EE247
Lecture 23

- ADC figures of merit
- Oversampled ADCs
  - Why oversampling?
  - Pulse-count modulation
  - Sigma-delta modulation
    - 1-Bit quantization
      - Quantization error (noise) spectrum
      - SQNR analysis
      - Limit cycle oscillations
    - 2nd order $\Sigma\Delta$ modulator
      - Practical implementation
      - Effect of various building block nonidealities on the $\Sigma\Delta$ performance

ADC Figures of Merit

- Objective: Establish measure/s to compare performance of various ADCs

- Can use FOM to combine several performance metrics to get one single number

- What are reasonable FOM for ADCs?

ADC Figures of Merit

$$FOM_1 = f_s \cdot 2^{ENOB}$$

- This FOM suggests that adding an extra bit to an ADC is just as hard as doubling its bandwidth
- Is this a good assumption?

Ref: R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, April 1999

Survey Data

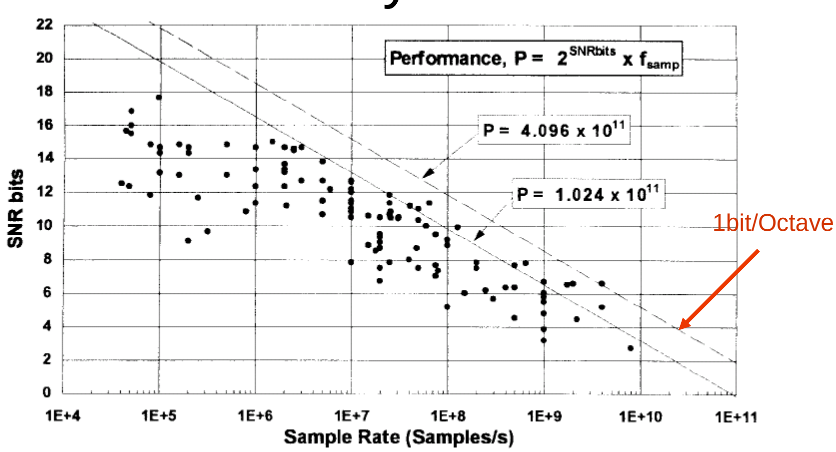

Ref: R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, April 1999

ADC Figures of Merit

$$FOM_2 = \frac{Power}{f_s \cdot 2^{ENOB}} \qquad [J/conv]$$

- Sometimes inverse of this metric is used
- In typical circuits power ~ speed, FOM_2 captures this tradeoff correctly
- How about power vs. ENOB?
  - One more bit 2x in power?

Ref: R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, April 1999

ADC Figures of Merit

- One more bit means...
  - 6dB SNR, 4x less noise power, 4x larger C
  - Power ~ Gm ~ C increases 4x
- Even worse: Flash ADC
  - Extra bit means 2x number of comparators
  - Each of them needs double precision
  - Transistor area 4x, Current 4x to keep same current density
  - Net result: Power increases 8x

ADC Figures of Merit

- FOM_2 seems not entirely appropriate, but somehow still standard in literature, papers

- "Tends to work" because:
  - Not all power in an ADC is "noise limited“
  - E.g. Digital power, biasing circuits, etc.

- Better use FOM_2 to compare ADCs with same resolution!

ADC Figures of Merit

$$FOM_3 = \frac{Power}{Speed}$$

- Compare only power of ADCs with approximately same ENOB
- Useful numbers:
  - 10b (~9 ENOB) ADCs: 1 mW/MSample/sec
    Note the ISSCC 05 example: 0.33mW/MS/sec!
  - 12b (~11 ENOB) ADCs: 4 mW/MSample/sec

10-Bit ADC Power/Speed

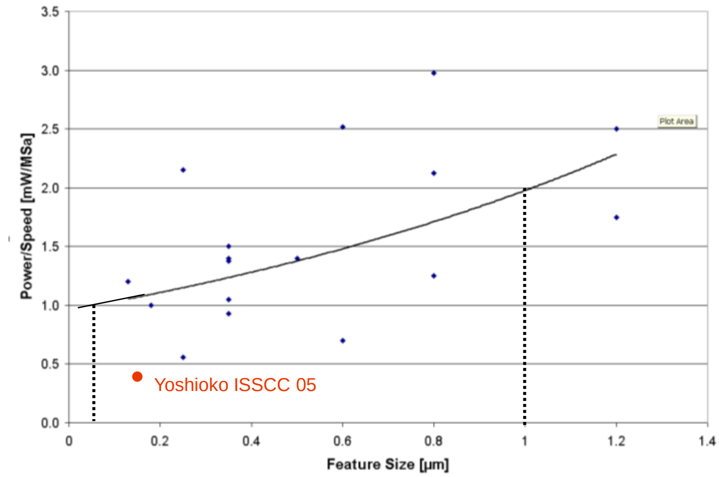

12-Bit ADC Power/Speed

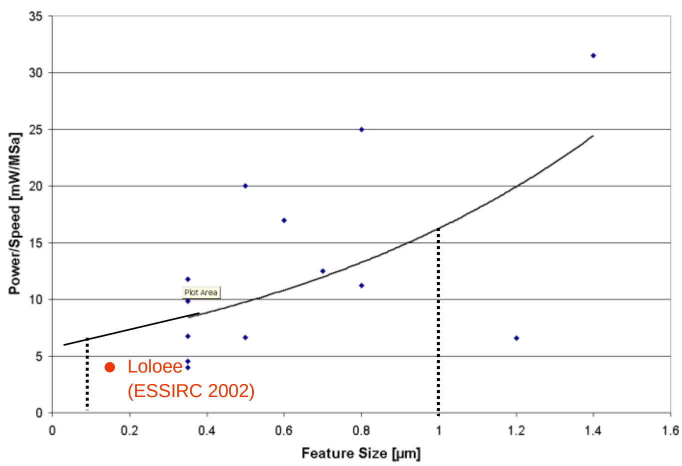

Performance Trend

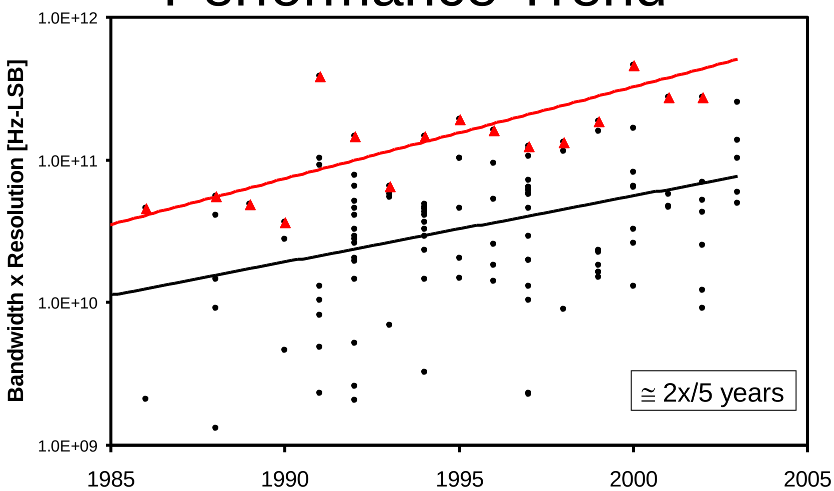

ADC Architectures

- Slope type converters
- Successive approximation
- Flash
- Interpolating & Folding
- Residue type ADCs
  - Two-step Flash
  - Pipelined ADCs
- Time-interleaved / parallel converter
- Oversampled ADCs

