


# EE247

## Lecture 22

- Pipelined ADCs
  - Combining the bits
  - Stage implementation
    - Circuits
    - Noise budgeting
- Figures of merit (FOM) and trends for ADCs
  - How to use/not use FOM
- Oversampled ADCs

### Summary

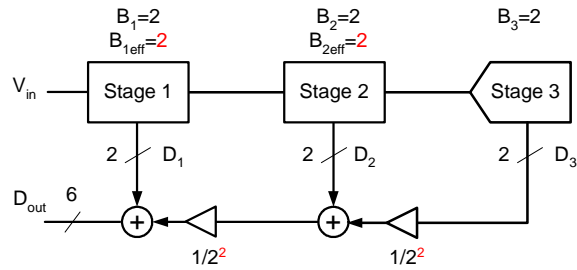
#### Pipelined A/D Converters



- Cascade of low resolution stages
  - Stages operate concurrently- trades latency for conversion speed
  - Throughput limited by speed of one stage ® Fast
- Errors and correction
  - Built-in redundancy compensate for sub-ADC inaccuracies
  - Digital calibration compensates:
    - Inter-stage gain inaccuracy
    - Sub-DAC error
    - Inter-stage gain nonlinearities

# Combining the Bits

- Example: Three 2-bit stages, no redundancy




$$D_{out} = D_1 + \frac{1}{4} D_2 + \frac{1}{16} D_3$$

# Combining the Bits

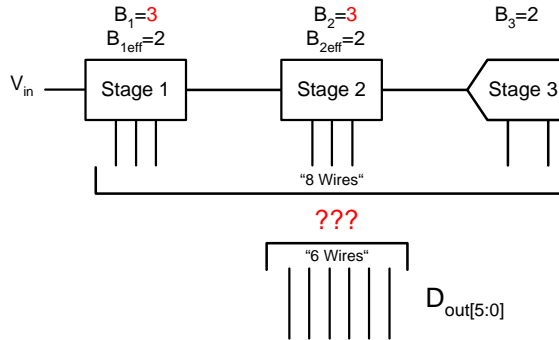
$D_1$     **XX**  
 $D_2$        **XX**  
 $D_3$             **XX**  
 -----  
 $D_{out}$     **DDDDDD**

- Only bit shifts
- No arithmetic circuits needed



# Combining the Bits

- Example: Three 2-bit stages, one bit redundancy in stages 1 and 2

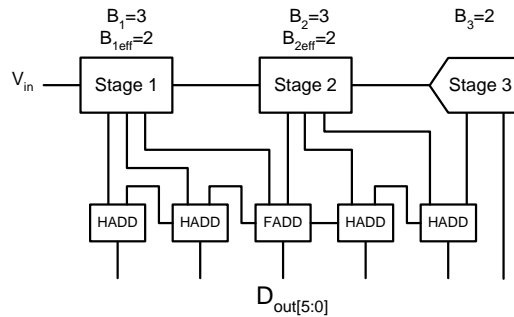


# Combining the Bits

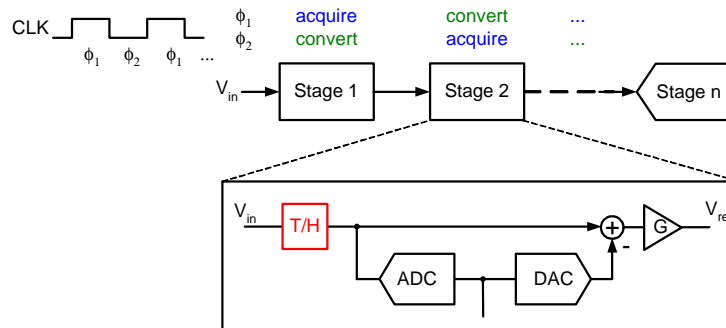
$$D_{out} = D_1 + \frac{1}{4}D_2 + \frac{1}{16}D_3$$

$D_1$     **XXX**  
 $D_2$     **XXX**  
 $D_3$     **XX**  
 -----  
 $D_{out}$     **DDDDDD**

- Bits overlap
- Need adders




## Stage Implementation




- Each stage needs T/H hold function
- Track phase: Acquire input/residue from previous stage
- Hold phase: sub-ADC decision, compute residue

## Stage Implementation




- Usually no dedicated T/H amplifier in each stage (Except first stage – why?)
- T/H implicitly contained as passive samplers in stage building blocks

