Upsampling Induced Distortion?

I don't know who coined the word 'upsampling' but it seems a good guess that it was prompted by wanting a term similar to 'oversampling', but which implied that the process was subtly different. Unfortunately, the word is now used by many manufacturers in their sales literature and it isn't always clear that they are using the same definition. So what is the vital distinction?

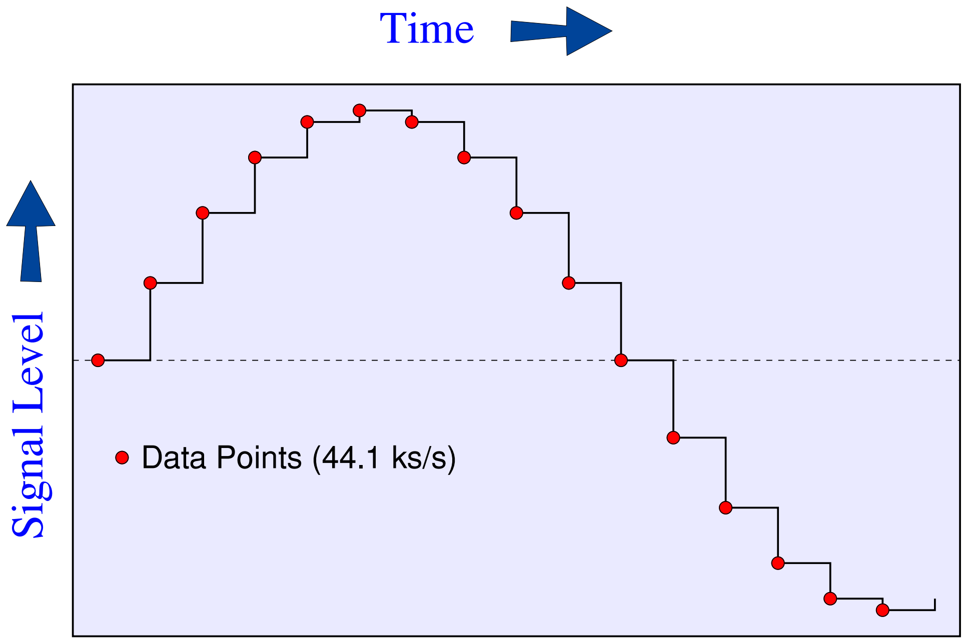

Fig 1a Sampled waveform with simple sample and hold

Oversampling DACs have been employed in Compact Disc Audio (CD-A) players right from the first domestic consumer models. If we use a simple ×1 DAC with no following low-pass filtering we get an output where the desired audio waveform is accompanied by spectral images of this at higher frequencies. To illustrate this we can start with a simple example where the sampled waveform was a sinusoid.

Figure 1a shows a brief snippet from our waveform. Each sample is indicated by a red circle. The assumed analogue output voltage level from the DAC is represented by a continuous black line. Here, for the sake of example, a sample-and-hold process is used to determine the level in between sample instants. This introduces a 'staircase distortion' effect. In terms of the power spectrum this distortion shows up as a series of unintended high frequency components that we can regard as images of the low-frequency components of the intended part of the output signal.

Figure 1b shows the power spectrum of this reconstructed waveform. As we might expect from the jagged shape shown in Figure 1a, the intended waveform sinusoid is accompanied by a series of unintended components which are symmetrically disposed about harmonics of the sampling frequency. i.e. Given an intended signal component at a frequency, f, we find that the reconstruction also includes components at a series of frequencies, $Nf_S \pm f$, where N is any positive integer and f_S represents the sampling frequency for the CD-A data stream. These extra components above $f_S / 2$ are called *images* of the intended frequency component below $f_S / 2$.

Note that in Figure 1b, and with the other spectra shown in this article, a linear frequency scale has been used. This makes it easier to appreciate the patterns of the unwanted components which may arise.

