CERN

Measurement and Acquisition for High Precision Power Converter Current Regulation for Particles Accelerators

Digital Signal Processing for Regulation Purposes

Dr. Michele Martino
CERN

September 12th 2018
on behalf of TE-EPC-HPM

Introduction

- This part of the lecture is not going to cover all the issues related to acquisition and digital post-processing
- It will highlight only some "peculiar" aspects related to:
  - Closed loop regulation : Negative feedback
  - Application to Power Converters : some disturbances present at pathological frequencies due to switching operation and imperfect suppression of 50 Hz harmonics
- Approach in-between a canonical lecture and a "tutorial"
- Quantities such as SNR (Signal-to-Noise-Ratio) and similar will be defined as ratio of powers rather than ratio of rms amplitudes as done in [1],[2] etc. When in doubt always use them in dB and you will be ok ☺

$$SNR_{dB} = 10log_{10}\left(\frac{Signal\ Power}{Noise\ Power}\right) = 10log_{10}\left(\frac{Signal\ rms}{Noise\ rms}\right)^2 = 20log_{10}\left(\frac{Signal\ rms}{Noise\ rms}\right)$$

- Some recap and case studies details as "extra slides" for the ones interested

Current measurement for “control” applications

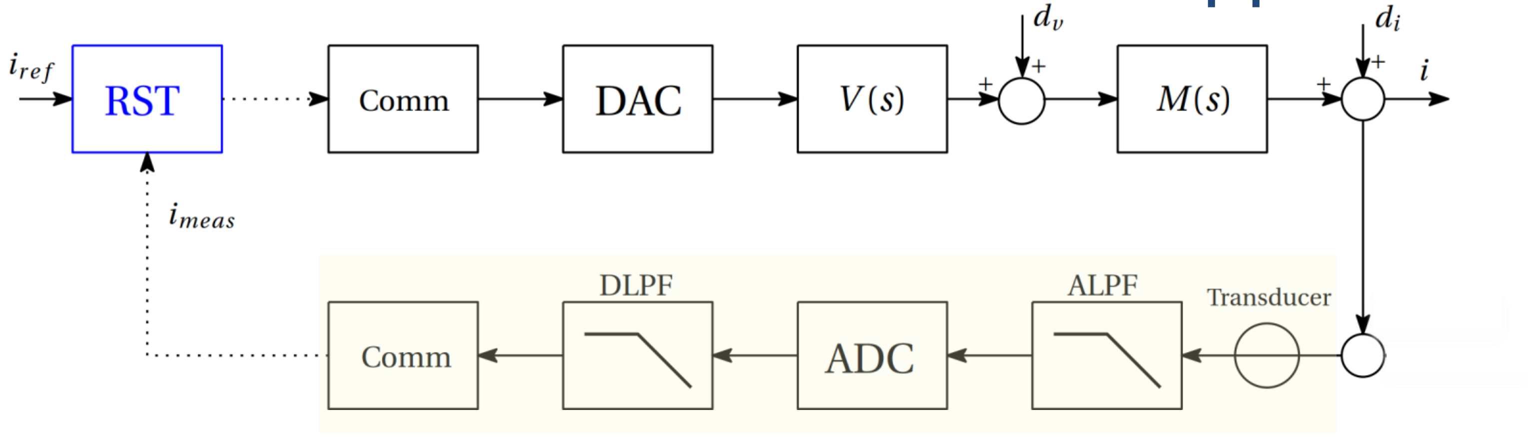

Delay in the measurement chain is highly detrimental for closed-loop stability

Even with optimized control techniques delay limits the achievable CLBW of the PC !

Fundamental trade-off of the measurement chain: precision vs speed !

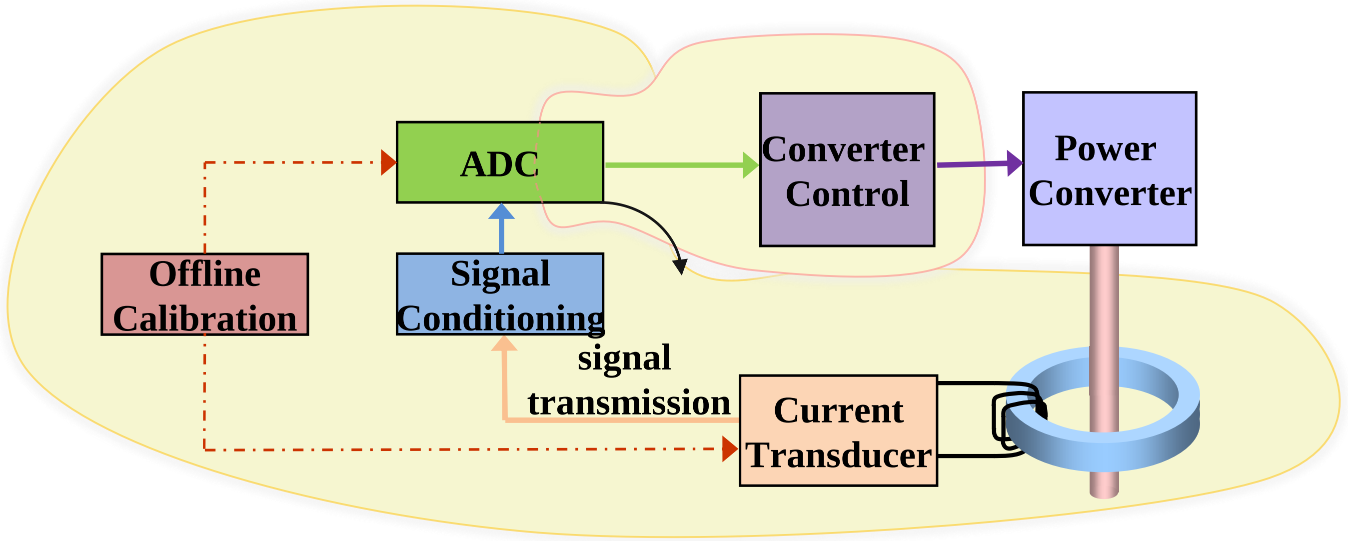

Current measurement for “control” applications

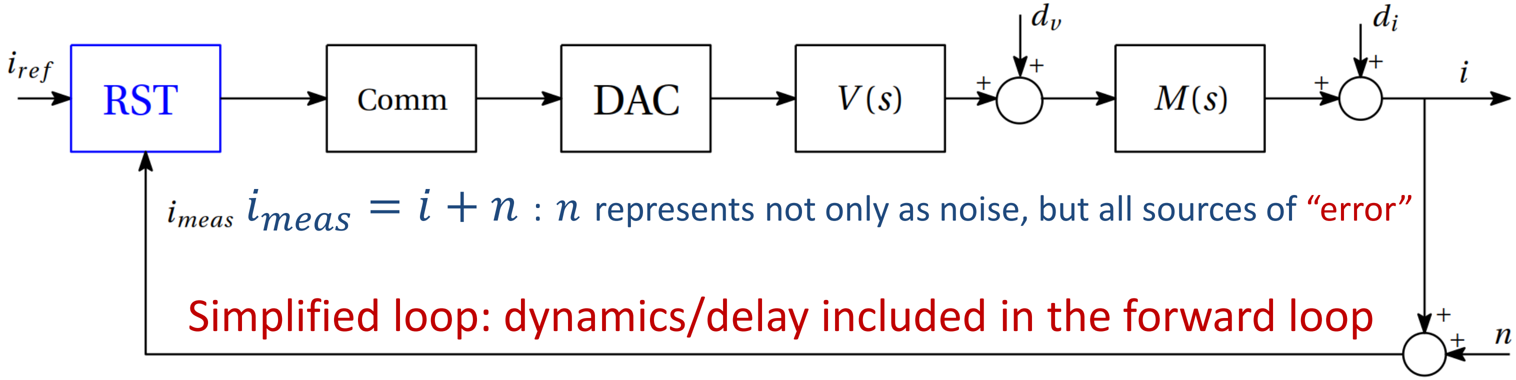

$i_{meas} = i + n$: n represents not only as noise, but all sources of “error”

Simplified loop: dynamics/delay included in the forward loop

Bode Integral for stable systems:

$$\int_0^{\infty} \ln \left| \frac{i(j\omega)}{d_i(j\omega)} \right| d\omega = 0$$

It can be seen as a “conservation law” …

… conservation of “dirt” ! [3]

You can reject disturbances usually in low frequency, but then you have to pay the price somewhere else, often where you don’t want it !

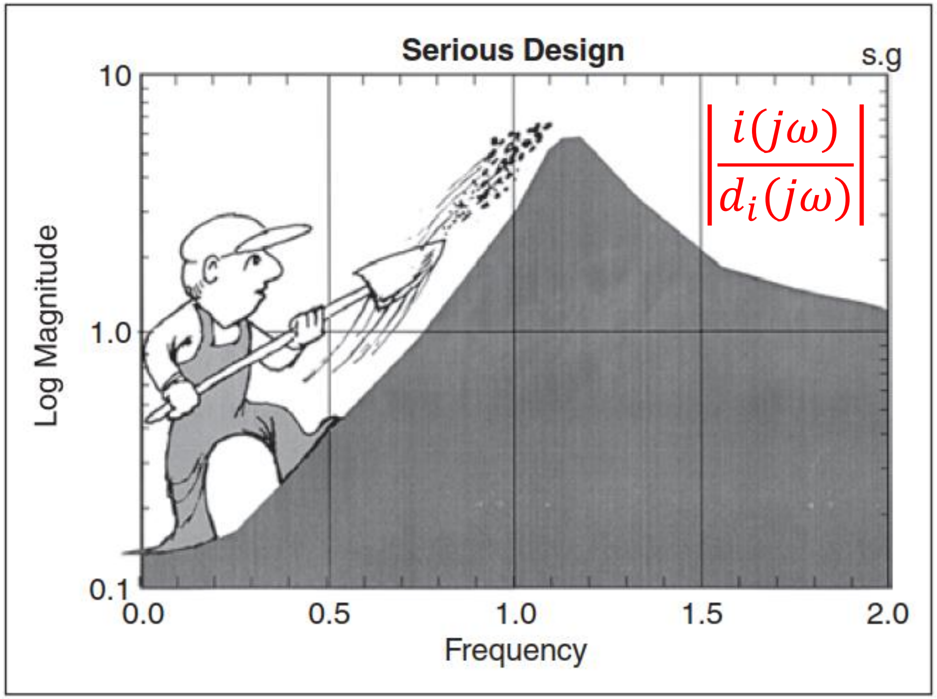

Current measurement for "control" applications

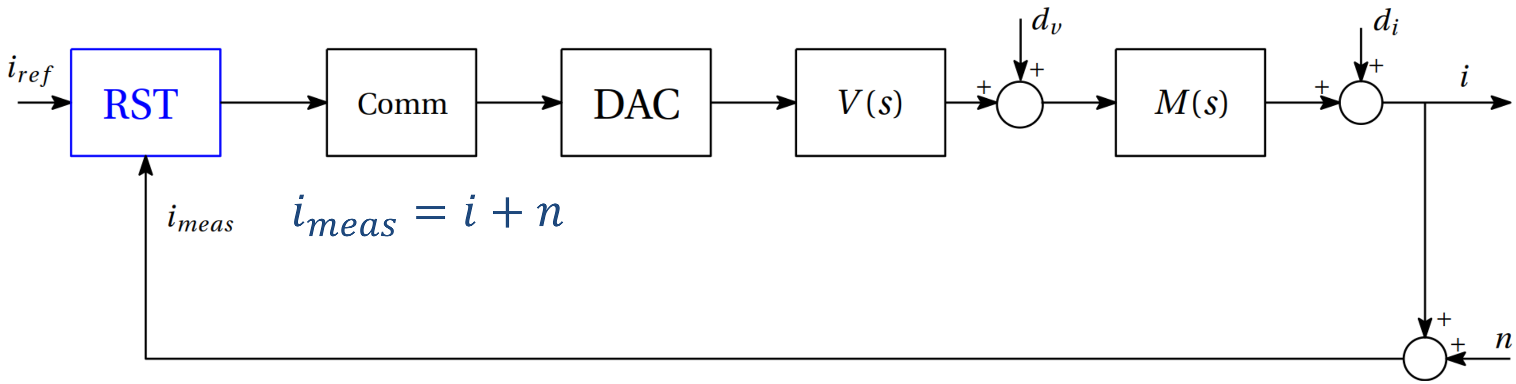

Bode Integral also sets important constraints for the measurement channel:

- in low frequency any "noise" will be "seen" by the current
  (and by the magnet hence the beam)
- in high frequency usually the "noise" is attenuated by the loop
- somewhere in the middle the "noise" is even amplified
  (by no more than 6 dB for a good design)

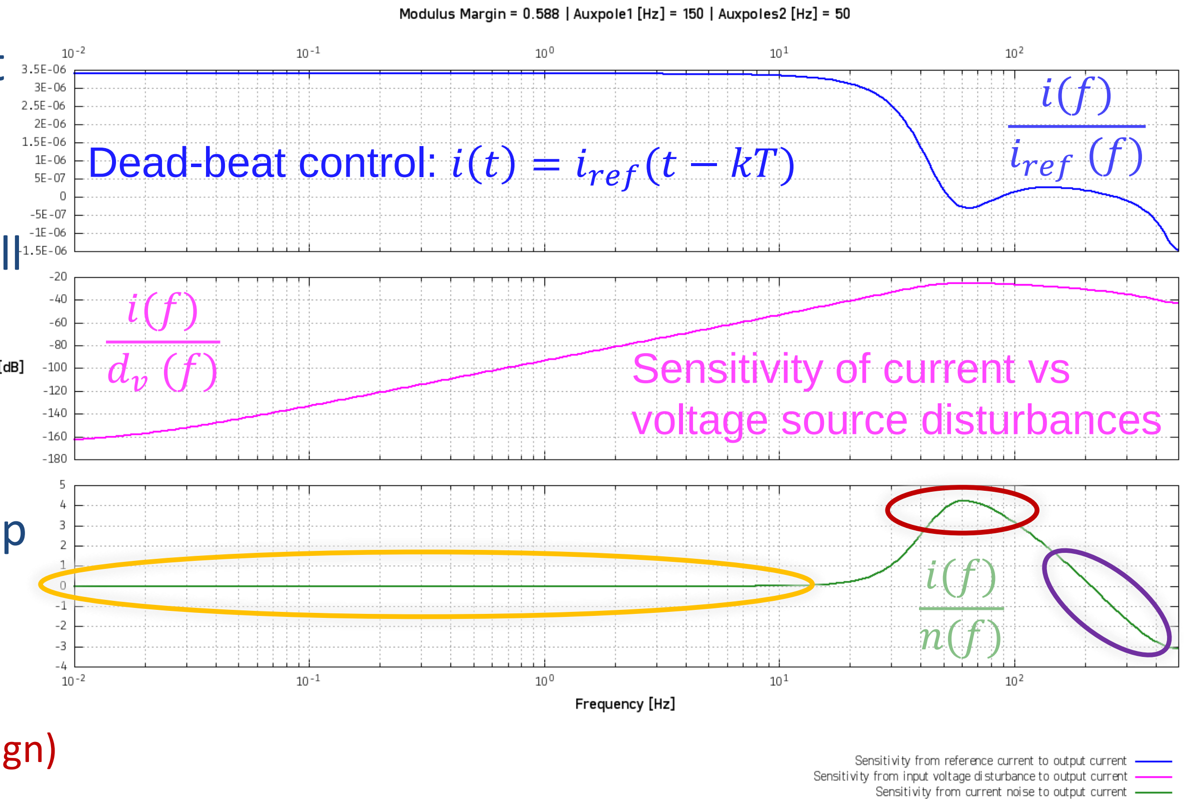

"Sampling basics"

Ideal Sampling

Time Domain	Frequency Domain
$x_a(t)$	$X_a(f) = \int_{-\infty}^{+\infty} x_a(t)\, e^{-j2\pi f t} dt$
$x_s(t) = \sum_{k=-\infty}^{+\infty} x_a(kT)\delta(t-kT)$	$X_s(f) = \frac{1}{T}\sum_{k=-\infty}^{+\infty} X_a(f - kf_s)$

Spectrum of a critically sampled strictly band-limited signal: no alias !

