ESE 568: Mixed Signal Design and Modeling

Lec 21: November 20, 2019

FOM, Performance Limits

ADC Figures of Merit

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

ADC Figures of Merit

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

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

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

Survey Data

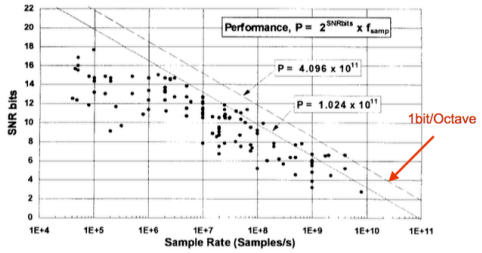

[Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Selected Areas Comm.*, April 1999]

ADC Figures of Merit

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

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

- Sometimes inverse of this metric is used
- In typical circuits power ~ speed
  - FOM2 captures this tradeoff correctly
- How about power vs. ENOB?
  - One additional bit = 2x in power?

ADC Figures of Merit

- In a circuit that is limited by thermal noise, each additional bit in resolution means...
  - 6dB SNR → 4x less noise power → 4x bigger C
  - Power ~ Gm ~ C increases 4x

ADC Figures of Merit

- In a circuit that is limited by thermal noise, each additional bit in resolution means...
  - 6dB SNR → 4x less noise power → 4x bigger 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 maintain current density
  - Net result: Power increases 8x

ADC Figures of Merit

- FOM_2 is inappropriate for comparing ADCs that are limited by matching or thermal noise
  - Still the most widely used FOM in publications...
- "Tends to work" because not all power in an ADC is noise limited
  - E.g. Digital power, biasing circuits, etc.
- To better capture the case of noise limited circuits, one could use 2^{2*ENOB} in the numerator of FOM_2...

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

ADC Figures of Merit

$$FOM_3 = \frac{Power}{Conversion\ Bandwidth}$$

- Compare only power of ADCs with approximately same SNR or SNDR (ENOB)
- Useful numbers (~state-of-the-art):
  - 10b(~9ENOB)ADCs:1...2mW/MHz
  - 12b (~11 ENOB) ADCs: 6...8 mW/MHz

Which FOM?

- Different limiting mechanisms
  - Matching: Power grows 8x per added bit
  - Thermal Noise: Power grows 4x per added bit
  - Process CV^2: Power grows 2x per added bit