## Stage Implementation




- Multiply-DAC-subtract function can be lumped into a single switched capacitor circuit
- "MDAC"

## 1.5 Bit Stage Implementation



## 1.5 Bit Stage Implementation


$$V_o = \begin{cases} \left(1 + \frac{C_s}{C_f}\right) V_i - \frac{C_s}{C_f} V_{ref} & \text{if } V_i > V_{ref}/4 \\ \left(1 + \frac{C_s}{C_f}\right) V_i & \text{if } -V_{ref}/4 \leq V_i \leq +V_{ref}/4 \\ \left(1 + \frac{C_s}{C_f}\right) V_i + \frac{C_s}{C_f} V_{ref} & \text{if } V_i < -V_{ref}/4 \end{cases}$$



Ref: A. Abo, "Design for Reliability of Low-voltage, Switched-capacitor Circuits," UCB PhD Thesis, 1999

## Stage Scaling


- Example: Pipeline using 1-bit<sub>eff</sub> stages



- Total input referred noise power:

$$N_{tot} \propto kT \left[ \frac{1}{C_1} + \frac{1}{4C_2} + \frac{1}{16C_3} + \dots \right]$$


## Stage Scaling



$$N_{tot} \propto kT \left[ \frac{1}{C_1} + \frac{1}{4C_2} + \frac{1}{16C_3} + \dots \right]$$

- If we make all caps the same size, backend stages contribute very little noise
- Wasteful, because Power  $\sim$  Gm  $\sim$  C

## Stage Scaling




$$N_{tot} \propto kT \left[ \frac{1}{C_1} + \frac{1}{4C_2} + \frac{1}{16C_3} + \dots \right]$$


- How about scaling caps down by 2<sup>2</sup>=4x per stage?
  - Same amount of noise from every stage
  - All stages contribute significant noise
  - Noise from first few stages must be reduced
  - Power  $\sim$  Gm  $\sim$  C goes up!

# Stage Scaling

Extreme 1: All Stages the Same Size



Extreme 2: All Stages Contribute the Same Noise




- Optimum capacitor scaling lies approximately midway between these two extremes

Ref: D. W. Cline, P.R Gray "A power optimized 13-b 5 MSamples/s pipelined analog-to-digital converter in 1.2um CMOS," JSSC 3/1996


# Stage Scaling

- Power minimum is "shallow"
- Near optimum solution in practice: Scale capacitors by stage gain
- E.g. for effective stage resolution of 1bit (Gain=2):







# Stage Scaling



stage 0 comparators    stage 0 sampling capacitors    stage 0 opamp  
stage 0 sampling switches



Ref: D. W. Cline, P.R Gray "A power optimized 13-b 5 MSamples/s pipelined analog-to-digital converter in 1.2um CMOS," JSSC 3/1996

## How Many Bits Per Stage?

- Many possible architectures
    - E.g.  $B_{1\text{eff}}=2, B_{2\text{eff}}=1, \dots$   
Vs.  $B_{1\text{eff}}=1, B_{2\text{eff}}=1, B_{3\text{eff}}=1, \dots$
  - Complex optimization problem, fortunately optimum tends to be shallow...
  - Qualitative answer:
    - Maximum speed for given technology
      - Use small stage resolution (large feedback factor)
    - Maximum power efficiency for fixed, "low" speed
      - Try higher resolution stages
      - Can help alleviate matching requirements in front-end
- Ref: Singer VLSI 96, Yang, JSSC 12/01

## Two State-of-the-Art Implementations

Reference	Yang (JSSC 12/2001)	Loloe (ESSIRC 2002)
Bits	12	12
Architecture	3-1-1-1-1-1-1-3	1-1-1-1-1-1-1-1-1-2
SNR	~70dB	~66dB
Speed	75MS/s	80MS/s
Power	340mW	260mW

## ADC Figures of Merit

- Objective: Want to compare performance of different ADCs
- Can use FOM to combine several performance metrics to get one single number
- What are reasonable FOM for ADCs?
- How can we use and interpret them?
- Trends?