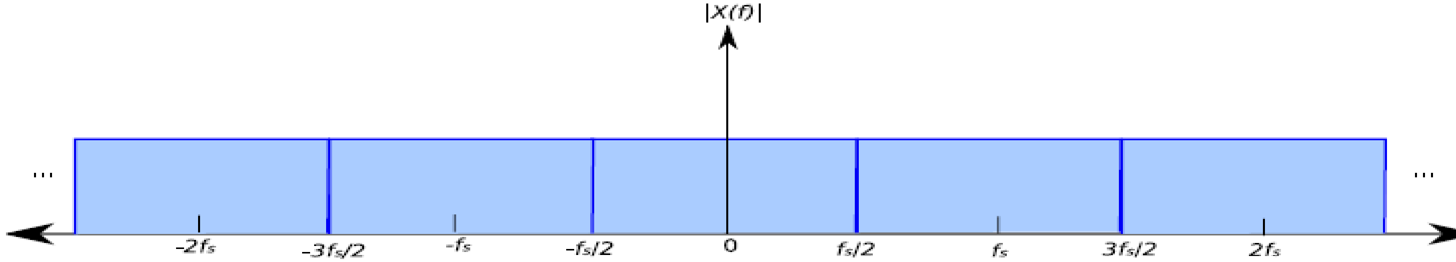

$f_s = \frac{1}{T}$ [4]

Signals whose bandwidth exceeds $f_s/2$ will be corrupted!

Noise is always present so strictly band-limited signals do not exist !

Anti-aliasing filtering always needed !

If the controller runs at f_c , what about sampling faster ?

Oversampling

Standard Nyquist Sampling

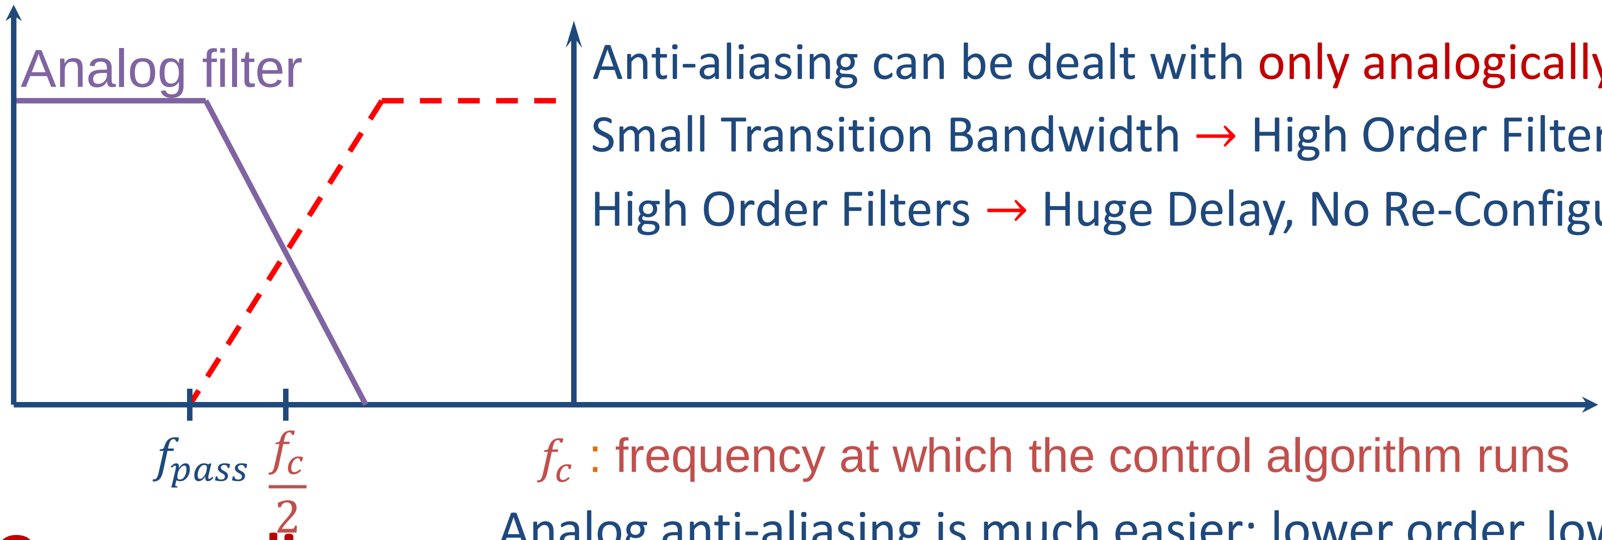

Anti-aliasing can be dealt with only analogically!

Small Transition Bandwidth → High Order Filter

High Order Filters → Huge Delay, No Re-Configurability ☹

f_c : frequency at which the control algorithm runs

Oversampling

Analog anti-aliasing is much easier: lower order, lower delay !

Anti-aliasing filtering can be shared between analog and digital!

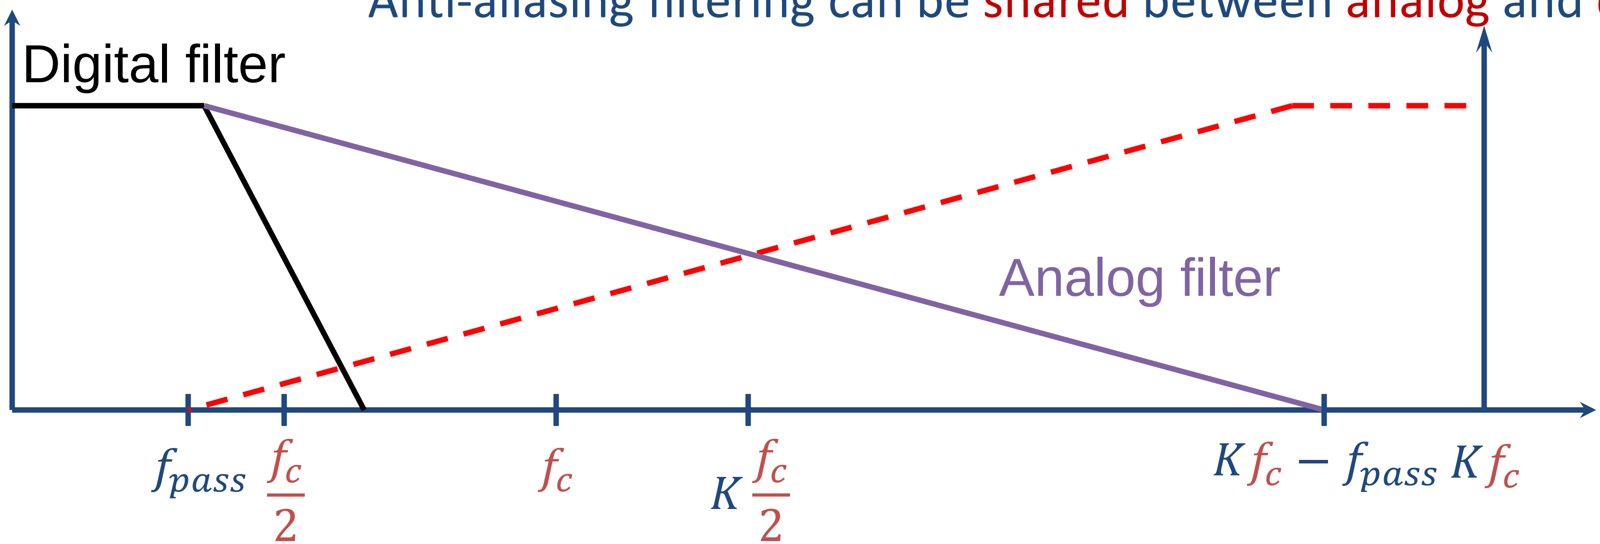

Oversampling

Sampling is not everything!

- **Quantization Process – a bit of theory**

  Quantization is a non linear process - vast and tricky subject (many experts in GMEE)

  Quantization of a signal can be seen as sampling its PDF (Probability Density Function)

  Fortunately an approximated model works very well ☺ (… almost all the time)

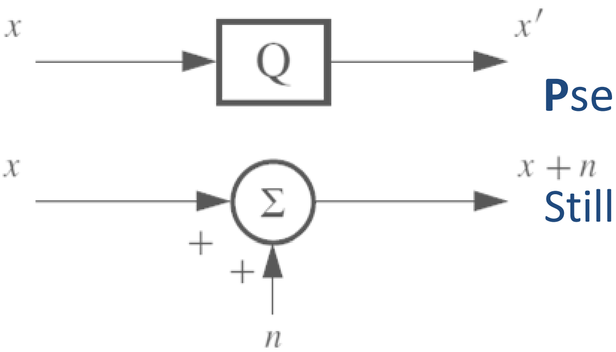

Pseudo **Q**uantization **N**oise model: $PDF(\acute{x}) \cong PDF(x+n)$ [5]

Still a bit too complicated → $\acute{x} \cong x+n$ now the model is linear ☺

$PQN\ n$: Uniform Statistical Distribution $\left[-\frac{\Delta}{2}, +\frac{\Delta}{2}\right]$

White

Independent of input signal x

- **Meaning :**

  Power of $PQN\ n$ does not change with sampling frequency $E[n^2] = \frac{\Delta^2}{12}$ (Uniform PDF)

  $$E[n^2] = \int_{-\infty}^{+\infty} PSD_n(f)df = \int_{-\frac{F_s}{2}}^{+\frac{F_s}{2}} PSD_n(f)df = \int_{-\frac{F_s}{2}}^{+\frac{F_s}{2}} \sigma^2 df = F_s\sigma^2$$ ideally (White Spectrum)

  By oversampling (increasing F_s) the PQN Power Spectral Density σ^2 is reduced !

Oversampling → Downsampling

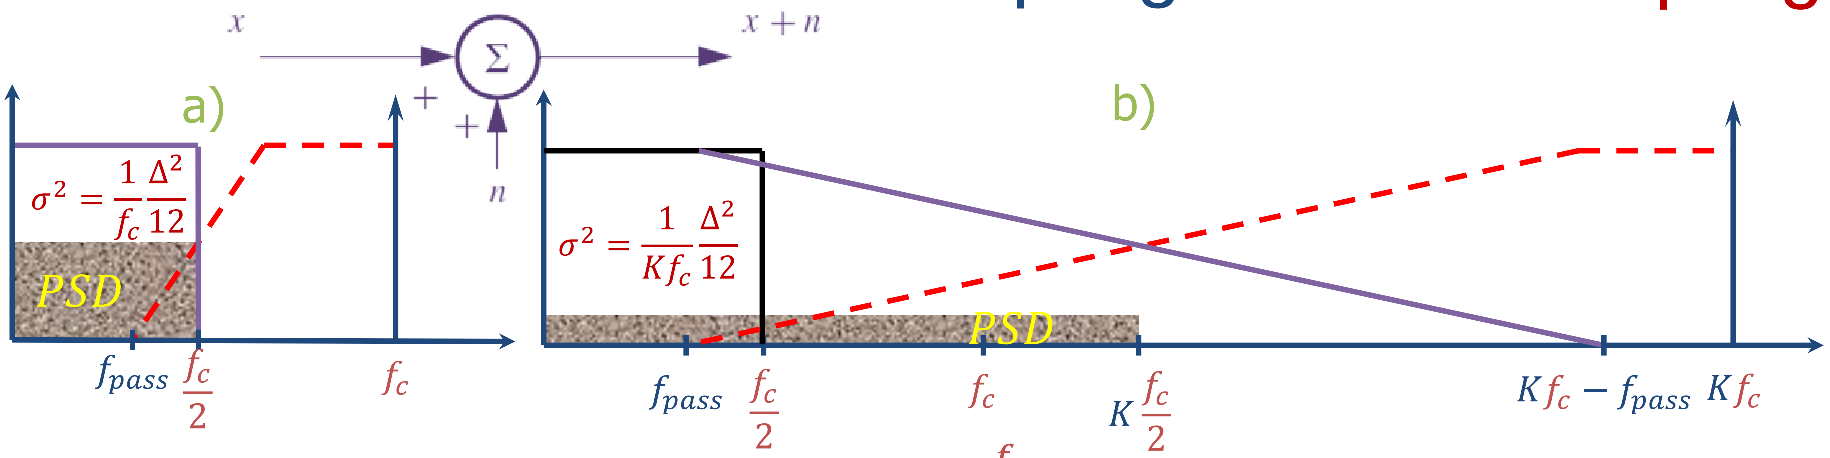

Let's assume ideal brick-wall lowpass filters at $\frac{f_c}{2}$ for both a) and b) and consider the :

Signal-to-Quantization-Noise Ratio $SQNR = \frac{E[x^2]}{E[n^2]}$ **:** Sampling Frequency Bandwidth

a) $SQNR = \frac{E[x^2]}{\frac{\Delta^2}{12} \cdot \frac{1}{f_c} \cdot f_c}$; b) $SQNR = \frac{E[x^2]}{\frac{\Delta^2}{12} \cdot \frac{1}{Kf_c} \cdot f_c}$ → $SQNR$ improves by a factor K

For "equal" lowpass filters in a) and b) no matter if analog or digital $SQNR$ still improves by K

For an analog unity gain lowpass: $\Delta_{OUT} = \Delta_{IN}\sqrt{\frac{\int_{-\infty}^{+\infty}|H^2(f)|df}{f_{sampling}}}$ so $\Delta \boldsymbol{SQNR_{bit}} = -\frac{1}{2}\log_2\left(\frac{\int_{-\infty}^{+\infty}|H^2(f)|df}{f_{sampling}}\right)$

For a digital unity gain lowpass: $\Delta \boldsymbol{SQNR_{bit}} = -\frac{1}{2}\log_2(\sum|h^2(n)|)$ since $\frac{\int_{-\infty}^{+\infty}|H^2(f)|df}{f_{sampling}} = \sum|h^2(n)|$

The digital lowpass filter can be also used to put notches at pathological frequencies before taking one sample out of K and using it in the digital control algorithm at rate f_c

The never-ending dispute: IIR vs FIR

FIR N^{th} order / $N+1$-taps	IIR Q^{th} order ($a_0 = 1$, $P \leq Q$)
$y[k] = \sum_{i=0}^{N} b_i\, x[k-i]$	$y[k] = \sum_{i=0}^{P} b_i\, x[k-i] - \sum_{j=1}^{Q} a_j\, y[k-j]$

A COMPARISON OF THE NONRECURSIVE (FIR) AND RECURSIVE (IIR) FILTERS OVER THE YEARS. COLORED BOXES INDICATE WHEN ONE TECHNIQUE OFFERS AN ADVANTAGE (IN GREY) OVER THE OTHER (IN ORANGE).

	1970		1980		NOW	
PROPERTY TRANSFER FUNCTION	**FIR ZEROS ONLY**	**IIR POLES AND/ OR ZEROS**	**FIR ZEROS ONLY**	**IIR POLES AND/ OR ZEROS**	**FIR ZEROS ONLY**	**IIR POLES AND/ OR ZEROS**
DESIGN METHODS FOR FREQUENCY SELECTIVITY	SUB-OPTIMAL USING WINDOWS	OPTIMAL ANALYTIC, CLOSED FORM	OPTIMAL USING ITERATIVE METHODS	OPTIMAL ANALYTIC, CLOSED FORM	OPTIMAL USING ITERATIVE METHODS	OPTIMAL ANALYTIC, CLOSED FORM
MULTIPLICATIONS/REGISTERS NEEDED FOR SELECTIVITY	MANY	FEW	MORE	FEWER	MORE	FEWER
CAN BE EXACTLY ALLPASS	NO	YES	NO	YES	NO	YES
UNSTABLE	NEVER	FOR POLES p_i, $\lvert p_i \rvert > 1$	NEVER	FOR POLES p_i, $\lvert p_i \rvert > 1$	NEVER	NEVER
DEADBAND EXISTS	NO	YES	NO	YES	NO	NO
CAN BE EXACTLY LINEAR PHASE	YES	NO	YES	NO	YES	YES (non-causal)
CAN BE ADAPTIVE			YES	DIFFICULT OR IMPOSSIBLE	YES	DIFFICULT OR IMPOSSIBLE
OPPORTUNITIES FOR PARALLELISM			MANY	SOME	MANY	MANY
HILBERT TRANSFORMER	INEFFICIENT	IMPRACTICAL BECAUSE NOT CAUSAL	INEFFICIENT	IMPRACTICAL BECAUSE NOT CAUSAL	INEFFICIENT	EFFICIENT

Charles M. Rader : The Rise and the Fall of Recursive Digital Filters – IEEE Signal Processing Magazine, Nov 2006 [6]

A "no match" against FIR: the 1-bit Σ-Δ

1-bit Σ-Δ ADC

- ❑ It is the digital filter that actually determines (most of) the ADC "precision"!

