Lecture 10
Nyquist-Rate ADCs

Trevor Caldwell
trevor.caldwell@awaveip.com

Lecture Plan

Date	Lecture (Wednesday 2-4pm)		Reference	Homework
2020-01-07	1	MOD1 & MOD2	PST 2, 3, A	1: Matlab MOD1&2
2020-01-14	2	MOD*N* + ΔΣ Toolbox	PST 4, B	2: ΔΣ Toolbox
2020-01-21	3	SC Circuits	R 12, CCJM 14	
2020-01-28	4	Comparator & Flash ADC	CCJM 10	3: Comparator
2020-02-04	5	Example Design 1	PST 7, CCJM 14	
2020-02-11	6	Example Design 2	CCJM 18	4: SC MOD2
2020-02-18	Reading Week / ISSCC			
2020-02-25	7	Amplifier Design 1		Project
2020-03-03	8	Amplifier Design 2		
2020-03-10	9	Noise in SC Circuits		
2020-03-17	10	Nyquist-Rate ADCs	CCJM 15, 17	
2020-03-24	11	Mismatch & MM-Shaping	PST 6	
2020-03-31	12	Continuous-Time ΔΣ	PST 8	
2020-04-07	Exam			
2020-04-21	Project Presentation (Project Report Due at start of class)			

Circuit of the Day: Preamp CM Range

- **How do we improve the CM range of a preamp?**

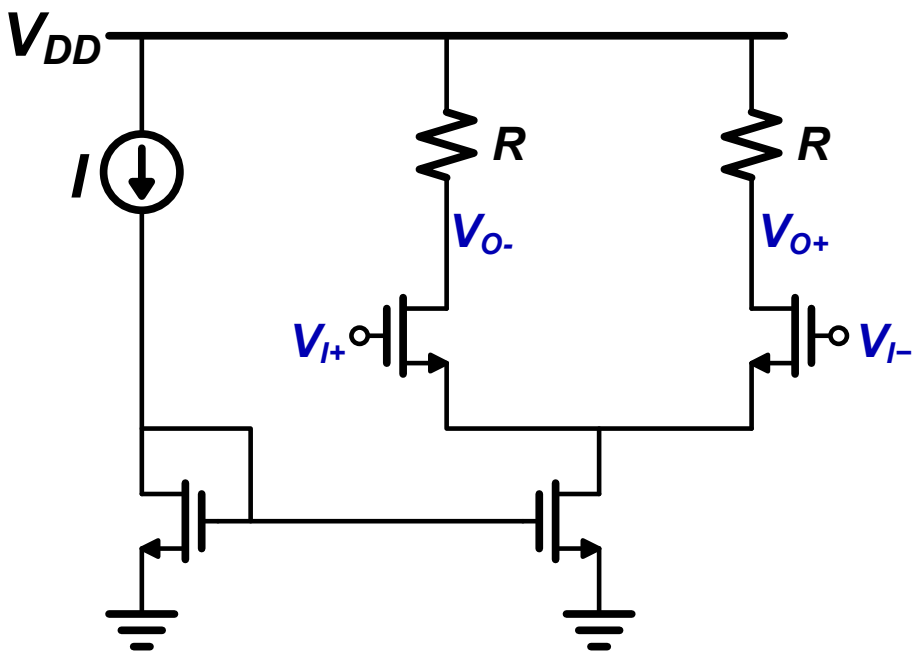

What you will learn…

- **ADC performance limitations**
- **Nyquist-Rate ADCs**
  - **Discuss 5 architectures**
  - **Look at SAR and Pipeline ADCs in more detail**
- **Figure of Merit (FOM)**
  - **Fundamental trade-offs in ADCs**
  - **Bandwidth, Resolution and Power**

ADC Performance

- **Performance Limitations in ADCs**
  - **SQNR, SNR and SNDR**
  - **Offset and Gain Error**
  - **Integral Nonlinearity (INL) Error**
  - **Differential Nonlinearity (DNL) Error**
  - **Sampling Time Uncertainty (Jitter)**

SQNR, SNR and SNDR

- **SQNR: Signal to Quantization Noise Ratio**

  For single-tone sinusoid with input V_P, LSB size V_{LSB}

$$SQNR = 10\log_{10}\left(\frac{V_P^2/2}{V_{LSB}^2/12}\right)$$

- **SNR: Signal to Noise Ratio**

  Typically includes thermal and quantization noise

$$SNR = 10\log_{10}\left(\frac{V_P^2/2}{\frac{V_{LSB}^2}{12} + \overline{V_T^2}}\right)$$

- **SNDR: Signal to Noise and Distortion Ratio**

  Typically includes all noise sources: thermal, quantization, distortion, etc.

Offset and Gain Errors

- **Look at transition points of input analog voltage**

  $V_{0...01}$ is transition from 0 to 1 in output code

- **Offset Error**

$$E_{off} = \frac{V_{0...01}}{V_{LSB}} - \frac{1}{2} LSB$$

- **Gain Error**

$$E_{gain} = \left(\frac{V_{1...1}}{V_{LSB}} - \frac{V_{0...01}}{V_{LSB}} \right) - \left(2^N - 2 \right)$$

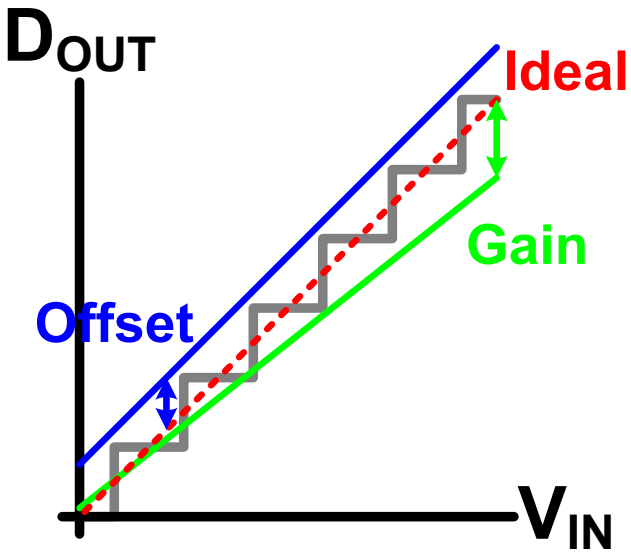

INL and DNL Errors

- **Remove Offset and Gain errors before finding INL and DNL error**

- **INL Error**
  - **Deviation from straight line**
  - **Accumulation of DNL Error**

- **DNL Error**
  - **Transition values should be 1 LSB apart**
  - **Error is deviation from 1 LSB**
  - **Differentiate INL Error**

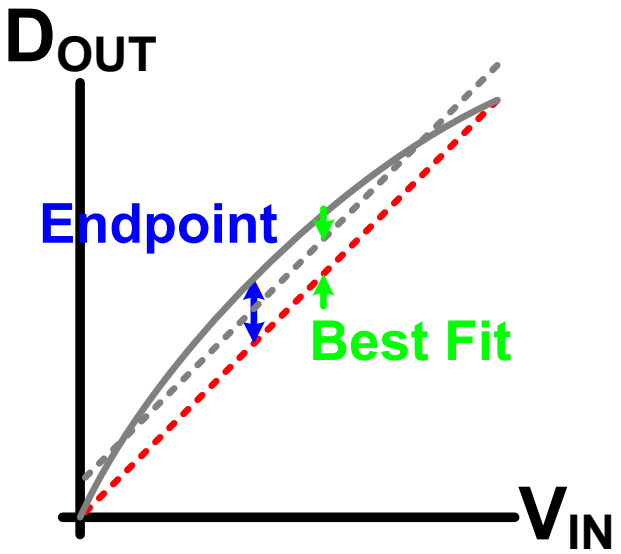

Sampling Time Uncertainty

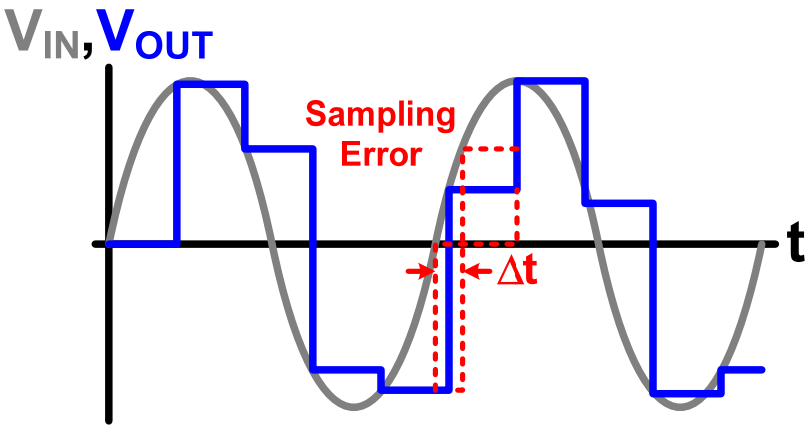

$$V_{IN} = V_P \sin(\omega_{IN} t)$$

- **Error occurs with uncertainty in sampling time**

  Worst case at maximum slope of signal

$$\frac{\Delta V}{\Delta t} = V_P \omega_{IN} \cos(\omega_{IN} t)$$

  Jitter error due to variation in sampling time

$$SNR_{jitter} = 20 \log_{10} \left(\frac{1}{\sigma(\Delta t)\omega_{IN}} \right)$$

Comparison: Nyquist vs $\Delta\Sigma$ ADCs

- **Larger bandwidth in Nyquist-rate ADCs**
  - **Oversampling costs bandwidth, increases accuracy**
  - **Nyquist ADC uses entire output spectrum**
- **Noise in $\Delta\Sigma$ primarily important in signal band**
  - **Larger out-of-band shaped noise gets filtered**
  - **All noise in a Nyquist-rate ADC is important**
- **Nyquist ADC has (typically) no memory**
  - **$\Delta\Sigma$ ADCs have memory; output is a function of many inputs**