Survey Data 1997-2004

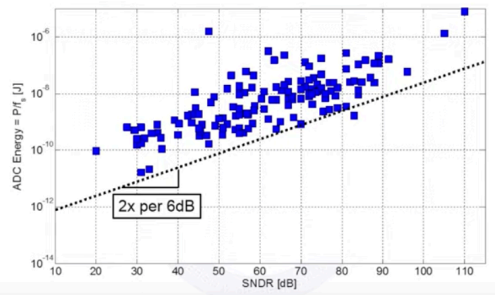

Survey Data 1997-2014

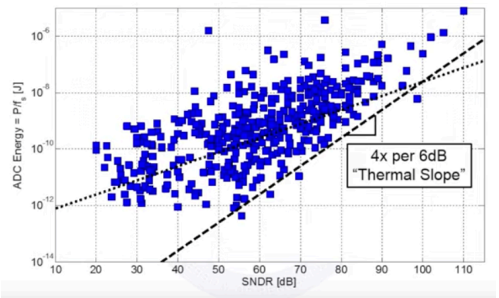

Popular FOM

- Walden FOM:
  - 2x per bit

$$FOM_W = \frac{P}{f_s \cdot 2^{ENOB}}$$

- Schreier FOM
  - 4x per bit
  - Ignores distortion

$$FOM_{S,DR} = DR + 10\log\left(\frac{BW}{P}\right)$$

- Schreier FOM
  - 4x per bit
  - Includes distortion

$$FOM_S = SNDR + 10\log\left(\frac{f_s/2}{P}\right)$$

Survey Data 1997-2014

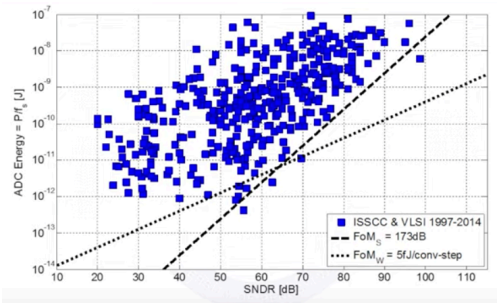

Fundamental Power Limits

- Sampling Power

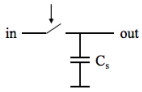

Noise voltage generated: $v_{nS}^2 = \frac{kT}{C_s}$

Maximum sine voltage with supply voltage V_{FS}:

$v_s^2 = \frac{V_{FS}^2}{8}$ gives dynamic range: $SNR = \frac{v_s^2}{v_{nS}^2} = \frac{V_{FS}^2 C_s}{8kT}$

Fundamental Limits

- Fundamental power limit for a class-B amplifier driving a single capacitor [Vittoz, ISCAS 1990]

$$P = 8 \cdot f_{sig} \cdot CV_{sig}^2 \qquad V_n^2 = \frac{k_B T}{C} \qquad SNR = \frac{0.5 \times V_{sig}^2}{V_n^2}$$

$$\boxed{\therefore P = 8k_B T \cdot SNR \cdot f_{sig}}$$

- Class-A power limit is π times higher

Switched Cap Circuits

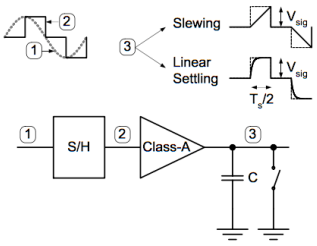

Case 1: 100% Slewing

$$I_{bias} = C \cdot \frac{dV}{dt} = C \cdot \frac{V_{sig}}{T_s/2} = 4 \cdot C \cdot V_{sig} \cdot f_{sig}$$

$$P = 2 \cdot V_{sig} \cdot I_{bias} \qquad SNR = \frac{0.5 \times V_{sig}^2}{k_B T / C}$$

$$\boxed{\therefore P = 16k_B T \cdot SNR \cdot f_{sig}}$$

Switched Cap Circuits

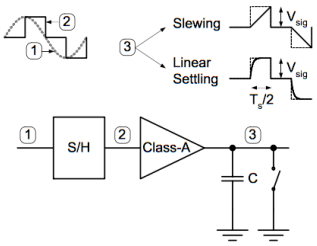

Case 2: 100% Linear Settling

$$I_{bias} = C \cdot \left.\frac{dV}{dt}\right|_{max} = C \cdot \left.\frac{d}{dt}\right|_{max} \left[V_{sig}\left(1 - e^{-t/\tau}\right)\right] = 4 \cdot C \cdot \frac{V_{sig}}{\tau}$$

Number of settling time constants: $N = \frac{T_s/2}{\tau}$

$$\therefore P = 16 \cdot N \cdot k_B T \cdot SNR \cdot f_{sig}$$

- Much worse
  - E.g. N=6.9 for settling to 0.1% precision

Reality Check

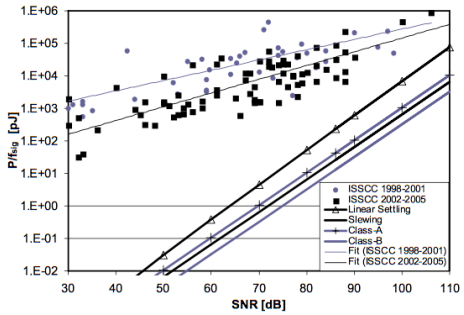

Discussion

- Orders of magnitude away from limits
- Slope of limit lines is much steeper than fit to experimental data
- What contributes to these large gaps?
  - Must keep in mind that ADCs are not just a single capacitor circuit...