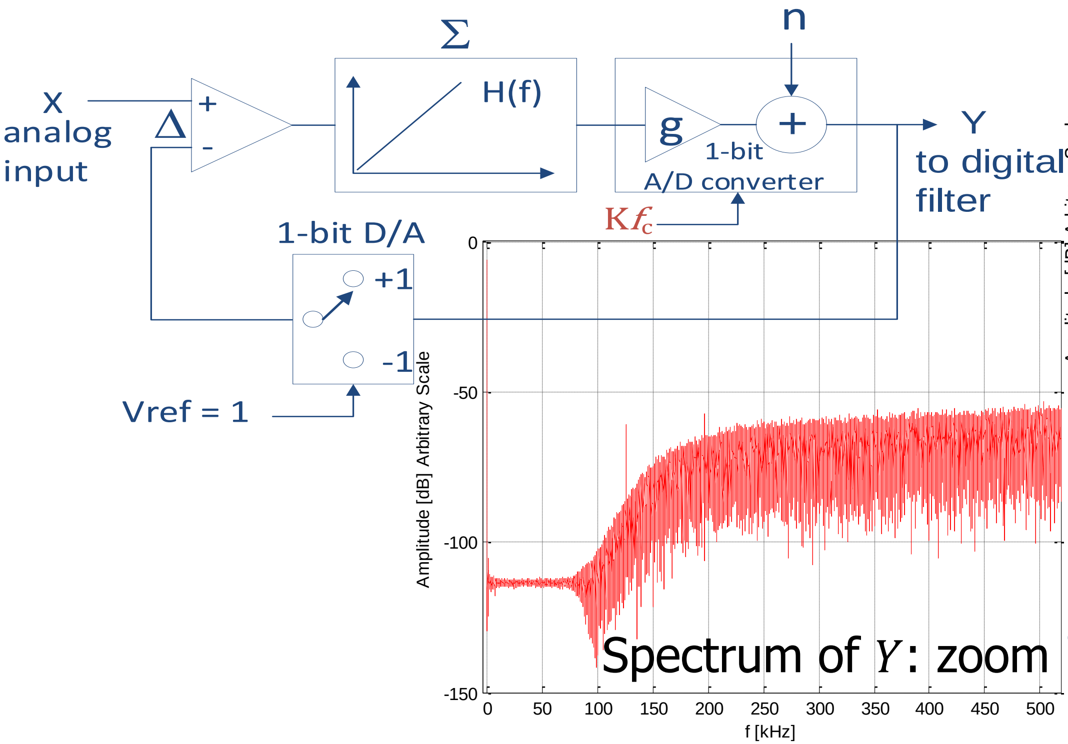

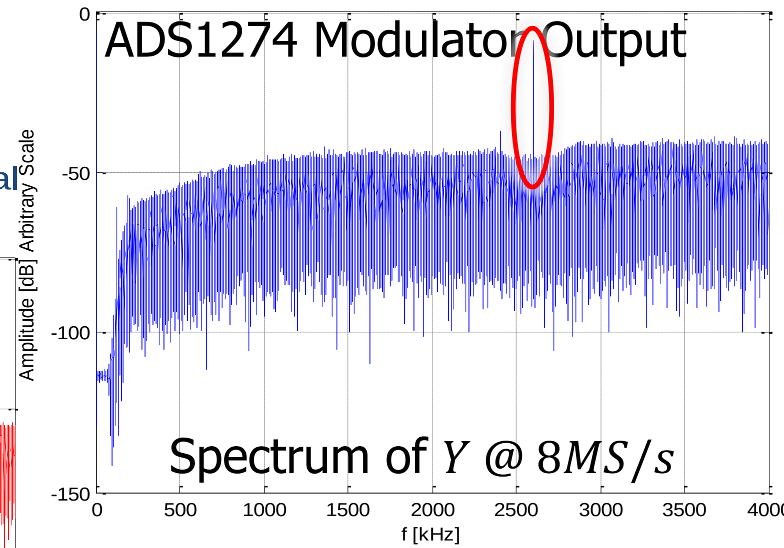

$$out[k] = \sum_{i=0}^{N} b_i\, y[k-i],\; y[k-i] = \pm 1$$

- ❑ $H(f)$ "shapes" the white, uniformly distributed noise n (PQN model is assumed)
- ❑ Only 1-bit means that no multiplication are needed for "downstream" FIR filters

Spectra in arbitrary dB scale shown only to illustrate the "noise shaping" behaviour of the Σ-Δ

Note that some "idle tones" can have amplitudes close to FS !

Specification of (digital) filters

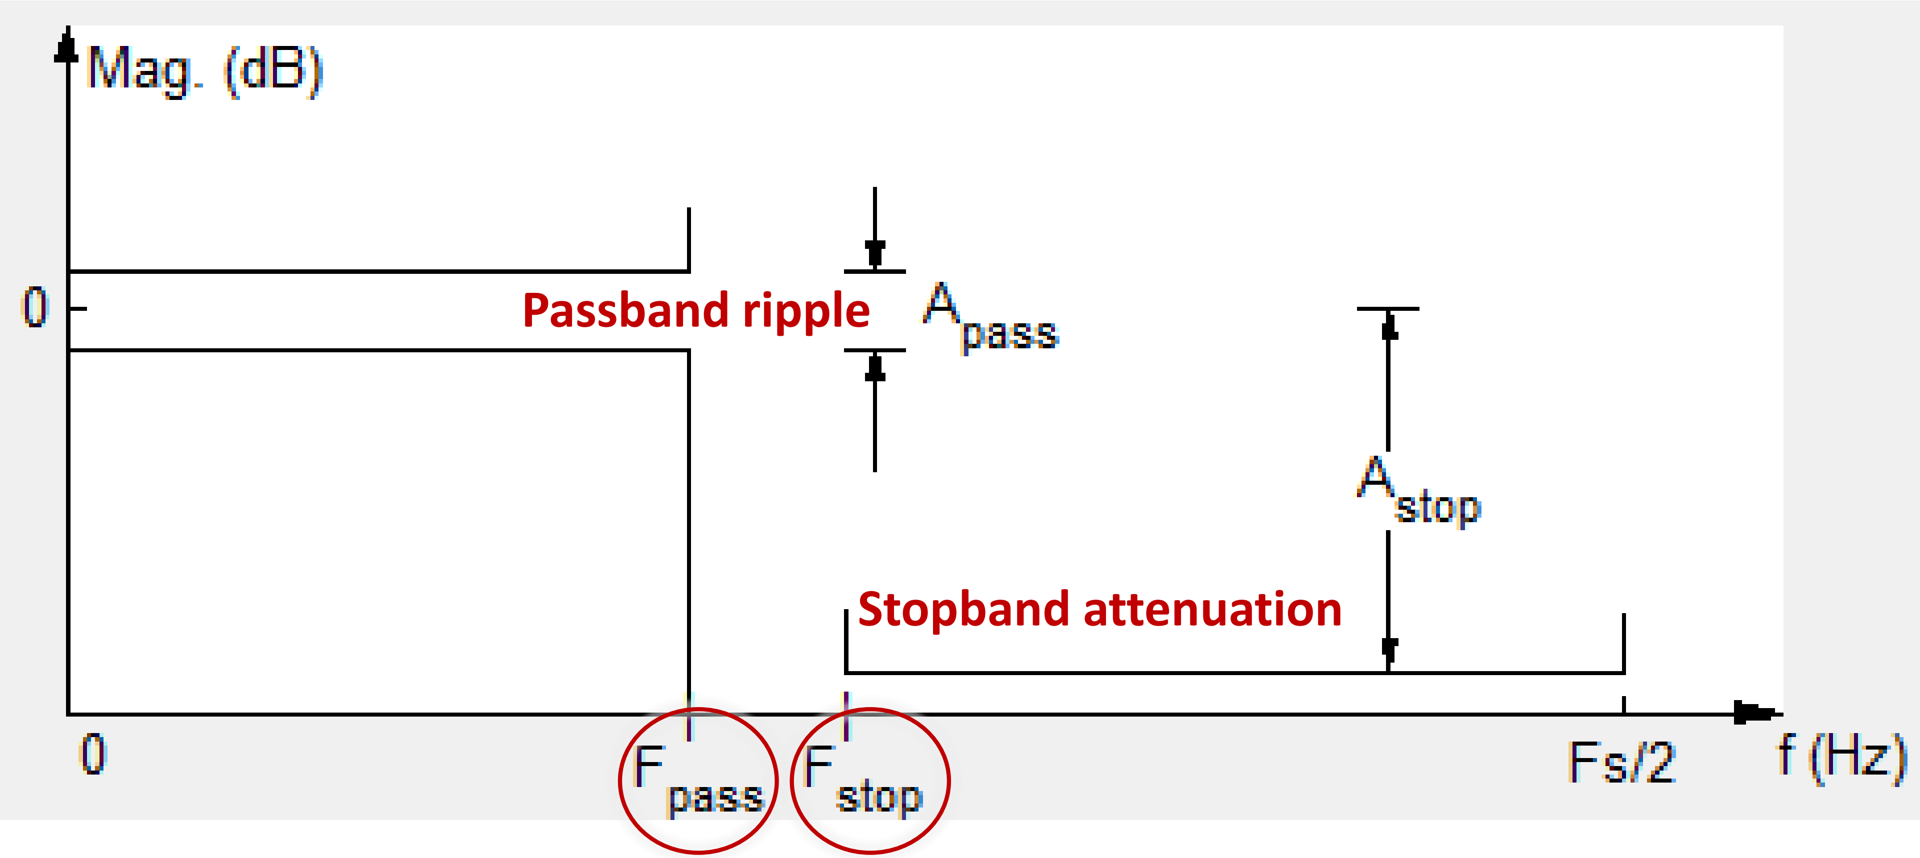

Specification of (digital) filters

From Precision to Filter Specs

Full scale FS is considered as the reference level $(\pm 10V\colon FS = 10V, FSR = 20V)$; part-per-million $ppm = 10^{-6}$ will always be referred to it; if a precision of $x\ ppm$ is required an ideal quantizer would then need to have a quantization step Δ such that the maximum "quantization error" would be s.t. $(|error| \leq) \frac{\Delta}{2} \leq x\ ppm$

- **Passband ripple** (beware of the implication of this requirement)

  In order to guarantee an "harmonic accuracy" of $x\ ppm$ all over the "useful band" the passband gain G should change less than the precision required $(FS - x\ ppm\ FS) \leq G(j2\pi f)FS \leq (FS + x\ ppm\ FS) \rightarrow$
  $(1 - x\ ppm) \leq G(j2\pi f) \leq (1 + x\ ppm) \rightarrow |Ripple_{max}|_{dB} = (1 + x\ ppm)_{dB}$

- **Stopband attenuation**

  For the stopband attenuation an estimation of the noise "level" to be rejected is required!
  In a relatively "silent" environment assuming FS noise components (worst case scenario) would result in heavy and sometimes **meaningless** over-specification → complexity, computational power, **delay**!

Specification of (digital) filters

From Precision to Filter Specs

- **Stopband attenuation (aim at regulation) - worst case scenario:**

$$SQNR_{dB}^{DC} = 10log_{10}\left({FS^2}/{\frac{\Delta^2}{12}}\right) = 20log_{10}\left(\frac{FS}{\Delta}\sqrt{12}\right) = 20log_{10}\left(\frac{\cancel{FS}}{x\ ppm\ \cancel{FS}}\sqrt{3}\right)$$

$\frac{\Delta^2}{12}$: (pseudo)quantization noise power of an ideal quantizer with quantization step Δ $(= 2\ x\ ppm)$

$$SQNR_{dB}^{DC} = 4.771 + 120 - x_{dB} \cong 124.8\ - x_{dB} \rightarrow A_{stop}\ =\ 124.8\ - x_{dB}$$

An extension of $SQNR$ for arbitrary shape waveforms is presented in [7] based on crest factor

- **Putting it (almost) all together**

$FS = 0\ [dB]$	$A_{pass}\ [dB]$	$A_{stop}\ [dB]$
$1\ ppm$	1.74 x 10^{-5}	124.8
$10\ ppm$	1.74 x 10^{-4}	104.8
$100\ ppm$	1.74 x 10^{-3}	84.8

As an example if the expected, or measured, noise "level" (rms amplitude) is 100 times smaller than the FS , then A_{stop} **should** be specified $40\ dB$ smaller than what is reported in the table.

So $10\ ppm$ precision can be achieved with only $64.8\ dB$ of attenuation in the stopband and so on. This will save computational power and most importantly: delay!

Specification of (digital) filters – some guidelines

From Precision to Filter Specs : balancing requirements

- **Passband and stopband frequencies:**
  - The easiest approach (minimum filter order - largest transition bandwidth):
    - End of the passband f_{pass} $=$ "alias-free" band
    - Beginning of the stopband f_{stop} $= f_c -$ "alias-free" band
  - The iterative approach:
    - Beginning of the stopband f_{stop} $= f_c -$ "alias-free" band ☺
      Replicas occurring around $\pm K\ f_c$ will not affect the desired precision in the "alias-free" band
    - Considerations for the choice of f_{pass}:
      - $f_{pass} \downarrow \rightarrow$ delay $\uparrow$
      - $f_{pass} \uparrow \rightarrow (f_{stop} - f_{pass}) \downarrow \rightarrow$ order $\uparrow$ ☹
    - f_{pass} should be chosen by trading off delay and filter complexity!

- **The "alias-free" band**
  - "alias-free" band $>$ $CLBW$ results in over specifying
  - "alias-free" band $\leq$ $CLBW$ should be considered as full precision might not be needed over the whole closed loop bandwidth, especially for very high precision applications

Minimum-phase FIR

FIR filters can have a nice linear phase, is that so important “in the loop” ?

$f_c = 50\ kS/s$, $F_s = 6.75\ MS/s$, $F_{pass} = 15\ kHz$, $F_{stop} = 45\ kHz$, $A_{pass} = 1\ dB$, $A_{stop} = 60\ dB$	
FIR Linear-Phase Minimum Order	FIR Minimum-Phase Minimum Order
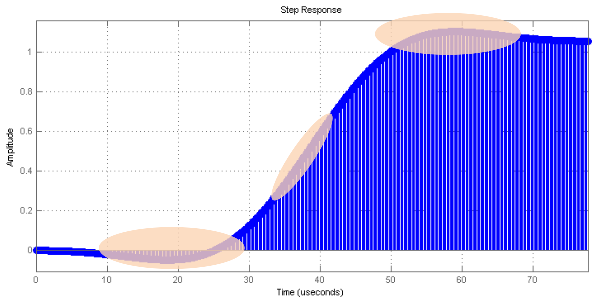	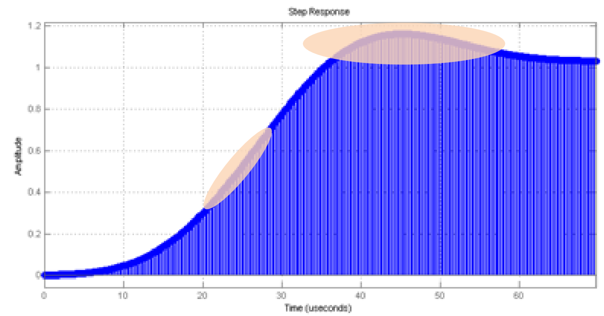

- Minimum-phase FIR filters: the delay in the passband (still approx constant) is significantly lower than that of a linear-phase having the same frequency constraints: $24.5\mu s$ vs $38.2\ \mu s$
- A minimum-phase FIR filter has additional advantages; the overall order of the design is less than that of a linear-phase FIR: 470 vs 525 coefficients which are less sensitive to quantization
- These are clear advantages, but minimum-phase FIR has a lot more overshoot! Is that a problem?