Nyquist ADCs

- **Several architectures to choose from:**
  1) **Flash**
  2) **Dual Slope**
  3) **Successive Approximation**
  4) **Cyclic**
  5) **Pipeline**

- **Architectures differ in speed, resolution, latency, power efficiency**

  Each architecture has a different way of converging on the most accurate digital representation

Comparator

- **Basic building block of an A/D converter**
  - **Acts as a 1-bit A/D converter**
- **Output amplifies difference between V_{IN} & V_{REF}**
  - **With a large gain, output is 'digital' at either the positive or negative supply rail**

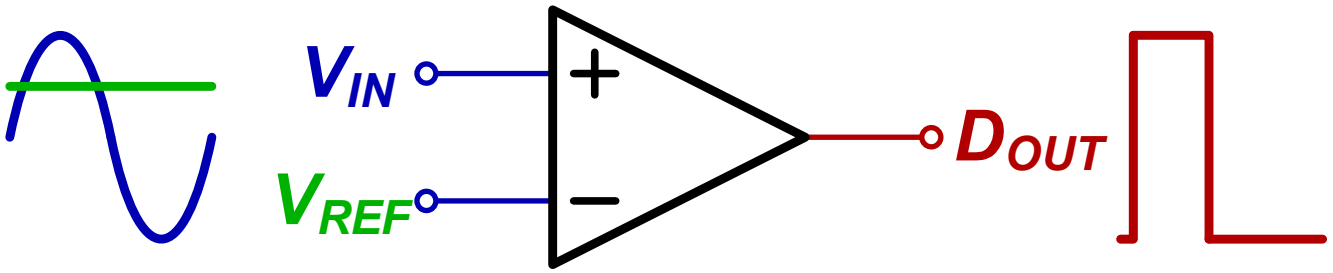

Comparator

- **Not an open-loop amplifier**

  Cascaded low-gain stages are faster for a given power

- **Latched comparator**

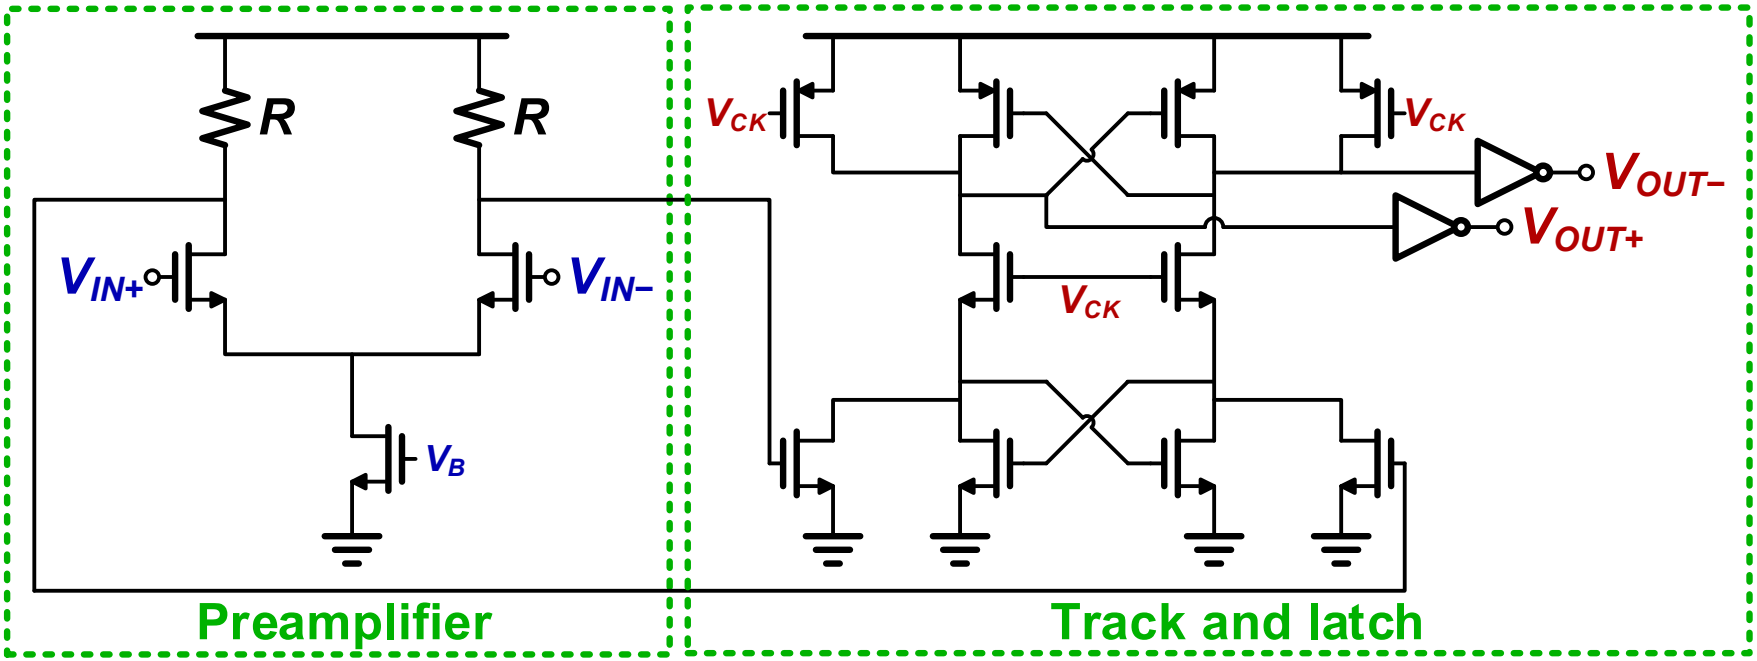

Architecture #1: Flash ADC

- Use 2^N-1 comparators to find the input level

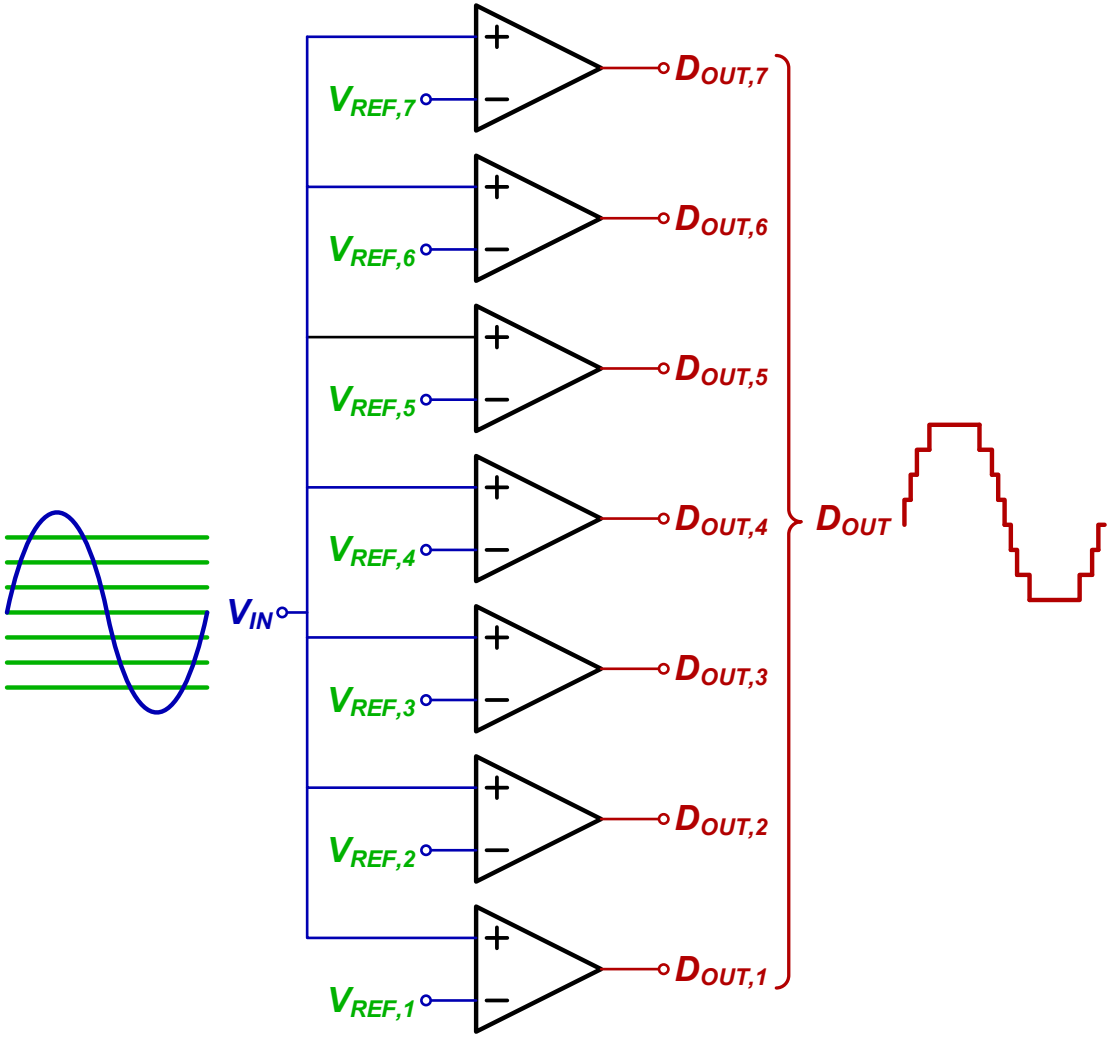

Architecture #2: Dual-Slope ADC

- **Only needs integrator, comparator and counter**
  - **Integrate V_{IN} for fixed time T**
  - **Subtract V_{REF} until V_{INT} reduces to zero**
  - **D_{OUT} proportional to time required to reduce V_{INT} to zero**

