

DIGITAL AUDIO, AN INTRODUCTION
Dan Lavry, Lavry Engineering, Inc.
Copyright 2007


Ideal AD followed by an ideal DA is as good as a perfect wire. The waveform into the AD is the same as the waveform out of the DA (we can allow for gain or attenuation). The fundamental goal and the proper base line is to be able to do just that, prior to sound alterations.

Is the concept of digital audio an approximation that we need to live with? Or does the concept offer a theoretically perfect solution, making the task of making it work an IMPLEMENTATION, not a conceptual problem?


While much talk to the contrary, the conversion process itself has no conceptual drawbacks.

Connecting the dots


Clearly, one can expect to take a complicated wave, pin down a few values and be able to claim much knowledge about the wave:




When taking very few sanp shots (sample), I we are left with insuficient information about the wave shape.




What can we do with such few samples? Not much. Let us try and connect them with a straight line to see how close we can get to the original wave:




Unfortunately, this is the way many non technical people in audio view digital audio. There is a wrong notion that we will get better approximation by taking "more and more dots", such as increasing the sample rate as shown below we will end up with better approximation:




Obviously we see a big difference between the original wave and the samples. Let us now change our process from RZ (return to zero) to NRZ (not return to zero). Each sample value is held until the next sample value:



The black NRZ signal (Steps) still does not look good enough to approximate the red Wave. Let us look at the difference between the Wave and the Steps signals by subtracting one from the other:




The above black difference wave looks like a very significant error, and again error is the difference between the original wave and the signal we called stepped. Intuition suggests that by taking more and more samples, the signal "Steps" signal we will get closer and closer to signal "Wave", which will of course make the signal "Diff" smaller and smaller. "Steet type common sense" sugets that digital representation (sampling) will never be good enough, but such conclusion is wrong. Here is an example when we do not need many points to describe a signal.. Let us begin with a DC signal Red):




A Obviously, adding samples (in blue) does not add any information in the DC case. A single sample is enough to describe the DC wave with zero error. While "simple" the above example has value. Lets take a look at a slow varyingt signal, sauch as 100Hz sin wave samplede at 44.1KHz:

Frequency :=




$$f = 1.58 \times 10^4$$

Points :=



The NRZ wave we call Steps looks like:



We can see that the error is always there, and the amplitude of the error grows as we increase the rate of change of the sampled signal. In other words, fast changes cause large errors. This fact is indisputable. But, if we can take the signal we call "Steps" and correct its shape by removing the error from it, we will end up with the original signal "Wave".

Steps = Wave + Error,

Removing error by subtraction yield:

Steps - Error = Wave + Error - Error = Waves.

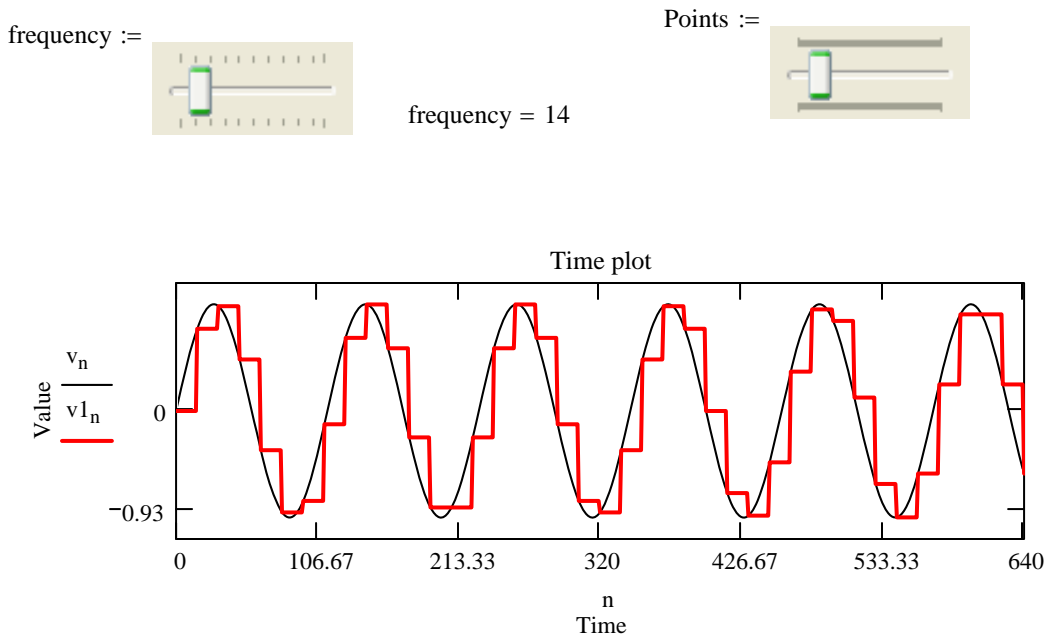
How can we "remove the steps" from the sampled wave? In other words, how can we remove the errors? We can do so if we conform to Nyquist Theorem, which says that we can do so as long as we sample faster than twice the highest frequency contained in the signal. In fact, conceptually, there is nothing to be gained by sampling faster (taking more "dots").