Minimum-phase FIR for "regulation"

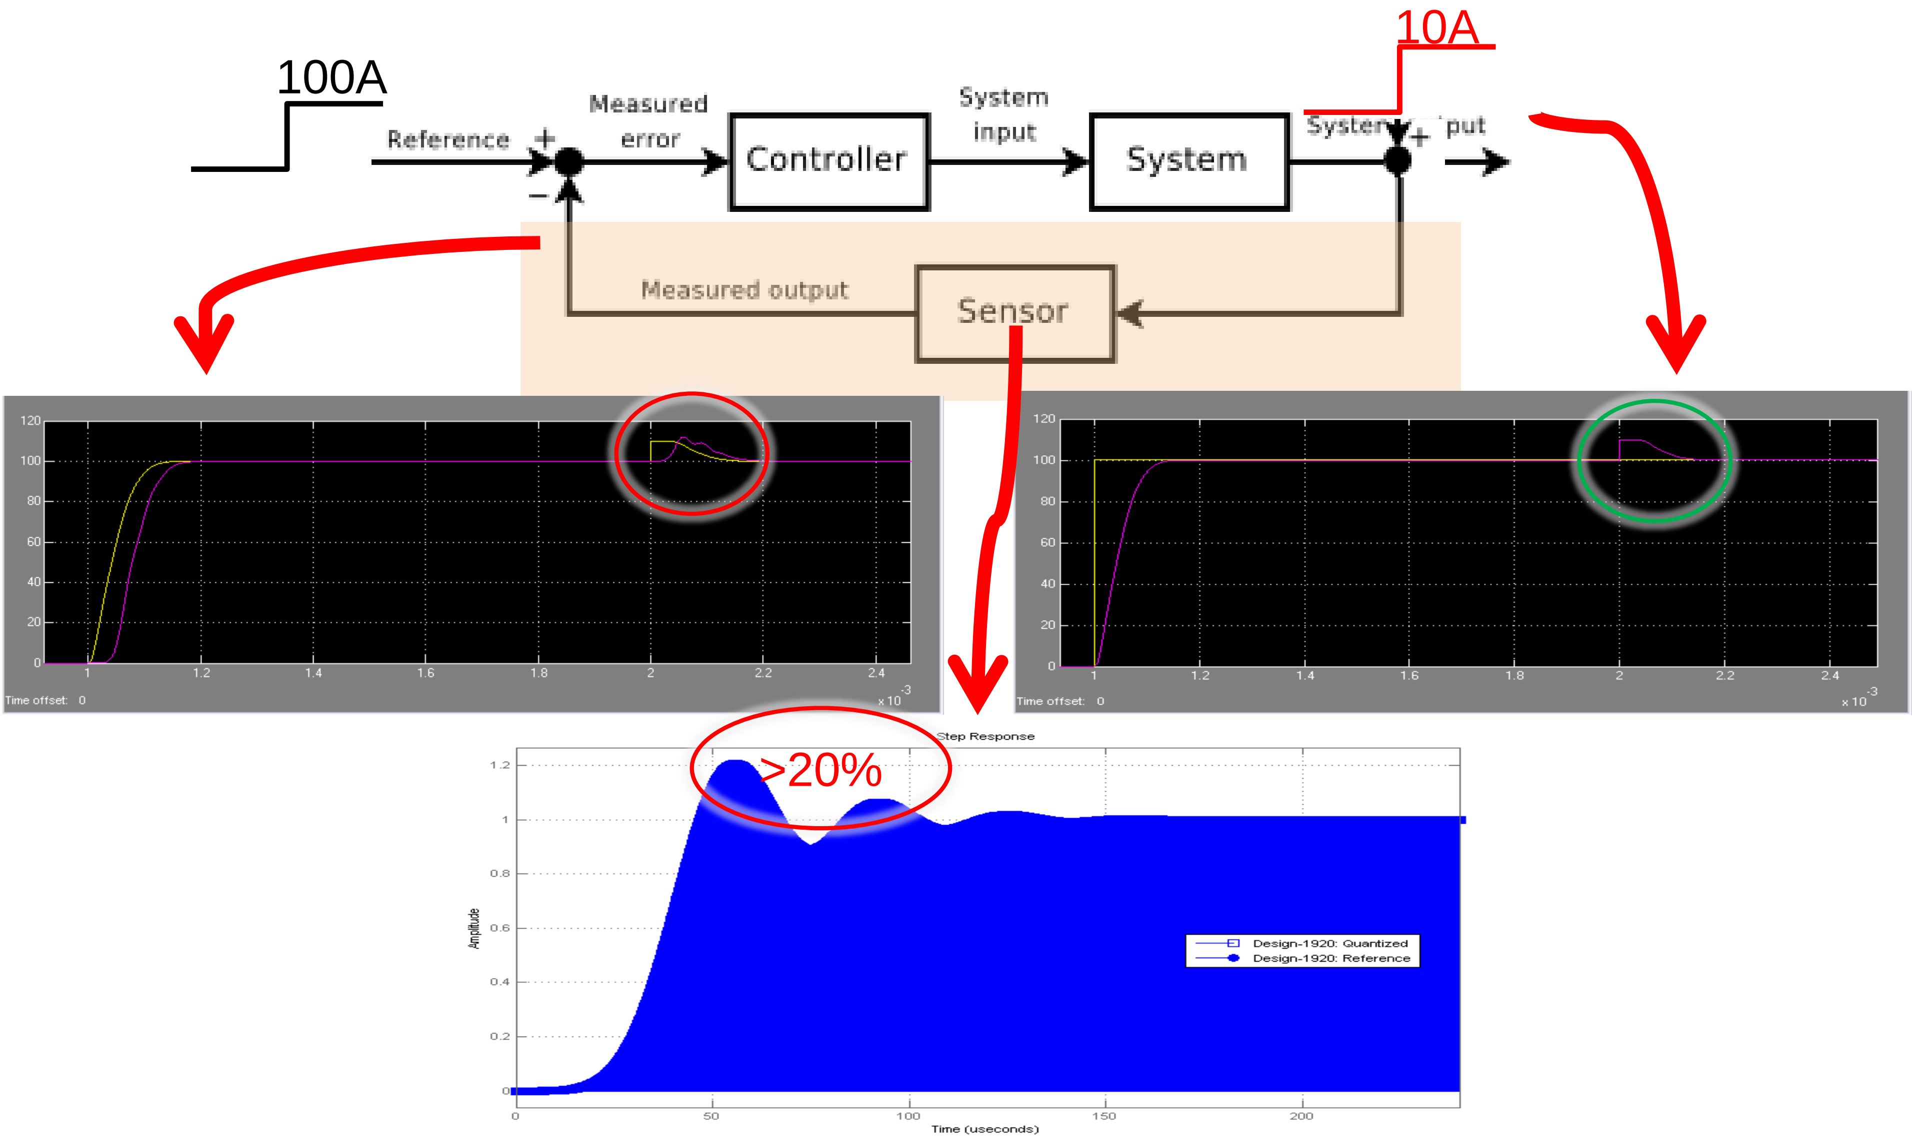

Minimum-phase FIR

FIR filters can have a nice linear phase, is that so important "in the loop" ?

$f_c = 50\ kS/s$, $F_s = 6.75\ MS/s$, $F_{pass} = 15\ kHz$, $F_{stop} = 45\ kHz$, $A_{pass} = 1\ dB$, $A_{stop} = 60\ dB$	
FIR Linear-Phase Minimum Order	FIR Minimum-Phase Minimum Order
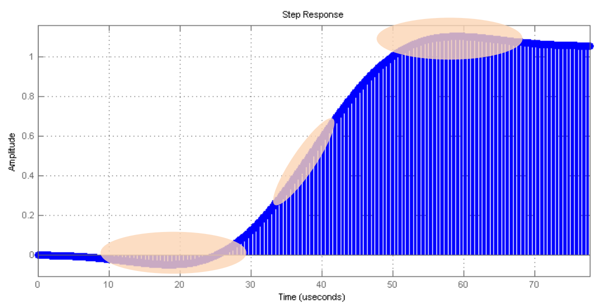	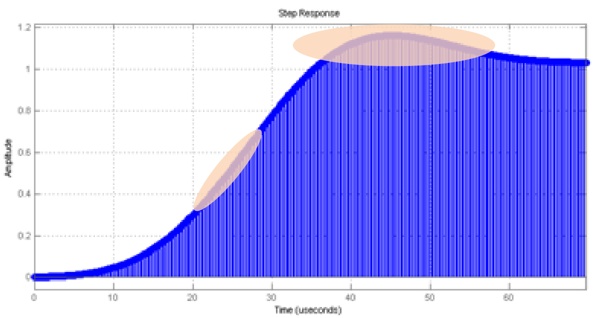

- Minimum-phase FIR filters: the delay in the passband (still approx constant) is significantly lower than that of a linear-phase having the same frequency constraints: $24.5\mu s\ vs\ 38.2\ \mu s$
- A minimum-phase FIR filter has additional advantages; the overall order of the design is less than that of a linear-phase FIR: $470\ vs\ 525$ coefficients which are less sensitive to quantization
- These are clear advantages, but minimum-phase FIR has a lot more overshoot! Is that a problem?
- Actually not if the filter is part of the measurement chain of a control loop!
- Minimum-phase FIR can readily be used in single-stage or multistage decimators… That needs some tricks: $F_s/f_c = K = 135$, as both $470/135$ and $525/135$ are not integers!

Minimum-phase FIR downsample: practical tricks

Given the specifications a minimum-order minimum-phase FIR with N coefficients can be calculated by means of, as an example, MATLAB - FDATOOL with Generalized Equiripple Algorithm

- If $N \le K$ then a new filter can be designed with $N_F = K$; the output at frequency f_c can then be generated by simply taking one samples out of K
- If $N > K$ then a new filter has to be designed with $N_F = \lceil N/K \rceil K$; in order to guarantee the minimum delay $N_{INTERLEAVE} = \lceil N/K \rceil$ filters need to be interleaved

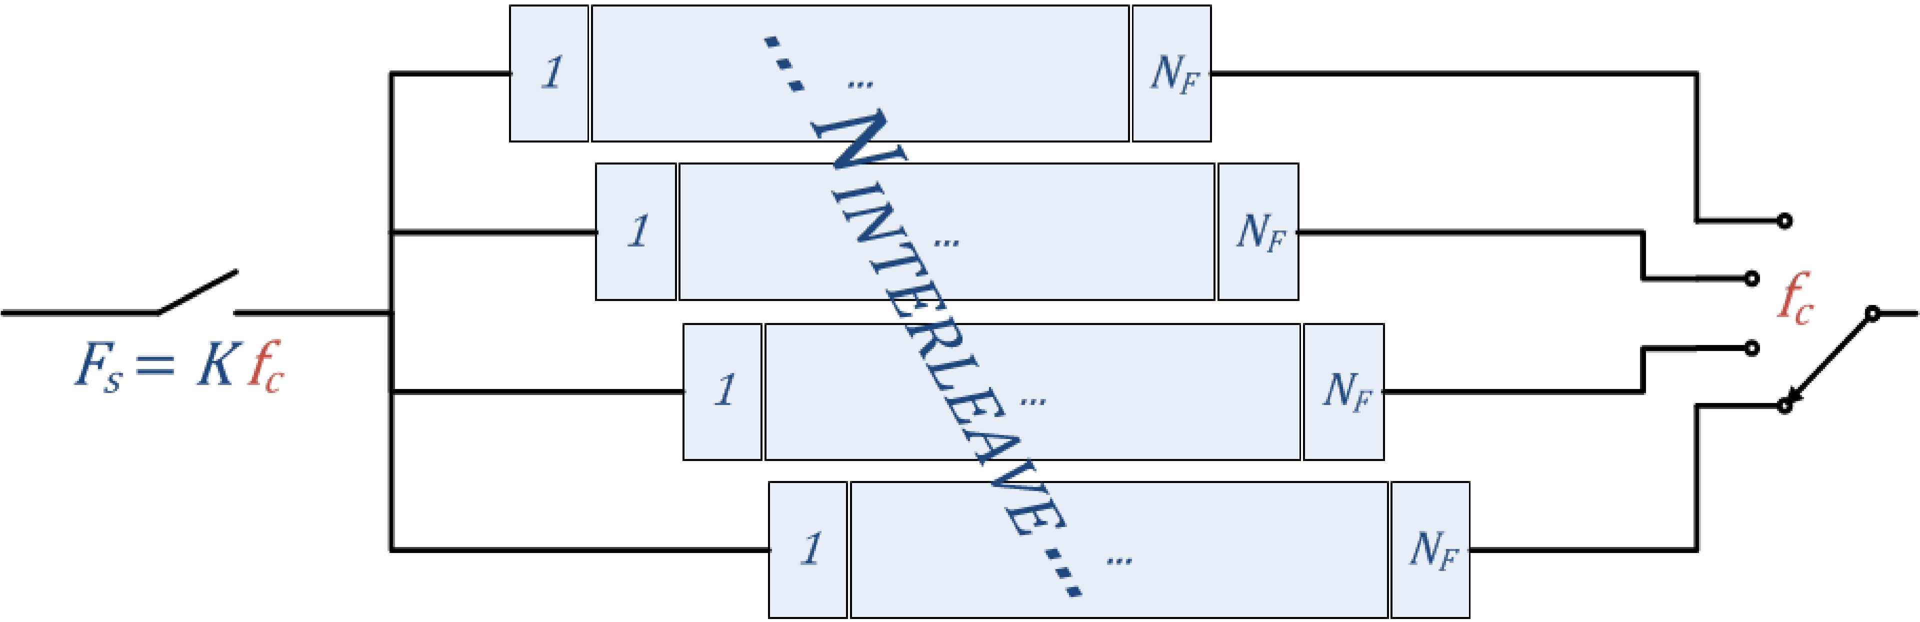

A case study for a 10 ppm Σ-Δ at 50 kS/s

Let's assume we need to extract the signal out of the ADS1274 bitstream with $f_s = 8MS/s$

Designed filter for $CLBW = 5kHz$ **with** $f_c = 50kS/s$ (oversampling factor $K = 160$)