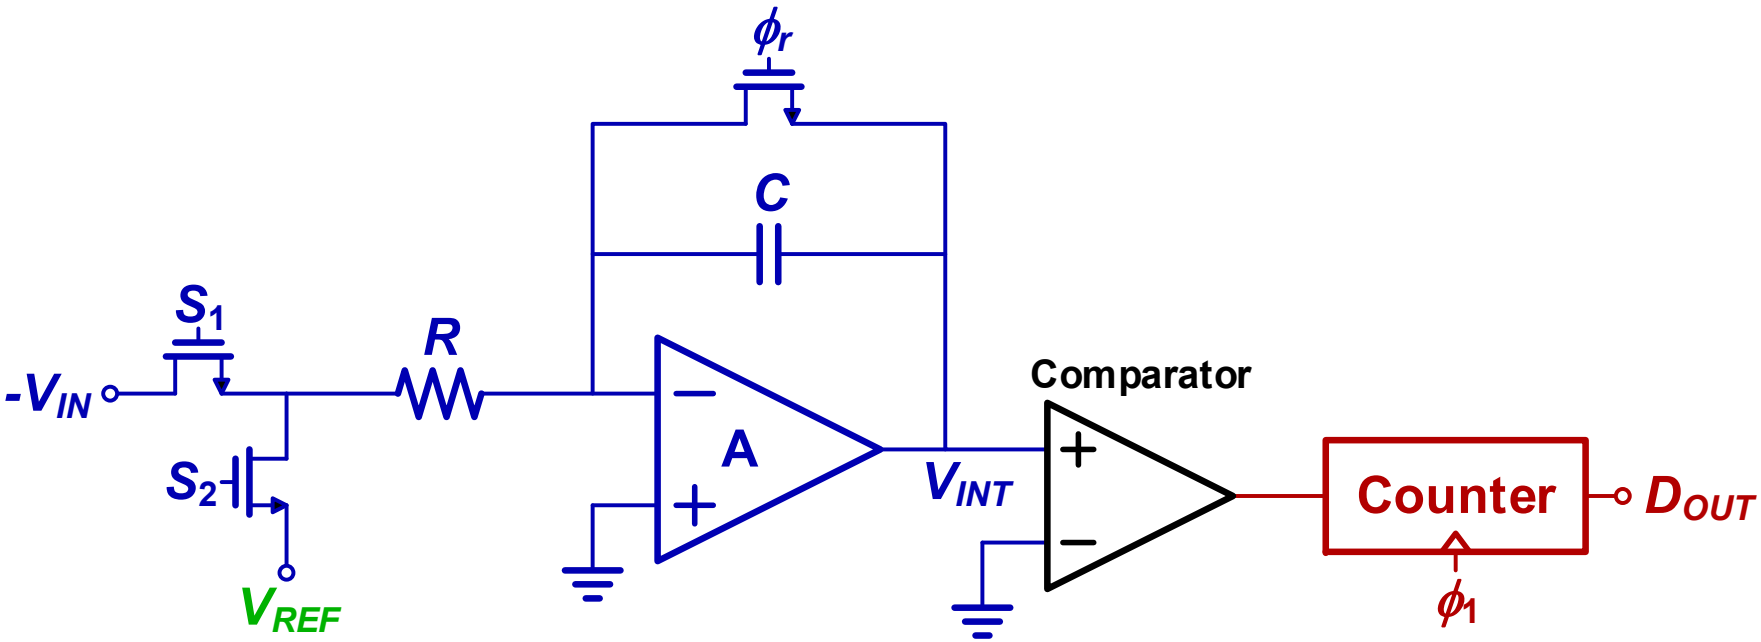

Architecture #2: Dual-Slope ADC

- **Example of integrator output curves**
  - **Slower but more accurate than flash**

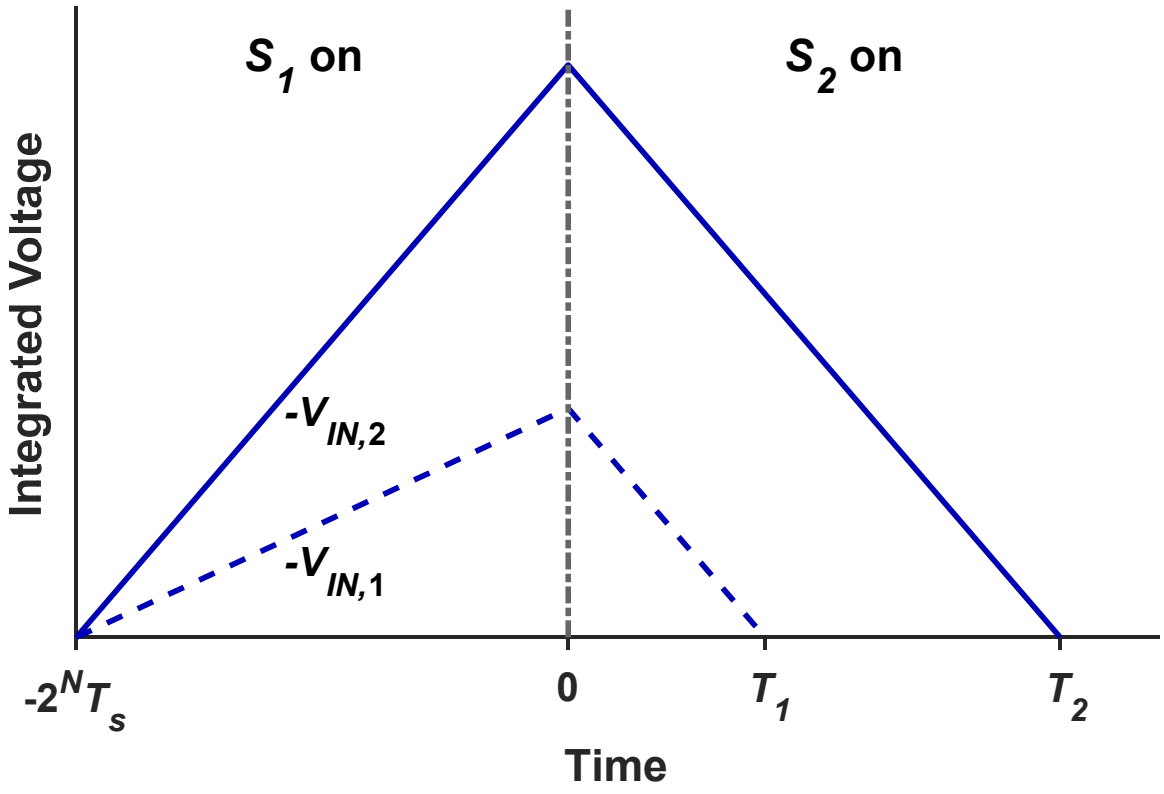

Architecture #3: SAR ADC

- **Binary search algorithm**

  D_1=1, $D_{2...n}$=0, check $V_{D/A}$ against V_{IN}
  If $V_{IN} > V_{D/A}$, D_1=1
  If $V_{IN} < V_{D/A}$, D_1=0
  Now try D_2=1, carry on…

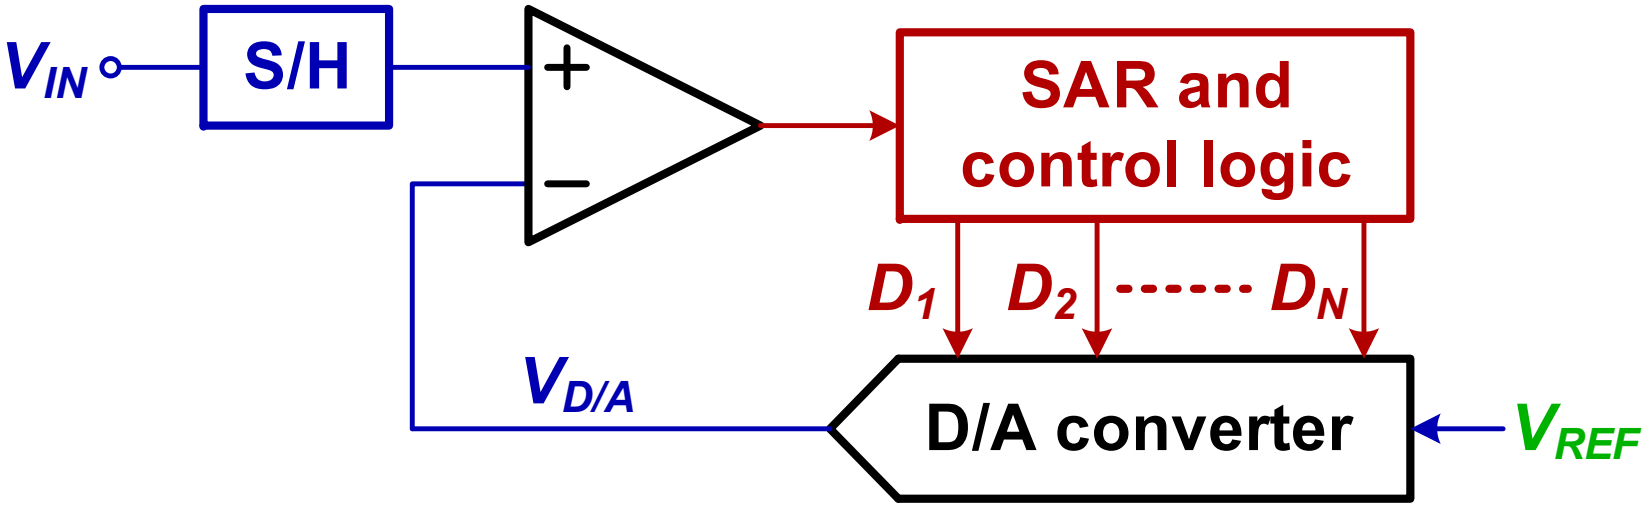

Architecture #4: Cyclic ADC

- **Binary search algorithm**
  - **Doubles the 'residue' every clock cycle**
  - **Compares against the same reference every cycle**

