Oversampled ADCs

- One-bit quantization
  – Quantization noise shaping
- A first-order, 1-Bit sigma-delta modulator
  – The name … sigma-delta, delta-sigma, ΣΔ, ΔΣ, …
  – Time domain model
  – Small-signal model
  – Oversampling

Oversampling

- Nyquist rate ADCs
  – Sample at f_s around 2x bandwidth
  – Resolution set by number of decision levels of quantizer
- Oversampled ADCs
  – Sample at f_s >> bandwidth (16 … 500x)
  – Use few quantization levels (typical 1-Bit)
  – Employ DSP to reduce quantization error

One-Bit Quantization

- Let's examine some properties of one bit random sequences
  - Values in the sequence are constrained to be either +1 or –1
- By picking +1's and –1's at random (using MATLAB's rand.m random number generator), we generate a sequence with zero mean
  - The sequence values model outputs of a hypothetical 1-Bit, 1MHz ADC
  - Nothing stops us from doing a DFT of this sequence …

Zero Mean 1-Bit DFT

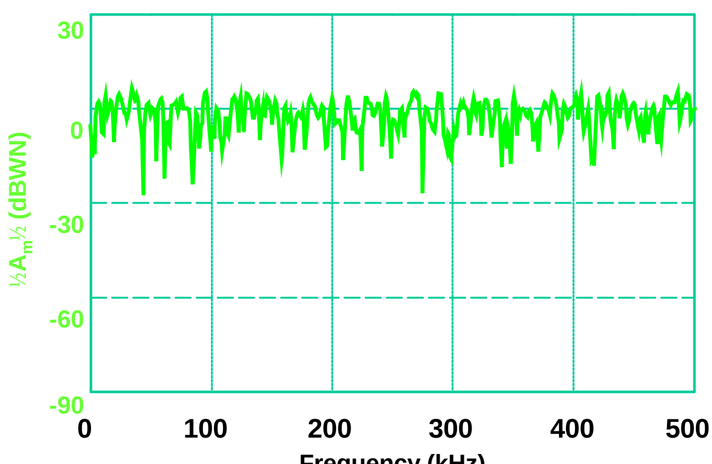

One-Bit DFTs

- For 1-Bit DFT plots, we'll use a different normalization scheme
  - The dBWN (dB White Noise) scale sets the 0dB line at the noise/bin of a random +1, -1 sequence
  - From the energy theorem,

$$\sum_{n=0}^{N-1} |a_n|^2 = N = \frac{1}{N}\sum_{m=0}^{N-1} |A_m|^2 \Rightarrow |A_m| = \sqrt{N}$$

$(+1)^2$ or $(-1)^2 = 1$

Zero Mean 1-Bit DFT

average of 30 spectra

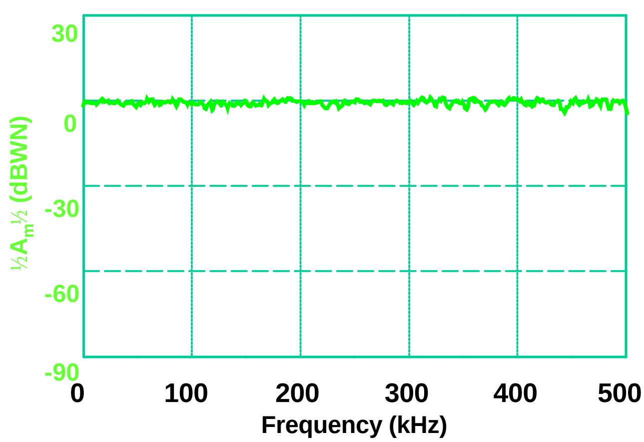

Non-zero Mean Sequences

- Of course, 1-Bit sequences can represent dc inputs from –1 to +1
- For example, an average value of +1/11 can be generated by a sequence with
  - 5/11 probability of –1
  - 6/11 probability of +1
- Let's look at DFTs of a non-zero mean sequence …