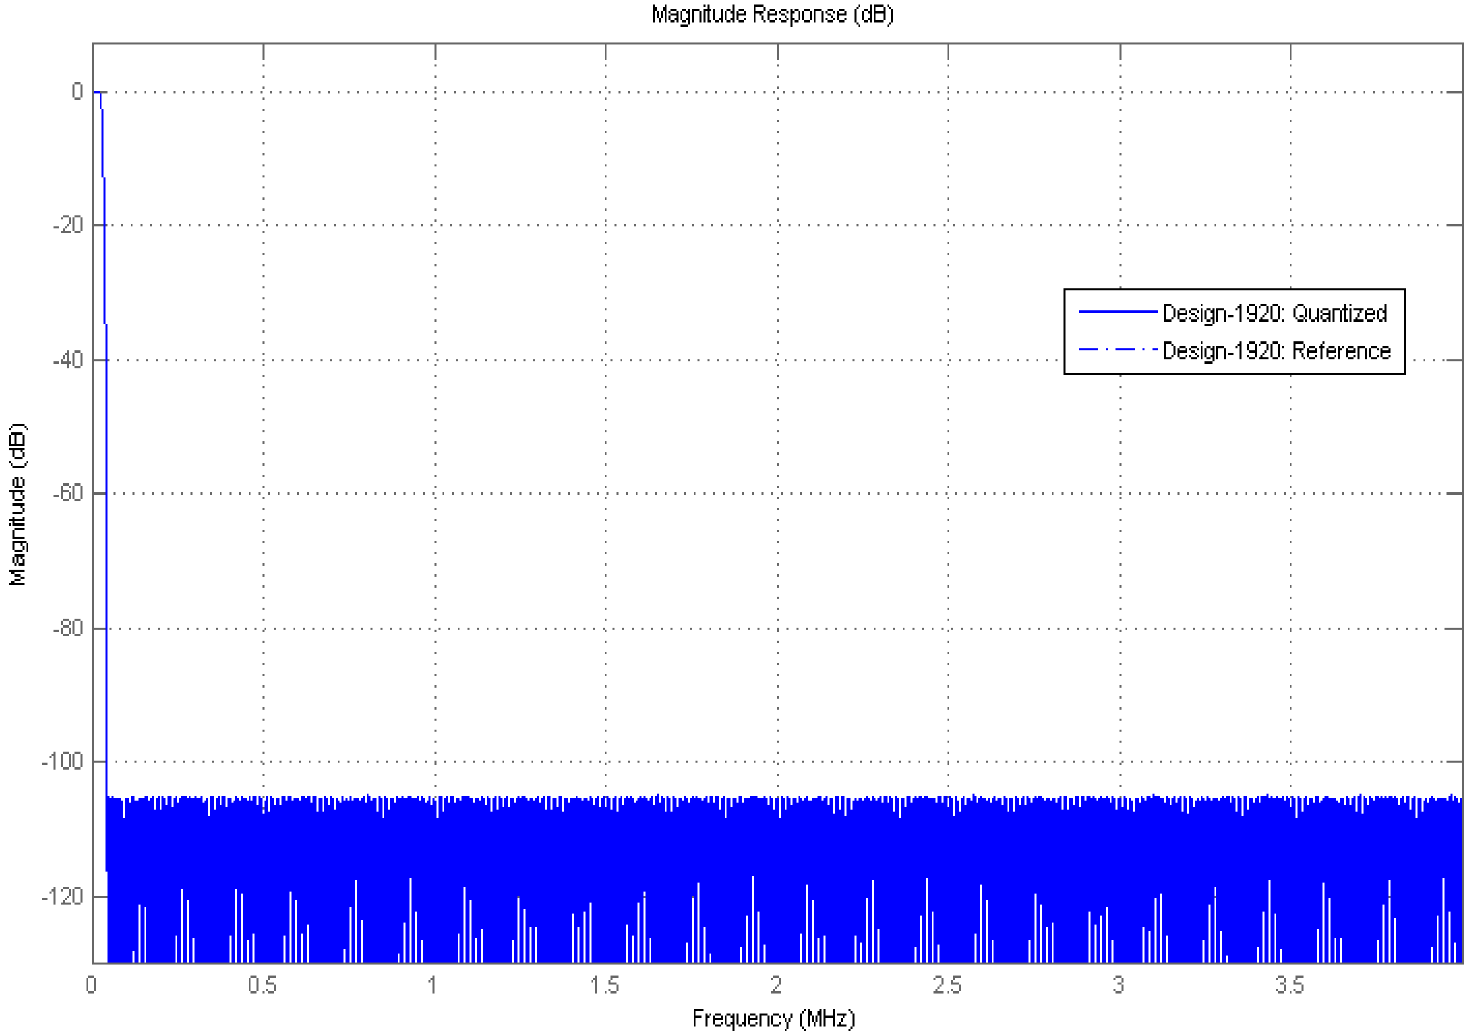

A case study for a 10 ppm Σ-Δ at 50 kS/s

Designed filter for $CLBW = 5kHz$ **with** $f_c = 50kS/s$ (oversampling factor $K = 160$)

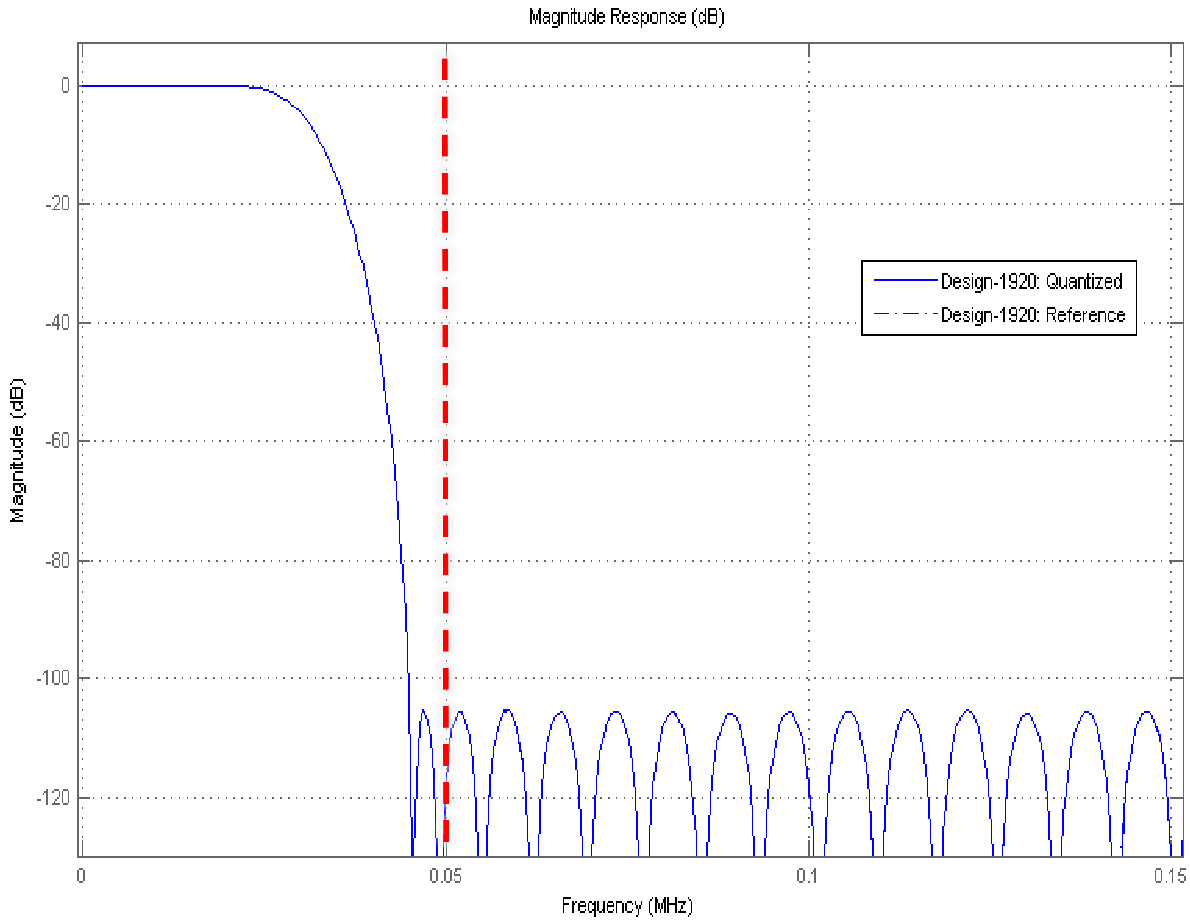

A case study for a 10 ppm Σ-Δ at 50 kS/s

Designed filter for $CLBW = 5kHz$ **with** $f_c = 50kS/s$ (oversampling factor $K = 160$)

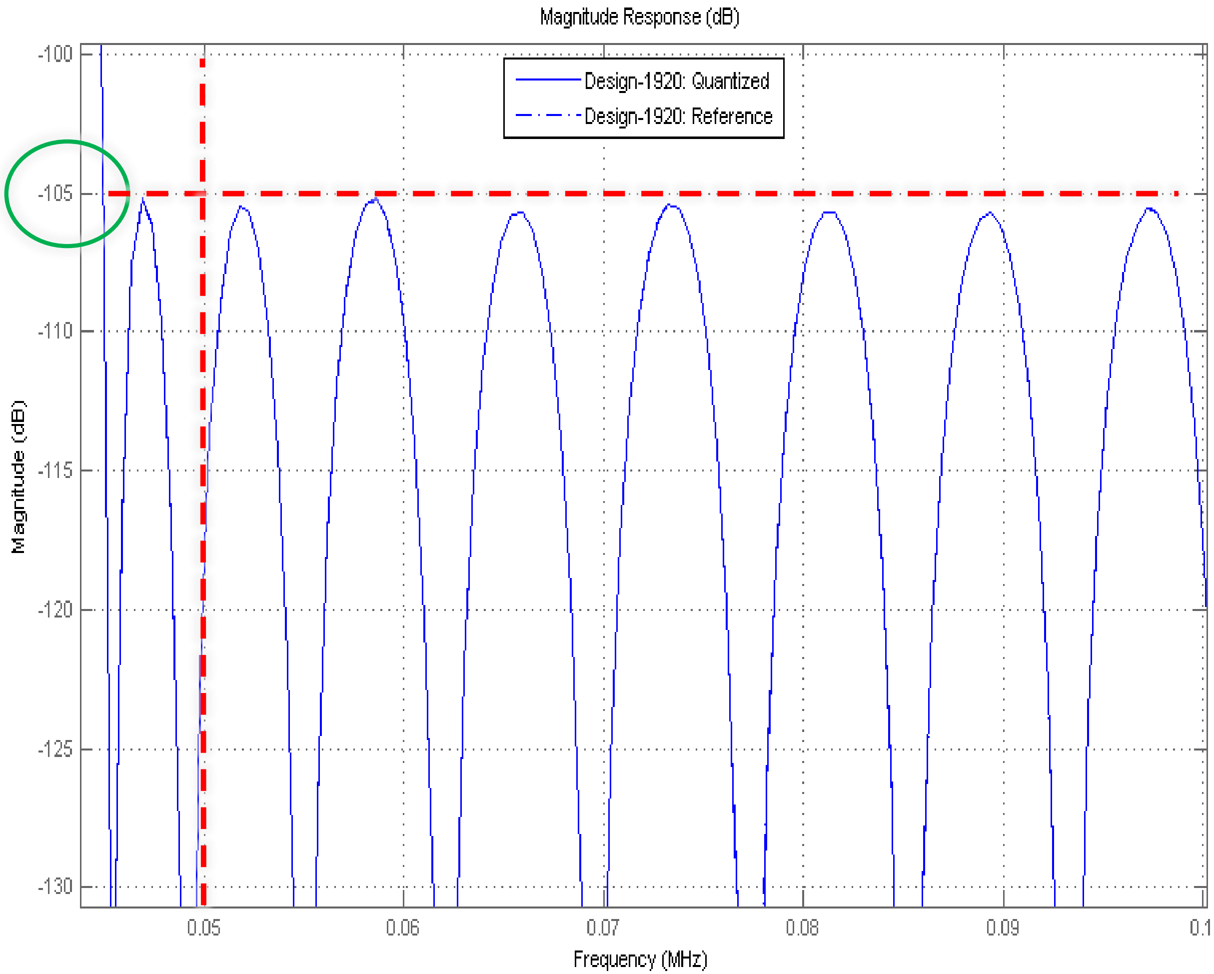

A case study for a 10 ppm Σ-Δ at 50 kS/s

Designed filter for $CLBW = 5kHz$ **with** $f_c = 50kS/s$ (oversampling factor $K = 160$)

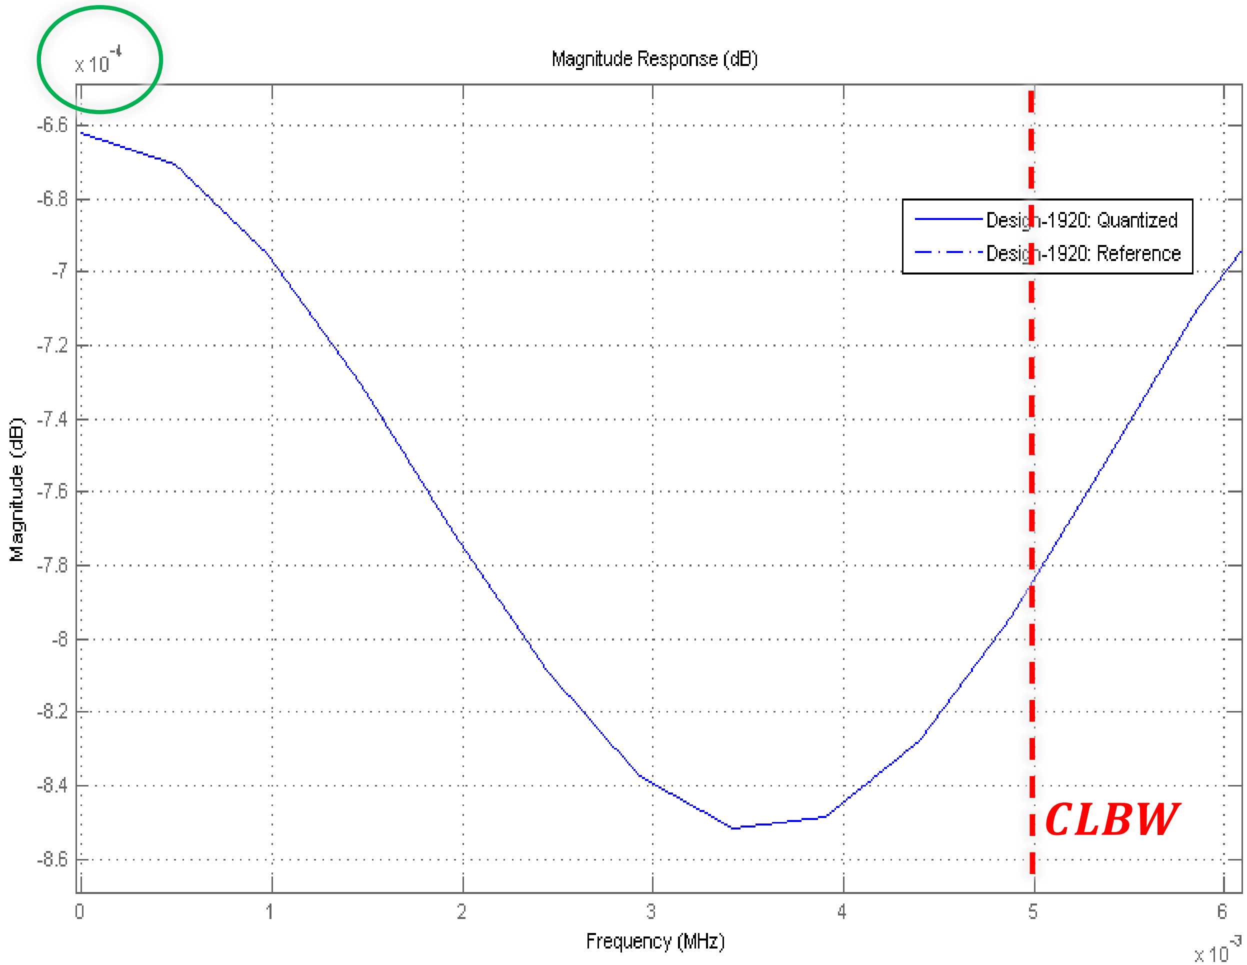

What can go wrong : an example at 6.75 MS/s

What can go wrong when applying the filter to the bitstream out of the ADS 1274 modulator ?