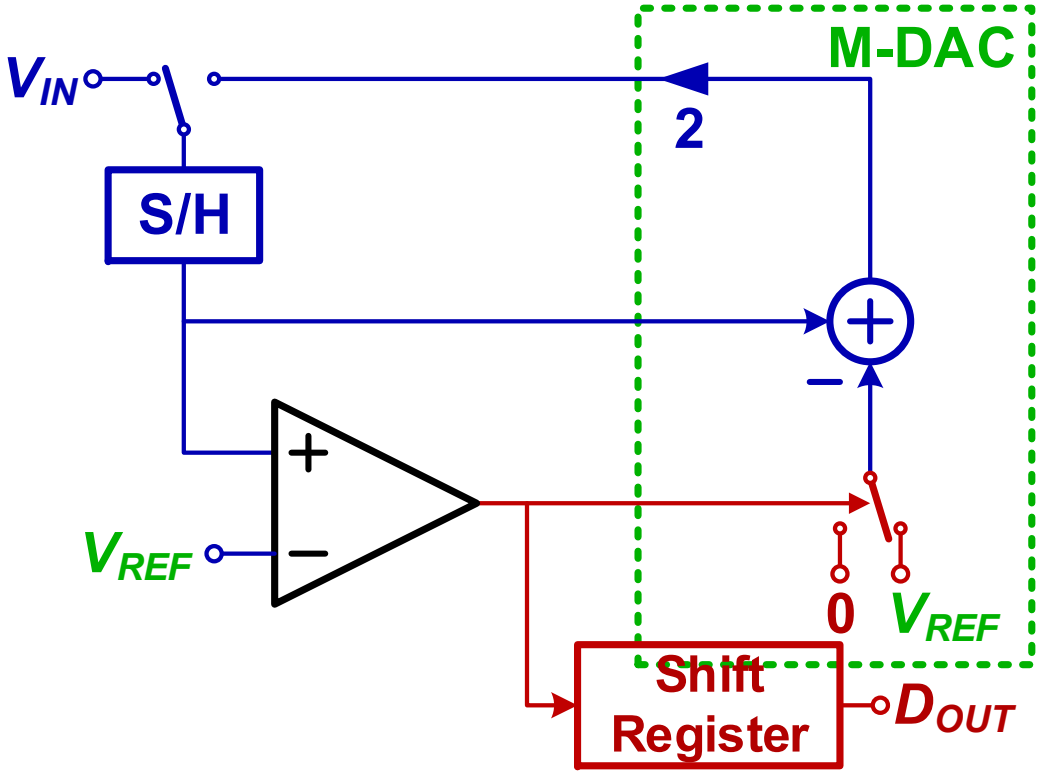

Architecture #4: Cyclic ADC

- **Multiplying DAC (M-DAC)**
  - **Multiplies DAC output with a gain of 2**
  - **Typically switched-capacitor circuits are used**

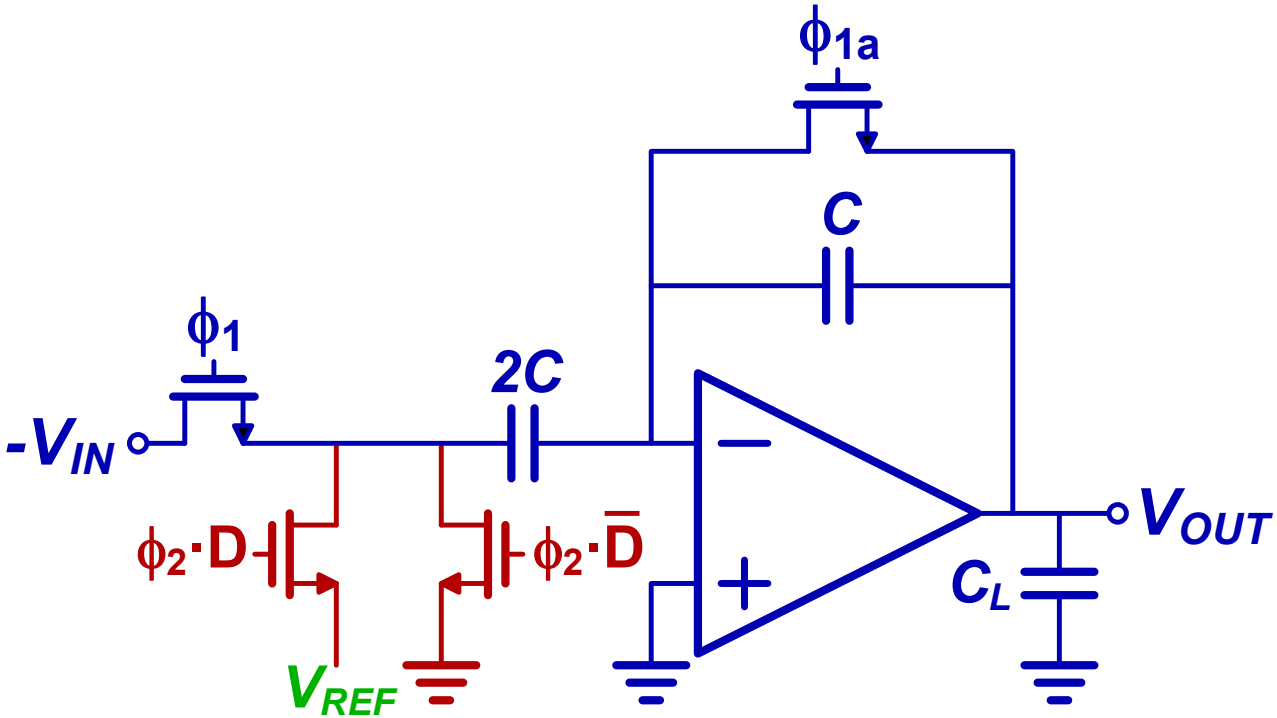

Architecture #5: Pipeline ADC

- ‘Unwrapped’ Cyclic ADC
  - Stages are cascaded rather than re-used
  - Allows higher speed/throughput
  - Increases overall area of converter

- 1-bit per stage design
  - N stages for N-bit resolution

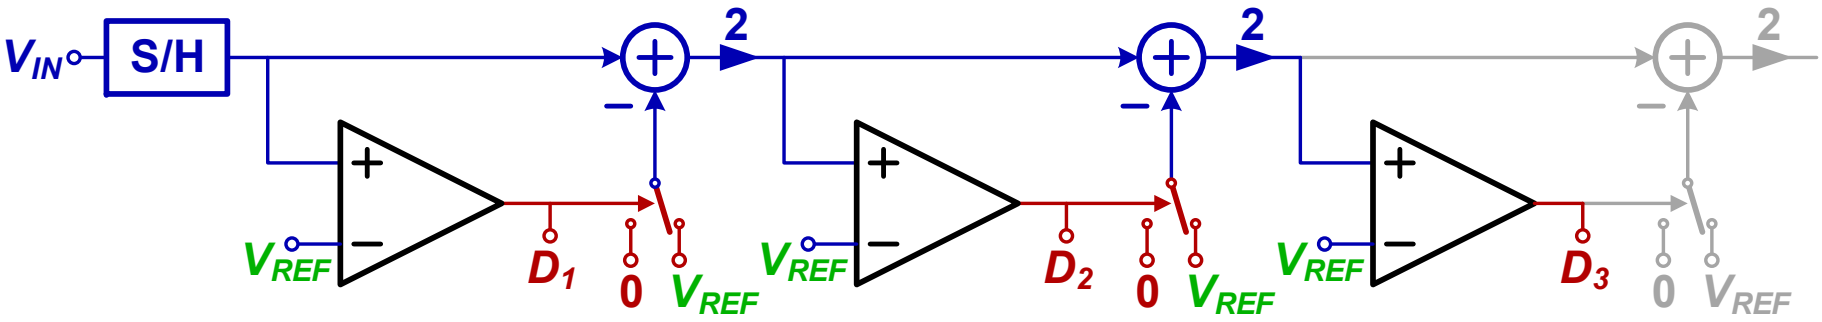

Motivation for SAR ADCs

- **Traditionally used for high SNR lower speed applications (<1 MS/s)**
  - **Industrial sensors**
  - **Power line measurements**
  - **Audio**
  - **Medical imaging**

- **Currently used for high-speed medium-res ADCs**
  - **Up to 1GS/s, 6-10 bit resolution, time-interleaved**
  - **Very power efficient**
  - **Do not require high-gain/linearity amplifier**
  - **Well suited to smaller process nodes**

DAC-based SAR ADC

- How do we implement the binary search?

