Sampling

Although sampling is most widely known for its importance to audio convertors, it shows up in a lot of other places too. JOHN WATKINSON explores.

IN THE LOOSEST sense, sampling is any process in which discrete measurements are taken instead of a continuous process. In audio, we are accustomed to samples being taken with respect to the time axis at a fixed frequency, and we become very nervous, even jittery, if the time between successive samples is not exactly the same. It is important to realise that not only can samples be taken with respect to other axes, such as distance, but there is also no fundamental need for the time or space between them to be constant, although it often is for simplicity.

In television, images are sampled into lines and in digital television the lines are sampled into picture elements or pixels. The sampling is spatial and the sampling rate is measured in samples per mm or samples per picture height. The temporal sampling rate is obtained by multiplying the spatial sampling frequency by the scanning velocity.

The first description of sampling theory is usually attributed to Shannon although, in the former Soviet Union, the same findings were independently made by Kotelnikov. Interestingly, they were both pre-dated by Whittaker, who worked in a shipyard. It is well known that the shape of a ship's hull is described by a set of lines. These are a set of cross sections that you would get by putting a ship through a giant bread slicer. Each cross section is a sample and what Whittaker wanted to know was the minimum number of samples that would accurately describe the shape. Clearly an aliased ship would be a source of embarrassment.

Actually shipbuilding offers a perfect analogy for audio reconstruction. The ships lines are infinitely narrow, just as audio samples represent a vanishingly small instant. The finished ship is shaped by creating a surface that smoothly joins together all of the lines. This is analogous to the audio reconstruction filter that joins the tops of the samples to obtain the waveform. We never see zero-order-hold ships looking as if they were made out of Lego, but we see plenty of incorrect digital audio explanations showing these quite erroneous staircases. It's a sad reflection on the audio industry that they know more about sampling in shipyards.

The sampling rate needed must be at least twice the bandwidth of the signal to be sampled. In the case of a signal going from zero Hz to some maximum, we would use low-pass filters before sampling and for reconstruction, whereas if the signal is band-limited, we would use bandpass filters. We could sample a signal whose spectrum lay between 10 and 11kHz using a sampling rate of only 2kHz.

In audio compression, we frequently split the audio spectrum into sub bands. MPEG uses 32 sub bands. After band splitting, the overall sampling rate has not risen at all. Each band is expressed using 1/32 of the original sampling rate. The bands can be recombined using a set of bandpass filters.

If we ignore these filtering requirements, the result will be aliasing. This is a word that doesn't translate; it appears to be limited to English. Aliasing causes the spectrum of the original signal incorrectly to be represented. In digital audio, aliasing is usually a bad thing. In other applications, it's very useful.

The mathematics of sampling are essentially the same as those of signal multiplication, which radio engineers call mixing. This should not be confused with audio mixing which is, of course, additive. When two signals are multiplied, the resulting spectrum contains new frequencies representing the sums and the differences of the original frequencies. It is the difference frequencies, or lower sidebands, that have the most interesting effects. In a system having a fixed sampling rate, if the input frequency goes up, the frequency of the lower sideband goes down.

In the sampling oscilloscope (not to be confused with the digital storage oscilloscope) an input signal of some appallingly high frequency is sampled at another appallingly high frequency so that the lower sideband, or difference, frequency is low enough to display on the tube.

In analogue tape recorders, the contact between the tape and the heads is imperfect and subject to rapid variations. This has the effect of amplitude modulating the audio waveform. The result is that any frequency in the original audio is accompanied by sidebands due to the modulation. Despite the fact that the unwanted signal added to the sound is audio waveform dependant, the term used is modulation noise.

In digital audio we divide the frequency range from zero up to the sampling rate in half. The input spectrum lives in the bottom half, whereas in the top half the same spectrum, having been reversed as if in a mirror, comes down from the sampling rate. Normally these two halves are kept separate by low-pass filters, but in the absence of such filters, if the input frequency rose beyond half the sampling rate, the frequency of the lower sideband would fall below it and we would have aliasing. If the input frequency continued to rise until it reached the sampling rate, the lower sideband frequency would come down to zero.

The stroboscope is a positive application of aliasing. The flashes of light are very short and sample the scene. By flashing a light at the same speed as a periodic mechanism, the motion can be arrested. Essentially the motion is being sampled at its own frequency, and the lower sideband is at zero Hz. Up-market vinyl disc turntables often had stroboscopic markings on the rim. Using a neon light connected to the household AC supply, the speed could be checked. Later turntables had quartz-crystal-locked motors and only retained the stroboscopic markings out of tradition. However, these were useful for measuring the stability of the AC supply. The strobe lights gradually became more powerful and were instead used to illuminate Sharon and Tracy dancing around their handbags.

A slight reduction of the flashing rate of a strobe causes the lower sideband to have a negative frequency, so the motion appears to reverse. This reversed motion phenomenon is often seen on films where the images contain spoked wheels. The spoke passing frequency is simply too high for the inadequate frame rates that are used.

In spectrum analysis, we often wish to know how much energy exists in a large number of different frequency bands. The approach of making one band-pass filter for each band is too expensive. Instead we can use a single fixed filter, typically a low-pass. We take the unknown signal to be analysed and multiply it by another known frequency and see if the result will pass through our single filter. If the input signal is close to the known frequency, the lower sideband will have a low frequency and can thus pass through out LPF. If the process is repeated with a set of known frequencies, the spectrum can be obtained.

A simple extension of this process is where we take a set of audio samples and multiply them by a set of samples representing a sine wave of known frequency, which is called a basis function. If the input frequency is the same as that of the basis function, the samples will be taken at the same point on each cycle and the sum of all the samples will be non-zero. For all other frequencies this will not happen and the sum of the samples will be zero. The only difficulty is if by chance we sample on the zero-crossings of the input signal and also get zero. To overcome this, we can use two basis functions, one a sine wave and the other a cosine wave. Thus if one basis function falls on the zero crossings, the other will fall on the peaks. If we call the sum of the samples in each case a coefficient, the ratio of the coefficients will reveal the phase of the input signal.

This sounds fairly simple, but it is actually how a discrete Fourier transform (DFT) works.

The other approach to the problem is to time reverse or mirror the input samples and place them directly before the original samples to obtain a block of double length. This mirroring will cause any sinusoids in the original block to be inverted in the mirrored samples and cancel out, whereas any cosinusoids will double. This is how the discrete cosine transform (DCT) works.

I can hardly raise the subject of sampling without being asked my opinion on the high audio sampling rates that are now supposed to be necessary. As no-one has shown that sampling theory is wrong and as no one has shown that the bandwidth of human hearing has suddenly increased, the case doesn't seem proven. I checked with the shipyard: 'Do you now use four times as many lines as before to build hulls?' You can imagine that the answer was in the negative. As I said earlier, shipyards know more about sampling than the audio industry. ■