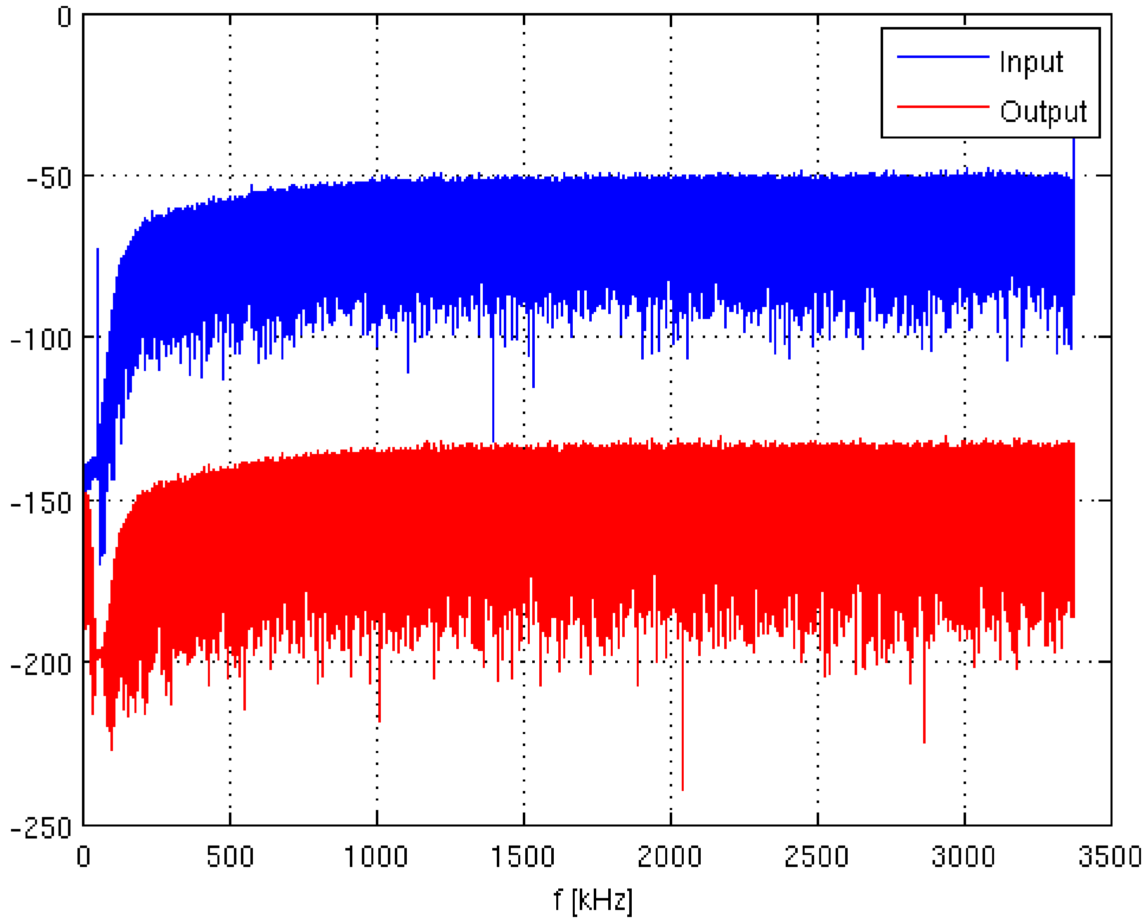

Note: the filter used is not the one presented beforehand – dB reported on arbitrary scales

What can go wrong : an example at 6.75 MS/s

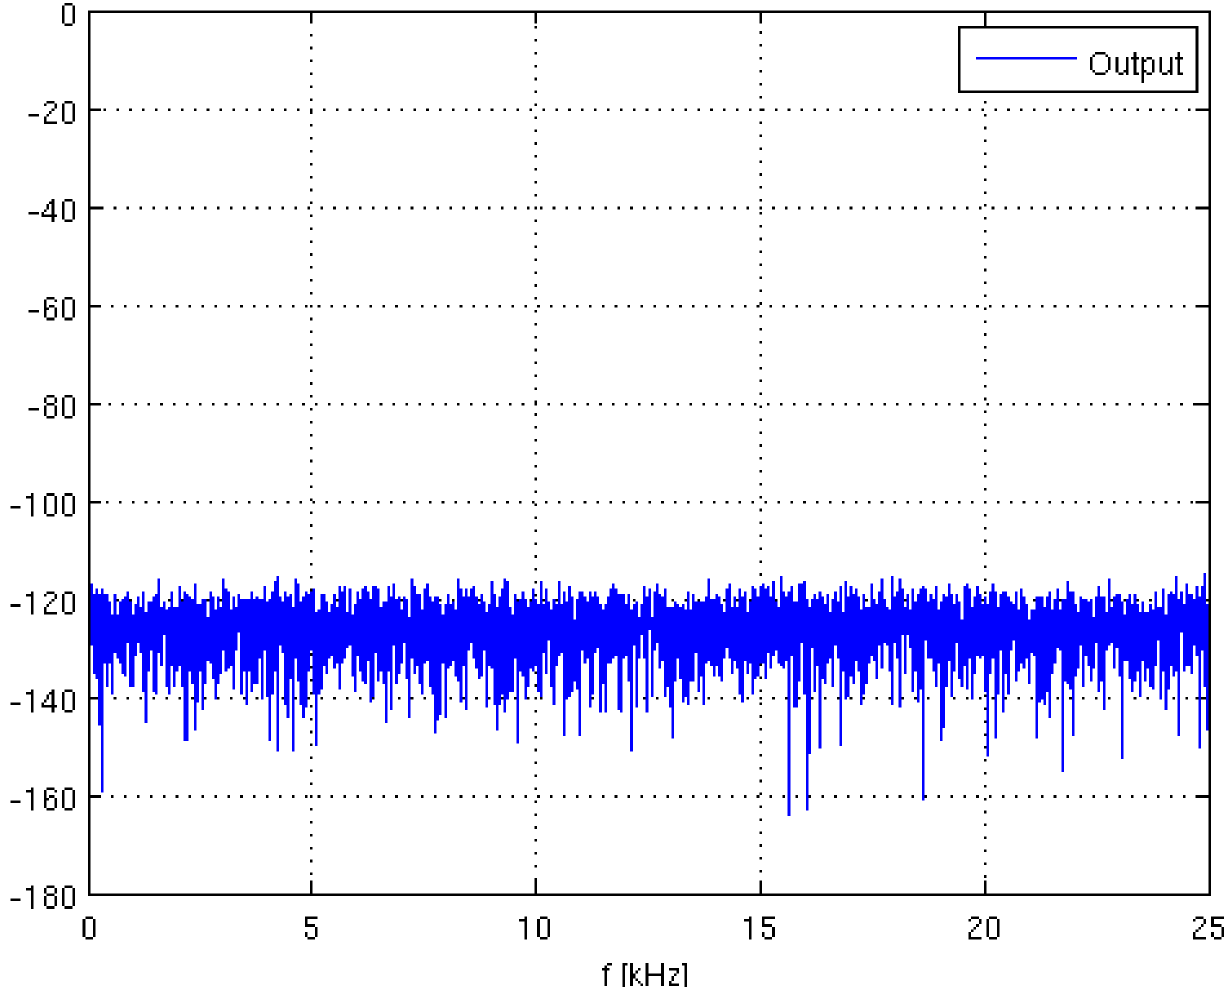

Note: the filter used is not the one presented beforehand – dB reported on arbitrary scales

The final result (after downsampling) looks very nice (it looks quite "white") isn't it?

What can go wrong : an example at 6.75 MS/s

Let's have a look at the histogram of the output:

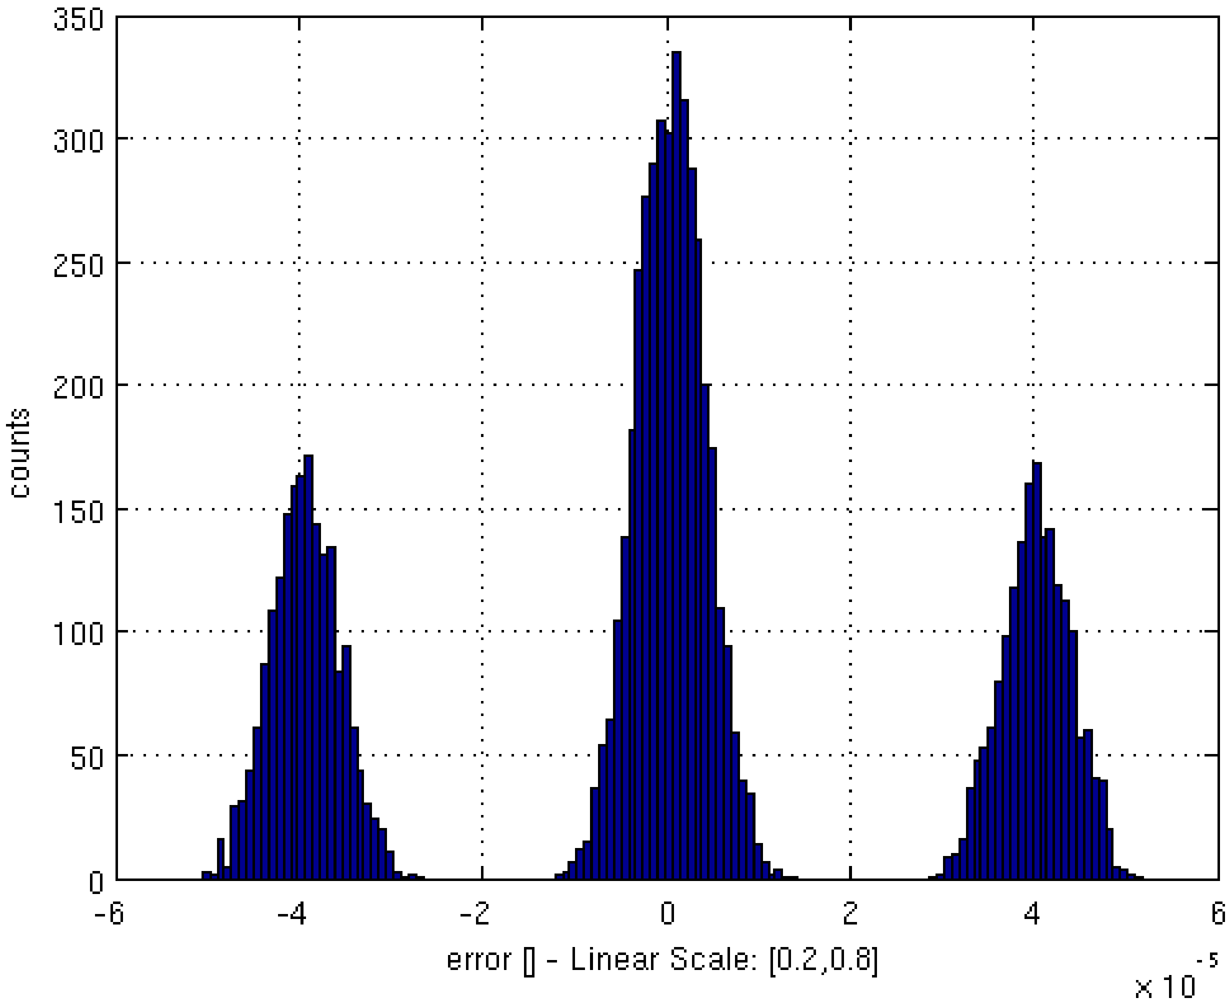

Absolute scale unimportant – assume a Full Scale of 1V

What can go wrong : an example at 6.75 MS/s

Let's have a look at the impulse response:

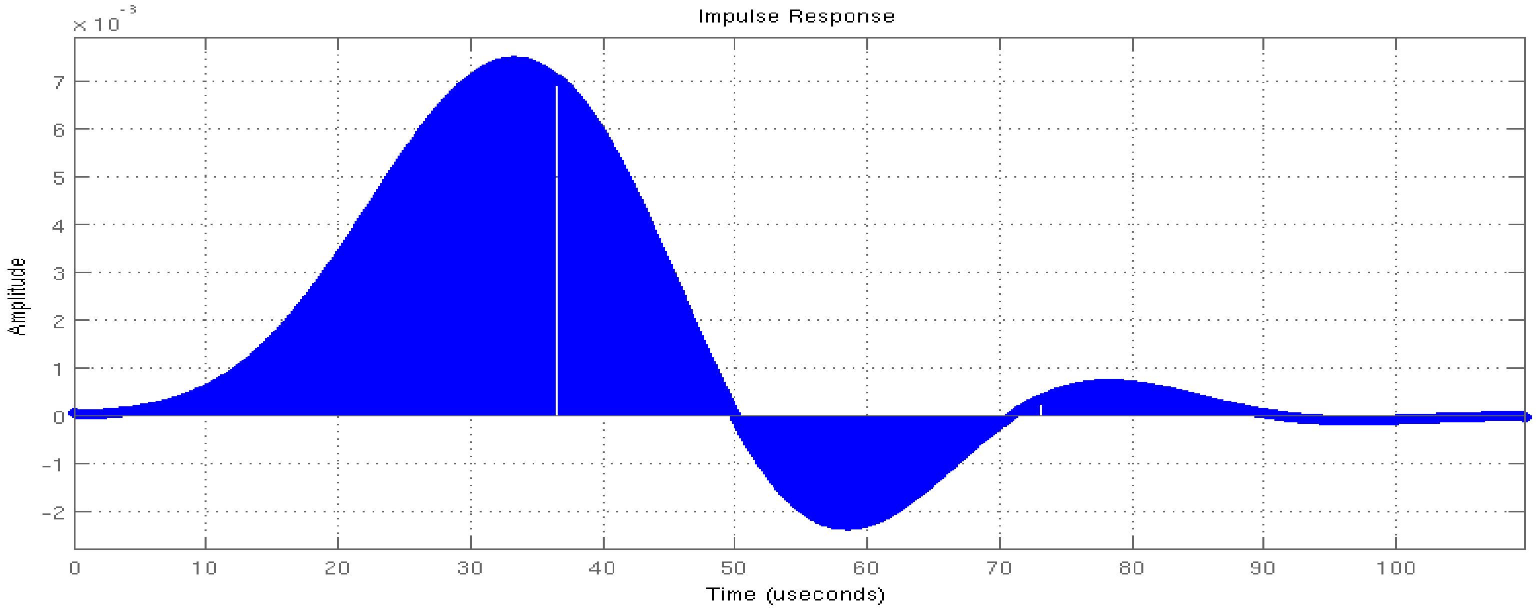

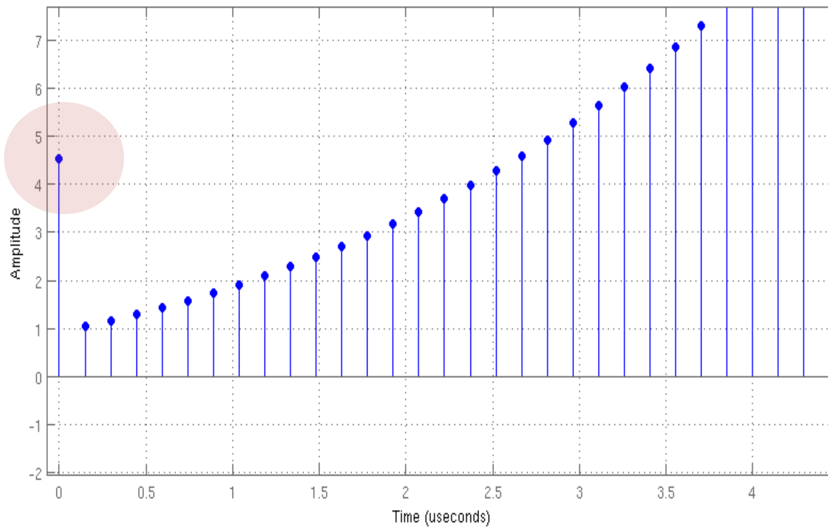

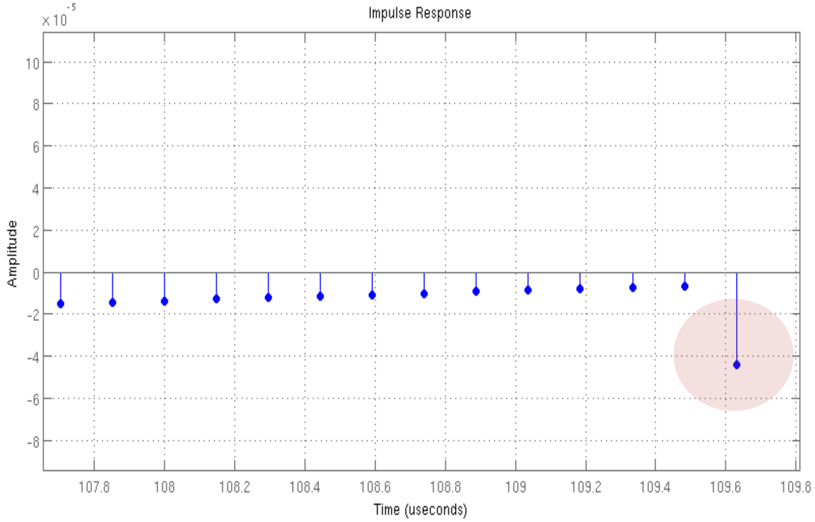

Back to the case study for a 10 ppm Σ-Δ at 50 kS/s

Actual implementation by means of 12 $\left(=\frac{1920}{160}\right)$ interleaved filters: first and last taps set to 0