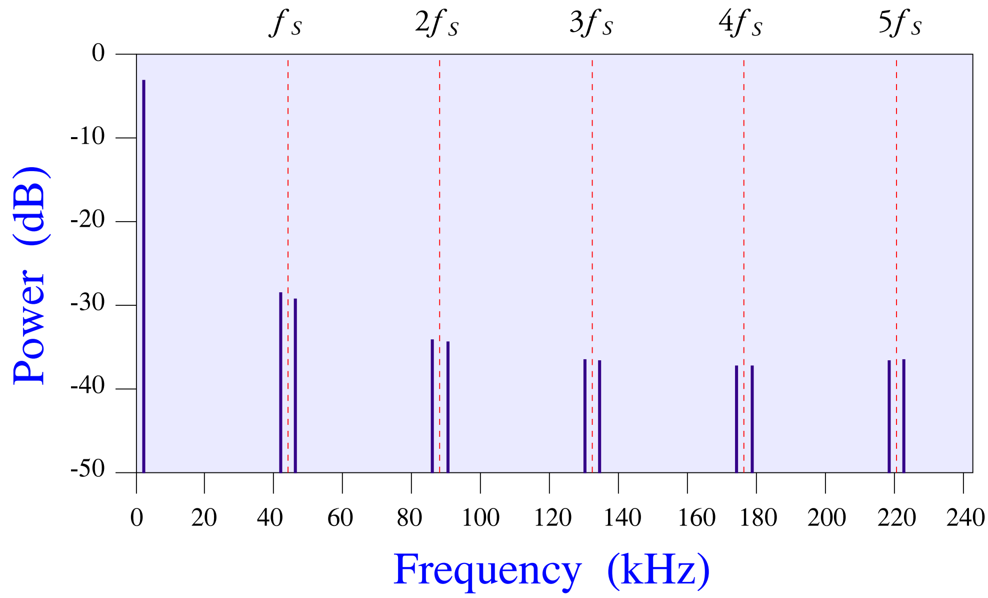

Fig 1b Power Spectrum of signal used for Fig 1a

This problem is well known. The standard method for dealing with it is to apply suitable low-pass filtering to suppress the unwanted ultrasonic frequency components. This filtering can be applied by passing the output from the DAC through an analogue low-pass filter. In principle at least, this can remove the staircase distortion, and suppress the unwanted high frequency image components. In practice, however, simply relying on such analogue filters means we may require very high performance analogue circuits in order not to compromise the performance implied by the dynamic range, etc, of the CD-A system.

An alternative approach for removing these unwanted image components is to perform some of the filtering in the digital domain by using an oversampling method. This replies upon assuming that the series of sampled values is what information scientists call a *complete record*. From information theory this means they tell us the level of the original audio waveforms at **all** instants within the duration of the recording – even those in between the sampled instants. We can therefore take the initial data and use it to generate the sampled levels we would have obtained if we'd taken extra samples at intermediate instants. (In fact, the same assumption is also implicit in the use of an analogue filter during the reconstruction process.)

The first generation Philips chipset for CD-A players employed an ×4 oversampling arrangement based upon the use of a Transverse Digital Filter. Figure 2a illustrates the effect of employing such a system to create oversamples. It is important to bear in mind that – provided that the samples *were* collected in accord with the sampling theorem – these new values are not 'guesses' or 'best fit' interpolations. They recreate the actual sample levels we would have obtained if we had taken extra samples at these instants when recording the original waveform. For this reason they came to be called 'oversamples' because actually taking them would have been overdoing things for the chosen audio bandwidth so far as information theory is concerned.

Figures 1 and 2 use the same initial waveform sample values. By comparing Fig 1a with Fig 2a we can see that oversampling has two obvious effects. The 'staircase' now has 'steps' of shorter duration. They also tend to be smaller in amplitude. This immediately implies that the overall amplitude of this unwanted staircase distortion will tend to be reduced by the oversampling process. In general terms we can tend to expect oversampling by a factor of N to reduce the mean power of the staircase distortion by a factor of N^2.

