

Mismatch-Immune Successive-Approximation Techniques for Nanometer CMOS ADCs

by

Nicholas Andrew Clark Collins

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering)
in The University of Michigan
2017

Doctoral Committee:

Professor Michael P. Flynn, Chair
Professor Zhong He
Associate Professor David D. Wentzloff
Associate Professor Zhengya Zhang

Nicholas Andrew Clark Collins

ncolins@umich.edu

ORCID iD: 0000-0002-9138-9212

© Nicholas Andrew Clark Collins 2017

I dedicate this dissertation to my family.

ACKNOWLEDGEMENTS

There are many people that have helped me complete this work, directly and indirectly. I am truly grateful for everyone that has been a friend or colleague throughout this journey.

I would not have become a Michigan Engineer without first attending Washtenaw Community College. I am grateful to all of my WCC professors, especially Bill Everin, Mohammed Abella, George Kapp, and Vlad Hlasny.

At Michigan, I first thank professors Michel Maharbiz and Fred Terry. Prof. Maharbiz taught my introductory circuits course that solidified my direction of study, and gave me my first research appointment. Prof. Terry was a supportive academic advisor during my transfer to Michigan from Washtenaw Community College, and recommended me to Prof. Flynn who became my Ph.D. advisor.

I wholeheartedly thank my advisor Prof. Michael Flynn for his unfailing generosity. Prof. Flynn gave me an opportunity I am forever grateful for.

Professor Dave Wentzloff has been extremely supportive of my efforts to complete my Ph.D. He has created a professional environment where integrity and honesty are rewarded and cultivated. I equally thank Professor Benton Calhoun of UVA.

Critical assistance and good company was always shared by Fran Doman and Beth Stalaker. Your support can never be appreciated enough. This work would not have been possible without assistance from the EECS Departmental Computing Organization (DCO), most especially Joel van Laven. This work relies heavily on CAD tools and Joel was often a life raft in a storm.

I have been lucky to have worked with many excellent engineers and friends in Prof. Flynn's group, most especially Jorge Pernillo, Andres Tamez, David Lin, Li Li, Hyo-Gyuem Rhew, Hyungil Chae, Mohammad Gharamani, Jeff Fredenburg, Aaron Rocca, Jaehun Jeong,

John Bell, Chunyang Zhai, Batu Dayanik, Yong Lim, Sunmin Jang, Kyojin Choo, Daniel Weyer, Adam Mendrela, and of course Lu Jie. I am also thankful to have had the opportunity to work with friends from other circuits groups, especially Dave Fick, Matt Fojtik, Nate Roberts, Osama Khan, and Mohammad Faisal.

Finally I am most grateful for the unyielding support from my family. I am extremely privileged to have been born to intelligent and loving parents who value education and culture of all kinds. They rewarded me for being interested in and excited about anything I attempted to do, and still encourage me to this day. Now that I have my own children, I strive to be a fraction of the parents they are.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vii
LIST OF TABLES	ix
ABSTRACT	x
CHAPTERS	
1 Introduction	1
1.1 Analog-to-Digital Converters: Context and Background	1
1.2 ADC Architectures	2
1.2.1 Flash ADCs	3
1.2.2 Pipeline ADCs	4
1.2.3 Sigma-Delta ADCs	5
1.2.4 SAR ADCs	6
1.3 Process Scaling and ADC Design	8
1.3.1 Process Scaling Overview	9
1.3.2 How Process Scaling Affects ADC Architectures	11
1.4 Recent SAR Techniques and Innovations	12
1.4.1 Power Reduction	13
1.4.2 Increased Speed and Conversion Bandwidth	17
1.4.3 Comparator Noise Mitigation	22
1.4.4 Mismatch Mitigation	23
1.5 Contributions of This Work	23
2 Completely Reconfigurable Capacitor DAC	25
2.1 Extending Scaling Benefits to Medium and Higher Resolution ADCs	25
2.1.1 Mismatch in Traditional DACs	25
2.1.2 Linearity Metrics in ADCs and DACs	27
2.1.3 How Element Mismatch Affects Linearity	28
2.2 Overcoming Mismatch Constraints Through Reconfigurability	29

2.2.1	Canceling Mismatch via Reconfigurability	29
2.2.2	Reconfiguration Techniques	30
2.3	Completely Reconfigurable DAC Details	35
2.3.1	Array Structure and Operation	35
2.3.2	Cell Structure	40
2.3.3	Brief Note Regarding Novelty	47
2.3.4	Other Contributors to This Work	48
3	A Mismatch-Immune 12-bit SAR ADC	49
3.1	ADC Design Considerations	49
3.2	Implementation of Recent Innovations	50
3.2.1	Comparator Design	50
3.2.2	Asynchronous Timing	52
3.2.3	DAC Switching Method	52
3.2.4	Flexible Transfer Function	55
3.3	Other Design Considerations	55
3.3.1	Voltage References	55
3.3.2	Sampling Technique	56
3.4	Measurements	60
3.4.1	Die Photo	60
3.4.2	Measurement Method	61
3.4.3	Linearity Enhancement Results Via Reconfiguration	66
4	Conclusion	72
4.1	Summary	72
4.2	Future Work	74
	BIBLIOGRAPHY	76