Design Space Partitioning

- High SNR
  - Complexity ~1 (e.g. first integrator in sigma-delta ADC)
  - Limited by thermal noise
- Medium SNR
  - Complexity ~Bits (e.g. pipelined ADC)
  - Partly limited by thermal noise
- Low SNR
  - Complexity ~2Bits (e.g. flash ADC)
  - Limited by matching, quantization noise

High SNR: SC-Stage

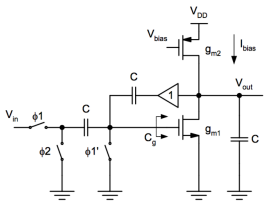

- Considerations
  - Noise is multiple of $k_B T/C$ (n_f)
  - Swing is only a fraction of V_{DD} (α)
  - Feedback factor (β)
  - g_m/I_D is upper bounded if slewing must be avoided

High SNR: SC-Stage

To avoid slewing: $\frac{g_{m1}}{I_{bias}} \le \frac{1}{\beta \cdot V_{sig}}$

$$\therefore P = 16 \cdot N \cdot n_f \cdot \frac{1}{\alpha} \cdot k_B T \cdot SNR \cdot f_{sig} \cdot max\left(1, \frac{1}{\frac{g_{m1}}{I_{bias}} \cdot \beta \cdot V_{sig}}\right)$$

- Graph on following slide shows result assuming
  - n_f=5, α=2/3, β=0.5, onset of slewing

High SNR: SC-Stage

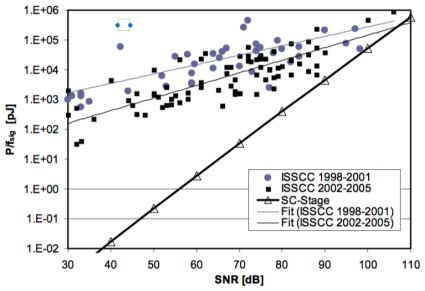

- Very close to experimental data at high SNR!

Medium SNR

- Pipeline ADC using SC stages
  - Partially limited by thermal noise

Medium SNR: Pipeline ADC

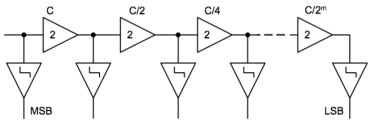

- Theoretical near optimum power scaling
  - Scale capacitance by gain of preceding stage
  - Stage 1 consumes half of total power
  - Adding one bit means power goes up 4x
- Caveat
  - Usually impractical to scale capacitors down to $C/2^m$
  - State-of-art: 10bits ~2mW/MHz, 12bits ~8mW/MHz

Medium SNR: Results

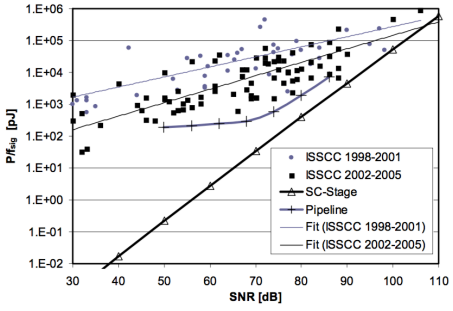

Low SNR

- Power of matching limited class-B circuit [Kinget, CICC 1996]

$$P = 24 \cdot C_{ox} \cdot A_{VT}^2 \cdot f_{sig} \cdot \left(\frac{V_{sig,rms}}{3 \cdot \sigma_{Vos}}\right)^2$$

- Refined result for flash ADC, assuming
  - Class-A, 1/2LSB matching with 3σ-confidence, 2^B components, additional E_{dyn} per clock cycle, partial supply usage (α)

$$P = \left(12\pi \cdot \frac{1}{\alpha} \cdot C_{ox} \cdot A_{VT}^2 \cdot 2^{3B} + 2 \cdot E_{dyn} \cdot 2^B\right) \cdot f_{sig}$$