+1/11 Mean 1 Bit DFT

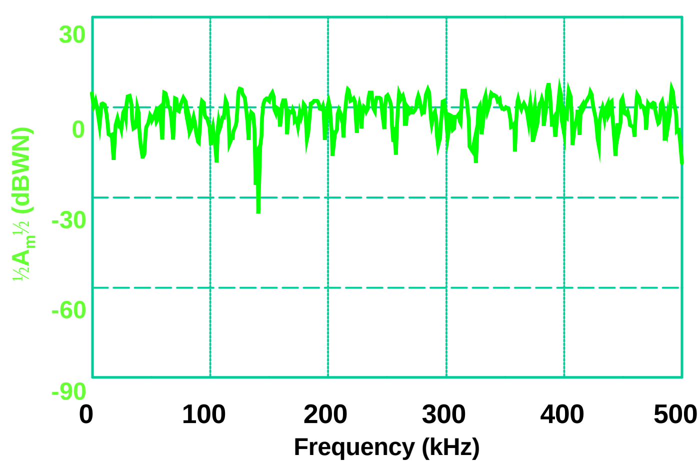

+1/11 Mean 1-Bit DFT

average of 30 spectra

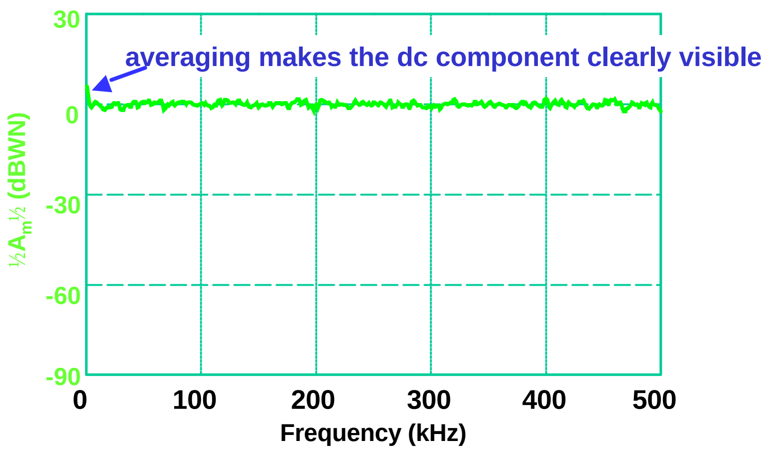

+1/11 Minimum Error Sequence

- No random number generator is required to produce the 1-Bit sequence which represents +1/11 with the minimum mean-squared quantization error
- This 11 term sequence averages to 1/11:
  [–1 +1 –1 +1 –1 +1 –1 +1 –1 +1 +1]
  – The sequence in []'s repeats
  – Its DFT follows …

+1/11 Minimum Error DFT

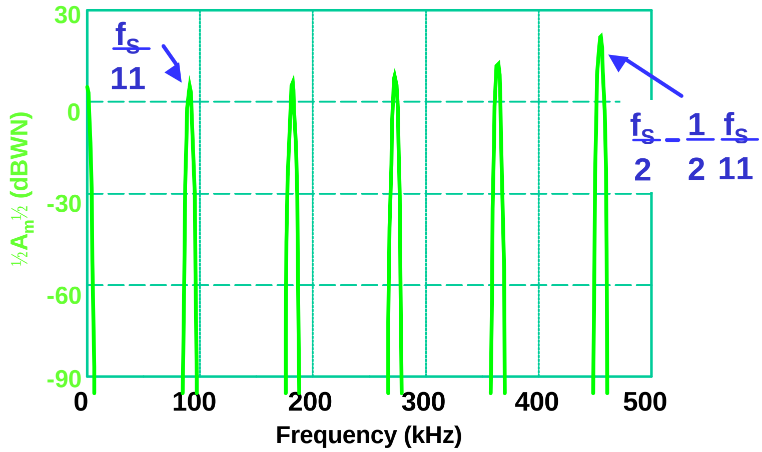

+1/11 Minimum Error Sequence

- Minimum error is periodic error with a period of $f_s/11$
  - Note that the fundamental term in this Fourier series is the smallest (a bit unusual)
- A quantization noise model is completely inappropriate for this type of sequence
- Let's deliberately increase the quantization error a little and attack its periodicity …

3 Pattern Options

- Generate another 1-Bit sequence by concatenation of the following sequences:
  - [–1 +1 –1 +1 –1 +1 –1 +1 +1]
  - [–1 +1 –1 +1 –1 +1 –1 +1 –1 +1 +1]
  - [–1 +1 –1 +1 –1 +1 –1 +1 –1 +1 –1 +1 +1]