LIST OF FIGURES

Figure

1.1	Bandwidth vs. Resolution for general ADC architectures, from Murmann [4].	3
1.2	Flash ADC architecture from [5]	4
1.3	Pipeline ADC architecture diagram from [5]	5
1.4	Second-order Sigma-Delta modulator diagram from [9]	6
1.5	Basic SAR architecture diagram from [5]	7
1.6	3b single-ended SAR conversion example	8
1.7	Traditional vs. Split MSB switching, from [25].	15
1.8	Comparator output decision time vs. input voltage.	18
1.9	Comparison of binary vs. non-binary search highlighting redundant search area, from [33].	20
1.10	Example of ideal vs. error tolerance in binary and two types of redundant switching schemes, from [32].	21
2.1	Example of 3-bit capacitor DAC assembled from unit element C_{UNIT}	26
2.2	Initial mismatched grouping (red) vs. reconfigured, corrected grouping (blue).	30
2.3	Before-and-after example of arbitrary layout-independent DAC reconfiguration.	31
2.4	Comparison of the 3 described reconfiguration algorithms at 0.1%, 1% and 3% unit mismatch.	33
2.5	Weighted fractional mismatch $(2^{i-1}\gamma_i)^2$ of each SAR grouping of a 12-bit differential DAC with 3% unit element mismatch across 4 different DAC configurations. The resulting ENOB and sum of $(2^i - 1)\gamma_i)^2$ of each configuration are noted in the legend.	34
2.6	Conceptual diagram of DAC and ADC system.	35
2.7	Detailed DAC and system diagram.	36
2.8	Simplified addresser schematic	38
2.9	Simplified 3-D diagram of 4x or 1x cell structure.	41
2.10	Simplified schematic of one DAC cell.	42
2.11	Schematic of individual memory bit cell.	43
2.12	Schematic of global bit vs. local bit match indicator cell.	44
2.13	Schematic of miniature state machine in each DAC cell.	44
2.14	Schematic of unitcap sub-DAC cell for final two bits of DAC resolution.	46

3.1	Schematic of dynamic comparator from [57].	51
3.2	Comparator timing diagram.	51
3.3	Asynchronous timing diagram of ADC.	53
3.4	DAC voltages during SAR cycle for output code 3103. Note the first 4 switching events are symmetric and differential, while the final 8 are set-and-down as in [27].	54
3.5	Conventional SAR DAC example. A single-ended 3-bit DAC samples signal V_{IN} onto the capacitors' bottom plates, and during conversion will switch between V_{REF} and GND . From [5].	57
3.6	Cell-level schematic and timing diagram for hybrid sampling method.	58
3.7	Die photograph of ADC.	60
3.8	Block diagram of test setup.	61
3.9	Simplified schematic of PCB used for ADC testing.	62
3.10	Photo of PCB used for ADC testing.	63
3.11	ENOB vs. DAC settling time settings on PCB <i>without</i> series resistance.	63
3.12	ENOB vs. DAC settling time settings on PCB <i>with</i> series resistance.	64
3.13	Noise floor, 2nd, and 3rd harmonics before and after noise averaging. The X-axis varies the total resolution of the ADC from an 11-bit ADC with a half-scale input to a 12-bit ADC at full-scale input.	65
3.14	Random scrambling algorithm in progress: ENOB vs. reconfiguration iteration.	66
3.15	ENOB vs. MSB used as non-binary group).	67
3.16	Noise-averaged ADC output spectrum (best measured): $F_{in} = 6\text{MHz}$, $F_s = 20\text{MHz}$	68
3.17	ENOB vs. Input Frequency and Sampling Frequency.	69
3.18	Measured ADC ENOB, optimized vs. default configuration, separated by test PCB number.	70
3.19	ENOB improvement histograms for PCB 1 and PCB 2.	71
3.20	ADC core power consumption (mW)	71
4.1	Energy vs. Resolution comparison from [4].	73
4.2	Bandwidth vs. Resolution comparison from [4].	73
4.3	Walden FoM vs. speed comparison from [4].	74
4.4	Schreier FoM vs. speed comparison from [4].	75

LIST OF TABLES

Table

1.1	Influence of scaling on MOS device characteristics.[14], [15]	9
1.2	Influence of scaling on interconnect characteristics. [14]	10
2.1	Default Binary-Weighted Array Programming	39
3.1	Sampling Techniques Summary	57

ABSTRACT

Mismatch-Immune Successive-Approximation Techniques for Nanometer CMOS ADCs

by

Nicholas Andrew Clark Collins

Chair: Michael P. Flynn

During the past decade, SAR ADCs have enjoyed increasing prominence due to their inherently scaling-friendly architecture. Several recent SAR ADC innovations focus on decreasing power consumption, mitigating thermal noise, and improving bandwidth, however most of those that use non-hybrid architectures are limited to moderate (8-10 bit) resolution. Assuming an almost rail-to-rail dynamic range, comparator noise and DAC element mismatch constraints are critical but not insurmountable at 10 bits of resolution or less in sub-100nm processes. On the other hand, analysis shows that for medium-resolution ADCs (11-15 bits, depending on the LSB voltage of the converter), the mismatch sizing constraint still dominates unit capacitor sizing over the $\frac{kT}{C}$ sampling noise constraint, and can only be mitigated by drawing increasingly larger capacitors.

The focus of this work is to extend the scaling benefits of the SAR architecture to medium and higher ADC resolutions through mitigating and ultimately harnessing DAC element mismatch. This goal is achieved via a novel, completely reconfigurable capacitor DAC that allows the rearranging of capacitors to different trial groupings in the SAR cycle so that mismatch can be canceled. The DAC is implemented in a 12-bit SAR ADC in 65nm CMOS, and a nearly 2-bit improvement in linearity is demonstrated with a simple reconfiguration algorithm.

