

The Bits In-Between

An EE's Guide to Survival Between Microphone and Voice Coil


Bruno Putzeys

Hypex Electronics, Grimm Audio, The Netherlands

On the occasion of the 123rd AES Convention, October 6, 2007

Digital Filters in AD/DA Converters

Sampling Theory's Basic Promise



A sampler flanked by low-pass filters with sufficient attenuation at $f_s/2$ does exactly the same as the low-pass filters alone.


Digital Filters in AD/DA Converters

The TOA Cue Fallacy

- “The ear can detect a 2 μ s Time-Of-Arrival difference”
 - Correct! (0.2° lateral shift in stereo image)
- “So we need 500kHz sampling”
 - Uhhh not quite...


Digital Filters in AD/DA Converters

The Promise in Practice




Digital Filters in AD/DA Converters

The Promise in Practice




Digital Filters in AD/DA Converters

The Promise in Practice




Digital Filters in AD/DA Converters

The Promise in Practice



Digital Filters in AD/DA Converters


The Promise in Practice



Digital Filters in AD/DA Converters

The Nonoversampling Fallacy


- “Digital Square Wave” test signal looks like this:



Digital Filters in AD/DA Converters

The Nonoversampling Fallacy


- After “NOS conversion” (=zero order hold) we get:



Digital Filters in AD/DA Converters

The Nonoversampling Fallacy

- Now insert a half-sample delay:




Digital Filters in AD/DA Converters

The Nonoversampling Fallacy

- Impulse Response becomes time-variant
- Fallacy was facilitated by the “Digital Squarewave” signals from test kit and test discs.


Digital Filters in AD/DA Converters

- The “Digital Step Function” reconstructs like this:




Digital Filters in AD/DA Converters

- Contrast with an actual band-limited step function



Digital Filters in AD/DA Converters

- Compare the spectra



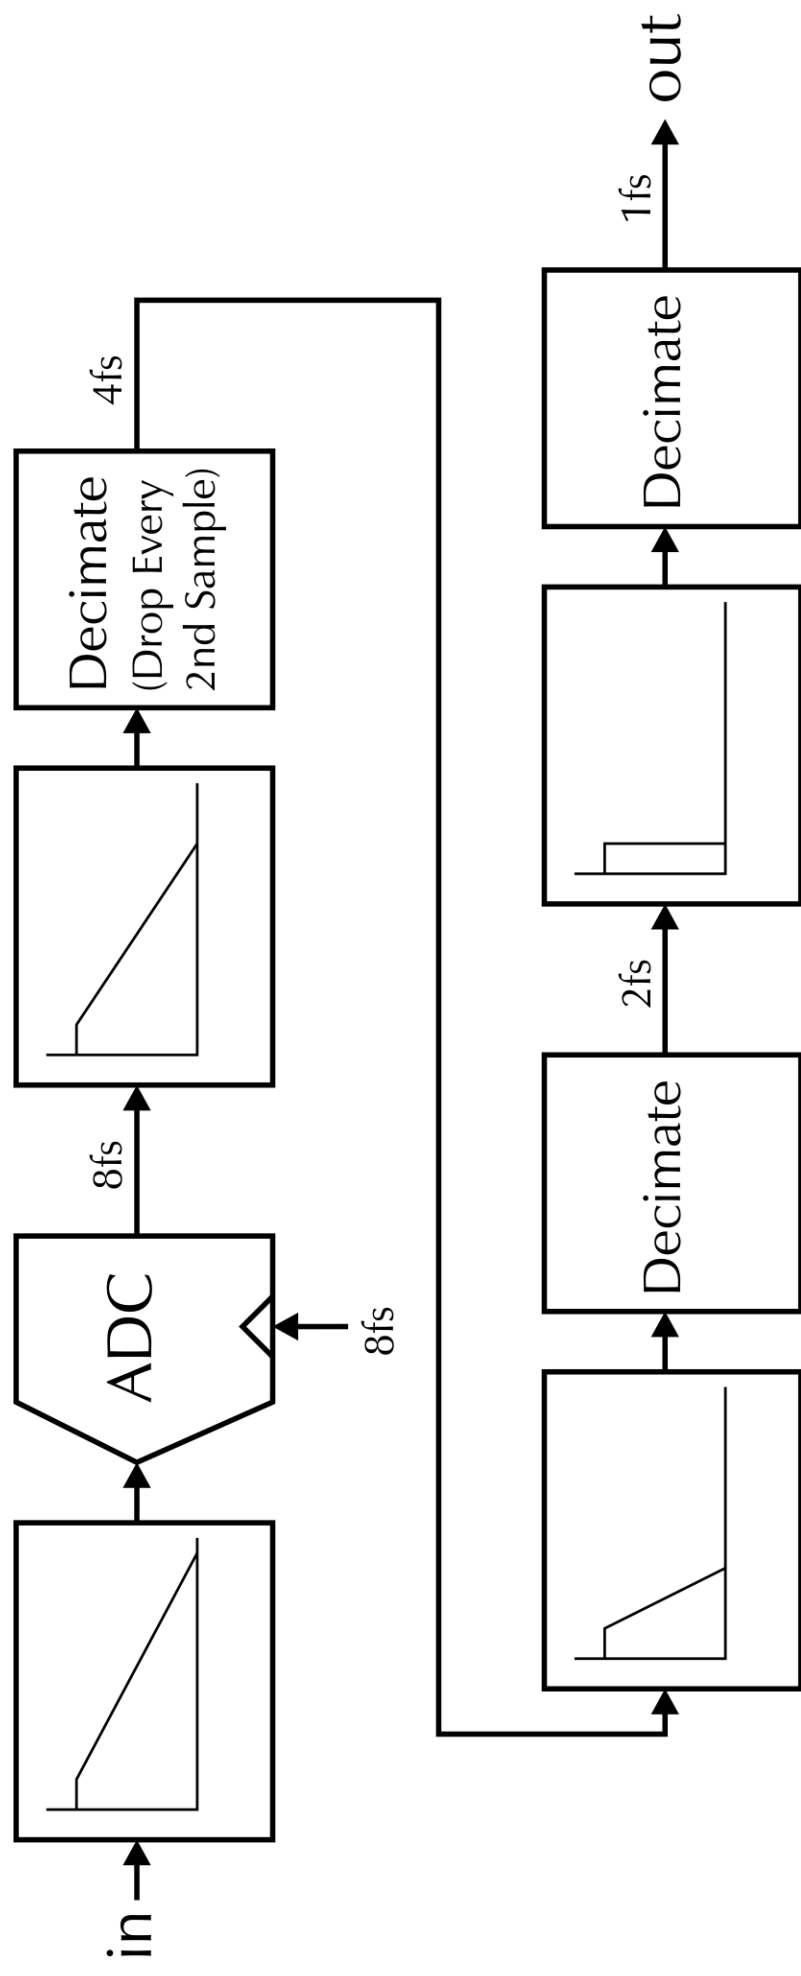
Dear test equipment designers: please provide a “true square wave” of arbitrary frequency

NOS Rundown

- NOS DAC may sound OK
 - We really don't notice much beyond 20k...
- NOS DAC sometimes sounds better than *same* DAC with digital filter
 - DAC in these experiments is invariably ladder type
 - Glitch contribution goes up with sampling rate
 - Latch signal passes through filter chip (increased clock jitter)
- None relate to impulse response


Digital Filters in AD/DA Converters

- Antialias filtering in contemporary ADC's is mostly done digitally, in a "decimation chain"




Digital Filters in AD/DA Converters

Gabarith for 8fs -> 4fs filter stage



Digital Filters in AD/DA Converters

Gabarith for 4fs -> 2fs filter stage




Digital Filters in AD/DA Converters

A perfect candidate: The Half-Band filter

- Magnitude response is chosen symmetrical round 0.25fs and 0.5.
 - Stop band = 0.5fs - pass band
 - Stop band rejection = stop band ripple


Digital Filters in AD/DA Converters

Half-Band filter, Magnitude Response




Digital Filters in AD/DA Converters

Half-Band filter, coefficients



Digital Filters: Design Compromises


Gabarith for 2fs -> 1fs filter stage.



Oops.

Digital Filters: Design Compromises

Typical final stage in commercial converters



Cut & dried breach of Nyquist criterion!