Analog-to-Digital Converters

- Two categories:
  - Nyquist rate ADCs → $f_{sig}^{max} \sim 0.5xf_{sampling}$
    - Maximum achievable signal bandwidth higher compared to oversampled type
    - Resolution limited to max. ~14bits
  - Oversampled ADCs → $f_{sig}^{max} << 0.5xf_{sampling}$
    - Maximum possible signal bandwidth significantly lower compared to nyquist
    - Maximum achievable resolution high (18 to 20bits!)

The Case for Oversampling

Nyquist sampling:

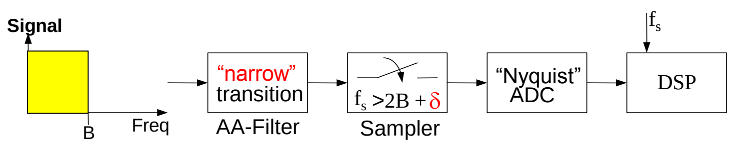

Oversampling:

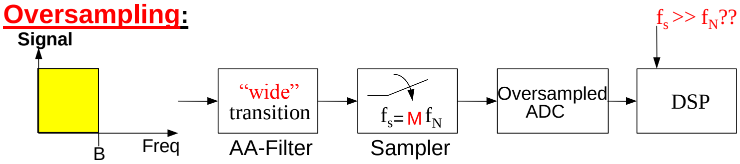

- Nyquist rate f_N ~2B
- Oversampling rate $M = f_s/f_N >> 1$

Nyquist v.s. Oversampled Converters
Antialiasing Requirements

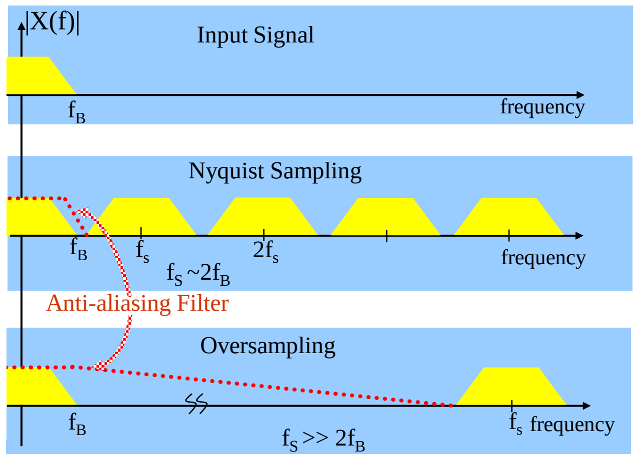

ADC Converters
Baseband Noise

- For a quantizer with quantization step size Δ and sampling rate f_s :
  - Quantization noise power distributed uniformly across Nyquist bandwidth ($f_s/2$)

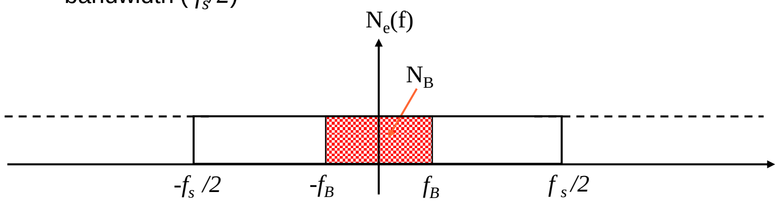

  - Power spectral density:

$$N_e(f) = \frac{\overline{e^2}}{f_s} = \left(\frac{\Delta^2}{12}\right)\frac{1}{f_s}$$

  - Noise is distributed over the Nyquist band $-f_s/2$ to $f_s/2$

Oversampled Converters
Baseband Noise

$$S_B = \int_{-f_B}^{f_B} N_e(f)\,df = \int_{-f_B}^{f_B} \left(\frac{\Delta^2}{12}\right)\frac{1}{f_s}\,df$$

$$= \frac{\Delta^2}{12}\left(\frac{2f_B}{f_s}\right)$$

$$where\ for\ f_B = f_s/2$$

$$S_{B0} = \frac{\Delta^2}{12}$$

$$S_B = S_{B0}\left(\frac{2f_B}{f_s}\right) = \frac{S_{B0}}{M}$$

$$where\ M = \frac{f_s}{2f_B} = oversampling\ ratio$$

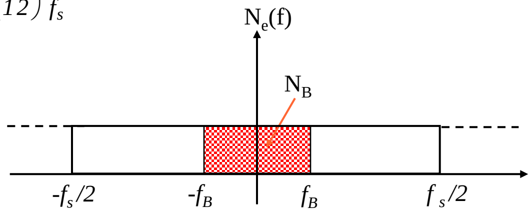

Oversampled Converters
Baseband Noise

$$S_B = S_{B0}\left(\frac{2f_B}{f_s}\right) = \frac{S_{B0}}{M}$$

$$where\ M = \frac{f_s}{2f_B} = oversampling\ ratio$$

2X increase in M
→ 3dB reduction in S_B
→ ½ bit increase in resolution/octave oversampling

To further increase the improvement in resolution:

- Embed quantizer in a feedback loop (patented by Cutler in 1960s!)
  →Noise shaping (sigma delta modulation)

Pulse-Count Modulation

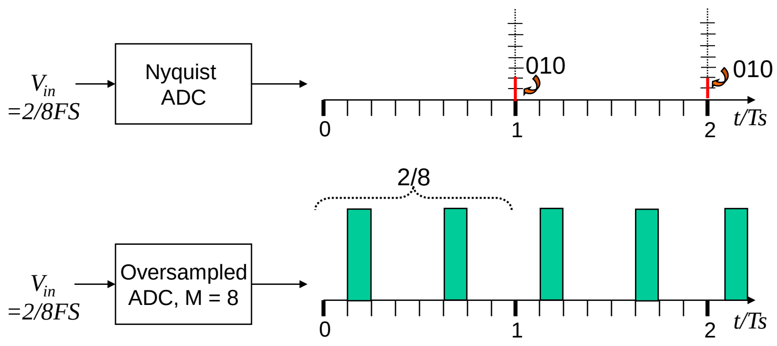

Mean of pulse-count signal approximates analog input!