CHAPTER 1

Introduction

1.1 Analog-to-Digital Converters: Context and Background

Analog-to-Digital Converters (ADCs) are omnipresent in modern (computing) society, because they are the boundary between the analog world and the digital world. The analog world represents the real-life continuous-time and continuous-voltage nature of electrical signals in the observable universe, whereas within the digital world data is represented with binary logic and discrete-time signals in computing systems. Every time a real world signal is measured for a digital system of any kind, somewhere in the signal path there is an ADC performing that measurement. ADCs work by comparing a real world signal against a reference signal, and represent that signal's relative value with a binary number. This process is called “quantizing” or “quantization.”

ADCs have three major attributes that determine their utility:

- **Resolution:** This is the number of “bits” of quantization the ADC performs, or the precision of measurement comparison.
- **Sampling Rate:** This is the frequency of measurement the ADC is capable of, which in turn affects the bandwidth of input signal that the ADC can measure.
- **Power Consumption:** The amount of power consumed by the ADC.

Examples of major applications for ADCs and their required attributes are:

- **Cable Modems:** High data-rate modulation schemes such as 256 and 1024-QAM require medium-high resolution ADCs with ample bandwidth at the expense of power consumption [1], [2].
- **Touchscreen Sensors:** Capacitive sensors (such as used in the iPhone and other mobile computing devices) require ADCs to detect the change in capacitance where one's finger is placed. These have lower bandwidth requirements but power consumption is critical in mobile applications [3].
- **Hard-Disk Drives (HDDs):** The read channels of HDDs require extremely high bandwidth and have historically been a market driver for state-of-the-art flash ADCs.
- **Audio:** The highest resolution ADCs (24 bits) are used for quantizing audio signals at relatively low (audio) bandwidth (20kHz - 192kHz).
- **Image Sensors:** CMOS Image sensors require multiple medium (10-14 bit) resolution ADCs.

1.2 ADC Architectures

Analog to digital converters can generally be categorized into four basic architectures: Flash, Pipeline, Sigma-Delta, and SAR (Successive Approximation Register). Each architecture is best suited for a particular balance between bandwidth and resolution. To illustrate this balance across the four architectures, Figure 1.1 plots the bandwidth vs. resolution of every ADC published at the International Solid-State Circuits Conference (ISSCC) between 1997-2016 as recorded in Murmann's *ADC Performance Survey* [4], by architecture. Hybrid architectures are not included, however time-interleaved examples of "pure" architectures are included, which on this plot are most of the highest bandwidth ADCs.

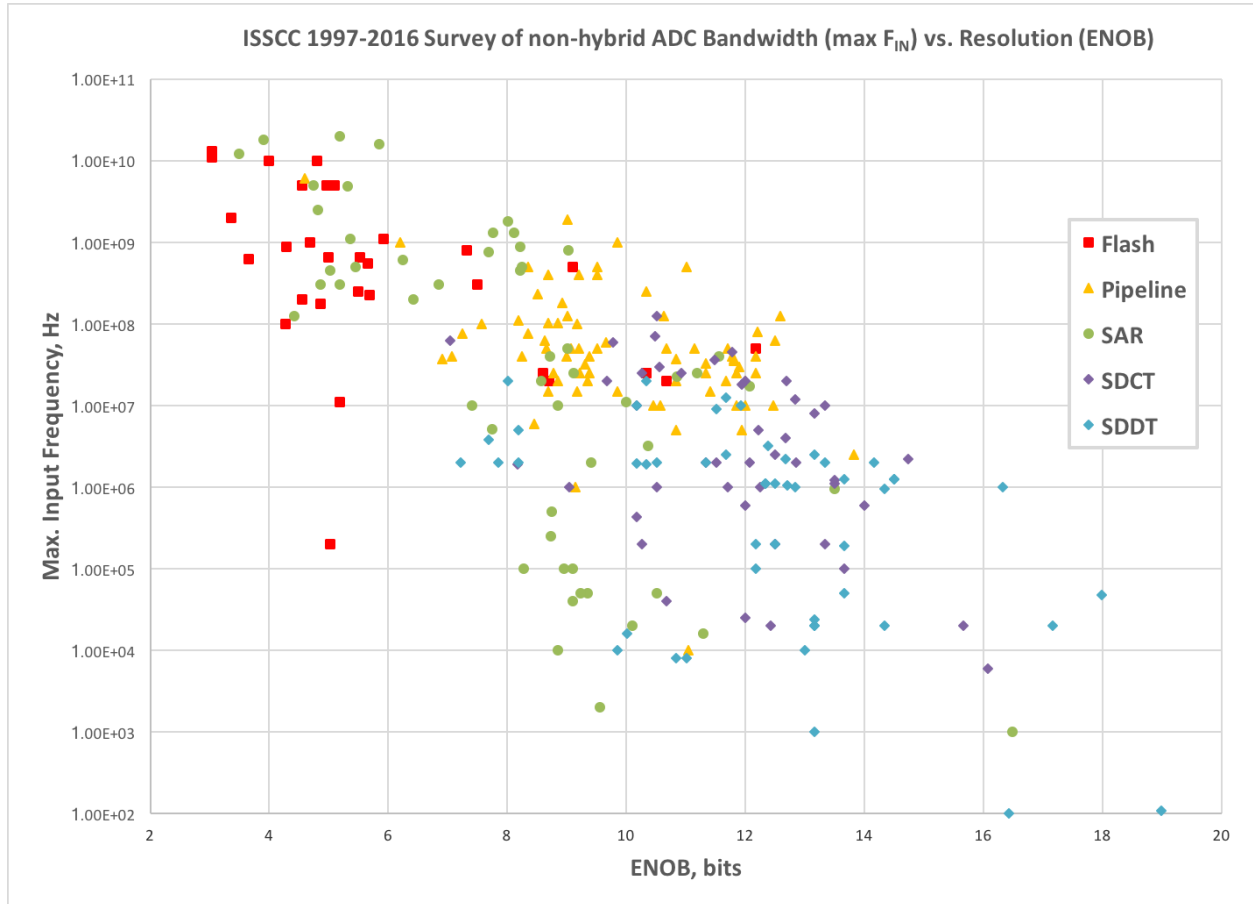


Figure 1.1: Bandwidth vs. Resolution for general ADC architectures, from Murmann [4].