- Selection of the above 9, 11, and 13 term sequences at random yields an average of +1/11 and the following DFT …

3 Pattern Options DFT

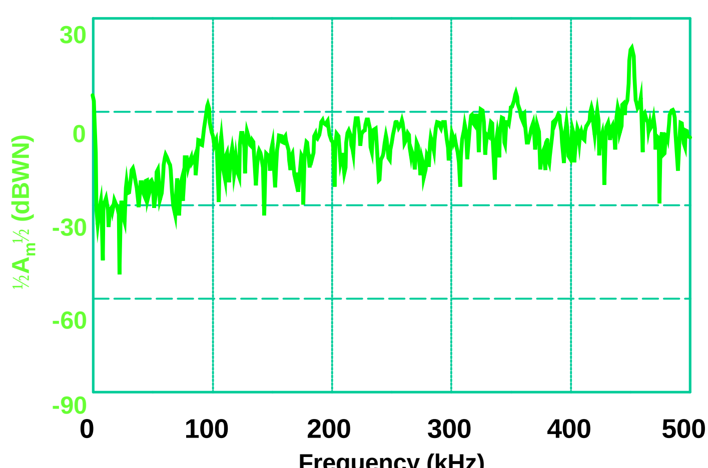

3 Pattern Options

- The randomized concatenation of longer patterns certainly breaks up periodic error
  - Some concentration of energy near $f_s/11$ and its harmonics (especially $5f_s/11$) is still visible
- Noise below 50kHz is significantly lower than that of the sample-by-sample random sequences of slide 9
  - The “noise shaping” obtained with 3 pattern options benefits low frequencies at the expense of increased quantization error at high frequencies

3 Pattern Options DFT

average of 30 spectra

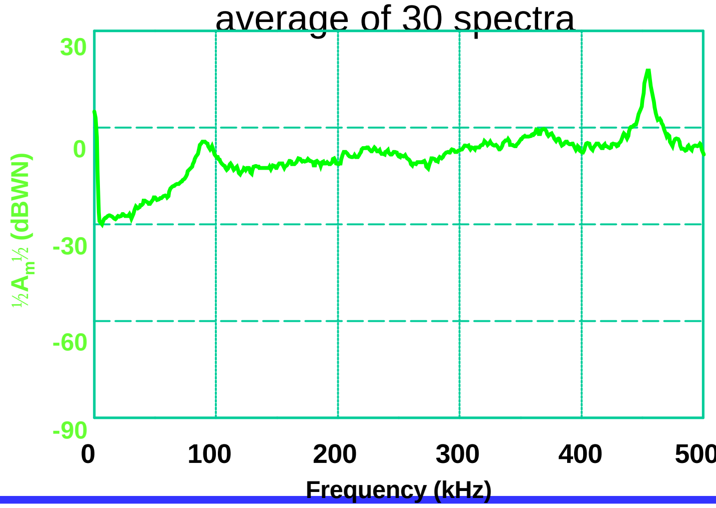

5 Pattern Options

- Let's add two more patterns to our set of pattern options:
  - [–1 +1 –1 +1 –1 +1 –1 +1 –1 +1 –1 +1 –1 +1 +1]
  - [–1 +1 –1 +1 –1 +1 +1]
- Selection of 7, 9, 11, 13, and 15 term sequences at random preserves the +1/11 average
  - Think of this process as a sort of time-variant dither
  - DFT follows …

3/5 Pattern Options

average of 30 spectra

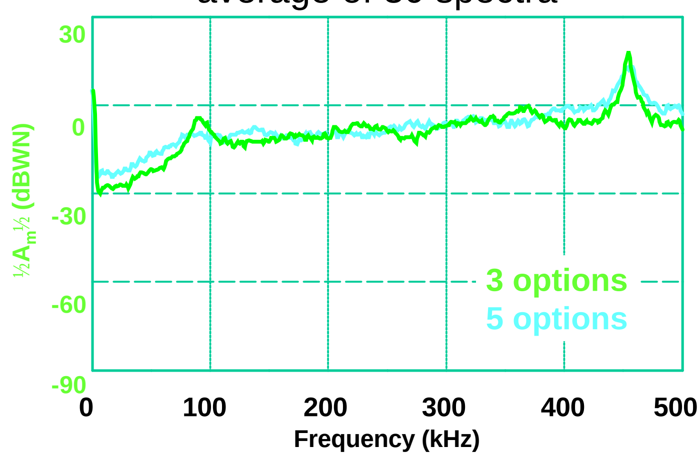

5 Pattern Noise Shaping

- Noise vs. frequency still follows the general upward tilt of the periodic error harmonics
  - 5 pattern options further reduce the $f_s/11$ bump
  - The $5f_s/11$ component remains large
- The model that quantization error is uncorrelated with the input signal becomes reasonable with only 5 pattern options
  - Such reasonableness is required for small-signal analysis of the sigma-delta modulator …