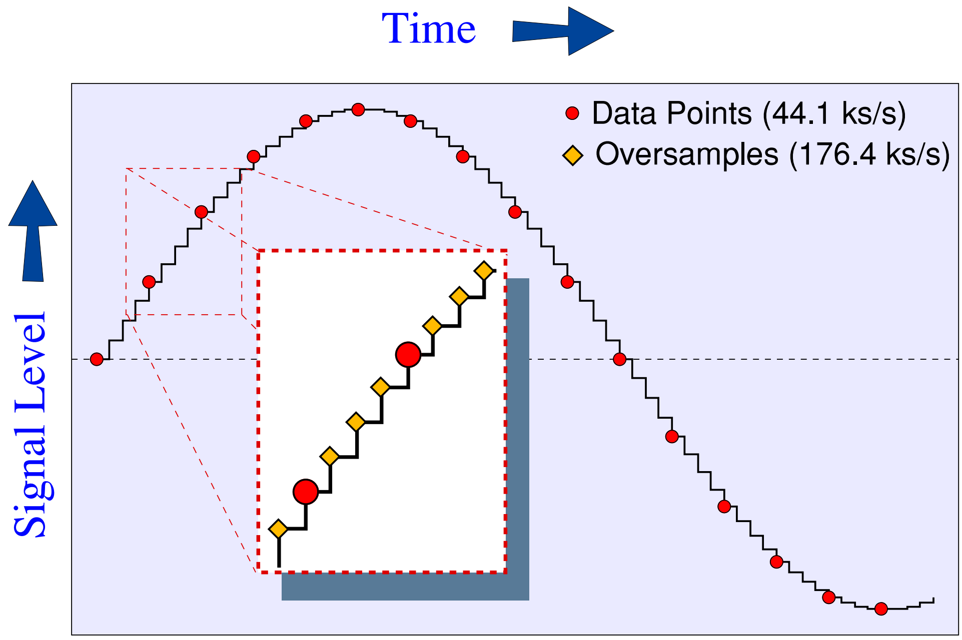

Fig 2a x4 Oversampled waveform with sample and hold

In effect, oversampling increases the sampling rate presented to the DAC, but without changing the bandwidth of the required output.

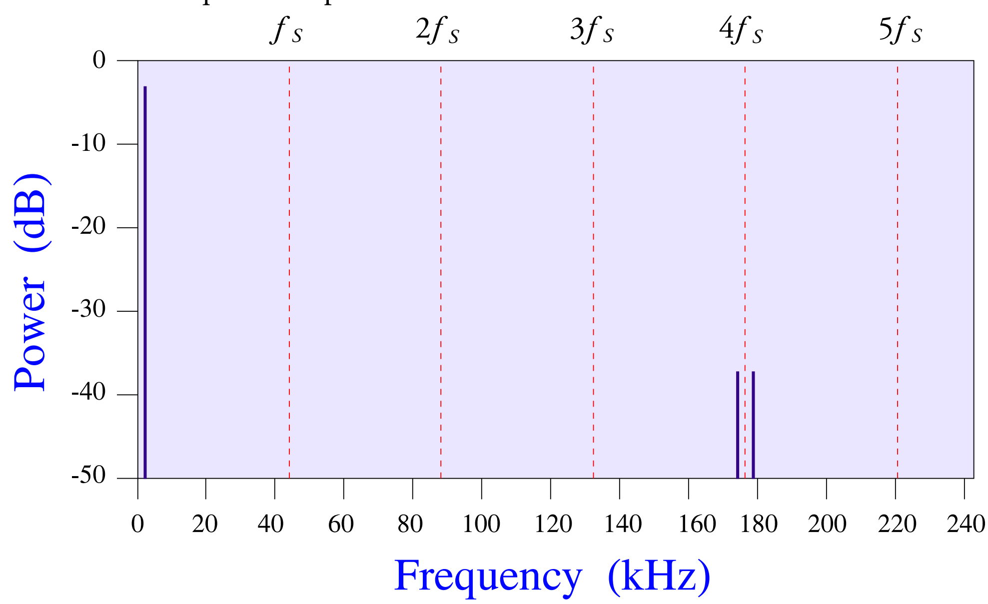

Fig 2b Spectrum of x4 oversampled signal used for Fig 2a

Figure 2b shows the spectrum of our example ×4 oversampled waveform. Comparing this with Figure 1b we can see that the ×4 oversampling has essentially removed the image components around the lower harmonics of f_S below $4f_S$. In general, oversampling by a factor of N will tend to remove the unwanted image components which would otherwise surround the harmonics below Nf_S. This makes it much easier to design a following analogue low-pass filter that can be used in domestic mass-market players and meet the specifications implied by CD-A, etc.

For the examples in this article I will assume we are starting with CD-A data (i.e. 44,100 samples/second) and that any oversampling, upsampling and digital filtering is based upon the use of Transverse Digital Filters (TDF). These assumptions are convenient for the purposes of explanation, but the conclusions we obtain should apply in more general circumstances.