1.2.1 Flash ADCs

The flash ADC architecture is best suited for high-bandwidth, low to medium resolution applications and a basic example is shown in Figure 1.2. Flash ADCs are what is referred to as a “Nyquist” ADC, meaning they adhere to the Nyquist-Shannon sampling theorem [6]. The Nyquist-Shannon theorem states that for a given sampling frequency (in our case F_s), a maximum signal bandwidth of $\frac{F_s}{2}$ can be sampled and reproduced.

An N -bit flash ADC achieves maximum bandwidth by performing the entire N -bit analog-to-digital conversion in only one operation for each period of the sampling clock F_s . To complete the entire conversion in one clock cycle, the traditional flash architecture uses 2^{N-1} comparators in parallel, each referencing a different voltage from a reference ladder. The comparator’s aperture time (similar to a $CLK \rightarrow Q$ of a flip-flop) determines the maximum sampling rate of the converter. The output of all the comparators is a thermometer-coded

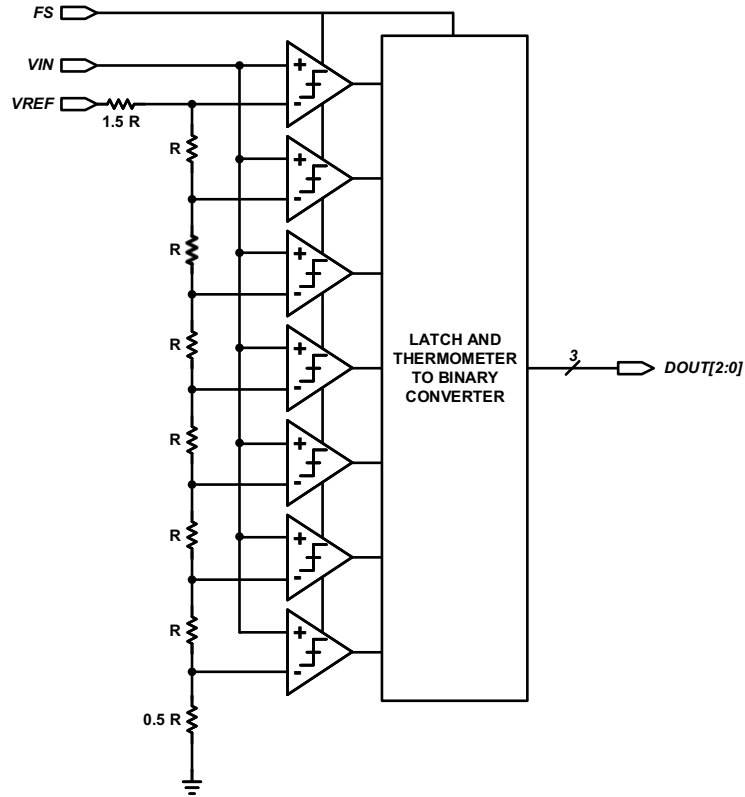


Figure 1.2: Flash ADC architecture from [5]

word that is converted to binary and sent off-chip. An excellent (and state-of-the-art when published) example of a flash ADC is [7].

While flash ADCs are excellent for lower-resolution applications, the exponential relationship between increasing resolution and number of comparators, as well as the challenge of creating a stable reference ladder at high-speed with resolution equal or better to the total converter present significant challenges to increasing flash architecture resolution beyond 6-7 bits.

1.2.2 Pipeline ADCs

The pipeline ADC architecture is best suited for moderate-medium resolution and medium to medium-high bandwidth. A conceptual diagram is shown in Figure 1.3. While generally not capable of achieving the same bandwidth as flash ADCs, the pipeline ADC achieves higher resolution than flash architectures by separating the conversion into multiple stages that operate serially; the output of one stage becomes the input to the next. It is this sim-




Figure 1.3: Pipeline ADC architecture diagram from [5]

ilarity to the digital circuit technique of the same name from which pipeline ADCs inherit their name.

Once the input is sampled by the first stage of the ADC and the initial partial conversion is made, the residue (or “leftover” un-quantized input) is amplified back up to the initial voltage range of the input and further quantized with following stages through the use of sub-ADCs and multiplying DACs with high-gain residue amplifiers. This is shown in Figure 1.3 where the input is sampled by a track-and-hold circuit and quantized by the first-stage ADC. The first stage ADC output is then fed to the first-stage DAC, whose output is then subtracted from the input and multiplied back up to full scale. In this way, the residue from the first stage becomes the input to the second stage. This process can be repeated as many times as desired, however total resolution is limited to the accuracy of the first stage; even if your first stage only quantizes the first 3 MSBs of a 10-bit converter, it still has to be 10-bit accurate [8].

As each stage may use an entire sampling clock period to complete, the pipeline architecture trades a small amount of bandwidth for a moderate increase in resolution compared with the flash architecture. Pipeline ADCs are a common choice for systems where 8-12 bits are required with high bandwidth. However while parallelism is the limiting factor for flash ADCs, the finite gain error from each stage’s residue amplifier and first-stage accuracy caveat are the key challenges facing design of pipeline ADCs.