Sigma- Delta Modulators

Analog 1-Bit ΣΔ modulators convert a continuous time analog input v_{IN} into a 1-Bit sequence d_{OUT}

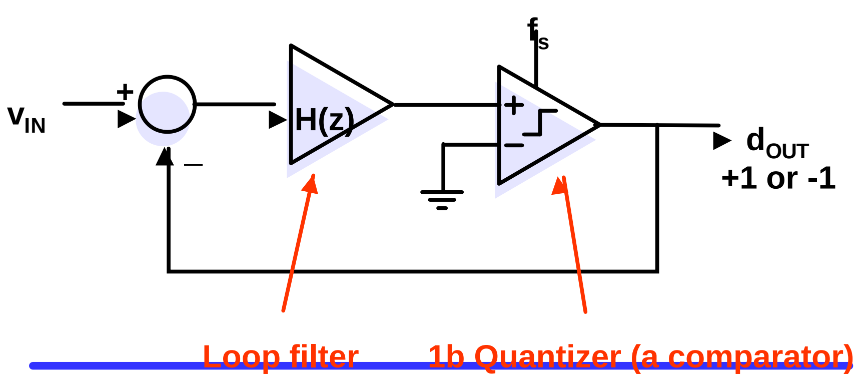

Sigma-Delta Modulators

- The loop filter H can be either a SC or continuous time
- SC's are "easier" to implement and scale with the clock rate
- Continuous time filters provide anti-aliasing protection
- Can be realized with passive LC's at very high frequencies

