Upsampling Upgrades Audio?

The terms ‘Oversampling’ and ‘Upsampling’ tend to be used quite often when people describe how devices like Audio CD players or DACs function. Since the first releases of Audio CD players back in the 1980s some models have said to be ‘better’ than others because they used these approaches – or not! However what are ‘Oversampling’ and ‘Upsampling’, and are there any practical differences between them that can affect the audible results?

Answering these sorts of questions in complete detail is a bit of challenge as the full answer requires delving into the details of complex technical mathematical areas like Information Theory, Digital Signal Processing, etc. So although I’ll try here to give an explanation I’ll try to pitch what I say to dodge some of the really hard sums.

Sampling and reconstruction

The well-established oversampling approach to consumer/home digital audio began with the first generation of Philips Audio CD players. Its still used by most domestic players, and in many other digital audio systems. Like much of conventional digital audio it is based firmly on the concept of what Information Theory calls a ‘Complete Record’. Quite simply this says that if we start with an analogue signal waveform then we can take a series of samples uniformly spaced in time that provide a recording of *all* of the details of the information carried in the original waveform. This can then be used to ‘reconstruct’ a duplicate waveform that carries exactly the same information. i.e. a ‘perfect’ reproduction of the original in terms of the details of the real original!

At first glance this idea tends to seem absurd to a practical engineer. This is because real-world engineered systems aren’t expected to be mathematically perfect, just “Good enough for the task”. However the above can be mathematically proved to be absolutely true in formal terms. The theoretical maths is water-tight. So where’s the catch?... we’ll come to that later on. But first, let’s see what it means in terms of digital audio.

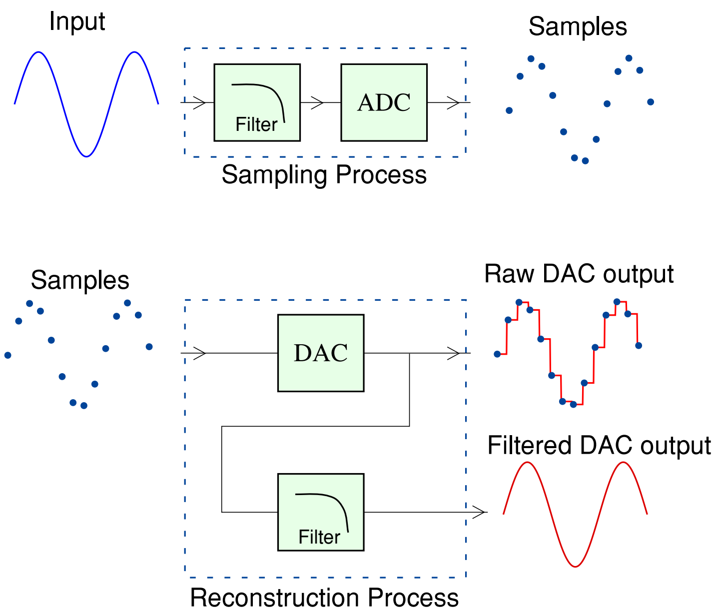

Figure 1 Sampling and reconstruction.

Figure 1 above shows a schematic diagram of an analogue signal waveform being converted into samples by an ADC (Analogue to Digital Convertor) and then this stream of digital values is used to reconstruct a fresh copy of the analogue waveform using a DAC (Digital to Analogue Convertor). Note the way I've added in a couple of (analogue) low-pass filters. The samples are taken at regular time intervals. e.g. For audio CDs the samples are taken at the rate of 44,100 samples per second. This is the value per channel. Since Audio CD is two-channel stereo this means 44,100 left-channel samples per second and 44,100 right-channel samples per second. However for the sake of clarity from now on I'll just consider one channel and we can assume that a similar process will be duplicated for the other.

The filter before the ADC is present to ensure that no frequencies above half the sample rate can reach the ADC. This is strictly required to obtain a Complete Record. When we come to reconstruct the required audio waveform from the series of sampled values we actually have to do two things. Firstly we need a DAC to obtain the original waveform levels at the sampled instants. Secondly, we then need to smoothly fill in between these values and re-create the correct original shape of the waveform everywhere *in between* these sampled instants. This requires what is called a 'reconstruction filter'. Early DACs used by engineers often used a 'sample and hold' process which produced a kind of 'staircased' version of the required waveform. The jagged steps of this staircase are then smoothed away using an analogue low-pass filter. This is the approach I've used to illustrate the process in Figure 1. The Sampling theorem tells us that if we use 'ideal' filters this process can reconstruct an exact duplicate of the original waveform.