Pulse-Count Output Spectrum

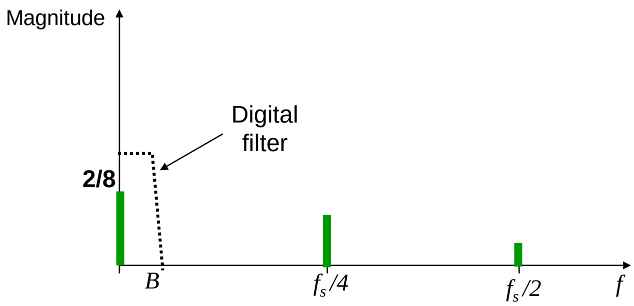

- Signal band of interest: low frequencies, f < B << f_s
- Quantization error: high frequency, B … f_s / 2
- Separate signal from Q error with digital low-pass filter!

Oversampled ADC
Predictive Coding

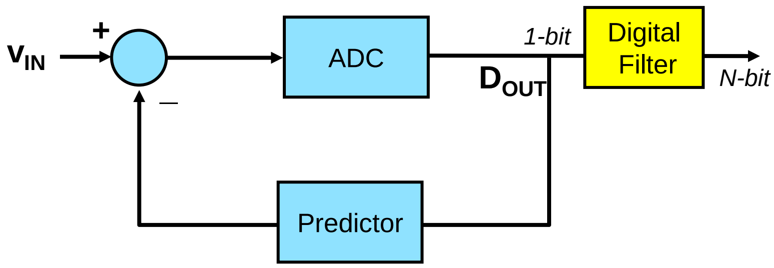

- Quantize the difference signal rather than the signal itself
- Smaller input to ADC → Buy dynamic range
- Only works if combined with oversampling
- 1-Bit digital output
- Digital filter computes "average" →N-bit output

Oversampled ADC

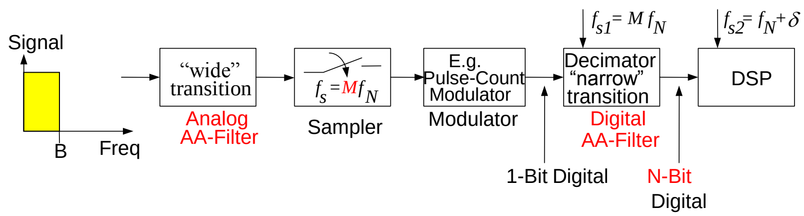

Decimator:
- Digital (low-pass) filter
- Removes quantization noise for $f > B$
- 1-Bit input @ $f_{s1}=Mf_N$, N-Bit output @ $f_{s2}=f_N+\delta$ (computes "average")
- Provides anti-alias filtering for DSP
- Narrow transition band, high-order (digital filters with high order consume significantly smaller power & area compared to analog filters)

Modulator
or Analog Front End (AFE)

- Objectives:
  - Convert analog input to 1-Bit pulse density stream
  - Move quantization error to high frequencies f >>B
  - Operates at high frequency $f_s >> f_N$
    - M = 8 … 256 (typical)….1024
    - Since modulator operated at high frequencies → need to keep analog circuitry "simple"

→ ΣΔ = ΔΣ Modulator

Sigma- Delta Modulators

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

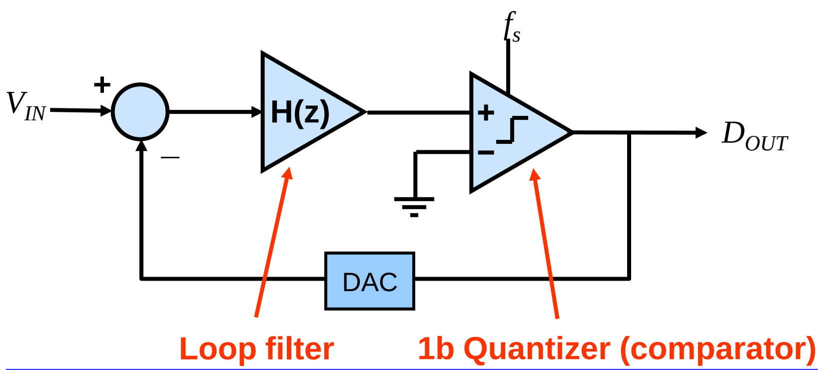

Sigma-Delta Modulators

- The loop filter H can be either switched-capacitor or continuous time
- Switched-capacitor filters are "easier" to implement + frequency characteristics scale with clock rate
- Continuous time filters provide anti-aliasing protection

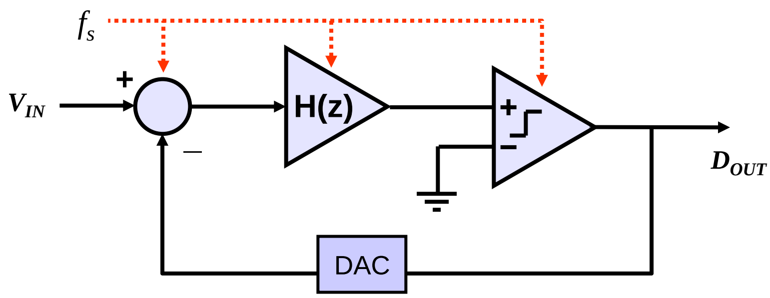

Oversampling A/D Conversion

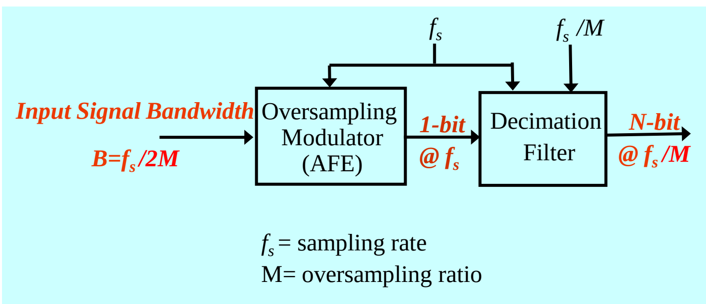

- Analog front-end → oversampled noise-shaping modulator
  - Converts original signal to a 1-bit digital output at the high rate of ($2BXM$)
- Digital back-end → digital filter (decimator)
  - Removes out-of-band quantization noise
  - Provides anti-aliasing to allow re-sampling @ lower sampling rate