Digital Filters: Design Compromises


$$0.4535 * 44.100\text{kHz} = 20.000\text{kHz}$$

Digital Filters: Design Compromises

Result

- Band between 0.4535fs and 0.5565fs suffers aliasing.
 - Only 12dB of attenuation at $f_s/2$. Signals with significant energy near $f_s/2$ are worst affected.
-
- Next slide: demonstration: $f_s=44.1\text{k}$. Square wave of $f\approx 3150\text{Hz}$ is fed into ADC. 7th harmonic aliases.

Digital Filters: Design Compromises




To human ears:

- TOA cues are affected for signals with significant HF.
 - Sibilants in choral music, wind and string instruments smear across the whole stereo image.
- Nearly no impact for panpot stereo.
 - Alias components are in phase across channels

Digital Filters: Design Compromises

How to Salvage a Burnt Steak

- Cut off the blackened bits.



Digital Filters: Design Compromises


Applicability of Steak-Salvaging filter.

- Use once in the entire record-replay chain
 - The rest of the chain may keep using halfbands.
- Check by ear
 - The 44.1kHz version has a sonic signature.
 - Weigh against improved imaging.

Digital Filters: Design Compromises


The Equiripple Filter

- “Just Enough” attenuation = minimum number of coefficients.
- Windowed Sinc filters roll off further inside the stop band.
- Unnecessary attenuation increases length.
- Example: **Equiripple, 75 taps**. **Windowed, 95 taps**.




Digital Filters: Design Compromises

In a halfband filter, ripple and attenuation are linked




Digital Filters: Design Compromises

Ripples are equal in amplitude and nearly equally spaced. Spacing $\approx 2/(\text{\#taps}+1)$




Digital Filters: Design Compromises

Let's define another linear-phase filter with nearly the same in-band response



Digital Filters: Design Compromises


Impulse response of that filter (exaggerated):



- Constant in-band ripple equates to echos at the ends of the filter.
- Amplitude of echos = stop band attenuation – 6dB
- Post-echo is certainly masked. Pre-echo possibly not.


Digital Filters: Design Compromises

Close-up of ripple of windowed sinc filter



Digital Filters: Design Compromises

Compare 2 halfband filters at 95 taps



Digital Filters: Design Compromises

Import on “digiphobia”

- Classic argument against digital: “pre-ringing”
 - Little serious evidence of audibility of pre-ringing outside the audio band exists.
 - Looks like a red herring
- 2 common impementation problems were identified
 - Aliasing and Pre-Echo
 - Audible deficiencies are linked to compromising.
 - Solved by better adhering to theory, not deviating further.
- Pre-ringing hypothesis is not needed!
 - You Hear What You Hear but it's Not What You Think.


Testing the Pre-Ringing Hypothesis

Testing audibility of brick wall filtering

- Use a 96kHz or 192kHz recording.
- Slice off 0.4535-0.5fs area.
- Test the following filters (never decimate):
 - 20kHz sharp-rolloff
 - 20kHz slow-rolloff
 - 40kHz sharp-rolloff


Testing the Pre-Ringing Hypothesis

Effect of Slow-Roll Filter after the fact...



Testing the Pre-Ringing Hypothesis

...reverses effect of sharp rolloff filters



(example: standard 96kHz AD/DA with slow LPF inserted)

Testing the Pre-Ringing Hypothesis



Should we put Slow-Rolloff filters in IC's?

- NO! Compounded SR filters amount to a brick wall.
- Only brick-wall filters are “idempotent” .
- Use brick-wall filters throughout and shape response only once.

Testing the Pre-Ringing Hypothesis

How About The Slow-Rolloff Filters in Chip XYZ?


- Intended to reduce latency, NOT improve sound quality



Testing the Pre-Ringing Hypothesis

The Phase-Optimised Filter

- Reduces pre-ringing at the expense of post-ringing



- Magnitude response is maintained
- Cost-effective implementation (IIR+short FIR at f_{sout})
- Reduces latency with minimal loss of sound quality

Testing the Pre-Ringing Hypothesis

Are phase-optimised filters a good thing?

- YES. Much better tradeoff between audio performance and latency.

Should phase-optimised filters be standard?

- NO. One pass may be inaudible but 2 passes? 10?
- “Improved sound quality” claim is based on pre-ringing hypothesis.