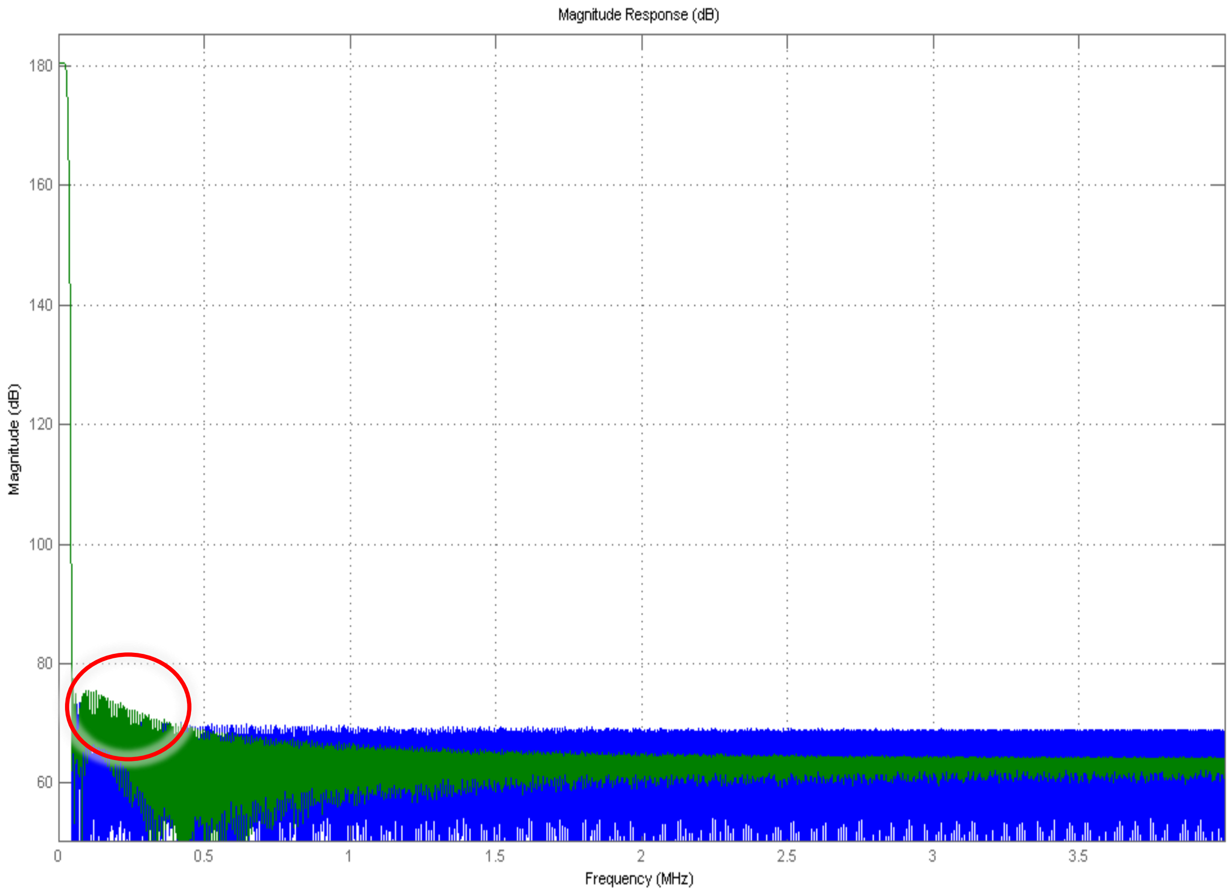

Back to the case study for a 10 ppm Σ-Δ at 50 kS/s

Actual implementation by means of 12 interleaved filters: first and last taps set to 0

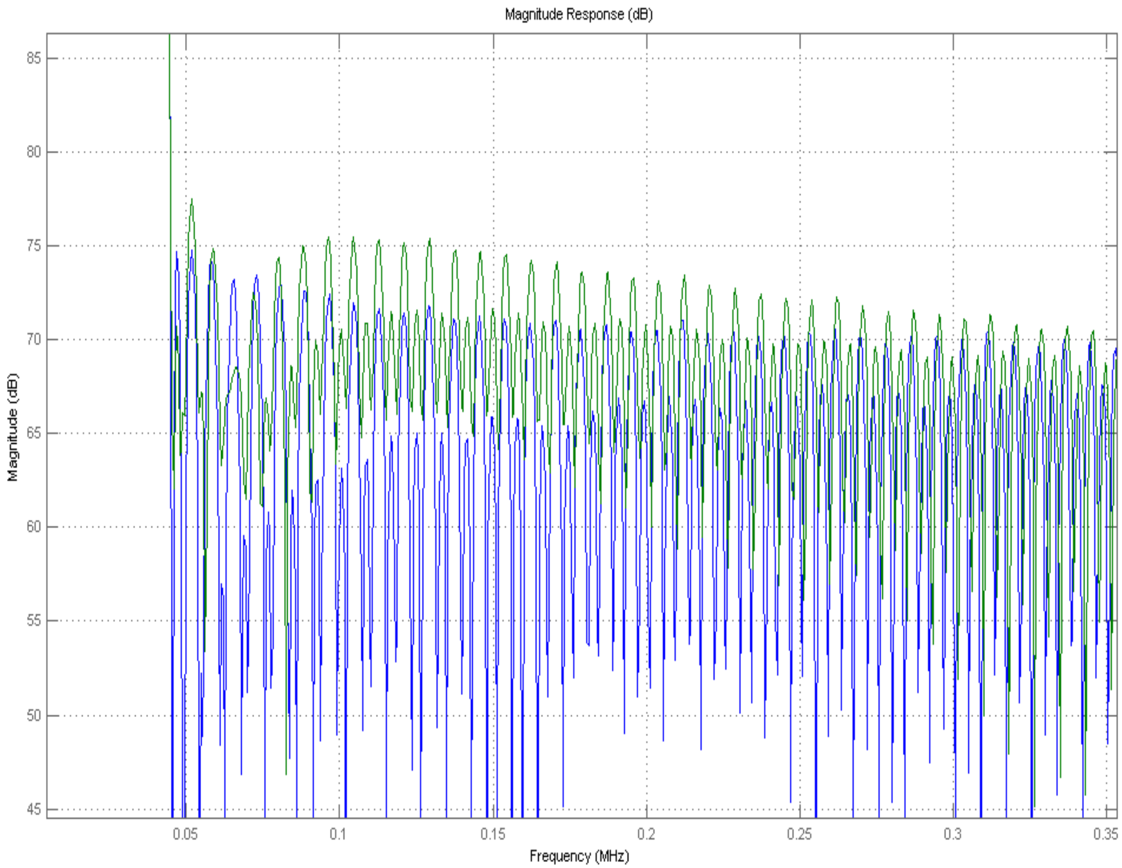

Back to the case study for a 10 ppm Σ-Δ at 50 kS/s

Noise distribution can be evaluated at a given frequency analysing the fit error of the 3-Parameter Sine-fit [8]

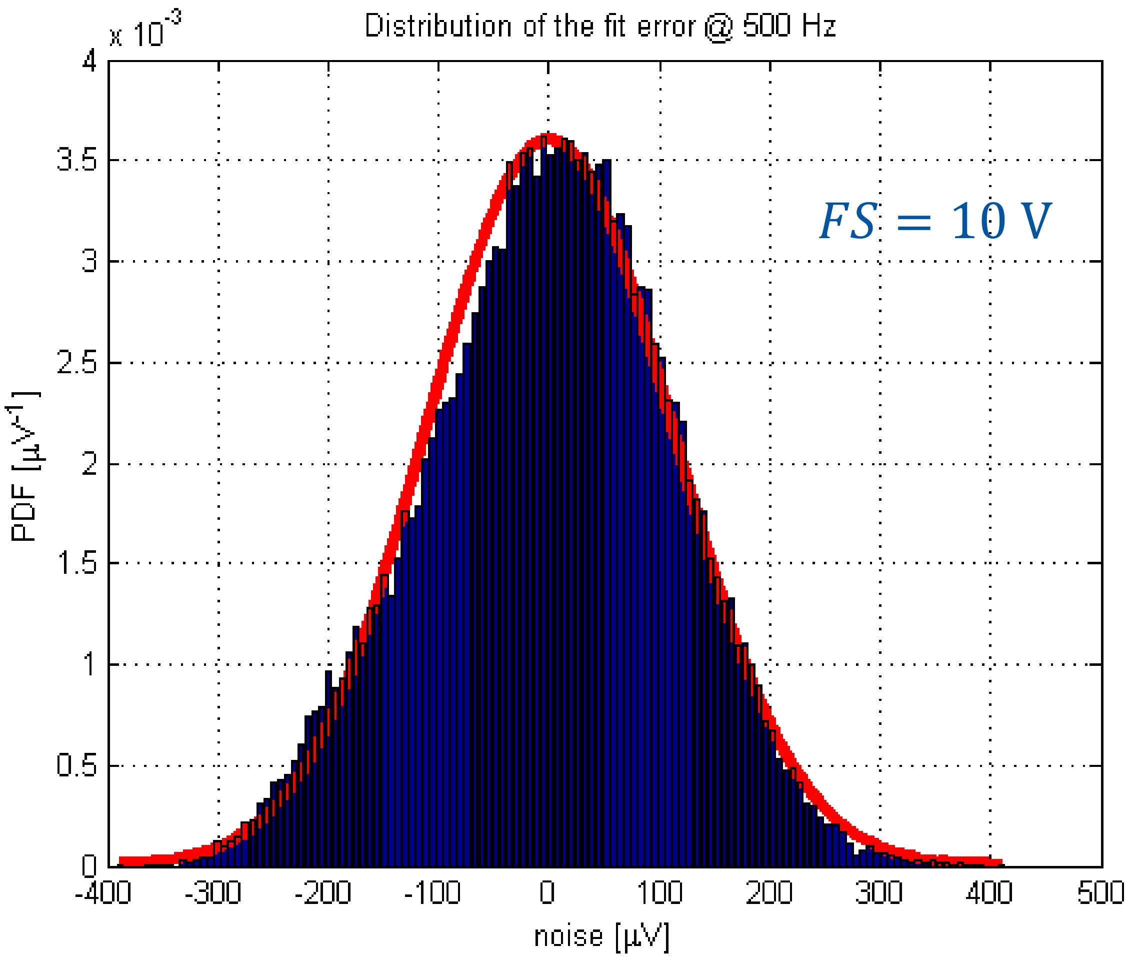

It turned out that clock had to be "corrected" by $-6.75ppm$ of the nominal 8 MHz for the data to make sense!

Bibliography

[1] IEEE Std 1241-2010 Standard for Terminology and Test Methods for Analog-to-Digital Converters

[2] IEEE Std 1057-2017 Standard for Digitizing Waveform Recorders

[3] G. Stein : Respect the unstable - IEEE Control System Magazine, 2003

[4] Digital Filters with MATLAB® : Ricardo A. Losada, 2009, The MathWorks, Inc.

[5] http://www.mit.bme.hu/books/quantization (by Bernard Widrow and Istvan Kollar)

[6] Charles M. Rader : The Rise and the Fall of Recursive Digital Filters – IEEE Signal Processing Magazine, Nov 2006

[7] R. Allan Belcher : ADC Standard IEC 60748-4-3 : Precision Measurement of Alternative ENOB Without a Sine Wave - IEEE Transactions on Instrumentation and Measurement, 2015

[8] Martino et. al : Analytical metrological characterization of the three-parameter sine fit algorithm - ISA Transactions, 2012

www.cern.ch

Oversampling - Recap

Standard Nyquist Sampling

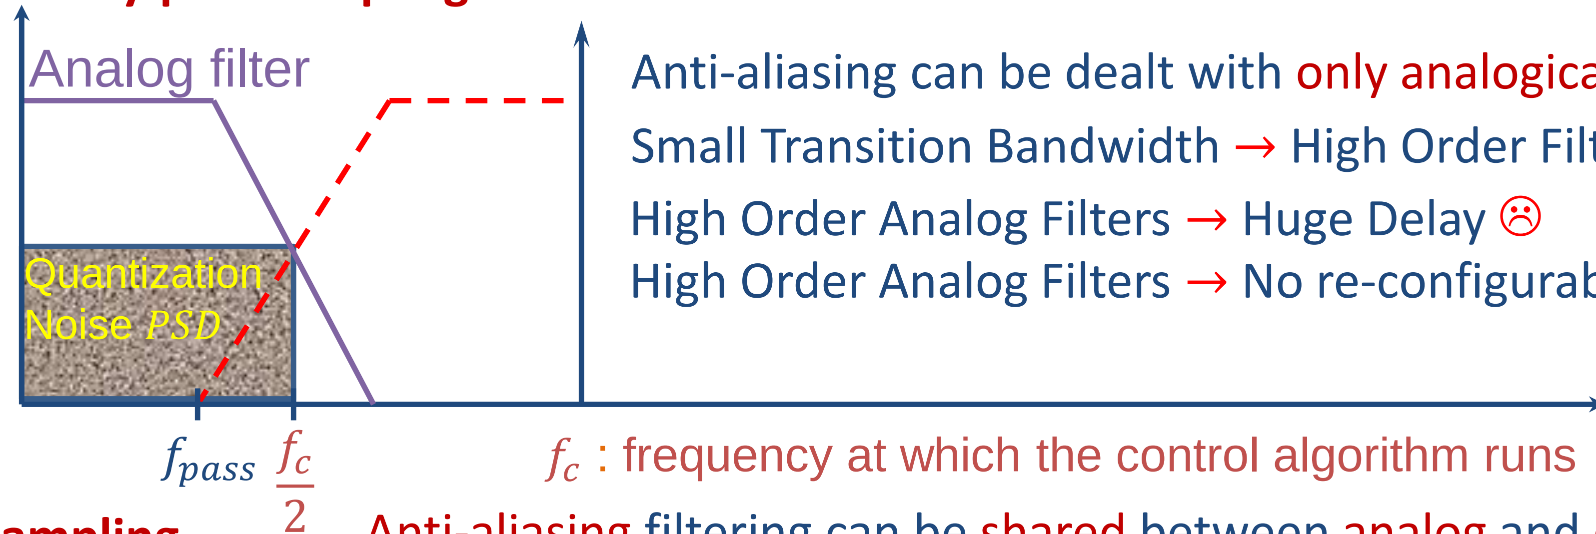

Anti-aliasing can be dealt with only analogically!

Small Transition Bandwidth → High Order Filter

High Order Analog Filters → Huge Delay ☹

High Order Analog Filters → No re-configurability ☹

f_c : frequency at which the control algorithm runs

Oversampling

Anti-aliasing filtering can be shared between analog and digital !