# ADC Figures of Merit

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

- This FOM suggests that adding a 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 J. Selected Areas Comm.*, April 1999

# Survey Data



Ref: R.H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Selected Areas Comm.*, April 1999

## ADC Figures of Merit

$$FOM_2 = \frac{f_s \cdot 2^{ENOB}}{Power}$$

- 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 J. Selected Areas Comm.*, April 1999

## ADC Figures of Merit

- One more bits means...
  - 6dB SNR, 4x less noise power, 4x bigger C
  - Power ~  $G_m$  ~ 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

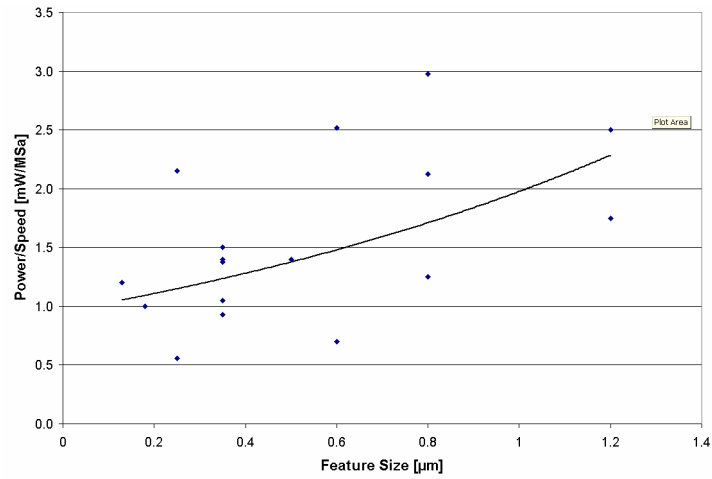
- $FOM_2$  seems inappropriate, but somehow still standard in literature, papers
- "Tends to work" because:
  - Not all power in an ADC is "noise limited"
  - E.g. Digital power, biasing circuits, etc.
- Avoid comparing different resolution ADCs using  $FOM_2$ !

## ADC Figures of Merit

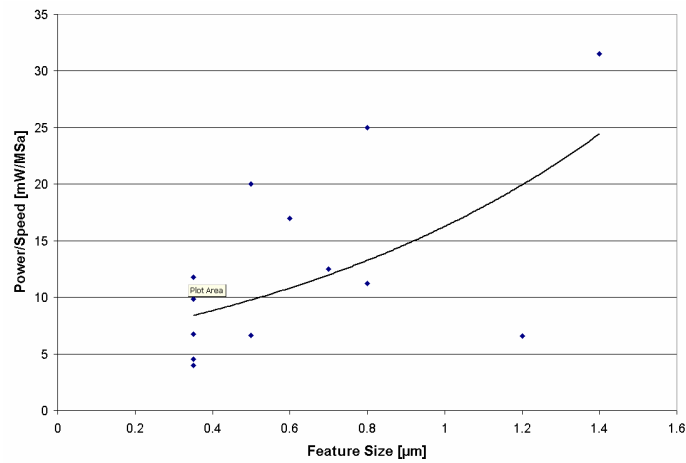
$$FOM_3 = \frac{Power}{Speed}$$

- Compare only power of ADCs with approximately same ENOB
- Useful numbers: (state-of-the-art):
  - 10b (~9 ENOB) ADCs: 1 mW/MSample/sec
  - 12b (~11 ENOB) ADCs: 4 mW/MSample/sec