Oversampling Benefits

- Almost no stringent requirements imposed on analog building blocks
- Takes advantage of the availability of low cost, low power digital filtering
- Relaxed transition band requirements for analog anti-aliasing filters
- Reduced baseband quantization noise power
- Allows trading speed for resolution

1st Order ΣΔ Modulator

1st order modulator, simplest loop filter → an integrator

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

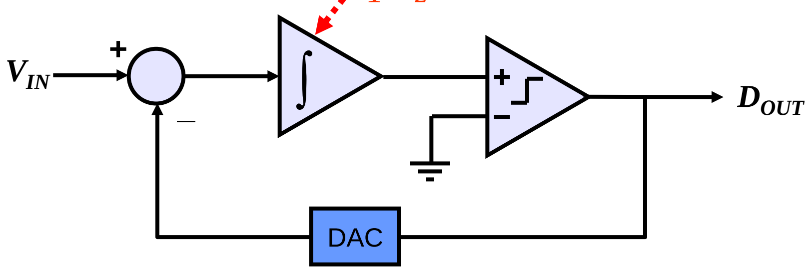

Note: Non-linear system with memory → difficult to analyze

1^{st} Order ΣΔ Modulator
Switched-capacitor implementation

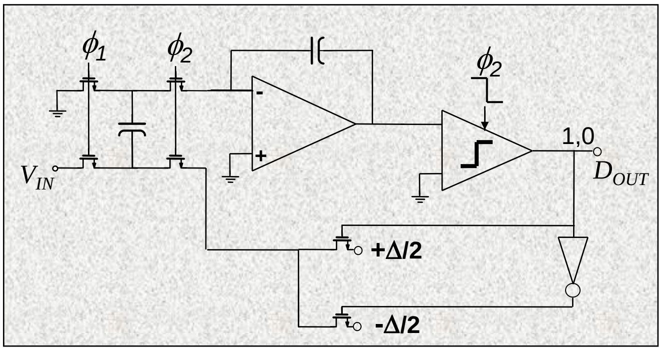

– Full-scale input range → Δ

– Note that Δ here is different from Nyquist rate ADC Δ (1LSB)

1^{st} Order ΔΣ Modulator

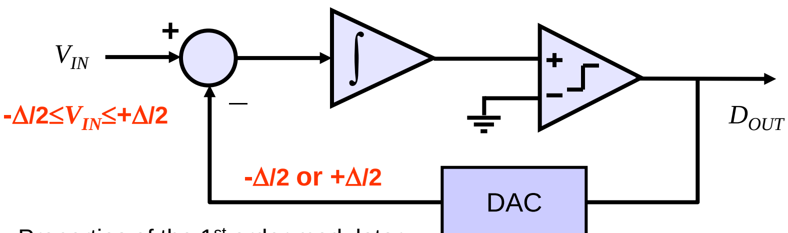

• Properties of the 1^{st} order modulator:

– Maximum analog input range is equal to the DAC reference levels

– 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}
→ To 1^{st} order, linearity is not dependent on component matching

– Alternative multi-bit DAC (and ADCs) solutions reduce the quantization error but loose this inherent monotonicity & relaxed matching requirements

1st Order ΣΔ Modulator

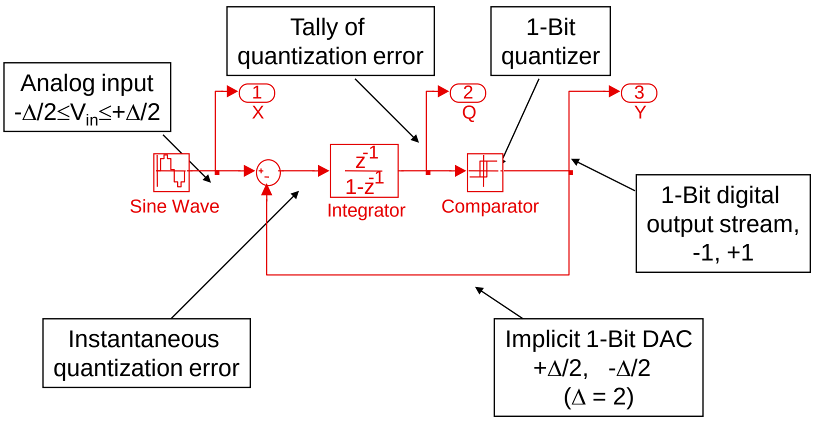

- M chosen to be 8 (low) to ease observability

1st Order Modulator Signals

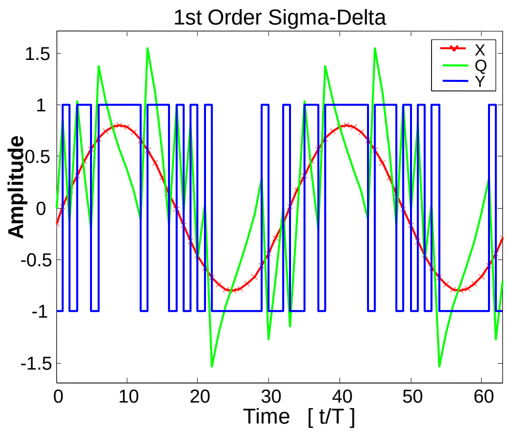

X analog input
Q tally of q-error
Y digital/DAC output

- Average of Y approximates X
- That is exactly what the digital filter does

$T = 1/f_s = 1/(M\ f_N)$
*M is usually >>8

ΣΔ Modulator Characteristics

- Inherently linear for 1-Bit DAC
- Quantization noise and thermal noise (KT/C) distributed over $-f_s/2$ to $+f_s/2$
  → Total noise within signal bandwidth reduced by 1/M
  → Required capacitor sizes $x1/M$ compared to nyquist rate ADCs
- Very high SQNR achievable (> 20 Bits!)
- To first order, quantization error independent of component matching
- Limited to moderate & low speed

1st Order ΣΔ Modulator Output Spectrum

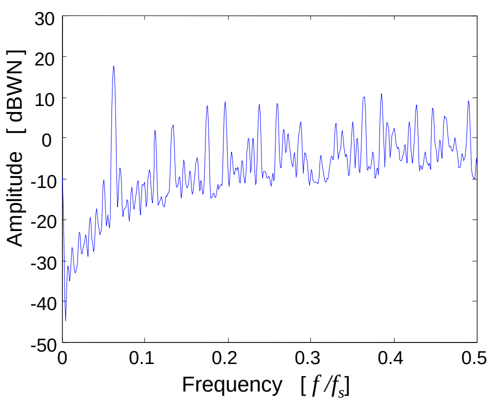

- Quantization noise definitely not white!
- Q-noise skewed towards higher frequencies
- Notice the distinct tones
- dBWN (dB White Noise) scale sets the 0dB line at the noise per bin of a random -1, +1 sequence