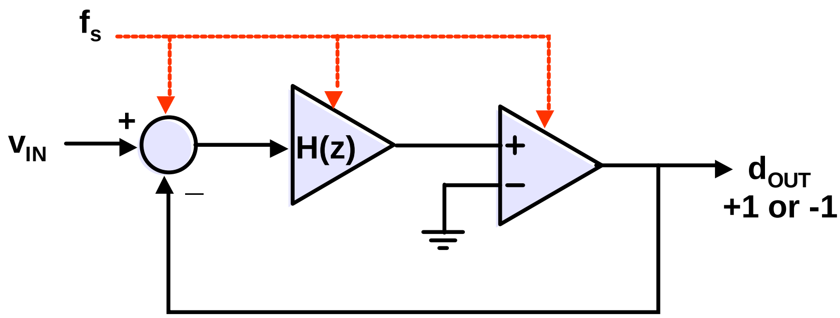

1st Order ΣΔ Modulator

In a 1st order modulator, the loop filter is an integrator

$$H(z) = \frac{z^{-1}}{1 - z^{-1}}$$

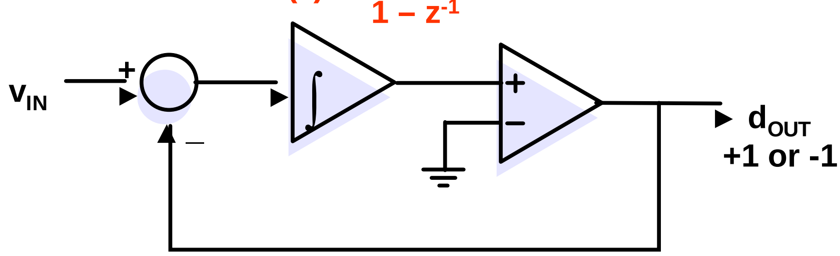

1st Order ΔΣ Modulator

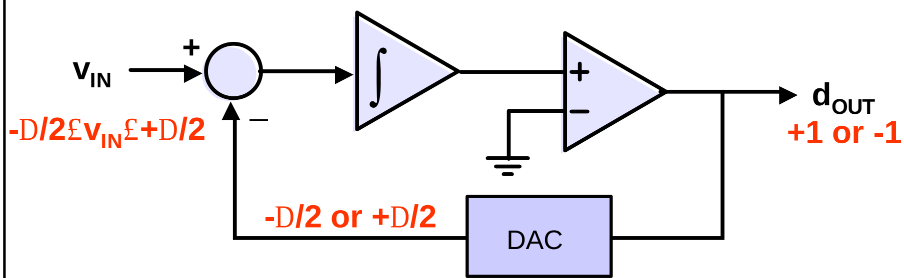

- Properties of the first-order modulator:
  - Analog input range is the d_{out} times the DAC reference
  - The average value of d_{OUT} must equal the average value of v_{IN}
  - +1's (or –1's) density in d_{OUT} is an inherently monotonic function of v_{IN}
    → linearity is not dependent on component matching
  - Alternative multi-bit DAC (and ADCs) solutions reduce the quantization error but loose this inherent monotonicity

Simulation

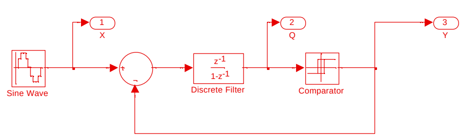

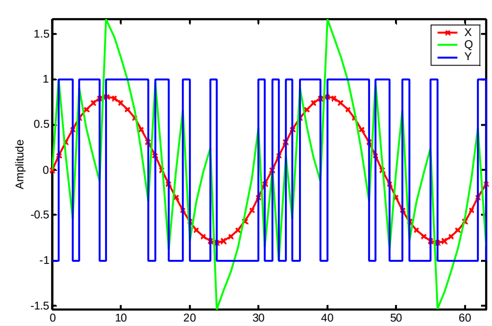

see L19_level1

1st Order ΣΔ, +1/11 dc Input

- Let's continue our +1/11 dc input example with the modulator sampling frequency of 1MHz

- A 1024 sample DFT plot appears on the following slide…

1st Order ΣΔ, +1/11 dc Input

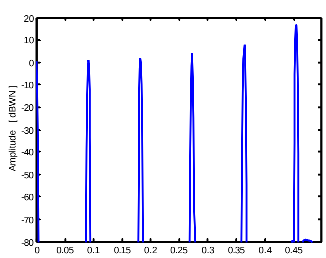

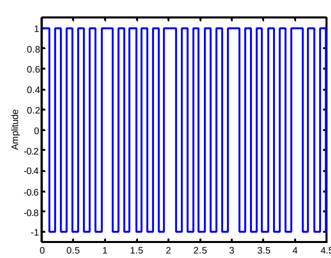

Looks familiar?

1st Order ΣΔ, +1/11 dc Input

- A check of the time samples confirms the obvious: the 1st order ΔΣ modulator produces the minimum error 11 term sequence:

  [–1 +1 –1 +1 –1 +1 –1 +1 –1 +1 +1]

- Of course, with this modulator model we can look at much more interesting inputs than dc…

1st Order ΣΔ, Sinewave Input

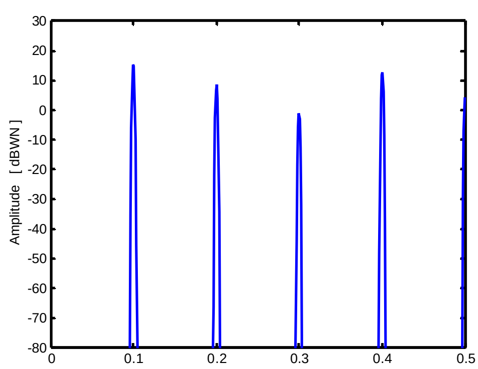

$v_{IN}(k) = 0.99\sin(2\pi\ 0.100001\ t)$

d_{OUT} = [+1 –1 +1 +1 +1 +1 –1 –1 –1 –1 …]

1st Order ΣΔ, Sinewave Input

- The modulator output for the undistorted sinewave input produces huge distortion, suggesting the need for dither
- We'll add a dither signal q at the comparator input:

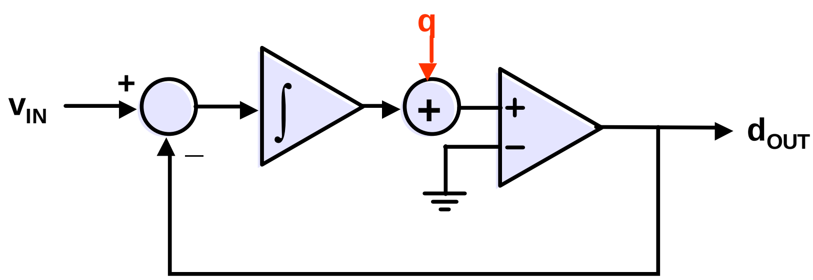

Dithered 1st Order ΣΔ

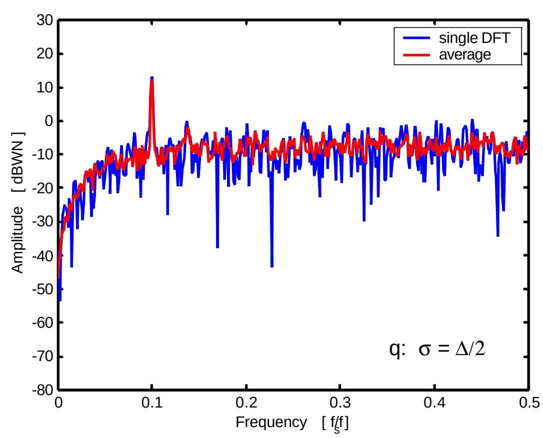

Dithered 1st Order ΣΔ

- 1Vrms Gaussian dither is difficult if not impossible to produce on mixed-signal ICs
  - On-chip digital circuitry adds excessive non-Gaussian interference to analog noise generators
- First-order modulators are too prone to limit cycles to be of much practical use
  - They do provide the basis for higher-order ΣΔ's

1st Order Noise Shaping

If q(k) and v_{IN}(k) are uncorrelated, we can compute the signal and noise transfer functions independently:

$$STF(z) = \frac{D_{out}(z)}{V_{in}(z)} \qquad NTF(z) = \frac{Q(z)}{V(z)_{in}}$$

How do we model the quantizer?

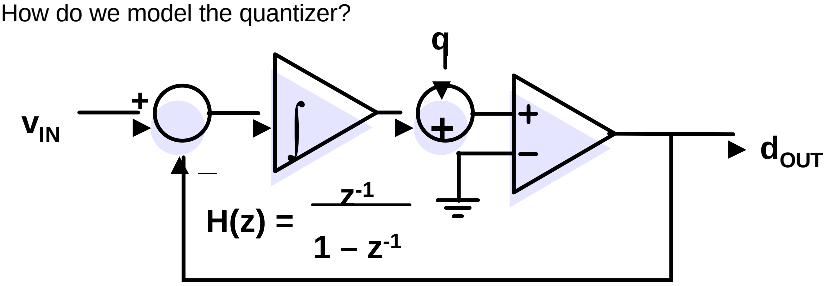

1^{st} Order Noise Shaping

- Nonlinear elements like comparators are frequently modeled by some sort of linearized “effective gain”, G
- One measure of effective gain that’s proven itself useful for ΣΔ analysis is:

$$G \ {}^{o}\ \frac{\textbf{rms value of the comparator output}}{\textbf{rms value of the comparator input}}$$

- The value of G depends on the input signal and can be determined with simulation
  - E.g. for the simulation in slide 30, G=0.7 (-3.1dB)

1^{st} Order Noise Shaping

- The input q models both dither, if added, and the input-referred noise of the comparator

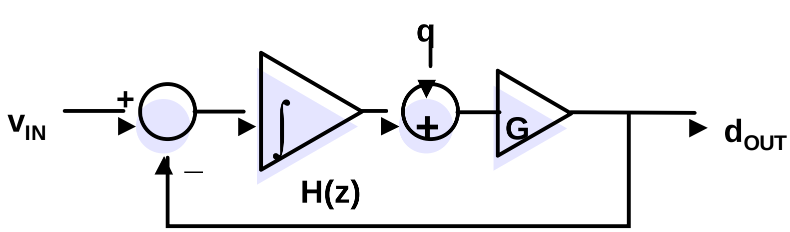

1st Order Noise Shaping

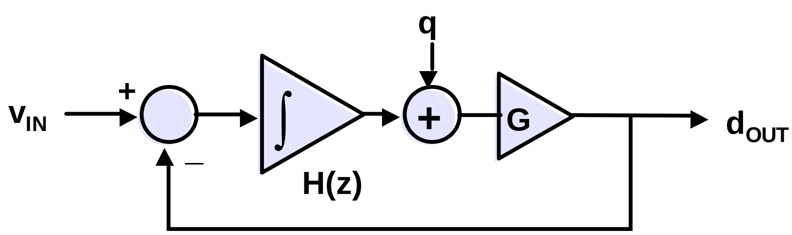

$$STF = \frac{D_{out}}{V_{in}} = \frac{GH(z)}{1+GH(z)} \qquad NTF = \frac{Q}{V_{in}} = \frac{G}{1+GH(z)}$$