# 10-Bit ADC Power



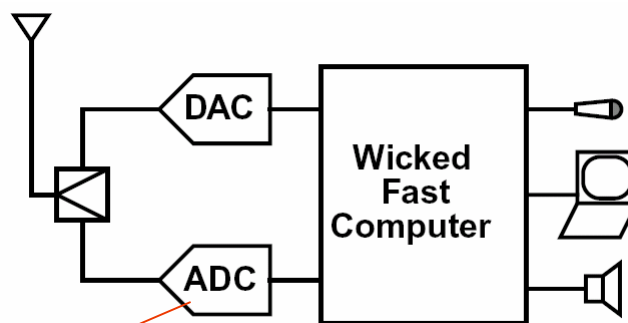
# 12-Bit ADC Power



# Performance Trend

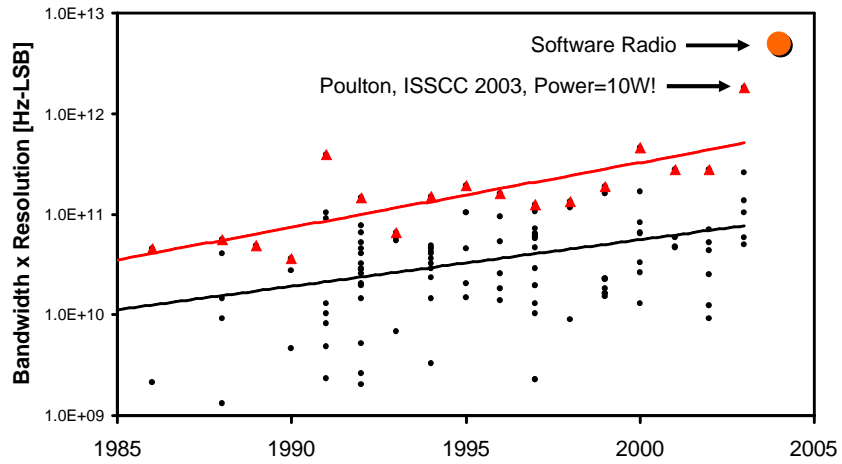


# The Dream Transceiver – "Software Radio"

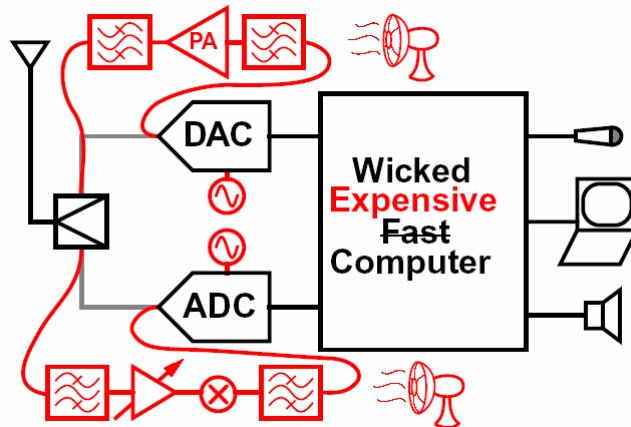


e.g.: SNR  $\cong$  100dB  
 BW  $\cong$  30MHz

# ADC for Software Radio



# Today's Transceiver...



Ref: Schreier, "ADCs and DACs: Marching Towards the Antenna," GIRAFE workshop, ISSCC 2003




# Oversampled ADCs


- Why oversampling?
- Pulse-count modulation
- Sigma-delta modulation
  - 1-Bit quantization
  - Quantization error spectrum
  - SQNR analysis

# The Case for Oversampling

## Nyquist sampling:

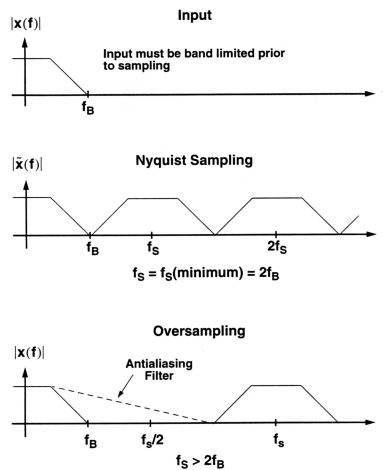


## Oversampling:



- Nyquist rate  $f_N = 2B$
- Oversampling rate  $M = f_s/f_N \gg 1$

## Oversampled Converters Antialiasing




## Oversampling Benefits

- Trades speed for resolution
- Relaxed transition band requirements for analog anti-aliasing filters
- Reduced baseband quantization noise power
- Utilizes low cost, low power digital filtering

## Oversampled Converters Baseband Noise

- For a quantizer with step size  $\Delta$  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 aliased into 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  $M = \frac{f_s}{2f_B} = \text{oversampling ratio}$

## Oversampled Converters Baseband Noise

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


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

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

Greater improvement in resolution:


- Embed quantizer in a feedback loop
  - Predictive (delta modulation)
  - Noise shaping (sigma delta modulation)

## Pulse-Count Modulation




Mean of pulse-count signal approximates analog input!