However when it comes to making a real system work the devil is in the details. It's one thing for me to draw a nice box on a diagram called 'filter'. It is something else to make real filters that meet the stringent demands of the Sampling Theorem!

In theory both of the filters in Figure 1 have to achieve a set or requirements. If r is the sample rate then they must pass though all frequencies below $r/2$ with exactly the same gain. Any time delay the filters impose on the signals must also be the same at all frequencies below $r/2$. i.e. The filters must have a perfectly 'flat' response below half the sample rate. i.e. for all frequencies below 22.05 kHz for a filter used for Audio CD sampling and reconstruction. The filters also have to *totally* reject anything at frequencies at or above 22·05 kHz. Alas, this is a very demanding set of requirements.

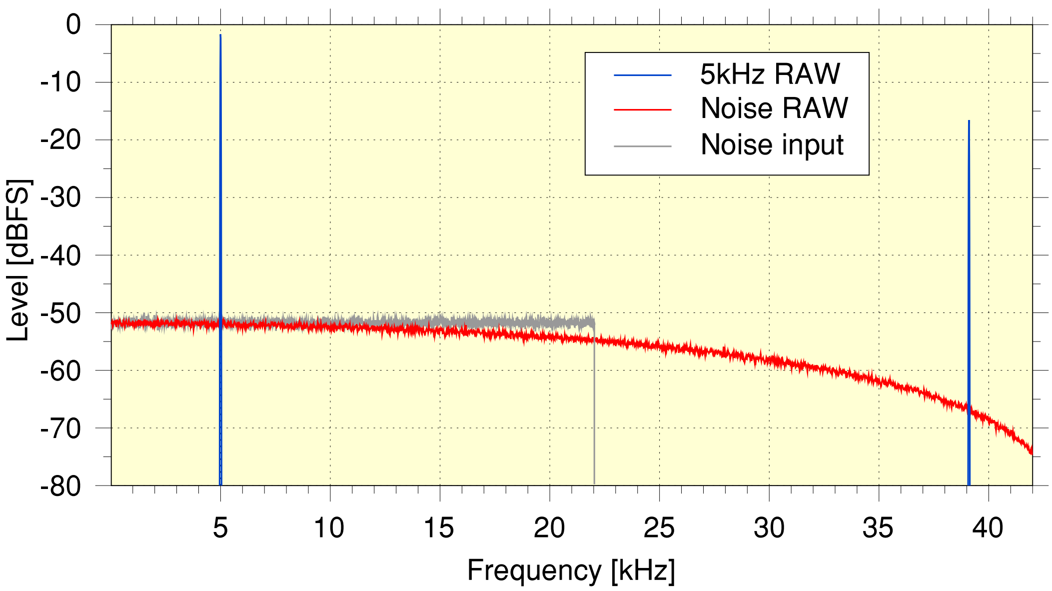

Figure 2 Examples of 'Raw' sample and hold output

Figure 2 illustrates the challenge facing the reconstruction filtering. Here I've used two examples as input waveforms/sample sequences. The first is a 44·1k sample-rate 5 kHz sinewave, shown by the blue lines. The 'staircase' output from the DAC circuit provides this, but adds an unwanted distortion component at 44·1 - 5 = 39·1 kHz that is only slightly lower in level. More generally, audio waveforms will contain a range of frequency components so the second example I've used is some sampled random noise with an (approximately) flat spectrum. This signal should have a spectrum that is as shown by the grey line. i.e. Note that because it was sampled at 44·1k samples/sec the filter used in front of the ADC must have removed *any* frequency components at or above 22·05kHz. However when we look at the spectrum of the raw staircased output from the DAC we get the result shown in red. The staircase distortion has added a great deal of unwanted noise above 22kHz! This is almost as high in level as the wanted waveform. So any reconstruction filter has to be very well designed and made if we wish to pass everything up to just under 22kHz whilst essentially eliminating this unwanted output at higher frequencies.

In theory we can build analogue filters that come quite close to being perfect. But they would tend to be very expensive, and their performance may deteriorate as the components in the circuitry age, etc. Fortunately, from the start Philips realised there was an alternative approach which made manufacturing high quality reconstruction filters more practical. Oversampling...

