides another method for averaging away differential non-linearity but without requiring multiple DACs. The ultrasonic image energy that a slow roll-off anti-imaging filter presents to the DAC can be thought of as a form of dither. As mentioned above, the image spectrum is folded which has the effect of de-correlating it from the baseband audio signal thus making it random for all practical purposes. To see how this works consider that for each original sample of the 44.1 kHz data stream there are normally 8 interpolation samples generated by the oversampling digital filter and presented to the DAC. If no ultrasonic energy is present, as is the case using a conventional sharp anti-imaging filter, the amplitude of each sample changes very little from one to the next within the 8-sample interpolation period. With the ultrasonic images present, as occurs with a gentle anti-imaging filter, these interpolated samples can and do change significantly and, for all practical purposes, randomly, from one sample to the next. We can think of the waveform presented to the DAC as consisting of the superposition of a slowly changing audio baseband signal plus a rapidly changing ultrasonic image signal. The ultrasonic image signal essentially frustrates any systematic interaction between the DAC's differential non-linearity errors and the slow-changing baseband audio signal, which would otherwise create audible non-linear distortion. Instead, the interaction of the ultrasonic dither with the DAC's differential non-linearity simply increases the noise floor of the DAC somewhat.

Eventually, the ultrasonic energy is filtered out, first by a low-order analog filter, then the power amplifier, the loudspeakers, the air in the room and finally by the ear. This post-DAC low pass filtering can be thought of as an averaging operation. Recall that, unlike integral non-linearity, differential non-linearity is a highly localized effect. Digital input codes differing by no more than a few LSBs generate differential DAC errors that are not statistically related to each other. Thus, each of the eight samples in the interpolation period can be thought of as representing eight separate DACs each having differential linearity errors that are statistically independent from one another. Their outputs are effectively averaged together by the low pass filtering provided by the post-DAC analog filtering.

What about simply adding a large amount of random ultrasonic dither to the digital signal presented to the DAC. Random ultrasonic dither would also reduce distortion due to differential non-linearity. Some digital filter IC products even include provisions for adding random ultrasonic dither. However, this feature is rarely enabled because the noise floor of the DAC is increased depending upon the amount of differential and integral non-linearity present in the DAC as well as the level of the ultrasonic dither selected. With upsampling, the level of the ultrasonic dither exactly tracks the level of the baseband audio at all times. Thus, the dither provided by upsampling is signal-dependent and not fixed like conventional dither. Auditory masking effects ensure that the increase in the noise floor is not perceivable. At low baseband audio levels, the ultrasonic dither provided by gentle anti-imaging filters is also low and so the noise floor is not raised significantly. At high audio baseband levels, the level of the ultrasonic dither provided to the DAC is also high which raises the noise floor but auditory masking ensures that this is inaudible. Note also that the signal-dependent ultrasonic dither provided by upsampling is large when most needed because the differential non-linearity of multibit DACs also tends to increase as the signal level increases.

Anecdotal evidence suggests that certain upsampling DACs sound better than others. This is to be expected based on the lack of understanding of why upsampling works, even by those who design upsampling DACs. In the light of the theory presented here, certain ultrasonic roll-off characteristics should provide a statistically more effective signal-dependent ultrasonic dither than other roll-off characteristics.   In particular, enough ultrasonic image energy must be present to sufficiently dither the DAC all the way to the Nyquist frequency at the output sampling rate. Otherwise, the benefits of upsampling to very high output sampling rates will not be fully realized. On the other hand, too much ultrasonic energy can raise the noise floor of the DAC enough to be audible. It may turn out that a rigorous mathematical analysis of upsampling would discover an optimum digital upsampling filter characteristic that maximizes the effectiveness of the signal-dependent ultrasonic dither provided to the DAC.

Final Thoughts

The sound quality of 44.1 kHz digital audio data can be dramatically improved by employing a “poor” oversampling digital anti-imaging filter having a slow roll-off in place of a “good” digital filter having a fast roll-off and a high stop band attenuation. It was shown that the ultrasonic images output by this “poor” filter is responsible for the improved sound quality, reducing certain forms of non-linear distortion such as that due to the differential non-linearity found in all DACs. There may very well be other, subtler, forms of non-linear distortion in DACs, which may also be reduced by signal-dependent ultrasonic dither.

In any case, there are certainly many other sources of non-linear distortion present in the signal chain. Some may question how such a small reduction in non-linear distortion due to differential non-linearity in DACs can be heard when much larger non-linear distortions are generated by loudspeakers, for example. The answer is that the non-linear distortions in question, like jitter-induced non-linearities, are uniquely digital in origin. Such digital distortions have no counterpart in the analog domain. It can be argued that human hearing is much more sensitive to certain digital forms of distortion as compared to the more common distortions of analog origin.   For example, it is widely recognized that very low levels of jitter are audible even in the presence of much larger levels of harmonic distortion generated by loudspeakers.