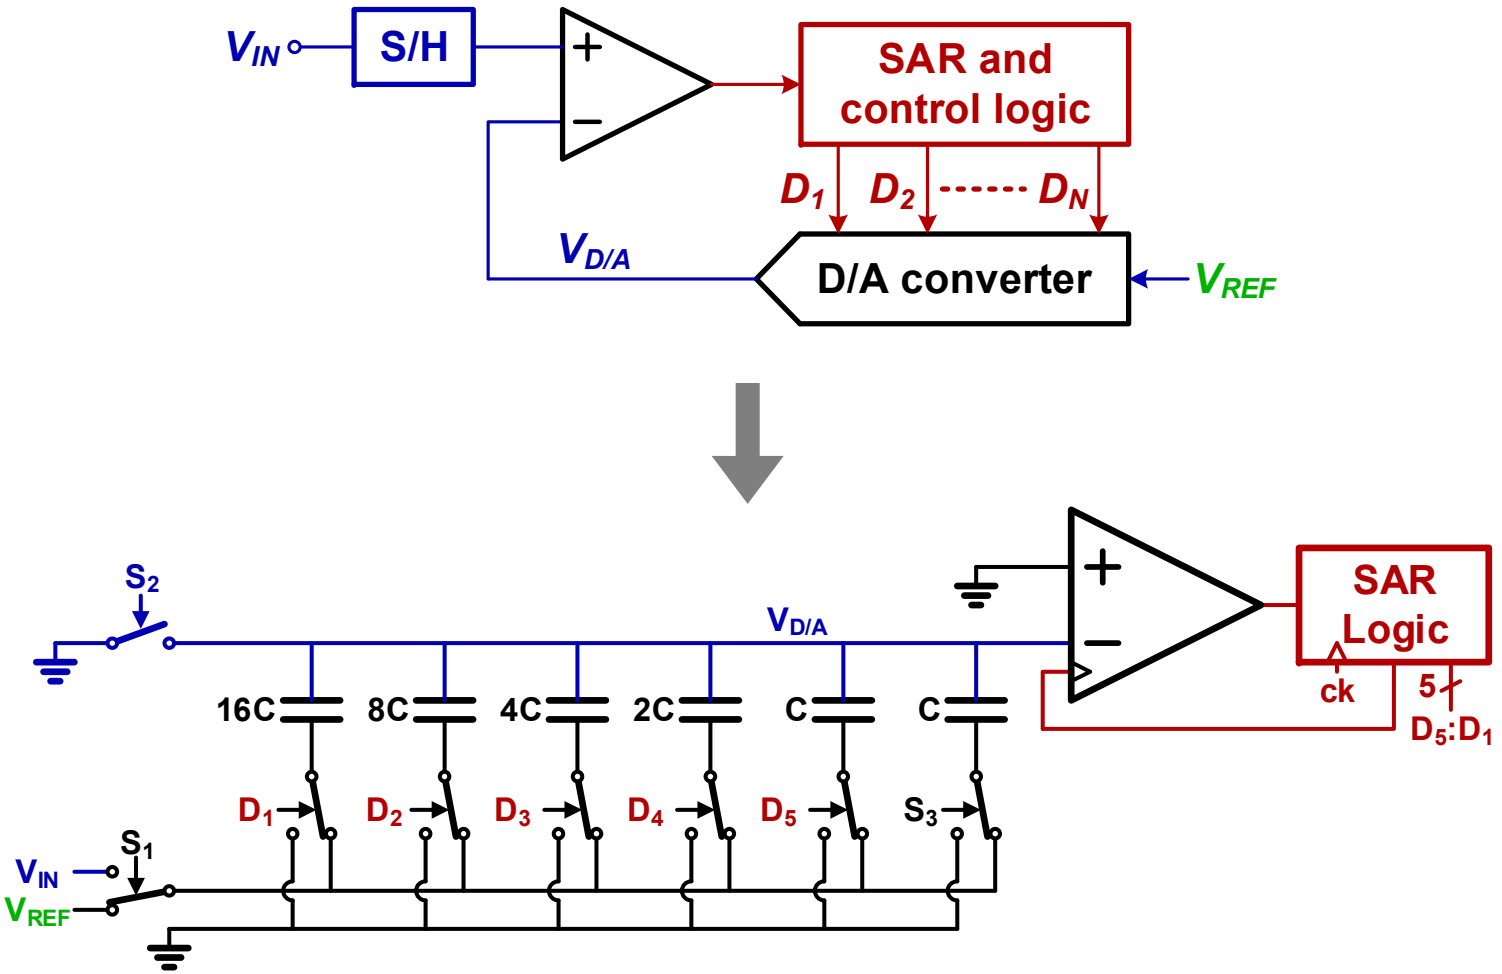

Unipolar Charge-Redistribution SAR

1. **Charge all caps to V_{IN} (S_1,S_2 on; S_3 off)**

   S_1,S_2 on; S_3 off

2. **Switch caps to ground, $V_{D/A}$=-V_{IN}**

   D_1 - D_5,S_1,S_3 on; S_2 off

3. **Switch S_1 to V_{REF}, start bit cycling (S_1 off)**

   Switch D_1 to V_{REF}, if $V_{D/A}$ negative then keep D_1 at V_{REF}

   Find V_{REF} weighting equivalent to V_{IN}

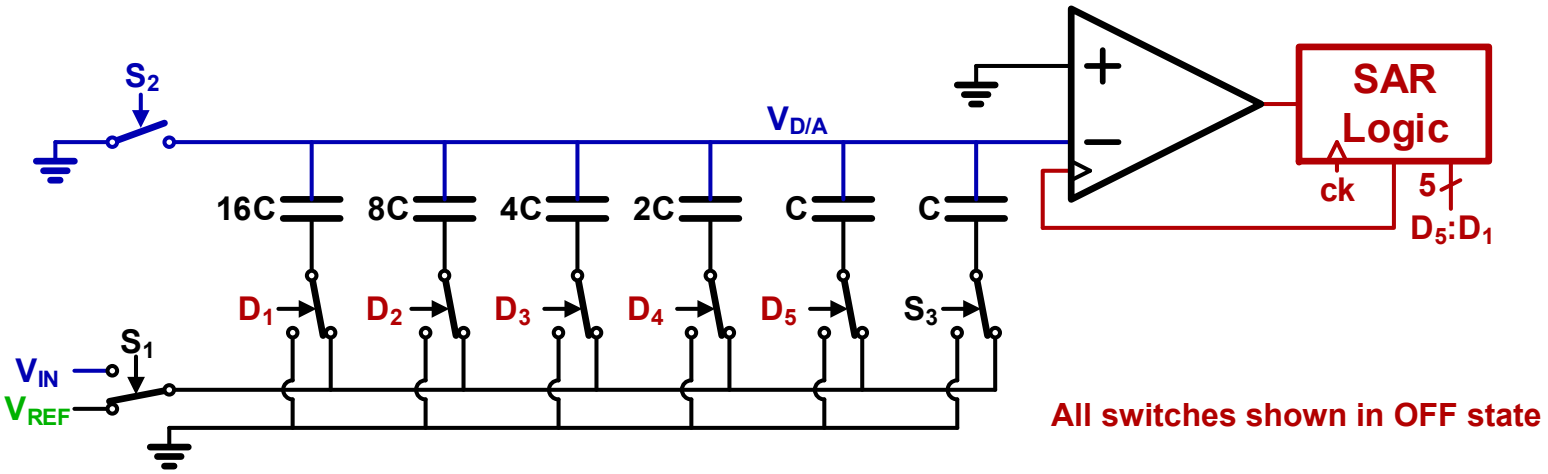

Unipolar Charge-Redistribution SAR

- Extra capacitor C required for exact divide by 2
- DAC capacitor array serves as Sample-and-Hold
- Switching sequence is parasitic insensitive
  - Parasitic capacitors attenuate the signal on $V_{D/A}$
  - Better to keep capacitor bottom plates on the V_{REF} side (not the comparator side)
- Signed conversion with added $-V_{REF}$ input
  - Third state for S_1: V_{IN}, V_{REF}, $-V_{REF}$
  - If $V_x < 0$ on step 2, proceed with V_{REF}
  - If $V_x > 0$ on step 2, use $-V_{REF}$, opposite comparison

SAR Algorithm

- Algorithm finds largest DAC value less than V_{IN}
- Every single bit acquired must be accurate to the resolution of the system
  - DAC linearity
  - Comparator Offset (for multi-bit)
  - S/H offset and linearity
- Comparator must respond to large changes
  - Large overdrive to 0.5 LSB in a single bit cycle period

SAR Algorithm

- **Bandwidth limitation**

  For N-bit resolution, N bit cycles are required, reducing the effective sampling frequency by N

- **Comparison nodes must be quiet before comparator triggers**

Speed Estimate of SAR

- **Two high speed paths**
  1) **Comparator Path: Comparator latching and resetting within clock period**
  2) **DAC Path: Comparator latching and DAC (RC switch network) resolving within period**

- **Comparator path**

  Discussed in Comparator lecture

  Requires fast regeneration time constant + strong reset

- **DAC Path (RC network)**

  Requires fast regeneration time constant in comparator

  Switches sized to reduce R in RC switch network (C sized for noise or matching)

Speed Estimate of SAR

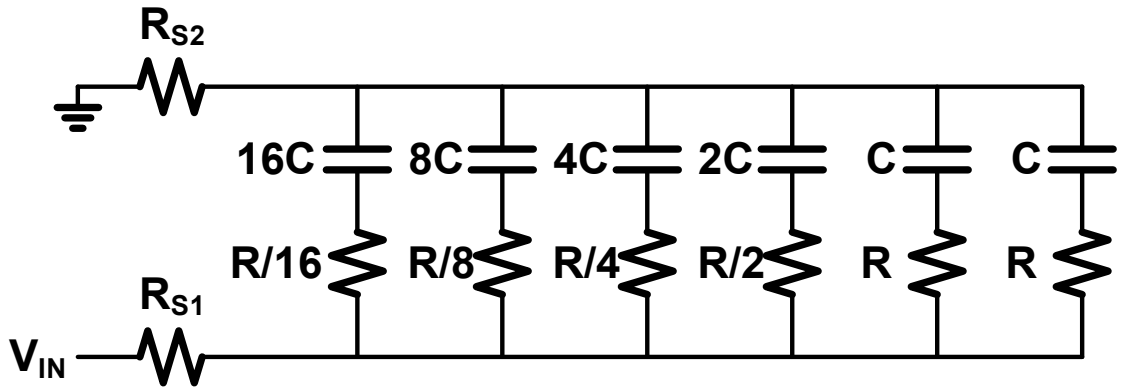

- **RC time constant of capacitor array + switches**
  - **Assume switches are sized proportional to capacitors**

$$\tau_{eq} = (R_{s1} + R_{s2} + R/2^N)2^N C$$

  - **For better than 0.5 LSB accuracy**

$$e^{-T/\tau_{eq}} < \frac{1}{2^{N+1}}$$

  - **Minimum value for the charging time T**

$$T_{DAC} > 0.69(N+1)(R_{s1} + R_{s2} + R/2^N)2^N C$$

Pipeline ADC

- Each stage quantizes the amplified quantization error signal from the previous stage
- Higher throughput, lower complexity