When H(z) is large, this is approximately 1.

1st Order Noise Shaping

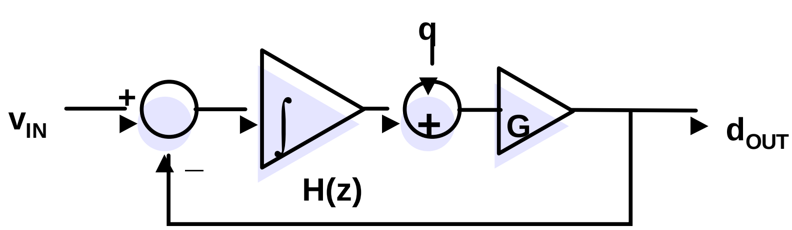

$$STF = \frac{D_{out}}{V_{in}} = \frac{GH(z)}{1+GH(z)} \qquad NTF = \frac{Q}{V_{in}} = \frac{G}{1+GH(z)}$$

For large H(z), this is »**0**.

Noise Transfer Function, NTF

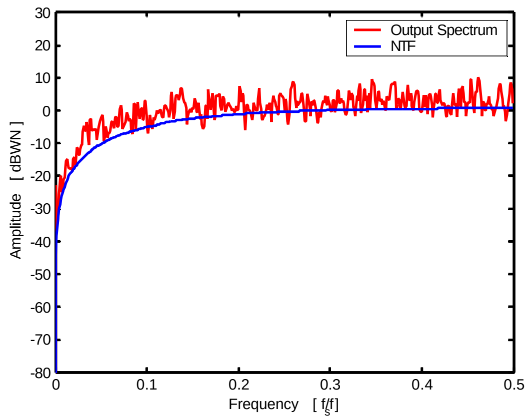

"Integrated" Noise

- Just as we did for thermal noise, let's look at the integrated noise at the output of the modulator
- In the discrete time case, noise integrals are summations, but the result is called "integrated noise" nonetheless

Noise Summations

- Integrated noise is computed from the energy theorem (N even):

$$a_{rms} = |A_0| / N \ , \ M=0$$

$$a_{rms} = \frac{1}{N}\sqrt{|A_0|^2 + \sum_{m=1}^{M} 2|A_m|^2} \ , \ 0<M<N/2$$

$$a_{rms} = \frac{1}{N}\sqrt{|A_0|^2 + |A_{N/2}|^2 + \sum_{m=1}^{N/2-1} 2|A_m|^2} \ , \ M=N/2$$

Integrated Noise

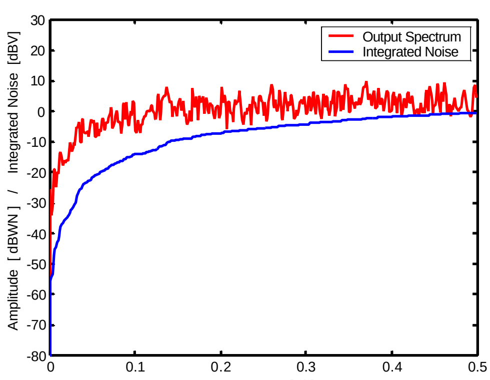

The total "noise" at the modulator output with no input sums to 0dB
→ This is consistent with a binary signal