Oversampling

The original Philips Audio CD players used an $\times 4$ oversampling system. They did this for a number of reasons. One was that it makes the reconstruction filtering much easier. For the sake of clarity, here I'll use a simpler example to show how – provided we have a set of samples which provide a complete record - we can tackle reconstruction by using *digital* filtering.

The implication of a series of sample values being a complete record is that the samples don't just tell us the actual levels at the sampled instants. They also give us all we need to know to be able to compute the signal levels at *any* other instant(s) in between the samples. This doesn't mean an 'interpolation' in terms of making a good 'guess'. It means recovering the *actual* instant-by-instant variations of the input that was sampled! Just like the concept of the Complete Record this idea seems startling at first, but is, mathematically, rigorously correct.

If you wade though the relevant mathematics you find that the value at any instant in between the given samples can be calculated on the basis of the Sampling Theorem using a sinc function. Those who are familiar with Fourier Transformations will realise this is the appropriate function in time because it goes with the ideal requirement outlined above of having a filter which is flat (in amplitude and phase) across the frequency range up towards $r/2$ and zero at and above that frequency. A sinc function in time is the FT of a flat, brick wall low-pass spectrum shape.

To illustrate the process we can look at an example and I'll give the maths for those who may find it useful. We can start by calling the interval between sampled values Δt, and the time where we want to generate a new sample t_0. The required sinc function can then be specified as

$$\text{sinc}\{z\} \equiv \frac{\sin\{z\}}{z} \qquad \text{where} \qquad z = \frac{\pi(t - t_0)}{\Delta t}$$

where when we are given a series of samples the value of the time t will change by Δt from each sample to the next.

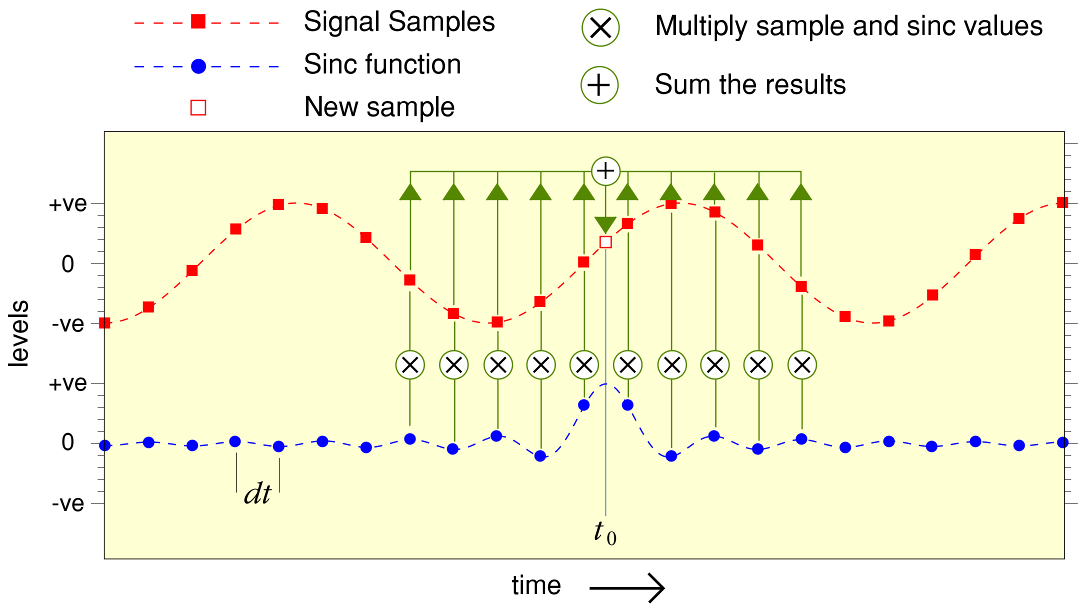

Figure 3 Digital, Time Domain oversampling

Figure 3 illustrates this process. In this case I've chosen the example of calculating the value at an instant mid-way between two particular given sample values. But the same general approach can be used for other instants.