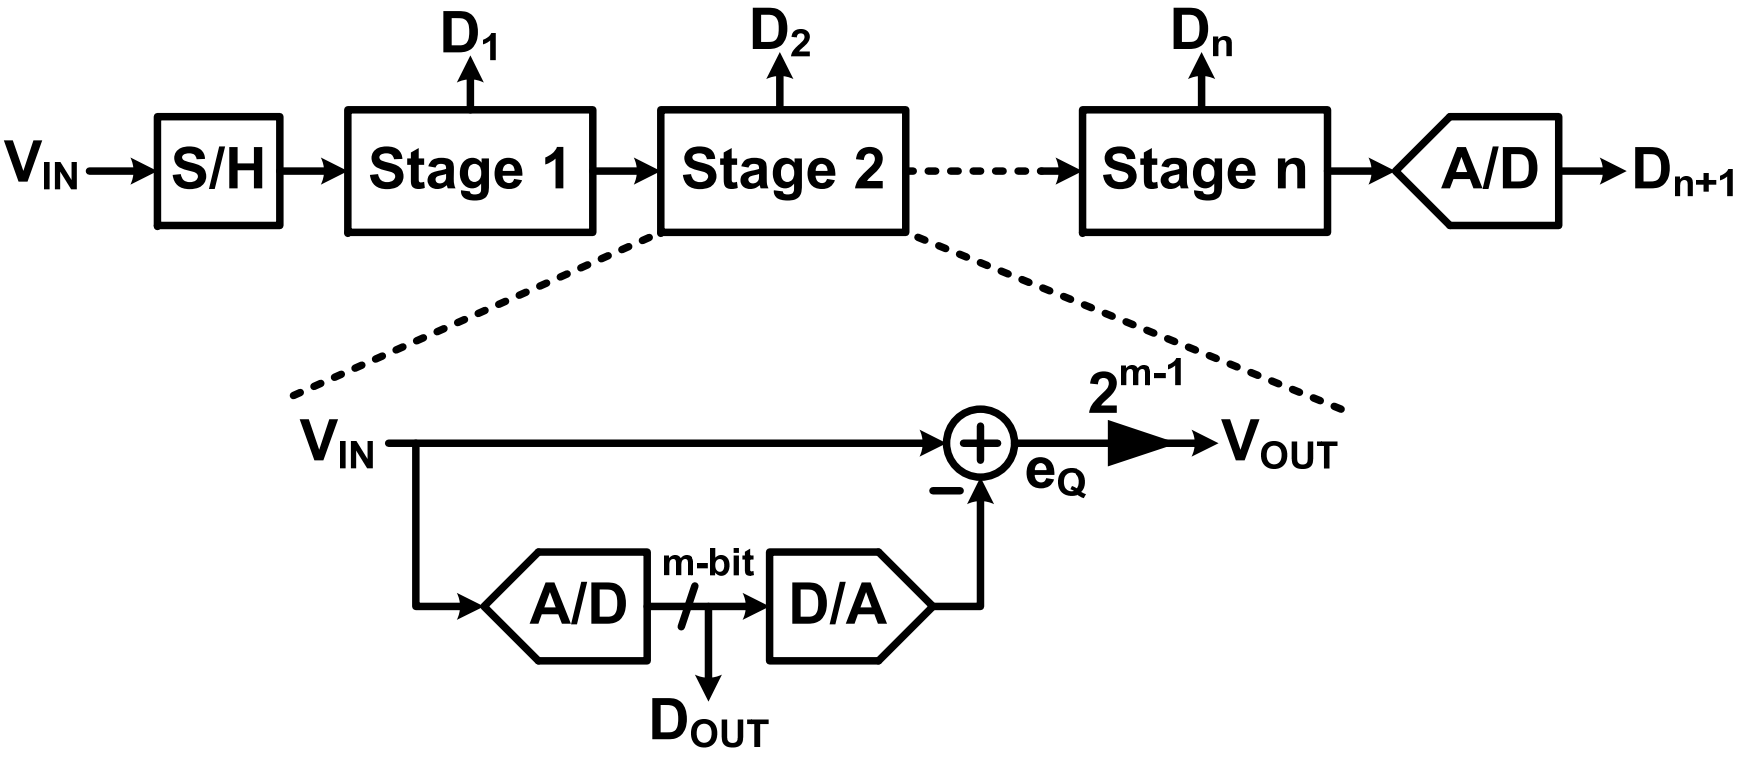

Digital Outputs

- **Add digital outputs of each stage with binary weighting**
  - **Assume some redundancy to correct offset errors**
  - **Inter-stage gain < 2^{m-1}**
- **Example: 10-bit ADC**

  8 stages with 1.5-bit per stage, 2-bit final stage
  - **$D_{1...8}$={-1,0,1}, D_9={-1.5,-0.5,0.5,1.5}**
  - **$D_{OUT}=D_1 2^8+D_2 2^7+D_3 2^6+D_4 2^5+D_5 2^4+D_6 2^3+D_7 2^2+D_8 2^1+D_9 2^0$**

  1024 output levels: 10-bit converter
  - **1-bit per stage x 8 stages + 2-bit final stage = 10 bits**

Gain and Quantization

- **Full Range: Output sensitive to A/D stage offsets**
  - **Offset saturates input to subsequent stage**

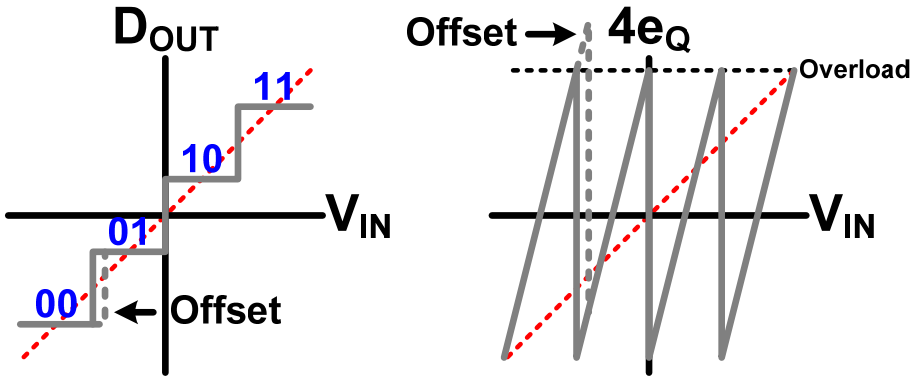

- **Redundancy: Smaller range, less offset sensitive**
  - **Fewer bits/stage, smaller gain to next stage**

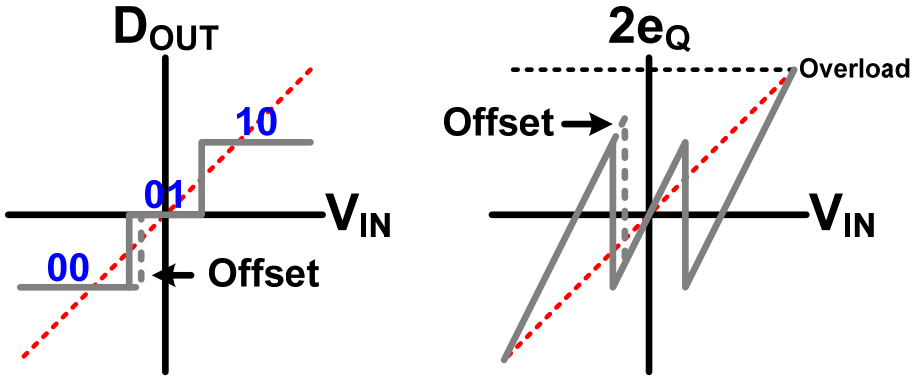

Error Sources

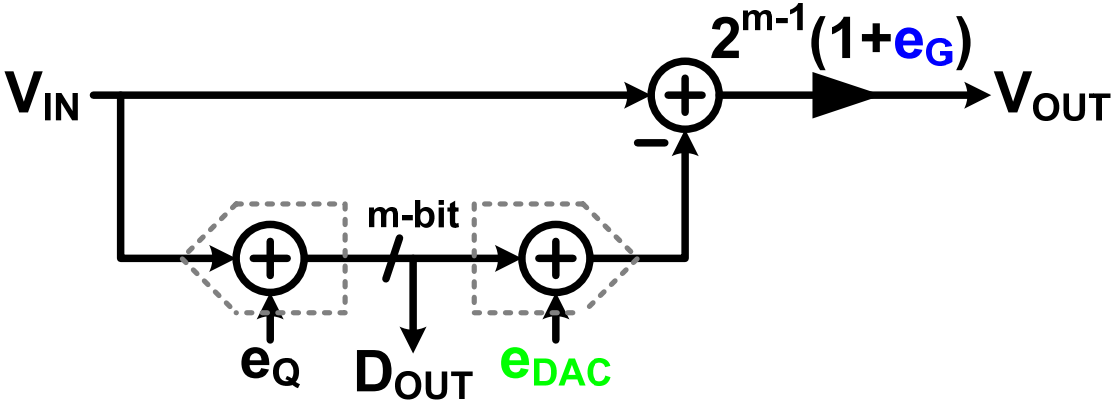

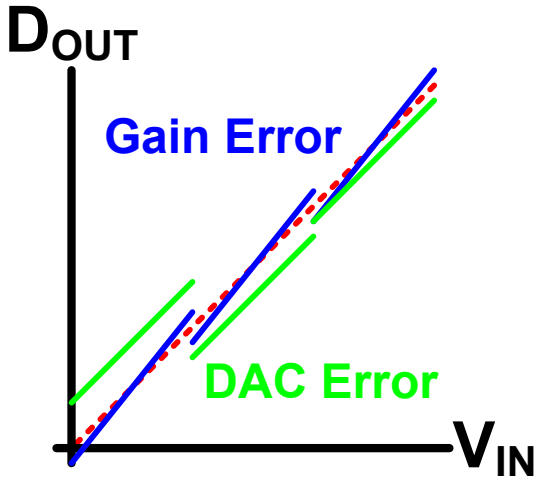

- **Sub-ADC error doesn't matter**
  - **DAC error and gain error cause discontinuity in transfer function**
  - **Signal dependent tones**