Integrated Output

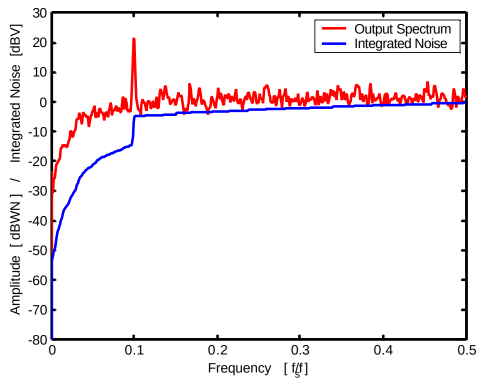

The total power still sums to 0dB

Only little "quantization noise" at low frequency

Oversampling and Noise Shaping

- ΣΔ modulators have interesting characteristics
  - Unity gain for the the input signal V_{IN}
  - Large attenuation of quantization noise injected at q
  - Much better than 1-Bit noise performance is possible if we're only interested in frequencies << f_s
- ADCs which sample their inputs at much higher frequencies than the Nyquist rate minimum are called "oversampling ADCs"

Oversampling and Noise Shaping

- Higher order loop filters consisting of several integrators provide much better noise shaping than 1st order realizations
- They are also less prone to limit cycles or need less dither
- If a 1-Bit DAC is used, the converter is inherently linear—independent of component matching
- References

J. C. Candy and G. C. Temes, “Oversampling Methods for A/D and D/A Conversion”, Oversampling Delta-Sigma Data Converters: Theory, Design, and Simulation, 1992, pp. 1-25.

S. R. Norsworthy, R. Schreier, and G. C. Temes, “Delta-Sigma Data Converters, Theory, Design, and Simulation,” IEEE Press, 1997.

Estimating Quantization Noise

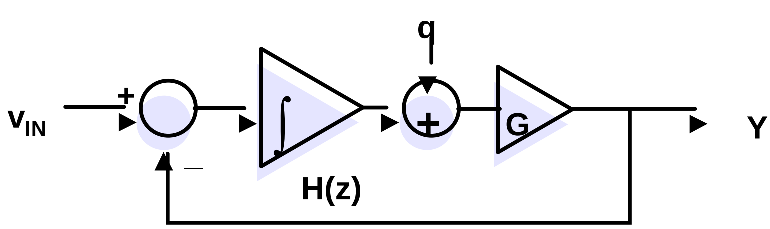

$$S_Q(f) = \frac{1}{f_s}\frac{\Delta^2}{12}$$

$$NTF(z) = \frac{G}{1+GH(z)}$$

$$\overline{S_Y} = \int_{-B}^{B} S_Q(f)\left|NTF(z)\right|^2_{z=e^{2\pi jfT}} df$$

Example: 1st Order Modulator

$$H(z) = \frac{z^{-1}}{1 - z^{-1}}$$

$$G = 1$$

$$NTF(z) = 1 - z^{-1}$$

$$\begin{aligned}
|NTF(z)|^2 &= NTF(z)NTF(z^{-1}) \\
&= (1 - z^{-1})(1 - z) \\
&= 1 - z^{-1} - z + 1 \\
&= 2 - 2\cos wT \\
&= 2\sin pfT
\end{aligned}$$

$$\begin{aligned}
\overline{S_Y} &= \int_{-B}^{B} S_Q(f) |NTF(z)|^2_{z=e^{2pjfT}} \, df \\
&\cong \int_{-f_s/M}^{f_s/M} \frac{1}{f_s} \frac{\Delta^2}{12} (2\sin pfT)^2 \, df \\
&\approx \frac{p^2}{3} \frac{1}{M^3} \frac{\Delta^2}{12}
\end{aligned}$$

Example: Dynamic Range

$$DR = \frac{\text{peak signal power}}{\text{peak noise power}} = \frac{\overline{S_X}}{\overline{S_Y}}$$

$$\overline{S_X} = \frac{1}{2}\left(\frac{\Delta}{2}\right)^2 \qquad \text{sinusoidal input,} \quad STF = 1$$

$$\overline{S_Y} = \frac{p^2}{3} \frac{1}{M^3} \frac{\Delta^2}{12}$$

$$DR = \frac{9}{2p^2} M^3$$

M	**DR**
16	33 dB
32	42 dB
1024	87 dB

Dynamic Range

- DR increases 9dB for each doubling of M
- 1st order modulators require very high M for >10-Bit resolution
  → higher order filters improve this tradeoff substantially
- Analysis is based on assumption that the quantization noise is "white"
  → not true in practice, especially for low-order modulators
  → practical modulators suffer from other noise sources also (e.g. thermal noise)
- Next time we'll design an oversampled audio ADC with better than 16-Bit resolution