1.2.3 Sigma-Delta ADCs

The Sigma-Delta (or $\Sigma\Delta$) architecture is most commonly used for high-resolution converters, and achieves the highest resolution of any converter through the use of oversampling




Figure 1.4: Second-order Sigma-Delta modulator diagram from [9]

and noise shaping [4], [10]. Oversampling sacrifices converter bandwidth for increased resolution. Unlike Nyquist ADC architectures, whose quantization bandwidth is half the sampling rate or $\frac{F_s}{2}$, sigma-delta ADCs instead provide only a small usable quantization bandwidth either at DC (in lowpass sigma-deltas) or at some intermediate frequency below the sampling clock (in bandpass sigma-deltas). Noise-shaping, made possible via oversampling and integrators “pushes” the quantization noise and other non-linearities outside of the conversion bandwidth of interest, thereby trading the traditional $\frac{F_s}{2}$ bandwidth for increased resolution. A block diagram of a second-order $\Sigma\Delta$ modulator is shown in Figure 1.4. Through use of high oversampling ratios (OSRs - many audio converters utilize OSRs of $128\times$ or even $256 \times F_s$), even a 1-bit quantizer can be utilized to achieve 16+ bits of resolution.

While increasing loop order and quantizer resolution improve the theoretical performance of $\Sigma\Delta$ converters, noise shaping is only effective on blocks in the forward path, so feedback DAC performance is often a limiting factor. In addition, integrator performance (finite gain error, slew rate, dynamic range) becomes increasingly challenging with higher loop orders [9].

1.2.4 SAR ADCs

The Successive Approximation Register (SAR) architecture is named as such because after the input voltage has been sampled, the ADC uses a DAC in feedback to perform a recursive search on the input, *i.e.* successive approximation. A simplified block diagram of a single-ended SAR ADC is shown in Figure 1.5.

Figure 1.6 shows a 3-bit example of a binary-weighted SAR algorithm. The algorithm in this example (and most SAR ADCs) is “binary weighted” because the search space is reduced

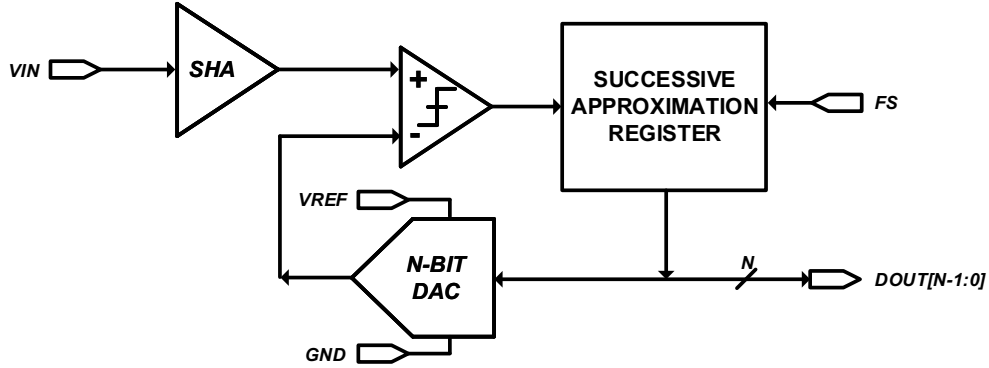


Figure 1.5: Basic SAR architecture diagram from [5]

in half between each successive step. In binary logic, each step determines a decreasing power of two. For the first decision the input (red line) is compared with a DAC voltage of $\frac{1}{2} \times V_{REF}$, and the result of that comparison is a boolean “true” or a “1”, confirming that $V_{IN} > \frac{V_{REF}}{2}$. The second decision will either add or subtract $\frac{1}{4} \times V_{REF}$ to the DAC from the previous value depending on the result of the first decision. Because the first result is a “1”, $\frac{1}{4} \times V_{REF}$ is added, and the comparison of V_{IN} and $\frac{3}{4} \times V_{REF}$ results in a logical “false” or “0”, confirming that $V_{IN} < \frac{3}{4} \times V_{REF}$. For the third and final comparison the search space is again reduced by a power of two, and $\frac{1}{8} \times V_{REF}$ is subtracted from the DAC resulting in a comparison of V_{IN} and $\frac{5}{8} \times V_{REF}$ which again yields a “0” result from the comparator. The 3-bit output code of this conversion is “100”, which from the available 3-bit code space of “000” to “111” is the 5th out of 8 codes.

While the iterative nature of the SAR algorithm trades overall conversion speed for resolution (although to a lesser degree than sigma-delta converters), it should be noted that it requires a minimum amount of analog circuitry. While the example in Figure 1.5 shows a separate SHA (Sample-and-Hold Amplifier), DAC, and comparator, most modern SARs incorporate the SHA functionality into the DAC by sampling the input directly onto the DAC. In time-interleaved or otherwise high-bandwidth applications where sampling bandwidth is a limiting factor, a standalone SHA is often used to ensure every interleaved ADC has the same sampling characteristics. In most SAR ADCs without standalone SHAs, this leaves the DAC and comparator as the sole analog circuits. The DAC requires analog design analysis to ensure its parameters will meet the overall desired performance of the ADC, however its control