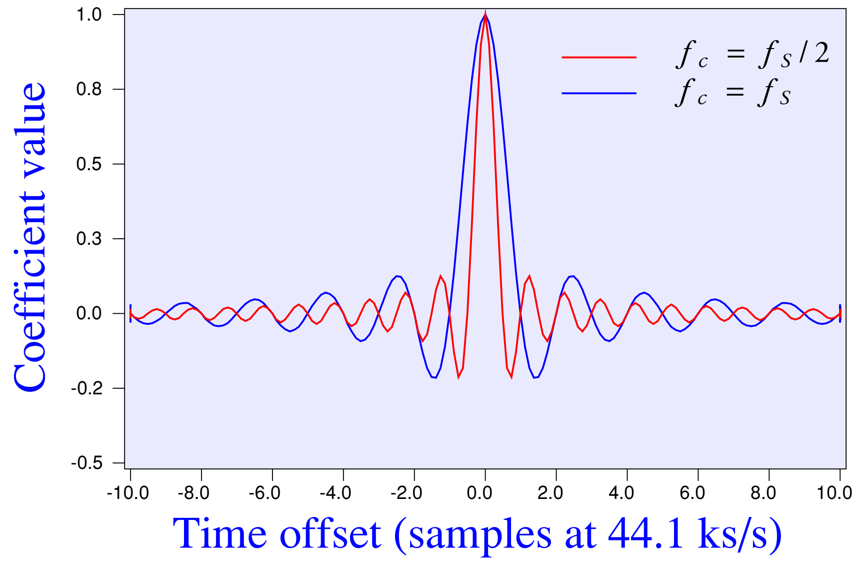

Fig 4 Impulse Responses used for following examples

Figure 4 shows two examples of the impulse response of oversampling low-pass filters. These patterns are essentially the same as the coefficient values we would employ in a TDF to obtain the desired filtering. In each case f_c represents the nominal cut-off frequency and the filter is therefore designed to pass any waveform components up to f_c, and to suppress components above this frequency. The blue line corresponds to a simple 'sinc' pattern filter whose low-pass bandwidth corresponds to half the sampling rate. The red line corresponds to a filter of similar shape, but with a nominal passband that extends up to the sampling rate.

The established practice for CD-A replay is that width and shape of the impulse response are determined by the requirements that, ideally:

- Within the required audio band up to 22kHz the relative amplitudes of the frequency components in the reconstructed waveforms should be the same as in the sampled original.
- Within the required audio band up to 22kHz the relative phases of the frequency components in the reconstructed waveforms should be the same as in the sampled original.
- There should be no output at frequencies above 22.05 kHz.

On this basis the example represented by the blue line in Figure 4 would be acceptable in normal practice, but the example shown in red would be regarded as having an excessive bandwidth.

The phase requirement given above has tended to prompt manufacturers to chose filters with a time symmetric impulse response. This avoids in-band phase distortion and should preserve the waveform shape.Impluse responses that approximate to a sinc function tend to give a reasonable trade-off between the in-band flatness of the frequency response and the level of ultrasonic image suppression.

Upsampling also employs a similar approach to generate 'extra' samples and hence present sample values to the DAC at a higher rate than was present in the original data. However the upsampled system does not necessarily place priority on satisfying all the above criteria. Instead, concern is directed towards the time-domain behaviour of the over/up-sampling action of the system. In particular, the designs may have a 'narrow' impulse response which implies a low-pass bandwidth whose cut-off does not satisfy the inequality, $f_c \leqslant f_S / 2$. We can therefore use the 'blue' impulse response as an example of an oversampling filter's impulse response, and the 'red' one as an example of an oversampling filter.

Detailed explanations of this argument can be found in the articles published by Mike Storey of *dCS*, and Keith Howard in *Hi Fi News*. From the work that has been published on such upsampling filters it may be seen that they have the effect of reducing the effective time-duration of the response to a transient event such as an impulse. It is then argued that this leads to improved perception of stereo audio information by humans. Various reviewers and others have tried listening to such upsampling DAC systems and concluded that they do sound different – and arguably 'better' – than the traditional oversampled DAC systems. The implication is then that this is due to the improvement in 'time resolution'. However, might any audible differences be due to something else?...

Nonlinearity

The audio consultant and writer, Richard Black, recently published an article on some of the possible problems which can arise due to imperfections in anti-aliasing filters. In particular, he examined intermodulation distortion due to the nonlinear behaviour of speaker units such as the HF units found in most domestic dynamic loudspeakers. His conclusions are mainly about the filters employed during the recording process. Doug Rife has also published a letter which suggests that imperfections of the reconstruction DAC may play a role in causing upsampling to lead to audible changes. He proposes that the image components may lead to a reduction in the level of such distortions, and that this may explain the perceived 'improvements'. However given the deterministic relationship between the image and intended components I have doubts that this would act as a satisfactory substitute for 'dither' as his letter suggests.