# Pulse-Count Spectrum




- Signal: low frequencies,  $f < B \ll f_s$
- Quantization error: high frequency,  $B \dots f_s / 2$
- Separate with 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 error for  $f > B$
- Provides most anti-alias filtering
- Narrow transition band, high-order
- 1-Bit input, N-Bit output (essentially computes “average”)

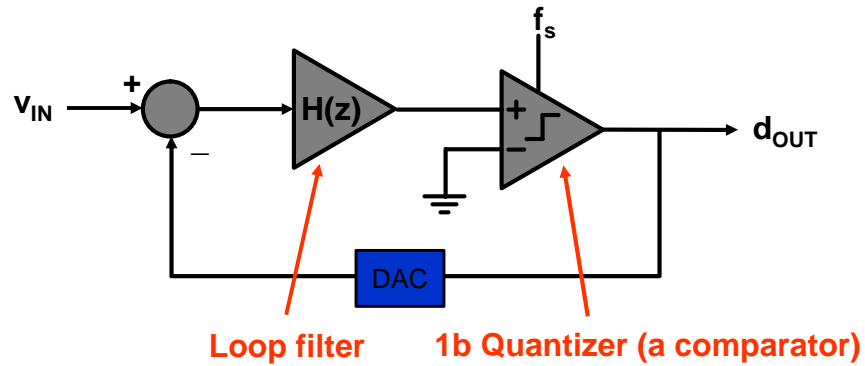
# Modulator

- Objectives:
  - Convert analog input to 1-Bit pulse density stream
  - Move quantization error to high frequencies  $f \gg B$
  - Operates at high frequency  $f_s \gg f_N$ 
    - $M = 4 \dots 256$  (typical)
    - Better be “simple”

→  $\Sigma\Delta = \Delta\Sigma$  Modulator

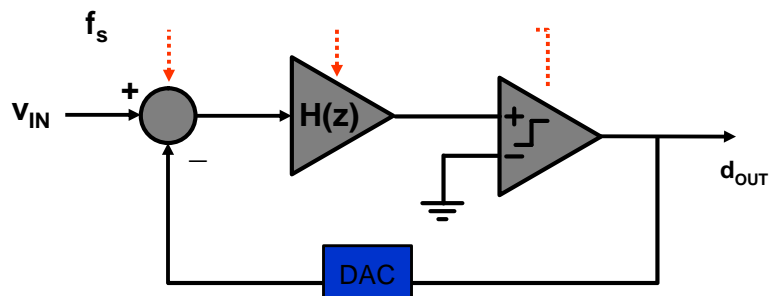
# Sigma- Delta Modulators

Analog 1-Bit  $\Sigma\Delta$  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




# 1<sup>st</sup> Order $\Sigma\Delta$ Modulator

In a 1<sup>st</sup> order modulator, simplest loop filter  $\rightarrow$  an integrator




# 1<sup>st</sup> Order $\Sigma\Delta$ Modulator

Switched-capacitor implementation



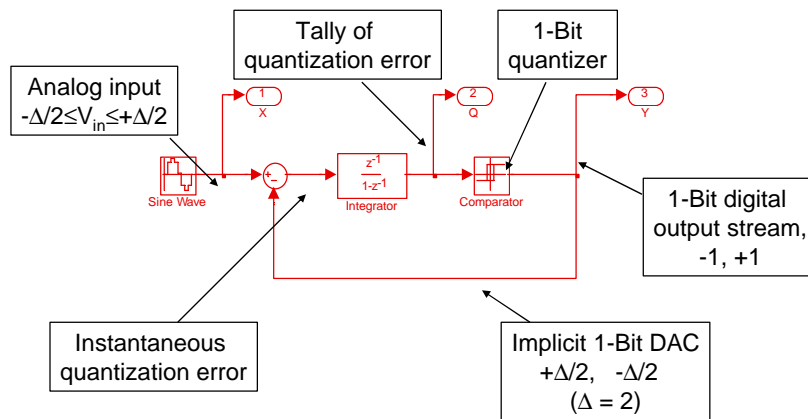


# 1<sup>st</sup> Order $\Delta\Sigma$ Modulator




- Properties of the first-order modulator:
  - Analog input range is equal to 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 lose this inherent monotonicity

# 1<sup>st</sup> Order $\Sigma\Delta$ Modulator



# 1<sup>st</sup> Order Modulator Signals



Mean of Y approximates X

$$T = 1/f_s = 1/(M f_N)$$

## $\Sigma\Delta$ Modulator Characteristics

- Quantization noise and thermal noise ( $KT/C$ ) distributed over  $-f_s/2$  to  $+f_s/2$ 
  - Total noise reduced by  $1/M$
- Very high SQNR achievable ( $> 20$  Bits!)
- Inherently linear for 1-Bit DAC
- Quantization error independent of component matching
- Limited to moderate to low speed