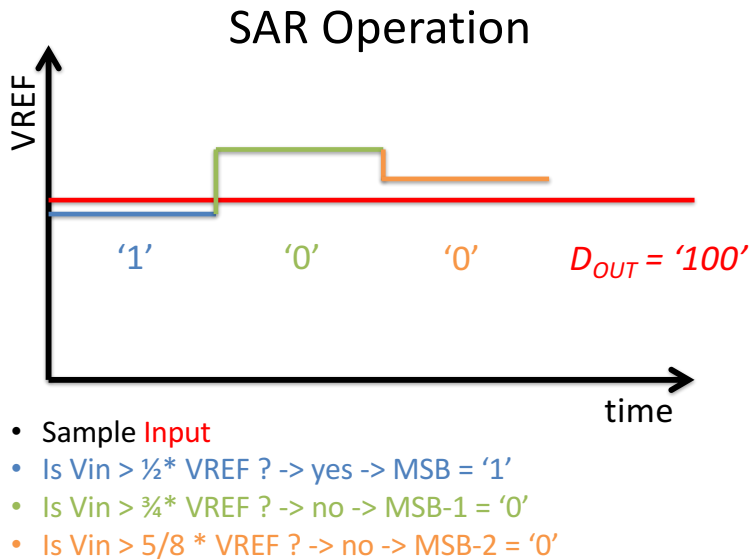


Figure 1.6: 3b single-ended SAR conversion example

circuitry is entirely digital. This leaves the comparator as the only remaining analog circuit in the SAR. It is this digital-dominant nature of the SAR ADC that makes it an inherently scaling-friendly architecture, as is discussed in Section [refsec:intro:scalingarchitectures](#).

1.3 Process Scaling and ADC Design

Process scaling is the shrinking of devices (transistors) in integrated circuits. What is commonly referred to as “Moore’s Law” [11], is Gordon Moore’s prediction of the number of devices on an integrated circuit doubling every two years due to advancements in semiconductor manufacturing processes and technology. This prediction held true roughly until 2012 when Intel’s CEO, Brian Krzanich, acknowledged that their pace of doubling was slowing down to approximately two and a half years instead of two [12]. This slowdown was due to state-of-the-art processes becoming increasingly difficult to develop. Modern Fin-FET processes feature small device sizes (gate lengths as small as 7nm) and several device and low-level metal layers now require double-pattern extreme-ultraviolet (EUV) lithography

Table 1.1: Influence of scaling on MOS device characteristics.[14], [15]

Parameter	Sensitivity	Constant Field	Constant Voltage
Scaling Parameters			
Length: L		$1/S$	$1/S$
Width: W		$1/S$	$1/S$
Gate oxide thickness: t_{ox}		$1/S$	$1/S$
Supply voltage: V_{DD}		$1/S$	1
Threshold voltage: V_{tn}, V_{tp}		$1/S$	1
Substrate doping: N_A		S	S
Device Characteristics			
β	$(W/L)(1/t_{ox})$	S	S
Current: I_{ds}	$\beta(V_{DD} - V_t)^2$	$1/S$	S
Resistance: R	V_{DD}/I_{DS}	1	$1/S$
Gate capacitance: C	WL/t_{ox}	$1/S$	$1/S$
Gate delay: τ	RC	$1/S$	$1/S^2$
Clock frequency: f	$1/\tau$	S	S^2
Switching energy: E	CV_{DD}^2	$1/S^3$	$1/S$
Switching power dissipation (per gate): P	Ef	$1/S^2$	S
Area (per gate): A		$1/S^2$	$1/S^2$
Switching power density	P/A	1	S^3
Switching current density	I_{ds}/A	S	S^3

[13].

1.3.1 Process Scaling Overview

Process scaling must be discussed in two contexts: *device* scaling, which refers to the shrinking of active devices on the silicon wafer (transistors), and *interconnect* scaling, which refers to the shrinking of metal interconnect.

Table 1.1, from data and analysis in [14], [15], shows how CMOS devices perform under constant field and constant voltage process scaling. Ideally, silicon manufacturers would like to follow constant field scaling, which improves performance but maintains the same power density. For large-scale systems such as CPUs, maintaining the same power density is important: since the Pentium 4 was released in 2000, processors have reached the maximum power (per chip) that can be supported given heat dissipation (*e.g.*, 125W). Perfect constant field scaling has not been achievable, due to increased leakage caused by reducing V_t , increased V_t variation, and V_t reduction via drain induced barrier lowering (DIBL). This results in an even greater power density, which designers are required to offset in other ways [15].

There are positive and negative consequences of device scaling for ADC (and other ana-

Table 1.2: Influence of scaling on interconnect characteristics. [14]

Parameter	Sensitivity	Scale Factor
Scaling Parameters		
Width: W		$1/S$
Spacing: s		$1/S$
Thickness: t		$1/S$
Interlayer oxide height: h		$1/S$
Die Size		D_c
Characteristics per Unit Length		
Wire resistance per unit length: R_w	$1/wt$	S^2
Fringing capacitance per unit length: C_{wf}	t/s	1
Parallel plate capacitance per unit length: C_{wp}	w/h	1
Total wire capacitance per unit length: C_w	$C_{wf} + C_{wp}$	1
Unrepeated RC constant per unit length: t_{wu}	$R_w C_w$	S^2
Repeated wire RC delay per unit length: t_{wr} (assuming constant field scaling)	$\sqrt{RCR_w C_w}$	\sqrt{S}
Crosstalk noise	w/h	1
Energy per bit per unit length: E_w	$C_w V_{DD}^2$	$1/S^2$
Local/Semiglobal Interconnect Characteristics		
Length: l		$1/S$
Unrepeated wire RC delay	$l^2 t_{wu}$	1
Repeated wire delay	$l t_{wr}$	$\sqrt{1/S}$
Energy per bit	$l E_w$	$1/S^3$
Global Interconnect Characteristics		
Length: l		D_c
Unrepeated wire RC delay	$l^2 t_{wu}$	$S^2 D_c^2$
Repeated wire delay	$l t_{wr}$	$D_c \sqrt{S}$
Energy per bit	$l E_w$	D_c / S^2

log circuit) designers. While the current gain (transconductance, or g_m) increases due to improved gate control of the electron channel (good), and gate delay decreases due to reduced channel length (or the device's unity gain frequency, ω_t increases, good), the output resistance (r_o) of each device decreases (bad), and the reduced drain-gate oxide-breakdown tolerance decreases the supported supply voltage (V_{DD} , bad), such that the self-gain and possible dynamic range of an analog circuit are reduced. Of equal if not greater importance to ADC designers are the negative effects that scaling has on transistor noise [16].