To explore this area, I wrote a set of computer models which took an input waveform (at CD-A rate) and both oversampled it by ×8 and upsampled it by ×8. The result is two representations of the input waveform. Each of these presents the DAC with nominal samples at a rate of 44,100 × 8 = 352,000 samples/second. The oversampled version employed a transverse filter which had a sinc pattern and a width intended to match a 22.05kHz output bandwidth. The upsampled version also employed a sinc pattern, but now narrowed to permit an out bandwidth of 2 × 22.05kHz. The filters chosen have impulse responses indicated by the red and blue lines in figure 4. I allowed the range of samples 'in scope' for any computation to extend for ±50 samples around the instant for which any output 'filtered' value was being computed. Since the interest in this article is on possible levels of intermodulation or anharmonic distortion I also chose to use as a test signal a waveform that contained two frequencies. This allows a fairer comparison of the levels of intermodulation distortion which may occur.

In conventional terms, the oversampled version should give good suppression of unwanted frequency images above 22.05kHz, but the upsampled version would permit images up to 44.1 kHz. The upsampled filter would produce an output such that any transients would have their energy dispersed over about half the range of times produced by the oversampled filter. Hence the upsampled version should have less 'time smear' but allow images. The output samples from these two filtering processes could then be passed though a nonlinear function that could apply a controlled amount of cubic distortion. This is designed to be a very simple and basic model of the nonlinearity of a loudspeaker unit (or of some types of amplifier such as some low-feedback valve designs which remain popular with some hi-fi enthusiasts).

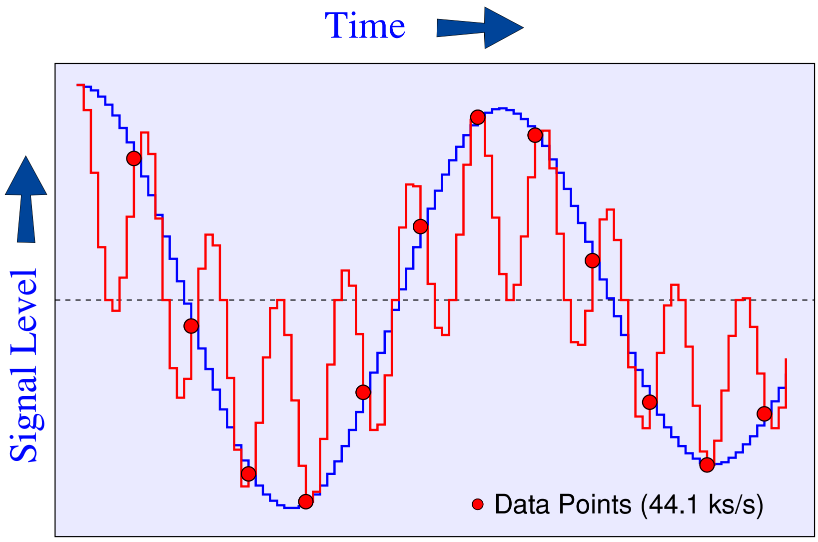

Fig 5 Examples of oversampled and upsampled waveforms

Figure 5 illustrates a short snippet from the two ×8 versions of the output created from the chosen test input series of CD-A rate sample values. The blue waveform represents what we obtain by ×8 oversampling with a sinc filter whose impulse response is as shown by the blue line in Figure 4. The red waveform represents what we obtain by ×8 oversampling using a sinc filter whose impulse response is shown by the red line in Figure 4. The red circles indicate the sample values at CD-A rate from which the over/upsampled waveforms are generated.