Offset Errors

- **Example: 3-bit ADC**

  1.5-bit stage followed by 2-bit final stage

  D_1={-1,0,1}

  D_2={-1.5,-0.5,0.5,1.5}

  $D_{OUT}=D_1 2^1+D_2 2^0$

- **With an offset error, the outputs are the same**

  Quantization error does not saturate the input to the subsequent stage

Offset Errors

- **Two Cases: No Offset; Offset present**

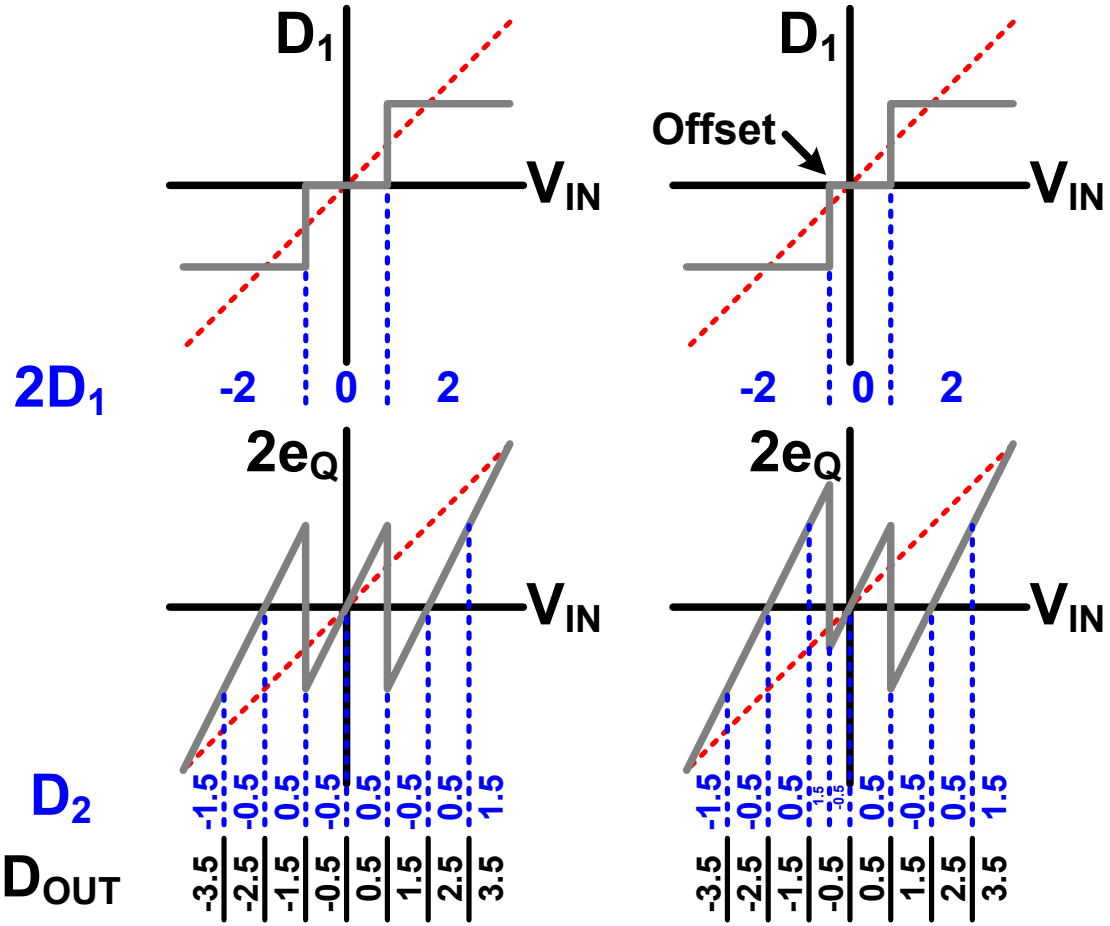

Sample Output Spectrum

- **What is the SQNR?**

  Rectangular Window, N=2048=2^{11}

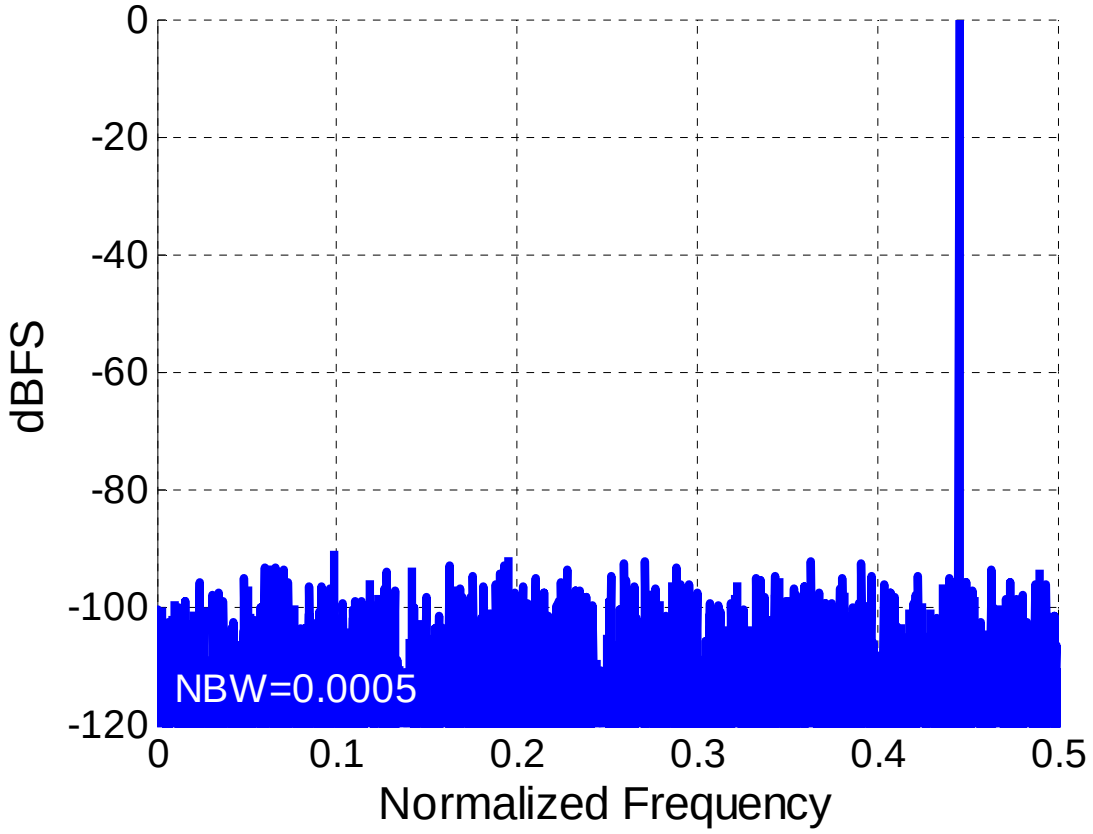

Fundamental Trade-offs between BW, DR and P

DR-P Trade-Off

- **To increase DR at the expense of P, put two ADCs in parallel and average:**

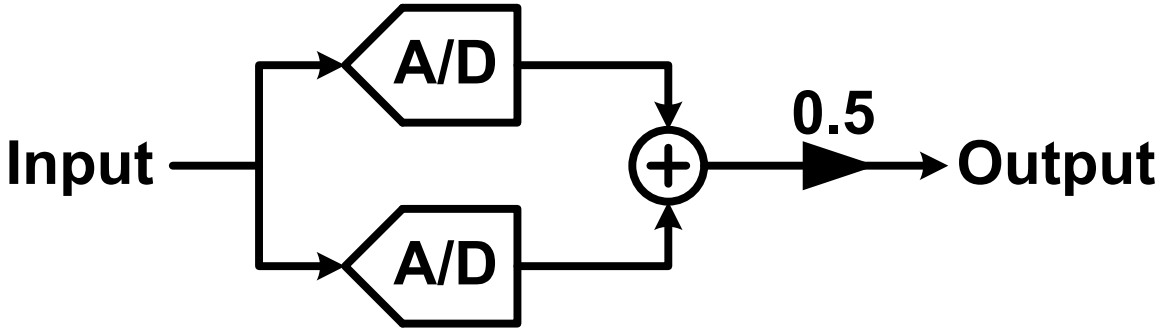

- **Reduces noise by a factor of √2 : DR ↑ 3 dB, but uses twice the power: P ↑ 3 dB**

  Assumes arithmetic requires no power, noise sources are uncorrelated and the source can drive two ADCs.