Table 1.2, also from [14],[15], shows how interconnect RC delay scales for local and global connections each generation [15]. While interconnect scaling severely limits power-dense systems such as CPUs, in general ADC designers benefit from the increased resolution of metal lithography allowing smaller and smaller passives to be created. Nearly every SAR ADC cited in this thesis uses custom-drawn Metal-Oxide-Metal (MOM) capacitors, and as minimum spacing and trace width decrease, ADC designers are better able to tune MOM

capacitors for system needs.

1.3.2 How Process Scaling Affects ADC Architectures

Process scaling affects the four major ADC architectures in different ways. This section briefly mentions key effects on each architecture.

Flash: While the shrinking of devices helps reduce the area penalty of the flash architecture’s inherent parallelism (2^{N-1} comparators), the increase in threshold voltage (V_t) variation carries either a severe calibration or severe over-sizing penalty. Two innovative examples of flash design in modern processes are [17] and [18]. In cases where devices are sized-up to overcome the increased variation in threshold voltage, the input capacitance and overall power consumption increase as well.

Pipeline: As previously discussed in Section `refsec:intro:pipeline`, the key limitation of the pipeline architecture is the first-stage accuracy requirement, and the critical block most affected is the residue amplifier. Scaling would potentially benefit every stage of a pipeline design except the first one. Recent state-of-the-art pipeline ADCs use alternative circuits instead of traditional op-amps as residue amplifiers, such as “ring amplifiers” [19], [20], and zero-crossing based circuits (open-loop amplifiers) [21].

Sigma-Delta: Being dependent on high-gain integrators, Sigma-Delta ADCs also suffer from scaling’s effect on dynamic range and reduced output resistance, however the increased speed enables higher oversampling ratios, and the nature of noise-shaping already mitigates the increased device noise from scaling. Bandpass sigma-deltas utilize the increased sampling speed capability to noise-shape an intermediate frequency (IF) instead of around DC as in traditional low-pass Sigma Deltas. An example of an innovative bandpass sigma delta is [22].

SAR: As the SAR architecture is mostly digital, SARs generally benefit from scaling, especially below 10 bits of resolution. Because scaling reduces the gate delay of digital logic gates, scaling will continually improve all speed-related issues with the exception of DAC settling time (further discussed in Section 1.4.2). A result of this scaling friendliness is that SAR ADC bandwidth has caught up to that of pipeline ADCs, as can be seen in the

bandwidth vs. resolution by architecture plot of Figure 1.1. For medium-higher resolution ADCs (11+ bits), the mismatch between DAC elements requires drawing larger capacitors than would otherwise be required to meet sampling noise requirements ($\frac{kT}{C}$ noise), and is further discussed in Section 2.1.3.

1.4 Recent SAR Techniques and Innovations

During the past decade, SAR ADCs have enjoyed increasing prominence due to their inherently scaling-friendly, mostly digital architecture. This section discusses key innovations that have improved the SAR architecture and have become increasingly common design techniques. These innovations fall into the following categories of improvements:

- **Power reduction:**

- Low-voltage operation [23],[24]
- Energy-efficient DAC switching [25],[26],[27],[23],[28],[24],[29]
- Tunable comparator performance [30]

- **Increased speed and conversion bandwidth:**

- Asynchronous operation [31],[32],[26]
- DAC settling time reduction
 - * Non-binary radix conversion [33], [34]
 - * Error-correction via extra decisions [32],[30],

- **Comparator noise mitigation:**

- Time-domain comparison or assistance [35], [36]
- Error-correction via extra decisions [32],[30],[37]
- Tunable comparator performance [30]

- **Mismatch mitigation:**

- Residue shaping [36]
- Tunable capacitors [38]
- Dithering [39]

1.4.1 Power Reduction

The power consumption vs. bandwidth and resolution tradeoff in SAR ADCs is captured by the following figure-of-merit (FoM), which is referred to in literature and performance surveys such as [4] as the “Walden FoM” [40].

$$Walden_{FoM}\left(\frac{J}{Conv. - step}\right) = \frac{P_{ADC}}{2 * F_{IN,MAX} * 2^{ENOB@F_{IN,MAX}}}$$

where

P_{ADC} = Total ADC power consumption

$F_{IN,MAX}$ = Maximum coherent input frequency in ADC conversion bandwidth

$ENOB@F_{IN,MAX}$ = The ADC’s ENOB at $F_{IN} = F_{IN,MAX}$

$\frac{J}{Conv. - step}$ = Resulting units are energy (Joules) per ADC code-step

(1.1)

Many of the works cited in this section achieved record-breaking Walden FoMs at their time of publication.