- Example: α=2/3, C_{ox}=15fF/μm^2, A_{Vt}=3mV·μm, E_{dyn}=60fJ (~10gates in 0.13 μm CMOS)

Low SNR: Result

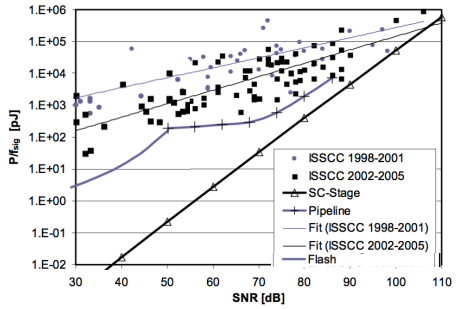

Discussion

- Missing factors
  - Biasing, reference generator
  - S/H hold front-end
  - Cost of "high gain"
  - I/O power, digital power … and more
- Shown results include only minor assumptions about technology
- Scaling brings some good, some bad news offsetting each other
  - CON: Lower V_{DD}, lower V_{swing}/V_{DD}
  - PRO: Lower E_{dyn}, higher f_t enables moderate/weak inversion operation with high g_m/I_D
- **Limit lines won't move much, unless someone hands us a new disruptive technology**

Future Opportunities

- More intelligent ADCs
  - Improved average power dissipation by adapting to instantaneous speed/resolution requirements
- "Sloppy" ADCs using significantly simpler circuits
  - Digital compensation of resulting non-idealities
  - Digital postprocessing is "free" at moderate to high SNR
- **Digitally Assisted ADCs**
  - **How many analog transistors do we really need?**

Conventional vs. Digitally Assisted ADC

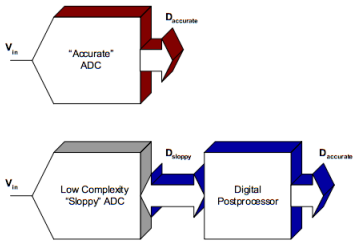

Digital Energy – Relative to ADC

- Example
  - Standard digital gate (NAND2) in 0.13 μm CMOS consumes about 6nW/Gate/MHz
    - Energy/Gate = 6fJ
  - State-of-the-art 10-bit Nyquist ADCs consume roughly 1mW/MSample/sec
    - Energy/Conversion = 1nJ
  - Energy equivalent number of gates (Assuming 10-bit ADC, 0.13 μm logic)
    - 1nJ/6fJ=166,666
- Up to several tens of thousand gates are "free" in terms of energy!

Digital Energy – Relative to ADC

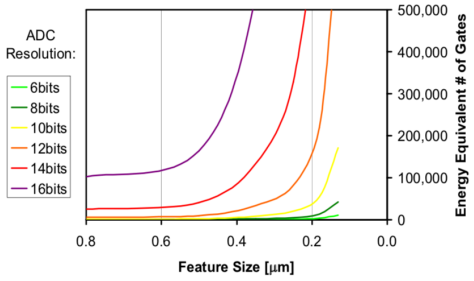

Example: Pipeline ADC

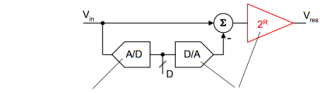

Redundancy helps tolerate large sub A/D errors [Lewis, 1987]

"Digital calibration" removes D/A and linear gain error by adjusting digital weights [Karanicolas, 1993]

- Bottleneck: Fast, highly linear gain element
- 50-70% of total pipeline ADC power is consumed by interstage amplifiers

Digitally Nonlinearity Compensation

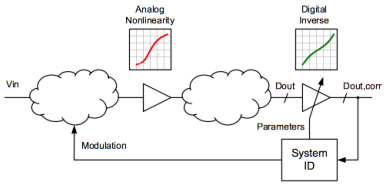

- Need system ID to determine optimum post distortion
- Possible to track variations over time without interrupting normal circuit operation