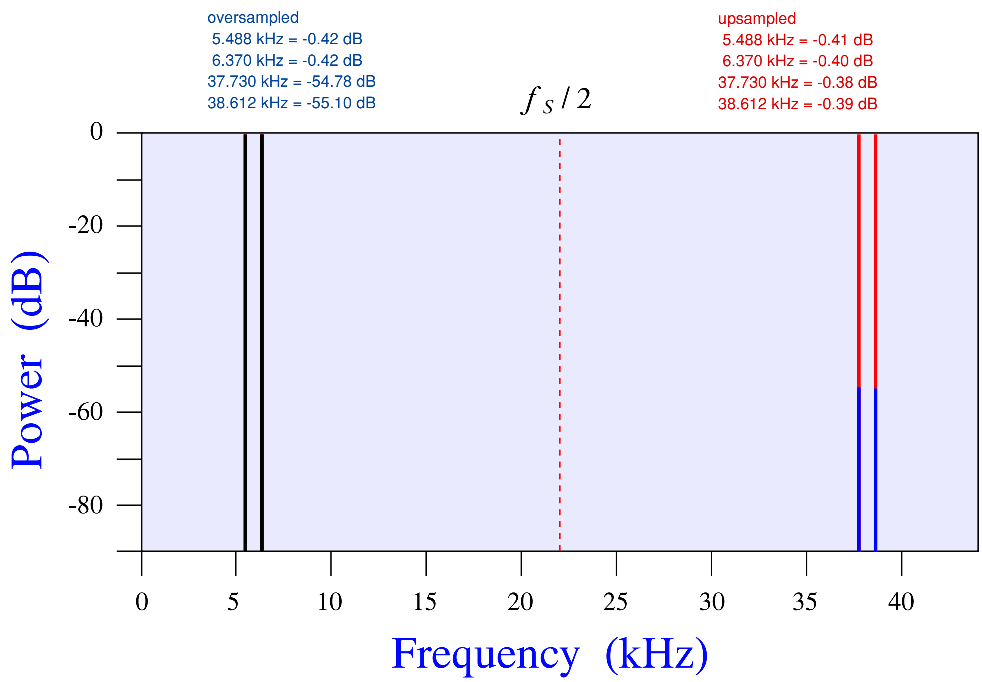

Fig 6 Spectra of two-tone test signals

Looking at Figure 5 we can see that both output waveforms give the same result at the initially sampled instants. However whereas the blue (oversampled) waveform does this with minimal overall error, the red (upsampled) version introduces unintended oscillations at frequencies above $f_S/2$. This effect also shows up clearly in the spectra of the reconstructed signal patterns as shown in Figure 6.

For the sake of Figures 5 and 6 (and some of what follows) I have arbitrarily chosen frequencies in the 5 - 7 kHz region. In Figure 6 – and in following spectra – the desired in-band components are represented by black bars. Unintended oversampling components by blue bars, and unintended upsampling components by red bars.

When the original sampled waveform contained components at two frequencies, f_a and f_b, we can expect the first image frequencies to be at $f_S - f_a$ and $f_S - f_a$. In the example chosen here we have $f_a = 5{\cdot}488$ kHz and $f_b = 6{\cdot}370$ kHz. This means we can expect the first images to occur (if present) at 38·612 and 37·730 kHz. The spectrum displayed in Figure 6 covers a wider range of powers than previous figures. As a result we can see that our chosen ×8 oversampling filter does actually produce a small amount of power at these frequencies. The reason for this is that the design is not ideal, and covers a finite range of times. This means it will 'leak' slightly and cannot perfectly suppress these unwanted images. In this case the resulting levels are about 55dB below the intended in-band signals which implies the ×8 oversampling filter produces a level of anharmonic imaging distortion (albeit well above 20 kHz) of the order of 0·3%. In a practical design this level would be further reduced by the use of a following analogue low pass filter.

The ×8 upsampling filter chosen for our examples has a cutoff frequency above these image frequencies. Hence the image component power levels are almost the same as those of the intended signals. However since these components are well above 20kHz it seems reasonable to decide that they are essentially inaudible. The question therefore becomes, "Does this matter?"...

To see why this may well matter we can now consider the effect of passing the reconstructed (filtered) signals through a stage that exhibits some static non-linear distortion. For the sake of example I have chosen a form of nonlinearity that can be represented by an expression of the form

$$v_o\{t\} = v_i\{t\} - \beta v_i^3\{t\}$$

where β is a positive value of reasonably small magnitude. i.e. we have introduced a small amount of compressive distortion of cubic form. This is a symmetric form of distortion. In reality, the distortion properties of any power amplifiers or loudspeakers that follow the DAC/filter will be quite complex. However this simple form of distortion will serve to illustrate the main points of interest in this article. It is chosen to be similar in form to that described by Richard Black.