Quantization Noise Analysis

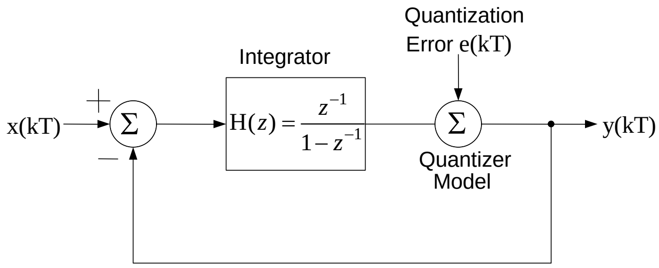

- Sigma-Delta modulators are nonlinear systems with memory → difficult to analyze directly
- Representing the quantizer as an additive noise source linearizes the system

Signal Transfer Function

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

$$H(j\omega) = \frac{\omega_0}{j\omega}$$

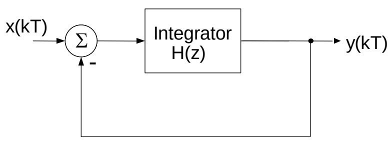

Signal transfer function
→ low pass function:

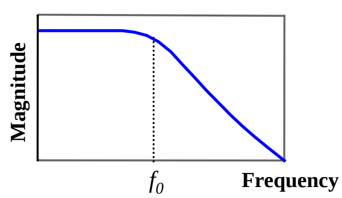

$$H_{Sig}(j\omega) = \frac{1}{1 + \frac{s}{\omega_0}}$$

$$H_{Sig}(z) = \frac{Y(z)}{X(z)} = \frac{H(z)}{1 + H(z)} = z^{-1} \quad \Rightarrow \quad \text{Delay}$$

Noise Transfer Function
Qualitative Analysis

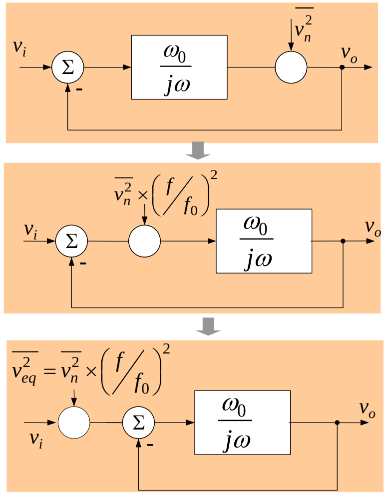

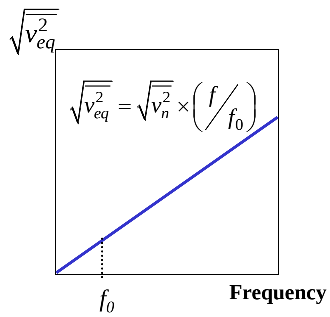

- Input referred-noise→ zero @ DC (s-plane)

1st Order ΣΔ Modulator
STF and NTF

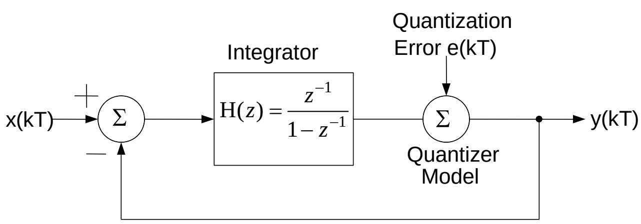

Signal transfer function:

$$\mathrm{STF} = \frac{Y(z)}{X(z)} = \frac{H(z)}{1+H(z)} = z^{-1} \quad \Rightarrow \quad \text{Delay}$$

Noise transfer function:

$$\mathrm{NTF} = \frac{Y(z)}{E(z)} = \frac{1}{1+H(z)} = 1 - z^{-1} \quad \Rightarrow \quad \text{Differentiator}$$

Noise Transfer Function

$$NTF = \frac{Y(z)}{E(z)} = \frac{1}{1+H(z)} = 1 - z^{-1} \qquad \text{set} \quad z = e^{j\omega T}$$

$$NTF(j\omega) = (1 - e^{-j\omega T}) = 2e^{-j\omega T/2}\left(\frac{e^{j\omega T/2} - e^{-j\omega T/2}}{2}\right)$$

$$= 2e^{-j\omega T/2}\, j\sin(\omega T/2)$$

$$= 2e^{-j\omega T/2} \times e^{-j\pi/2}\left[\sin(\omega T/2)\right]$$

$$= \left[2\sin(\omega T/2)\right]e^{-j(\omega T - \pi)/2}$$

where $T = 1/f_s$

Thus:

$$|NTF(f)| = 2|\sin(\omega T/2)| = 2|\sin(\pi f/f_s)|$$

Output noise power spectrum: $N_y(f) = |NTF(f)|^2 N_e(f)$

First Order ΣΔ Modulator
Noise Transfer Characteristics

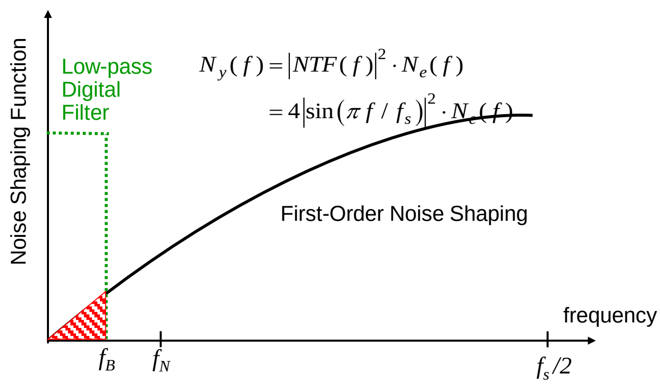

Key Point:
Most of quantization noise pushed out of frequency band of interest

Quantizer Error

- For quantizers with many bits

$$\overline{e^2(kT)} = \frac{\Delta^2}{12}$$

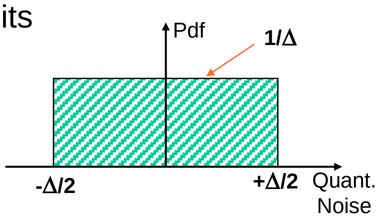

- Let's use the same expression for the 1-bit case
- **Use simulation to verify validity**
- <u>Experience</u>: Often sufficiently accurate to be useful, with enough exceptions to be careful