In effect we align the central ($z = 0$) peak of the sinc function with the time where we want to add a new sample. Then multiply each nearby waveform sample value by the sinc function's value at the time that sample occurred. We then add together these results to obtain the value for the new sample we are adding into the sequence. Figure 3 only shows this process as including the five sample values on each side that are nearest to the chosen instant. In theory we need to extend this process to include *all* the values we have in the sampled waveform sequence to get precisely the correct value. If we do this we get the correct result provided we started with a complete record any our maths is done with perfect precision. If we represent the samples as a series of K sample values, $x[1]$ ... $x[n]$ ... $x[K]$, taken at the instants $t[n] = n \times \Delta t$ then the full expression for calculating a 'new' sample value, $x\{t_0\}$ at the time t_0 will be

$$x\{t_0\} = \sum_{n=0}^{K} x[n] \ \mathrm{sinc}\left\{\frac{\pi(t[n] - t_0)}{\Delta t}\right\}$$

This new value was named an 'oversample' because it isn't needed by the Sampling Theorem to give us a complete record of the original waveform. If we were to now use the same approach to compute a new sample mid-way in between each existing adjacent pair of samples in the original sequence we would be generating an 'oversampled' output. Taking these samples along with the originals we then have a series that are closer together in time. If the original sequence of samples were at a sample rate of 44100 samples/sec then the result is the series of values we'd have obtained if we sampled the original waveform at 88200 samples/sec. The original set has been converted into an '× 2 oversampled' 88200 samples/sec version. (The word 'Upsampled' can also quite logically be used to describe this process as we have generated an equivalent output at a higher sample rate. However the term is often used to indicate a slightly different approach which I'll discuss later.)

Alas, in practice this ideal process may require a great deal of computation. For example, a 5-minute audio track at the Audio CD sample rate will contain over 13 million samples per channel. To get the correct result we'd ideally need to use *all* of these to work out *each* additional sample. Requiring over 100,000,000,000,000 calculations!

Fortunately the amplitude of the sinc function tends to fall away toward zero as we move away from its central peak. As a consequence we can, in practice, omit including sample points that are far away from t_0. Figure 3 shows the process using the 10 closest original samples to work out each new oversample. How many we choose is a trade-off between the accuracy of the results and the amount of computational work required. It becomes a matter of judgement and skill to get the most acceptable results. But the general rule is that: the more sample points we have 'in scope' either side of t_0 then the better the results.

One consequence of limiting the scope of the samples used for the calculation is that the frequency response of the results are altered. To deal with this it is usual to modify or 'Window' the function employed so that it isn't exactly a sinc function. To illustrate this we can look at an example.

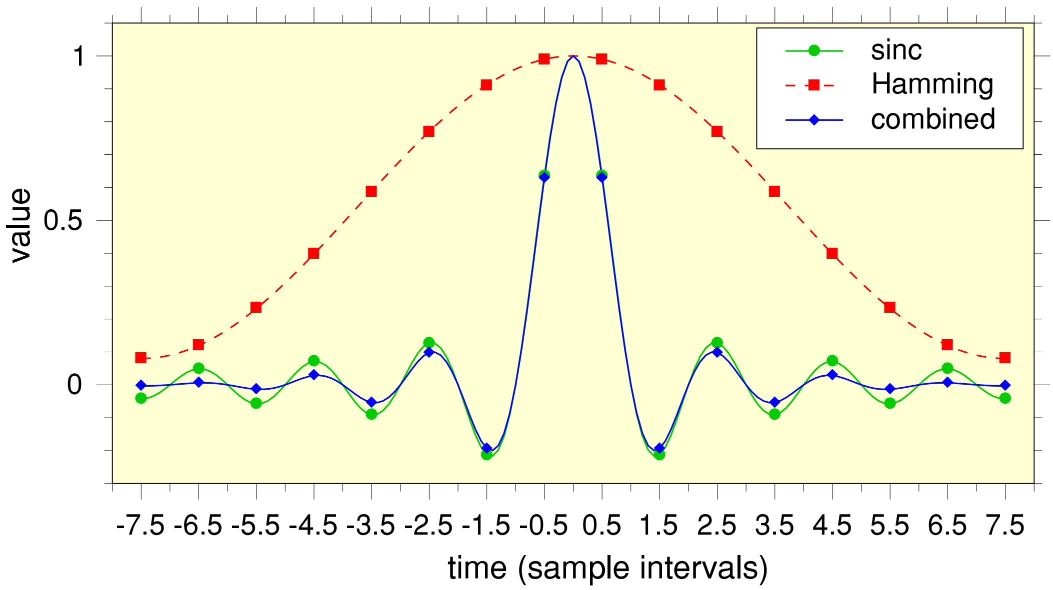

Figure 4 Reconstruction filter functions

Figure 4 shows an example for computing a new sample value midway between two existing samples by using 8 original samples on either side. The original samples occur at instants that are ± 0·5, 1·5, 2·5, etc sample intervals from the instant where the new sample is to be produced.