How is it possible? An examination of a sampled signal will reveal that all the energy content below half the sampling frequency (the Nyquist frequency) remains intact, and that all of the error is contained by frequencies ABOVE half the sampling rate (above the Nyquist frequency).

Therefore removal of all the high frequency content above Nyquist will remove the error completely, and we end up with the original wave. It may not be "street smart intuitive" but it is a fact, that the removal of the energy above Nyquist will yield the proper "connecting of the dots" with the perfectly correct curvature. We do not end up connecting dots with straight line approximations, or any other approximations. We are ending up with a perfect replica of the original.

But How can we remove high frequencies from a sampled signal? Lucky for us, electronics is highly evolved in the art of filtering which means passing and blocking frequencies. For audio, The AD needs to sample faster than the highest frequency of interest, which will generate the signal that is a sum of the original wave plus an error signal. The DA will replicate the signal we call Steps. Placing an appropriate low pass filter after the DA will reproduce the original wave. By appropriate I mean a filter that passes signals below Nyquist and blocks signals above Nyquist.


Therefore There is nothing conceptually wrong or missing from digital audio. All issues are implementation specific issues.






Let us introduce a low pass filter

Filter :=






frequency :=



frequency = 36

Points :=






Filtering to avoid aliasing, No oversampling, $f_s=44.1\text{KH}$:

High frequency energy may find its way into the audio range. Filtering is needed to avoid aliasing. There are some fundamental compromises to be made.

1. It is desirable to move the filter frequency higher, to avoid impact on the audio band
2. The higher the passband, the less anti aliasing protection
3. The SINC function (NRZ sampling) is an added consideration in terms of the upper audio pass band.

Order :=



Order = 24

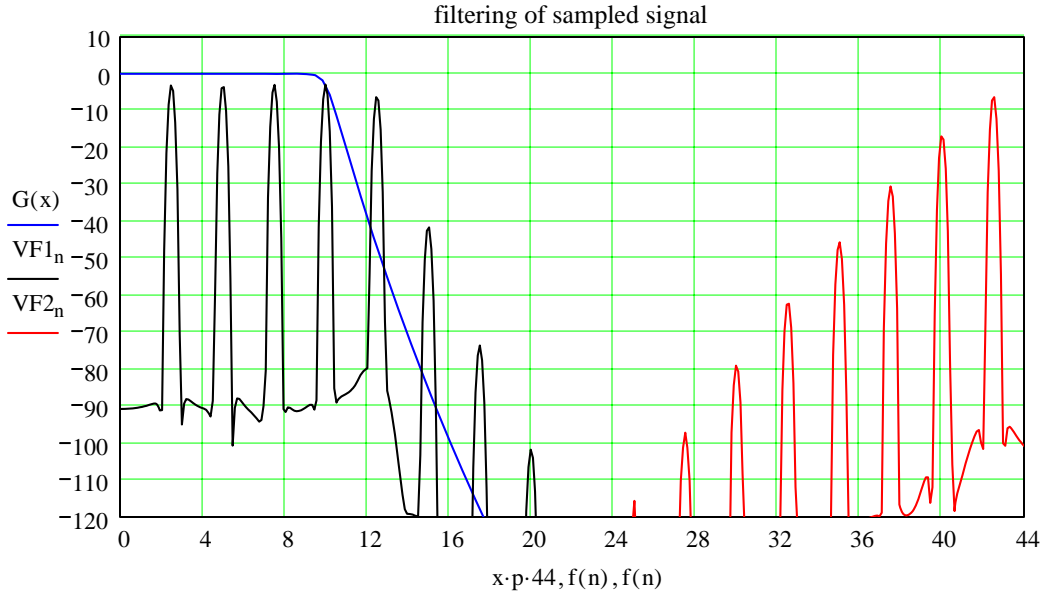
Cutoff :=



Cutoff = 10

High filter orders yield better alias rejection but increase circuit complexity and can increase noise and distortions

blue: filter curve
black: filtered audio
red: filtered aliased energy




Filtering to avoid aliasing, X2 oversampling, $f_s=88.2\text{KHz}$:

High frequency energy may find its way into the audio range. Filtering is needed to avoid aliasing. There are some fundamental compromises to be made.


1. It is desirable to move the filter frequency higher, to avoid impact on the audio band
2. The higher the passband, the less anti-aliasing protection
3. The SINC function (NRZ sampling) is an added consideration in terms of the upper audio pass band.

Order :=



Order = 23

Cutoff :=



Cutoff = 25

High filter orders yield better alias rejection but increase circuit complexity and can increase noise and distortions

blue: filter curve
black: filtered audio
red: filtered aliased energy