"Quantization" Noise PSD decreases of a factor K

$SQNR$ increases by a factor K, $SQNR_{bit}$ by $\frac{1}{2}\log_2 K$

Re-configurability now possible

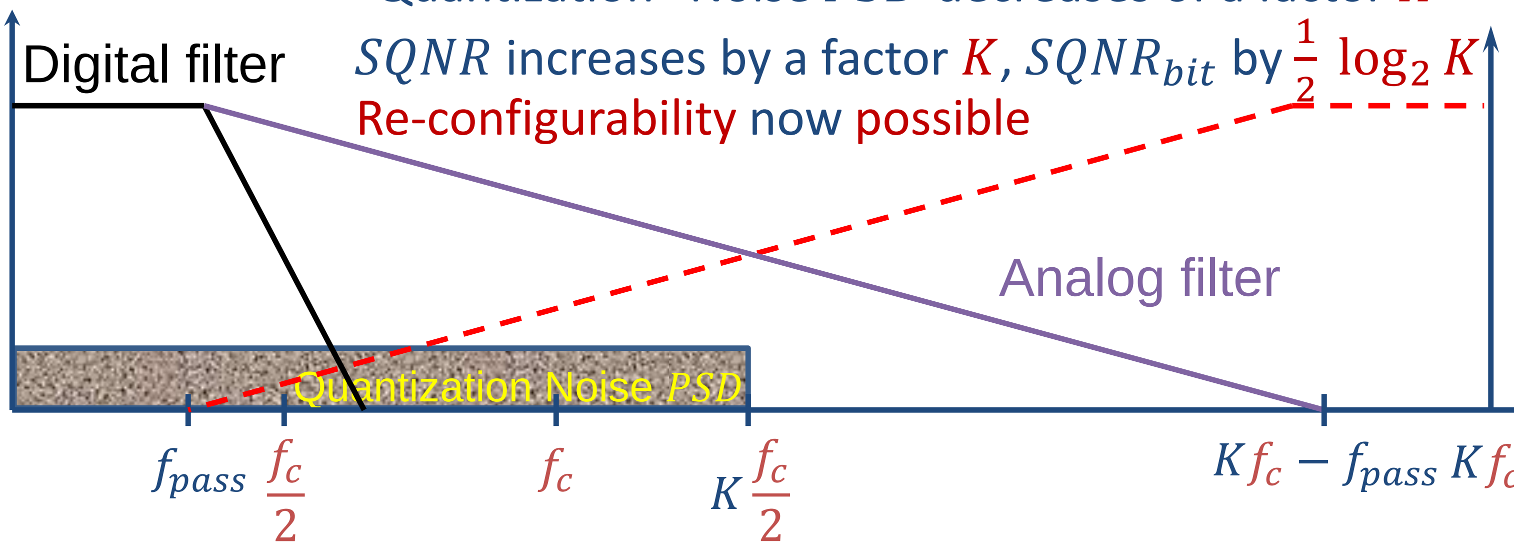

Specification of (digital) filters - Recap

Target precision of $x\ ppm$:

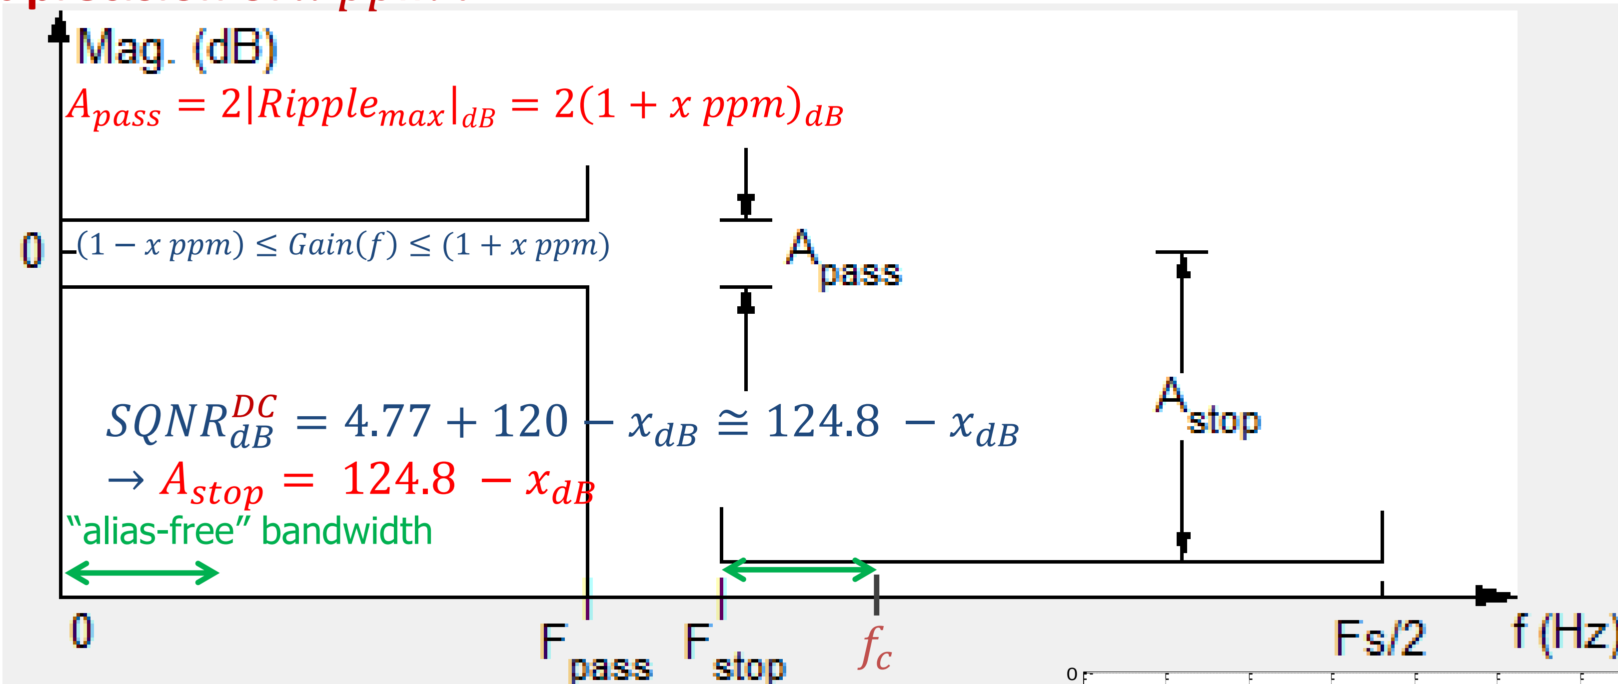

Idle tones can be close to FS so a conservative $SQNR_{dB}^{DC} = -20log_{10}(\sqrt{3}\ x\ ppm)$ has to be guaranteed by the filter!

In terms of equivalent resolution $ENOB^{DC} = \frac{SQNR_{dB}^{DC} - 4.77}{6.02}$

$F_{stop} = f_c$ - "alias-free" bandwidth

"alias-free" bandwidth $\leq CLBW$ of the PC : no overspecs

$F_{pass} \geq CLBW$: minimizing delay

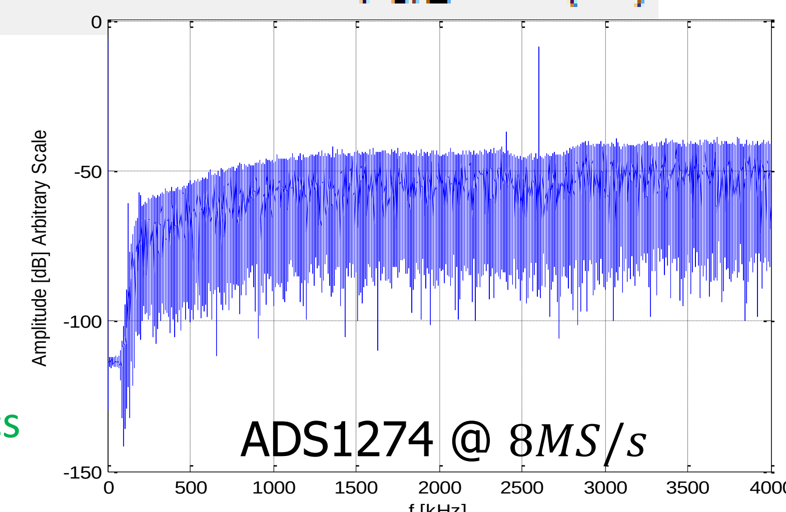

Additional slides

A case study for a 10 ppm Σ-Δ at 50 kS/s

AC performance: is it possible to reach the nominal flatness in the pass-band ?

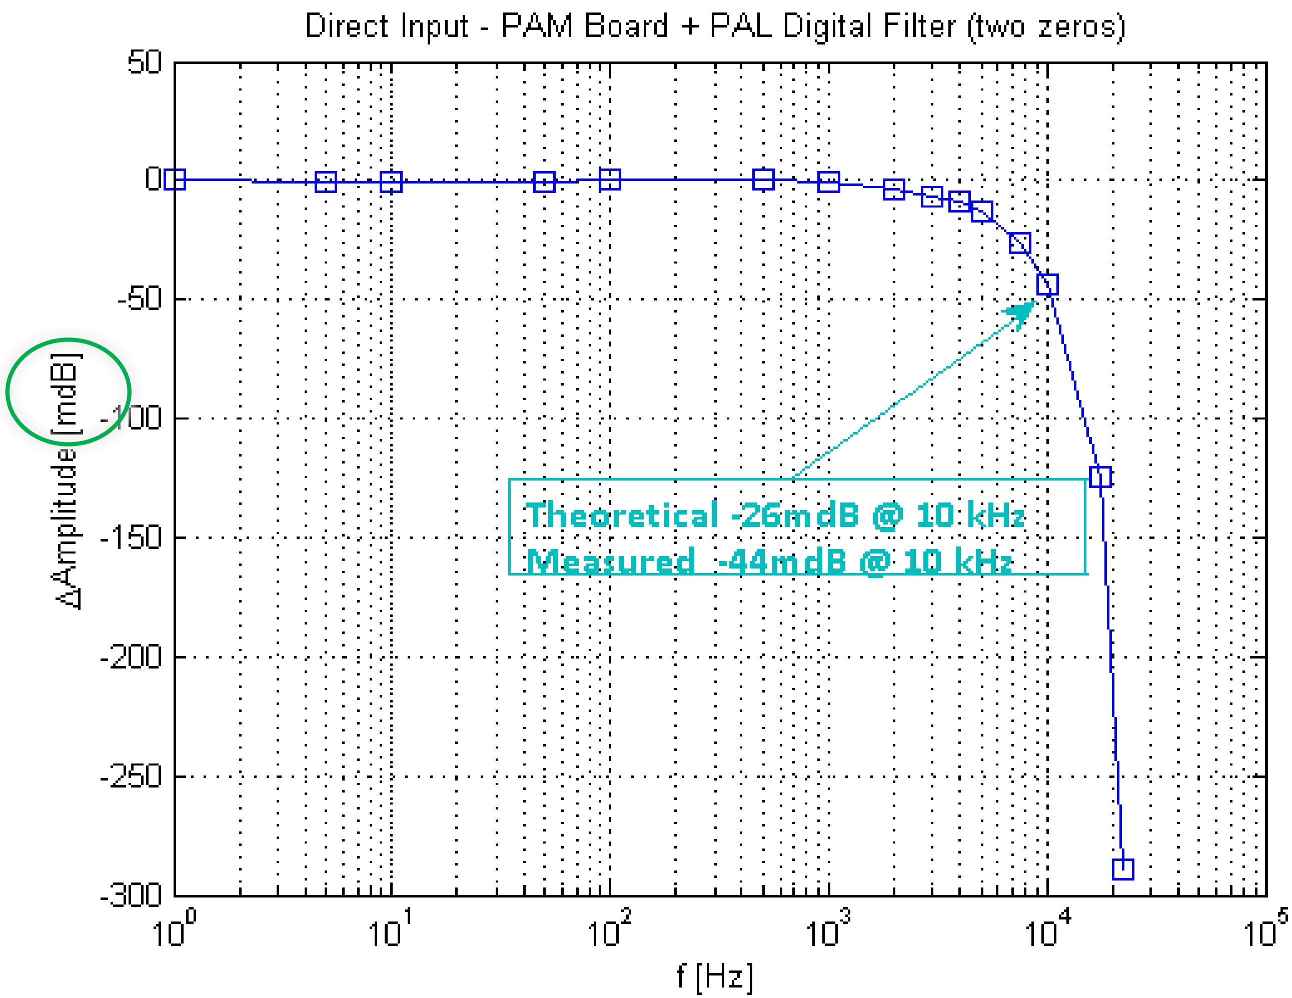

A case study for a 10 ppm Σ-Δ at 50 kS/s

- **AC performance: flatness in the pass-band**

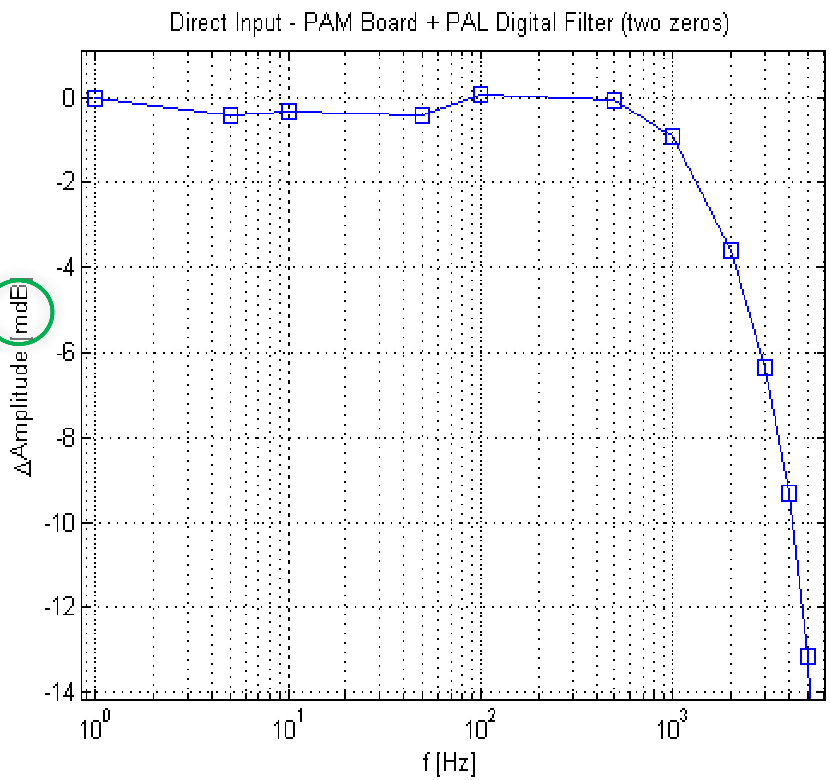

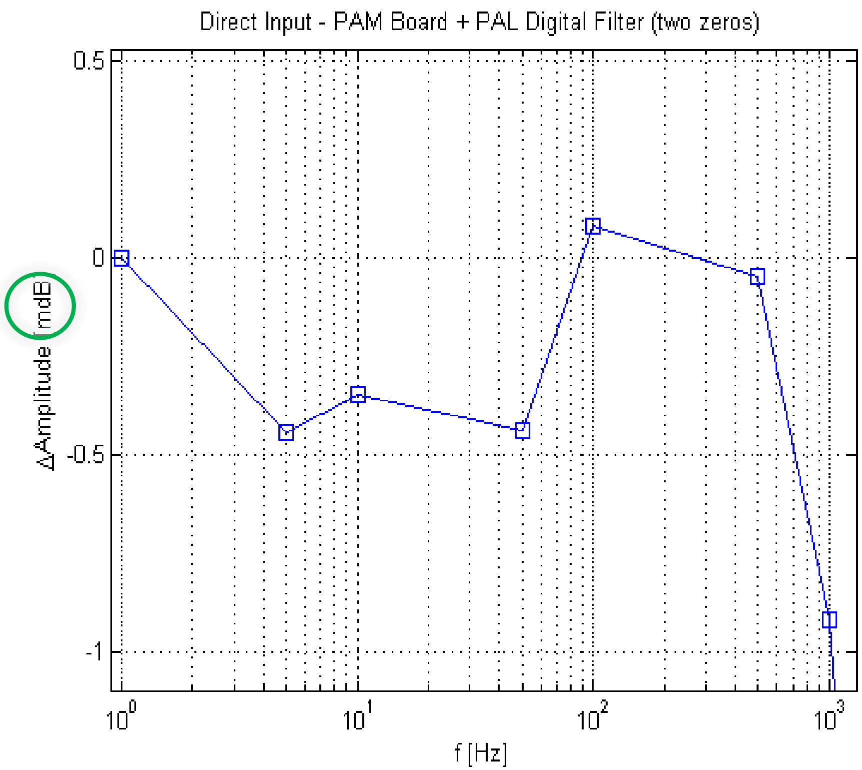

Theoretically it should be: $A_{pass}\,[dB] \leq 1.74 \times 10^{-4}$

- **Making hardware as flat as the digital filter or perform such measurements may turn out to be unfeasible or unworthy!**

A case study for a 10 ppm Σ-Δ at 50 kS/s

DC performance $FS = 1\,V$