Correction Concept

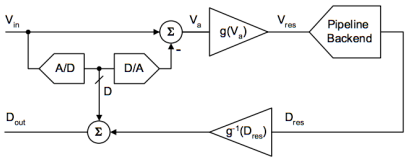

Digital Inverse

- V_{res} is digitized by pipeline backend
- Amplifier nonlinearity cancelled by digital inverse

Prototype

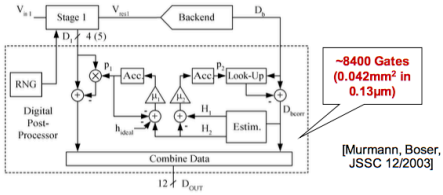

- Open-loop amplifier in first, most critical stage
- Statistics based system ID allows continuous parameter tracking
- Judicious analog/digital co-design

Measurement Results

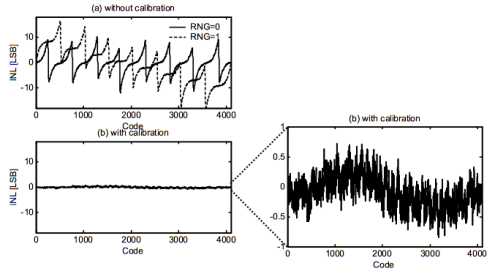

Stage1 Power Breakdown

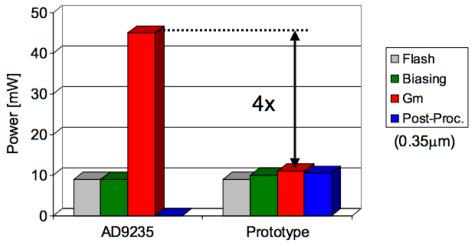

14.7 Li, ISSCC 2018

- A Signal-Independent Background-Calibrating 20b 1MS/s SAR ADC with 0.3ppm INL

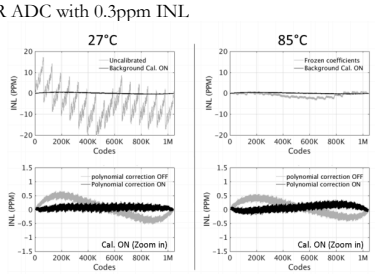

14.1 He, ISSCC 2018

- A 50MHz-BW Continuous-Time $\Delta\Sigma$ ADC with Dynamic Error Correction Achieving 79.8dB SNDR and 95.2dB SFDR

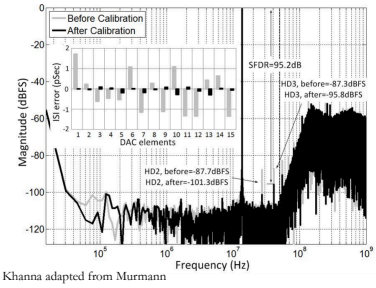

Future

- Today's ADCs are extremely well optimized
- For non-incremental improvements, we must explore new ideas in signal processing that tackle ADC inefficiency at the system level
  - Compressed sensing
  - Finite innovation rate sampling
  - Other ideas?

Big Ideas

- FOM effective to compare performance of ADCs
- Performance trends and limits
  - Moving into digitally assisted designs

Revisit: HW 1

- Read state-of-the-art ADC design publications (2019) and fill out EXCEL spec sheet
- Watch video of ADC performance trends (2014)

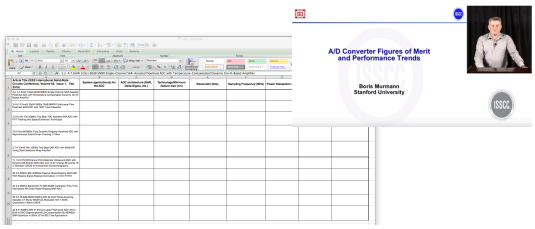

Admin

- Proj 2 out