First Order ΣΔ Modulator
Simulated Noise Transfer Characteristic

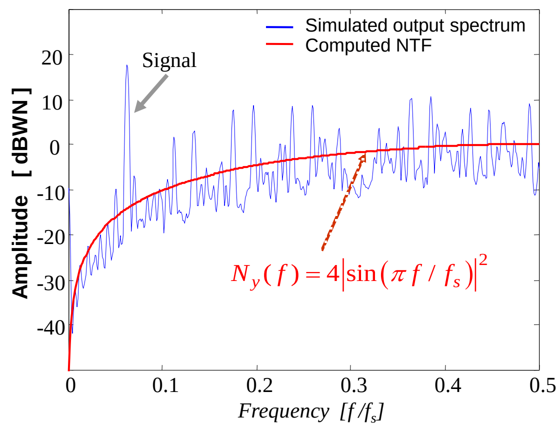

- Confirms assumption of quantization noise being white at insertion point
- Linearized model seems to be accurate

First Order ΣΔ Modulator In-Band Quantization Noise

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

$$|NTF(f)|^2 = 4|\sin(\pi f / f_s)|^2 \quad for \quad M >> 1$$

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

$$\cong \int_{-f_s/2M}^{f_s/2M} \frac{1}{f_s} \frac{\Delta^2}{12} (2\sin\pi fT)^2 df$$

$$\boxed{\rightarrow \overline{S_Q} \approx \frac{\pi^2}{3} \frac{1}{M^3} \frac{\Delta^2}{12}}$$

Total in-band quantization noise

1^{st} Order ΣΔ Dynamic Range

$$DR = 10\log\left[\frac{\text{full-scale signal power}}{\text{inband noise power}}\right] = 10\log\left[\frac{\overline{S_X}}{\overline{S_Q}}\right]$$

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

$$\overline{S_Q} = \frac{\pi^2}{3} \frac{1}{M^3} \frac{\Delta^2}{12}$$

$$\frac{\overline{S_X}}{\overline{S_Q}} = \frac{9}{2\pi^2} M^3$$

$$DR = 10\log\left[\frac{9}{2\pi^2} M^3\right] = 10\log\left[\frac{9}{2\pi^2}\right] + 30\log M$$

M	DR
16	33 dB
32	42 dB
1024	87 dB

$$\boxed{DR = -3.4dB + 30\log M}$$

2X increase in M➔9dB (1.5-Bit) increase in dynamic range

Oversampling and Noise Shaping

- $\Sigma\Delta$ modulators have interesting characteristics
  – Unity gain for input signal V_{IN}
  – Significant attenuation of in-band quantization noise injected at quantizer input
  – Performance significantly better than 1-bit noise quantizer possible for frequencies << f_s
- Increase in oversampling ($M = f_s/f_N >> 1$) improves SQNR considerably
  – 1st order $\Sigma\Delta$: DR increases 9dB for each doubling of M
  – To first order, SQNR independent of circuit complexity and accuracy
- Analysis assumes that the quantizer noise is "white"
  – Not entirely true in practice, especially for low-order modulators
  – Practical modulators are affected by other noise sources also (e.g. circuit thermal & 1/f noise)

1st Order $\Sigma\Delta$ Modulator Response to DC Input

- Matlab & Simulink model used
- Input → DC at 1/11 full-scale level

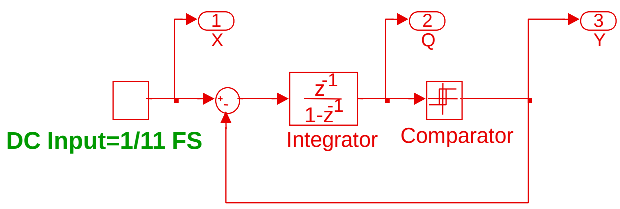

1^{st} Order ΣΔ
Response to DC Input

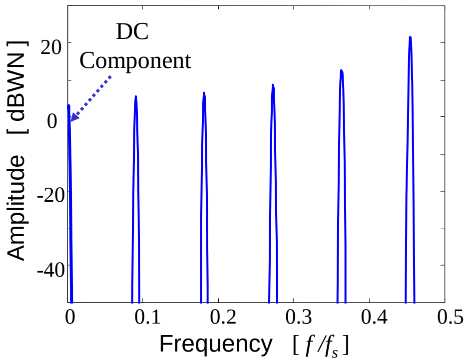

- DC input A = 1/11
- Output spectrum shows DC component plus distinct tones!!
- Tones frequency shaped the same as quantization noise
  → More prominent at higher frequencies

  → Seems like periodic quantization "noise"

Limit Cycle Oscillation

DC input 1/11 →
Periodic sequence:

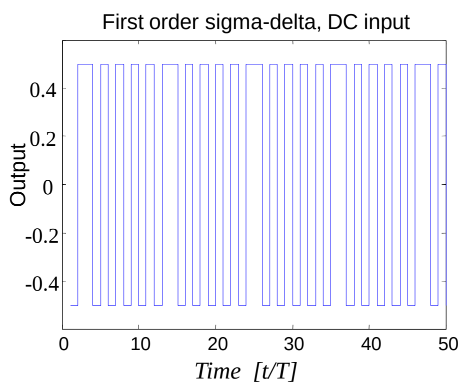

1	+1
2	+1
3	-1
4	+1
5	-1
6	+1
7	-1
8	+1
9	-1
10	+1
11	-1

→ Average =1/11

1^{st} Order $\Sigma\Delta$
Limit Cycle Oscillation

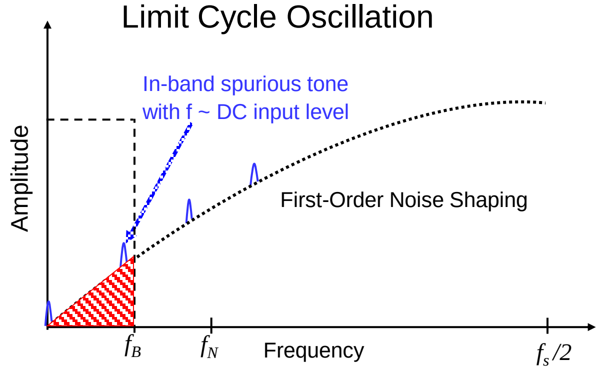

- Problem: quantization noise becomes periodic in response to low level DC inputs & could fall within passband of interest!
- Solution:
  - Use dithering (inject noise-like signal at the input): randomizes quantization noise
    - If circuit thermal noise is large enough→ acts as dither
  - Second order loop

1^{st} Order $\Sigma\Delta$ Modulator
Linearized Model Analysis

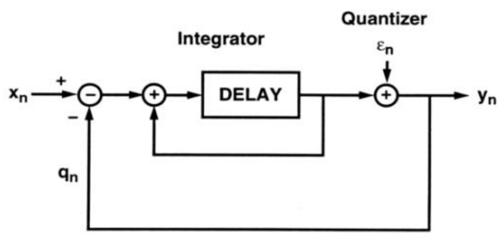

$$Y(z) = \underbrace{z^{-1}}_{\text{LPF}} X(z) + \underbrace{\left(1 - z^{-1}\right)}_{\text{HPF}} E(z)$$