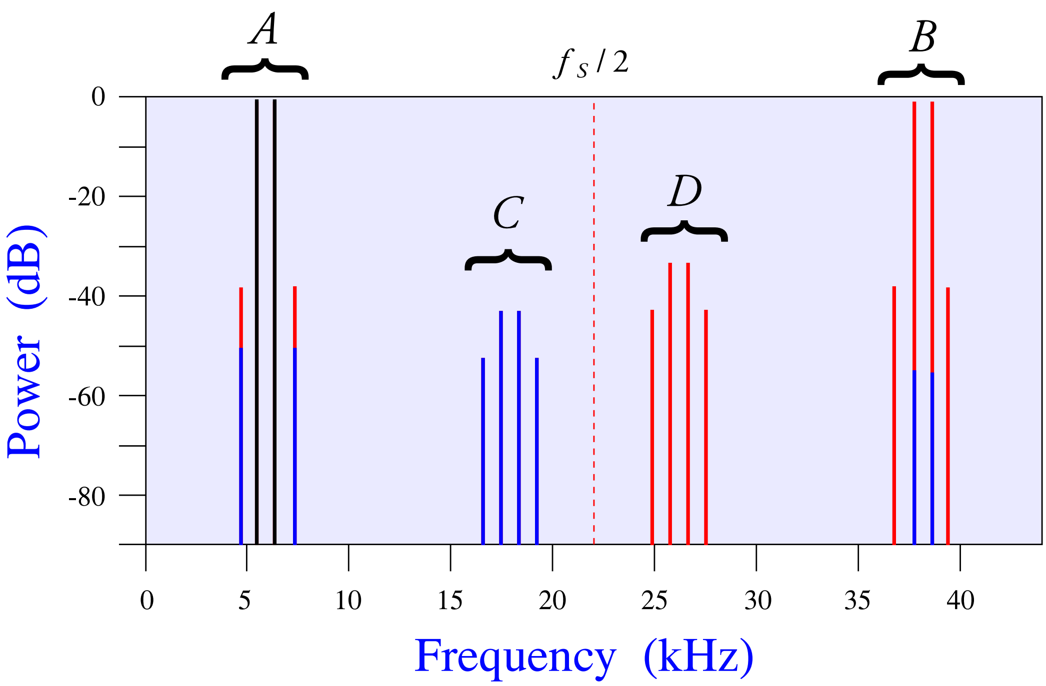

Fig 7 Spectra of distorted two-tone test signals

An example of the effect of this distortion is illustrated in Figure 7. For the sake of explanation I have identified the spectrum components in four labelled groupings.

- Group 'A' shows the intended signal components, f_a and f_b, (black) along with the third order distortion components at $2f_a - f_b$ and $2f_b - f_a$. These show that the upsampled (red) levels of these distortion components are around 10dB higher than for the oversampled (blue) components. Thus we can see that the upsampling may cause an increase in the level of the Intermodulation Distortion (IMD) at these frequencies when the signal passes though a nonlinear stage after the DAC/filter.
- Group 'B' represents the components at the image frequencies of group 'A'. These show that the upsampled filter produces these at a level around 50dB higher than the oversampled filter.
- Group 'C' represents the distortion components at $3f_a$, $3f_b$, $2f_a + f_b$, and $2f_b + f_a$. The levels of these seem essentially unaffected by the choice of oversampling versus upsampling filters.
- Group 'D' represents the components at the image frequencies of group 'C'. These components are effectively absent for the oversampling filter, but present at a level comparable to group 'C' when using the upsampling filter.

As a result of the above we can see that – when we compare the results of allowing nonlinear distortion to be applied after the up/oversampling DAC/filter – the result is that the upsampled version may lead to a higher level of overall intermodulation distortion than occurs with the oversampled version. The audibility of any components above 22 kHz is debatable. However we can now plot the results of varying the degree of nonlinearity encountered by the output from the DAC filter. and compare the two cases.