DR-P Trade-Off

- Can increase DR by 3 dB by reducing T by a factor of 2:

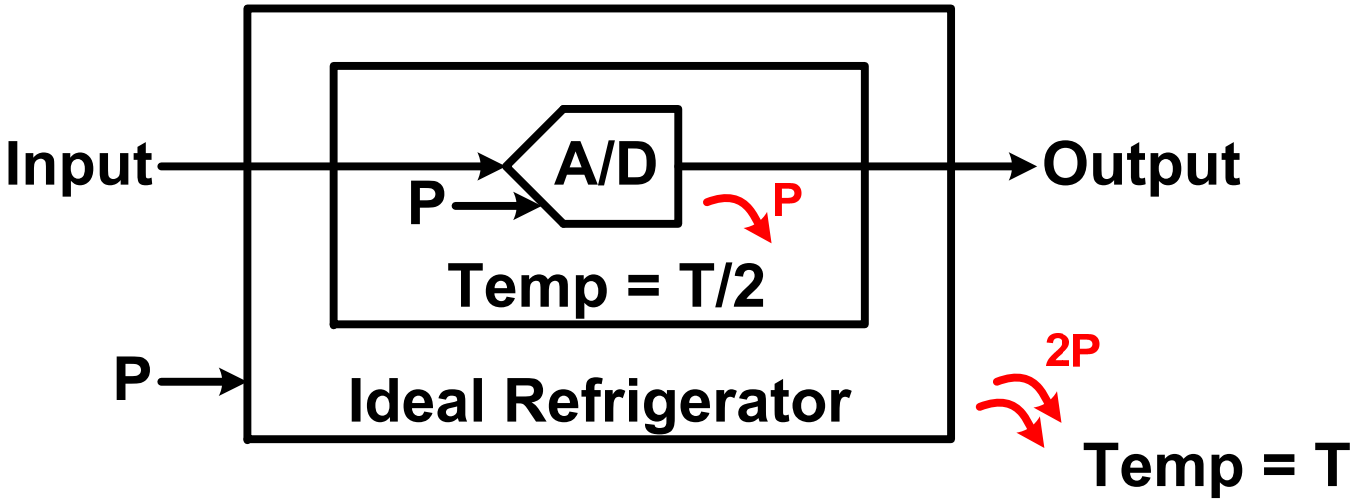

- But this also costs twice the power (based on thermodynamics)

DR-P Trade-Off

- **To reduce P at the expense of DR, "cut the ADC in half"**

  May not be practical if the ADC is already small, but if it can be done,
  P ↓ 3 dB & DR ↓ 3 dB

- **∴ For an ADC of some BW**

X dB in DR costs X dB in P,

or

DR (in dB) – 10log10(P) is a constant

What About BW?

- **Reducing BW by a factor of 2 increases DR by 3 dB but leaves P alone**
  - **Assumes the noise is white (distortion is not dominant)**
  - **Assumes digital filtering takes no power**
- **Time-interleaving two ADCs doubles BW and doubles P, but leaves DR unchanged**
  - **Assumes that interleaving is perfect (can be calibrated with no extra power)**
  - **I/Q processing also doubles BW & P @ same DR**
- **In these examples,**

  DR (in dB) + 10log10(BW/P) is a constant

Resulting FOM (Schreier FOM)

$$FOM = (DR)_{dB} + 10\log_{10}\frac{BW}{P}$$

- For a given FOM, factors of 2 in BW or P are equivalent to a 3 dB change in DR

- Should really include T, but since T is usually assumed to be 300K, omit it

  Steyaert et al.

$$FOM = \frac{4kT \cdot DR \cdot 2BW}{P}$$

Oversampled vs. Nyquist Rate

- For switched-capacitor design, assume power is dominated by kT/C noise
- Assume power is dominated by first stage
  - In $\Delta\Sigma$ ADCs, this is usually more accurate
  - In Pipeline ADCs, more noise is from the later stages
- Which is more power efficient?
  - For the same resolution and speed, oversampled operates N times faster, capacitor is N times smaller
  - $\therefore$ Power is the same
  - It is the 'little things' that make a difference
  - (biasing, clocks, stage sizing, OTA gain, architecture)

Common FOM (Walden FOM)

$$FOM = \frac{P}{2BW \cdot 2^{ENOB}}$$

- **Units of Joules per Conversion Step**

  P in J/s, BW in 1/s, 2^{ENOB} in conversion steps

- **Fundamentally incorrect:**
  1 extra bit costs 2x P, instead of 4x P

  1 extra bit means noise is reduced by 6dB, which requires 4 parallel ADCs, not 2

- **Favours lower resolution ADCs**

Common FOM (Walden FOM)

- **Example: OTA dominated by thermal noise**

SNR = Signal Power / Noise Power

Reduce noise power by 6 dB to get 1 more bit

Noise Power proportional to 1/C

To reduce noise power by 6 dB, or 4x, we need to increase C by 4

Keep V_{EFF} and BW (g_m/C) constant

Increase I_D and W by 4 (like 4 OTAs in parallel)

OTA power increased by 4

$\therefore$ 4x power required for 1 extra bit

Inconsistent with 'Common FOM'

FOM vs. BW (1990-2015)

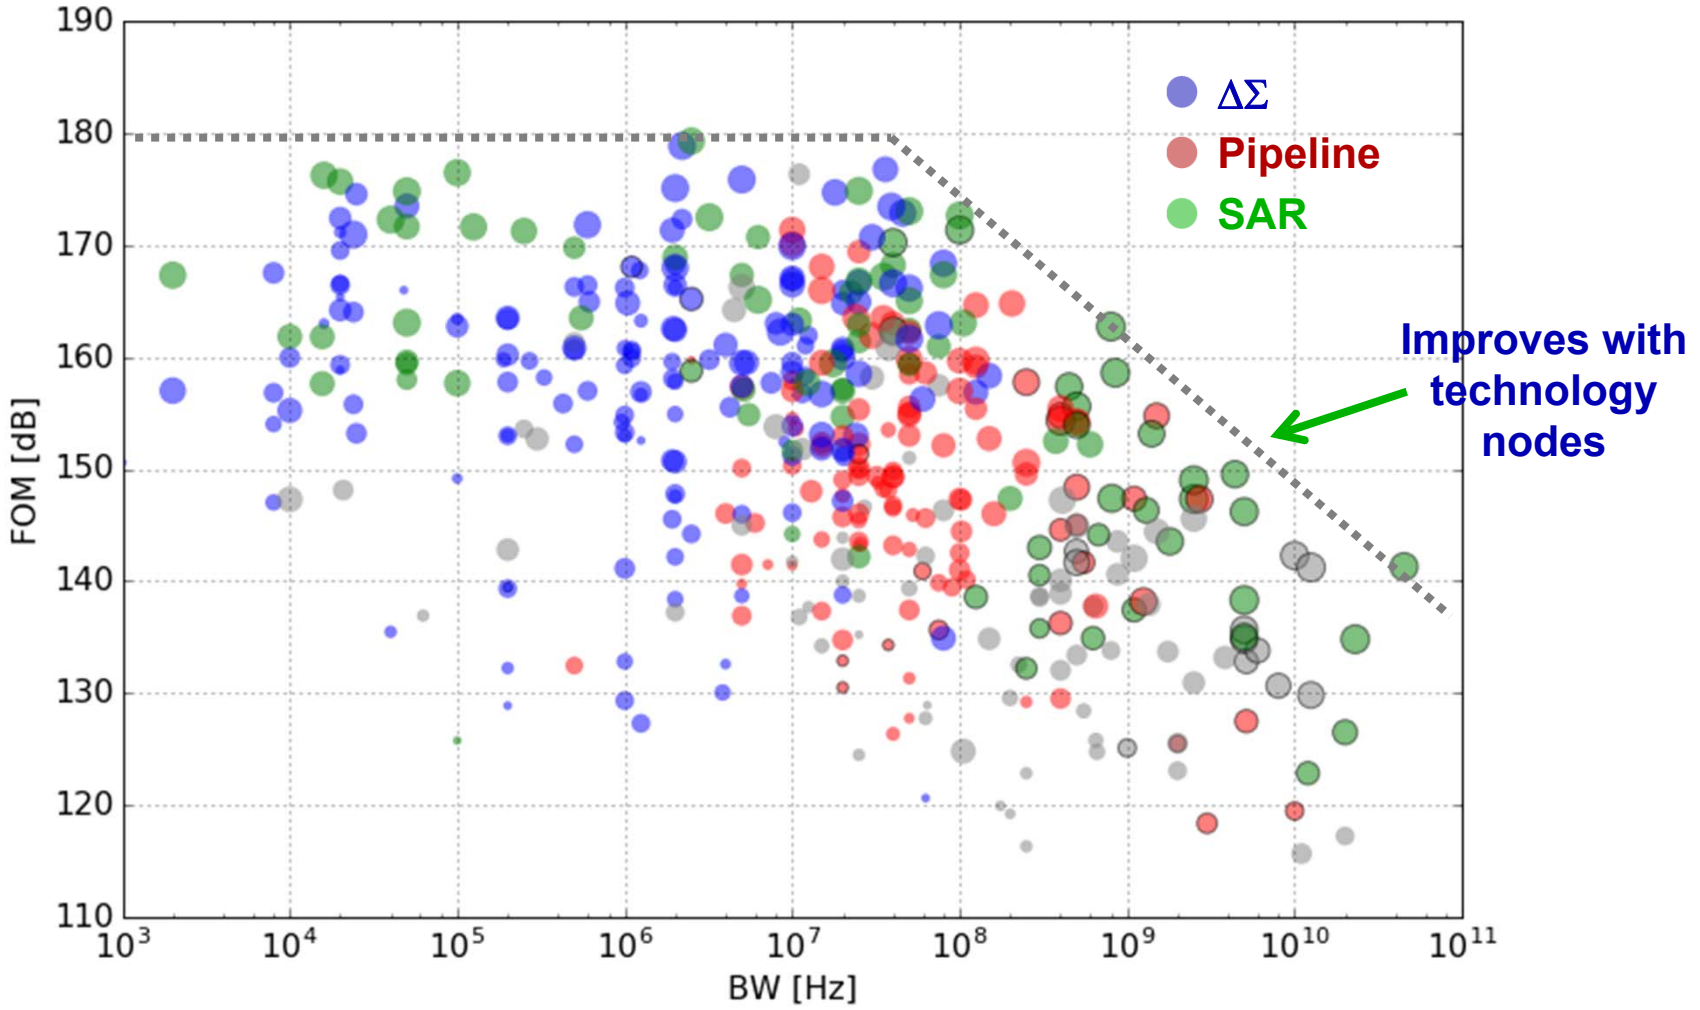

Circuit of the Day: Preamp CM Range

What You Learned Today

- **ADC performance limitations**
- **Examples of Nyquist-rate ADCs**
  - **Basic operation of Pipeline and Successive Approximation ADCs**
- **Fundamental trade-offs between BW, DR and P and ADC FOMs**