The green line shows the sinc function and the points indicate its value at the instants of the original samples. The broken red line shows a Hamming function which is often used for 'weighting' such a computation to try and get improved results. The green line is the Hamming-weighted sinc function. You can see that the effect is to reduce the values as we approach the sides of the 16-sample-wide window. Sample values further away are ignored when doing this calculation.

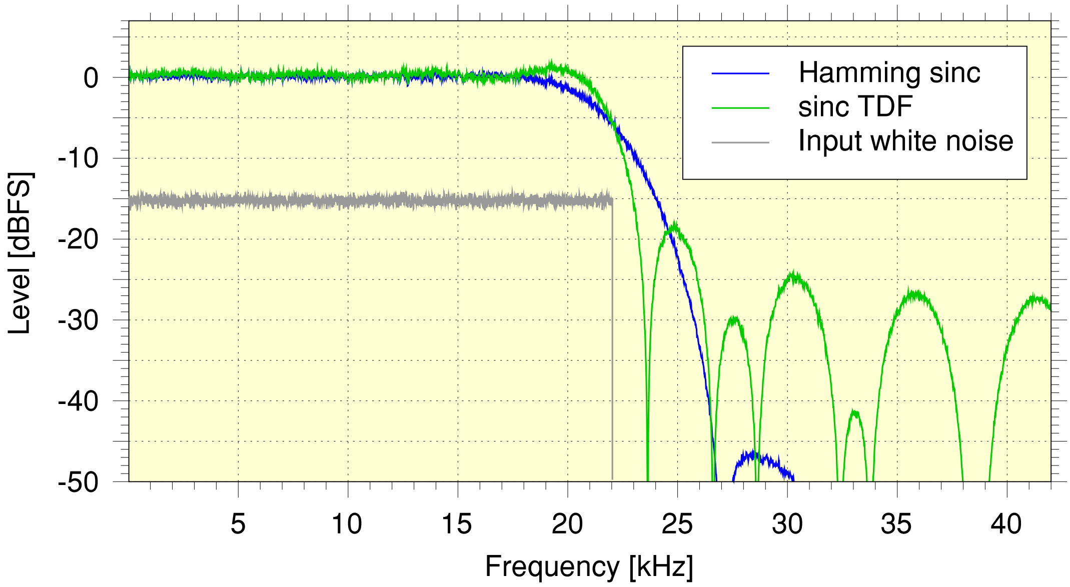

Figure 5 Spectra of the digital filtering results

Figure 5 shows the results of using the above method to × 2 oversample an input white noise 44·1k rate sequence of samples to an 88·2k rate output. The grey line shows the spectrum of the input white noise used for the test. If we simply converted this with a ‘sample and hold’ DAC of the kind described in Figure 1 the result would be a fairly high amount of unwanted added components above 22kHz, as illustrated in Figure 2. A very good analogue filter would then be needed to reduce these to a low level.

The Hamming weighted × 2 oversampling process using just a window 16 samples wide also produces some unwanted products above 22kHz. But these are at a much lower level than we’d get from the raw ‘staircase’ DAC. This means a following analogue filter can now be a simpler design because much of the work is already done. In practice the original Philips CD players used × 4 oversampling with a window using many more samples, so gave an even flatter response below 22kHz and rejected the unwanted ultrasonic rubbish more efficiently than our simple example. This made it relatively easy to provide a following analogue filter that achieved excellent results. More modern DACs often use even higher oversampling ratios and give better results, much closer to the theoretical ideal.

The important point to take from this is that the oversampling process can, itself, act as a form of digital filter. In this case it forms an important part of the reconstruction filter when we want an analogue output from a series of audio samples. In most CD players the main drive when designing and making such digital filters was to obtain performance as close to the ideal Sampling Theorem sinc function using as many input samples as possible. However this ideal was eventually challenged by an alternative approach. I’ll say more about that and the practical details in another article...

2500 Words
Jim Lesurf
25th Nov 2013