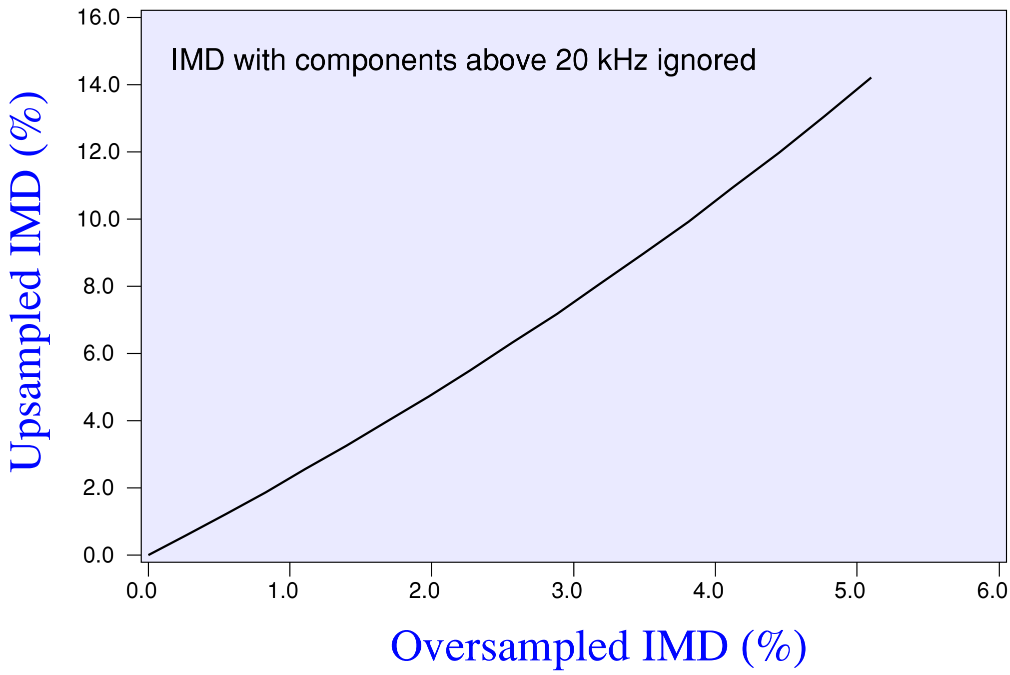

Fig 8 IMD levels produced by nonlinearity

Figure 8 shows a plot which compares the levels of IMD for the upsampled and oversampled filters as we vary the amount of subsequent nonlinearity. This plot shows that – even if we omit any components above 20 kHz from consideration on the assumption that they would make no audible contribution – that the upsampled filter tends to lead to a level of IMD that is typically a factor of two or three times higher than is the case when using the oversampled filter.

If we had considered high orders nonlinearities than cubic and/or a recorded waveform that contained components at frequencies above 5 - 7 kHz then the above picture might change slightly. In the example chosen above the Group ‘C’ components are at frequencies we can expect to be within the audible range, but those on Group ‘D’ are not. Higher values for f_a or f_b, and/or for the distortion order might mean that Group ‘D’ components would then have frequencies below 20 kHz. This may alter the perception of the presence of the distortion components as they would tend to have a distinctly anharmonic relationship with the intended waveform components. Under such circumstances even quite low levels of distortion products might lead to a perceptible alteration in the sound quality.

These results are interesting as they imply that the combination of an upsampling filter and some following nonlinearity may lead to a higher level of distortion than would be the case when using a conventional oversampling filter. The details of the distortion spectrum within the nominally audible frequency range may also be different for the two cases.

In practice, the details of a given commercial upsampling DAC/filter, and any subsequent nonlinearities, will usually differ from those assumed here. The test waveforms employed here are also much simpler than many of the musical and speech waveforms which will be encountered in practice. For all these reasons, the specific results shown in the Figures, etc, should not be taken as describing in precise detail the behaviour of any individual commercial system.

However the examples do indicate that it is possible that – in part at least – perceived differences in the audible results might perhaps be due in some cases to changes in the level and nature of the resulting IMD spectrum when reproducing signals with a number of spectra components. Indeed, this may also go some way to explaining why some observers perceive differences whereas others do not. It may be that in some cases the users are employing speakers or amplifiers that have an unusually low (or high!) level of nonlinearity, or where there are distinct differences in the response to the oversampled image components. Hence unless this possibility can be ruled out it remains unclear if any perceived differences are due to the ‘time smear’ or ‘pre ringing’ effects which are often claimed to be the cause.

Jim Lesurf
20, Priestden Place,
St Andrews
Fife KY16 8DW
jcgl@st-and.demon.co.uk

Example References

Anti-alias filters: the invisible distortion mechanism in digital audio?
Richard Black *http://www.museus.co.uk/aespaper.htm*

A suggested explanation for (some of) the audible differences between high sample rate and conventional sample rate audio material.
Mike Storey *http://www.dcsltd.co.uk/Papers.htm*

Theory of upsampled digital audio
Doug Rife *http://www.mlssa.com*

Upsampling Upheaval
Keith Howard *Hi Fi News march & april 2002*

17th Feb 2006
approx 3300 words