2^{nd} Order ΣΔ Modulator

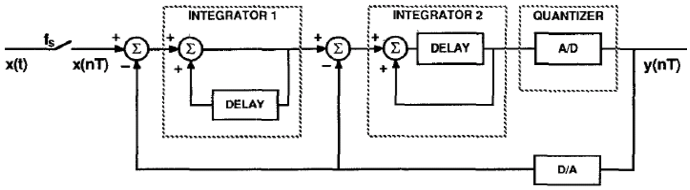

- Two integrators in series
- Single quantizer (typically 1-bit)
- Feedback from output to *both* integrators
- Tones less prominent compared to 1st order

2^{nd} Order ΣΔ Modulator Linearized Model Analysis

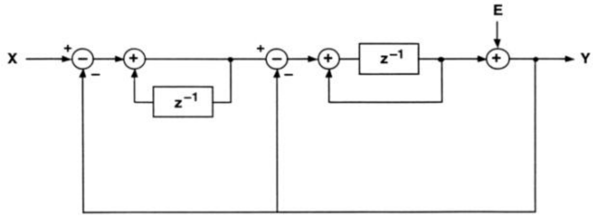

Using the delay operator z^{-1}: $Y(z) = \underbrace{z^{-1}}_{\text{LPF}} X(z) + \underbrace{\left(1-z^{-1}\right)^2}_{\text{HPF}} E(z)$

2^{nd} Order ΣΔ Modulator
In-Band Quantization Noise

$$NTF(z) = \left(1 - z^{-1}\right)^2$$

$$\left|NTF(f)\right|^2 =$$

$$= 2^4 \left|\sin(\pi f / f_s)\right|^4 \quad \text{for} \quad M >> 1$$

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

$$\cong \int_{-f_s/2M}^{f_s/2M} \frac{1}{f_s} \frac{\Delta^2}{12} \left(2 \sin \pi f T\right)^4 df$$

$$\approx \frac{\pi^4}{5} \frac{1}{M^5} \frac{\Delta^2}{12}$$

Quantization Noise
2^{nd} Order ΣΔ Modulator vs 1^{st} Order Modulator

$$\overline{S_Q} \approx \frac{\pi^2}{3} \frac{1}{M^3} \frac{\Delta^2}{12}$$

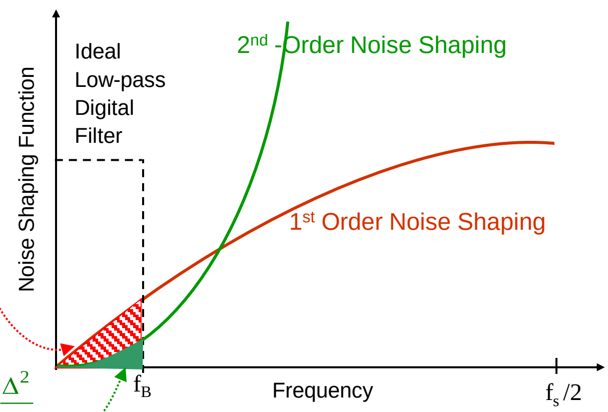

$$\overline{S_Q} \approx \frac{\pi^4}{5} \frac{1}{M^5} \frac{\Delta^2}{12}$$

2nd Order ΣΔ Modulator Dynamic Range

$$DR = 10\log\left[\frac{\text{full-scale signal power}}{\text{inband noise power}}\right] = 10\log\left[\frac{\overline{S_X}}{\overline{S_Q}}\right]$$

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

$$\overline{S_Q} = \frac{\pi^4}{5}\frac{1}{M^5}\frac{\Delta^2}{12}$$

$$\frac{\overline{S_X}}{\overline{S_Q}} = \frac{15}{2\pi^4}M^5$$

$$DR = 10\log\left[\frac{15}{2\pi^4}M^5\right] = 10\log\left[\frac{15}{2\pi^4}\right] + 50\log M$$

$$\boxed{DR = -11.1dB + 50\log M}$$

2X increase in M → $15dB$ (2.5-bit) increase in DR

2nd Order vs 1st Order ΣΔ Modulator Dynamic Range

M	2nd Order D.R.	1st Order D.R.	Resolution (2nd order - 1st order)
16	49 dB (7.8bit)	33dB (5.2bit)	2.6 bit
32	64 dB (10.3bit)	42dB (6.7bit)	3.6 bit
256	109 dB (17.9bit)	68.8dB (11.1bit)	6.8 bit
1024	139 dB (22.8bit)	87dB (14.2bit)	8.6 bit

– Note: For higher oversampling ratios resolution of 2nd order modulator significantly higher compared to 1st order

2nd Order ΣΔ Modulator Example

- Digital audio application
  - Signal bandwidth *20kHz*
  - Desired resolution *16-bit*

$$16-bit \rightarrow 98\,dB \quad \text{Dynamic Range}$$

$$DR_{2nd\ order\ \Sigma\Delta} = -11.1dB + 50\log M$$

$$M_{min} = 153$$

$M \rightarrow 256 = 2^8$ ($\rightarrow DR = 109dB$) two reasons:

1. Allow some margin so that thermal noise dominate & provides dithering to minimize level of in-band limit cycle oscillation
2. Choice of M power of $2 \rightarrow$ ease of digital filter implementation

→ Sampling rate *(2x20kHz + 5kHz)M = 12MHz* (quite reasonable!)

Limit Cycle Tones in 1st Order & 2nd Order ΣΔ Modulator

- Higher oversampling ratio → lower tones
- 2nd order tones much lower compared to 1st
- 2X increase in M decreases the tones by 6dB for 1st order loop and 12dB for 2nd order loop

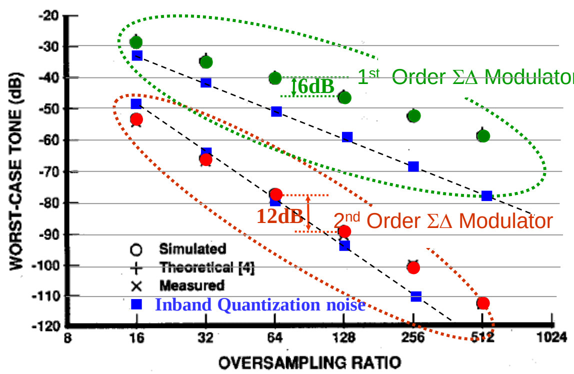

Ref: B. P. Brandt, et al., "Second-order sigma-delta modulation for digital-audio signal acquisition," IEEE Journal of Solid-State Circuits, vol. 26, pp. 618 - 627, April 1991.
R. Gray, “Spectral analysis of quantization noise in a single-loop sigma–delta modulator with dc input,” *IEEE Trans. Commun.,* vol. 37, pp. 588–599, June 1989.

ΣΔ Implementation
Practical Design Considerations

- Internal node scaling & clipping
- Effect of finite opamp gain & linearity
- KT/C noise
- Opamp noise
- Effect of comparator nonidealities
- Power dissipation considerations

Switched-Capacitor Implementation 2nd Order ΣΔ Nodes Scaled for Maximum Dynamic Range

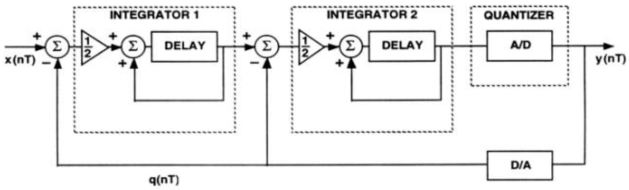

- Modification (gain of ½ in front of integrators) reduce & optimize required signal range at the integrator outputs ~ 1.7x input full-scale (Δ)
- Note: Non-idealities associated with 2nd integrator and quantizer when referred to the ΣΔ input is attenuated by 1st integrator high gain

  → The only building block requiring low-noise and high accuracy is the 1st integrator

Ref: B.E. Boser and B.A. Wooley, “The Design of Sigma-Delta Modulation A/D Converters,” IEEE J. Solid-State Circuits, vol. 23, no. 6, pp. 1298-1308, Dec. 1988.

2^{nd} Order ΣΔ Modulator
Example: Switched-Capacitor Implementation

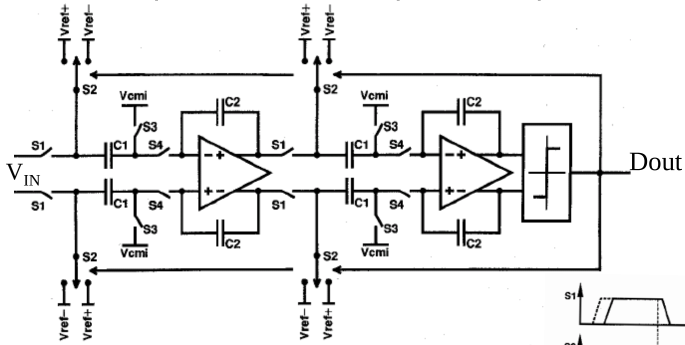

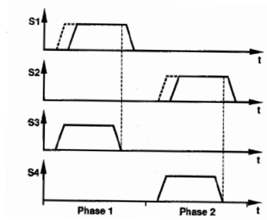

- Fully differential front-end
- Two bottom-plate integrators
- 1-bit DAC is made of switches and Vrefs

Switched-Capacitor Implementation 2^{nd} Order ΣΔ
Phase 1

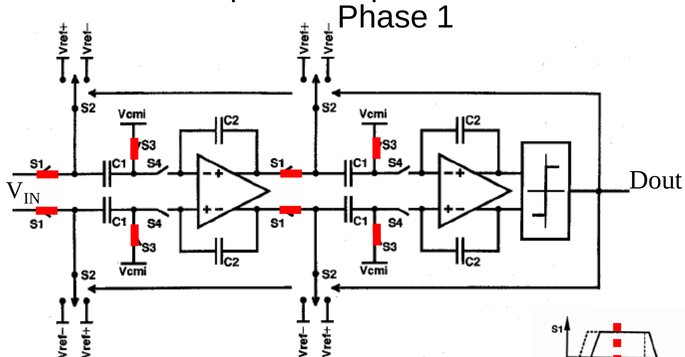

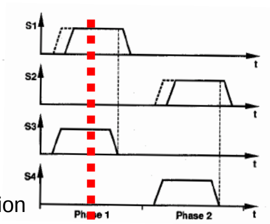

During phase 1:
- 1^{st} integrator samples Vin on 1^{st} stage C1
- 2^{nd} integrator samples output of 1^{st} integrator
- Comparator senses polarity of 2^{nd} intg. output → result saved in output latch
- S3 opens prior to S1 → minimize effect of charge injection

Switched-Capacitor Implementation 2nd Order ΣΔ Phase 2

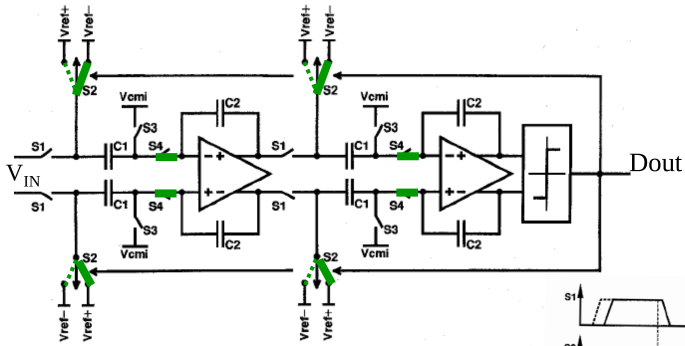

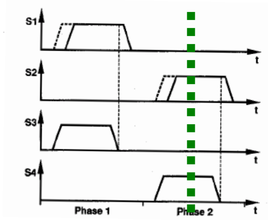

- Note: S2 connects integrator inputs to **+** or **–** Vref, polarity depends on whether Dout is **0** or **1**
- Input sampled during ϕ_1 – or + C1xVref transferred to C2 → DAC output subtraction & integration

2nd Order ΣΔ Modulator Switched-Capacitor Implementation

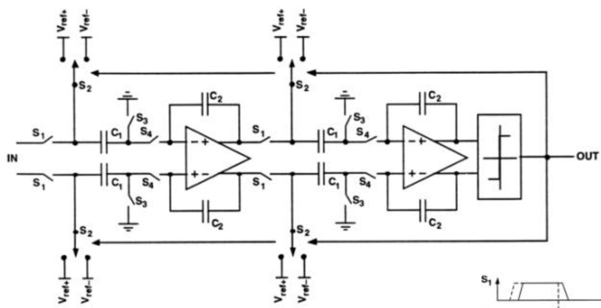

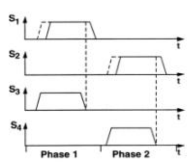

- The ½ loss in front of each integrator implemented by choice of:

$\boxed{C_2=2C_1}$ $\rightarrow f_0^{intg}=f_s/(4\pi)$