Low-voltage Operation

Most of the circuits in SAR ADCs, other than the comparator and the DAC switches, are digital. Power consumption in digital CMOS circuits is given by fCV^2 , where f is the frequency of operation (e.g. a clock for a flip-flop), C is the internal device and interconnect capacitance being charged and discharged during each operational cycle, and V is the supply voltage of the logic and the voltage to which the capacitance is being charged. Because a reduction in V results in a quadratic decrease in power consumption, many of the lowest-power SARs operate well below the supported supply voltage of their fabrication process. A

flexible supply voltage between 0.4-0.7V is utilized in [24], 0.35V is used in [23], while 0.45V is used in [29].

Energy-efficient DAC switching

The DAC switching algorithm determines how a successive approximation is performed on an input voltage during a SAR conversion. While there are several ways to achieve a linear transfer function, the choice of DAC switching scheme affects power consumption from the reference used by the DAC, as well as how comparator and reference errors affect the overall ADC performance. A large number of DAC switching schemes have been published between 2005-2010, however the first in this group was Ginsburg [25], and most of the subsequent works [26], [23], [28], and [24] are extensions of Ginsburg’s work and cite his publications. Ginsburg noted that most SARs use switching algorithms that charge and discharge more capacitance than necessary, thereby wasting reference energy.

Regardless of the switching technique, when a capacitor in a SAR DAC is switched “up” to V_{REF} or “down” to ground, the current consumed from the reference being connected is a function of the change in charge experienced by the switched capacitor. An example from Ginsburg is shown in Figure 1.7. An initial state is shown where C_2 is connected to V_{REF} while C_1 and C_0 are connected to GND (0V). Ginsburg achieves significant energy savings in the case where the SAR algorithm requires the DAC voltage to be reduced for the next decision. Ginsburg notes, instead of discharging the entire large capacitor C_2 and subsequently re-charging the next (smaller) capacitor C_1 , if C_2 were “split” into two equal parts, only half of C_2 would need to be discharged to create the desired DAC voltage change without any subsequent charging of C_1 . The end result of Ginsburg’s paper is that by “splitting” *one* binary weighted DAC of 2^N units into *two* DACs of 2^{N-1} units, and *switching the least amount of capacitance necessary at any time*, DAC energy consumption from reference switching can be reduced by 36.5%, averaged across all ADC codes. Thus, Ginsburg’s switching scheme is often referred to as “split-MSB” switching. The only negative consequence from Ginsburg’s technique is additional routing complexity, however because there are no additional references and the logic is trivial, there is a minimal penalty.

The next switching scheme that is often cited and replicated is Liu’s “set-and-down”

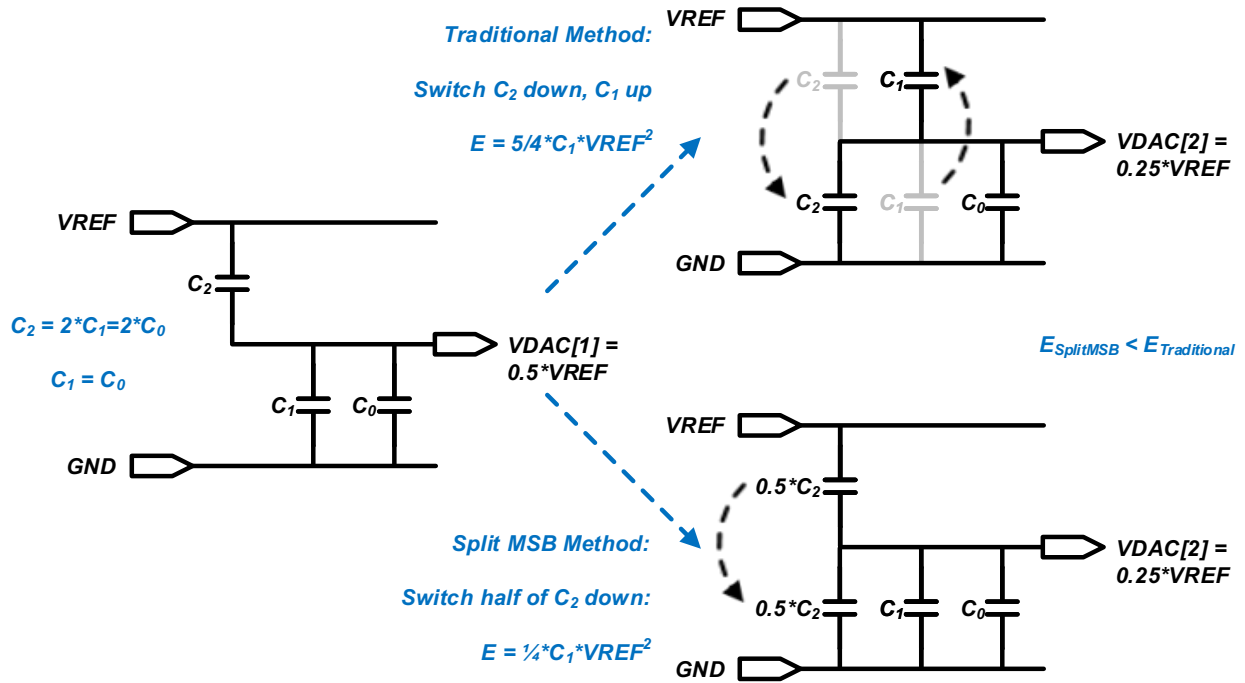


Figure 1.7: Traditional vs. Split MSB switching, from [25].

switching method, first described in [26] and adjusted in [27]. Liu goes farther than Ginsburg in reducing switching events by abandoning differential DAC switching altogether, in favor of a monotonically decreasing scheme. In Liu’s initial work [26], all DAC capacitors (in a differential 10-bit ADC) are sampled against V_{REFP} and the DAC voltage on either side continually decreases for the entire SAR conversion. This results in a 70% energy reduction over Ginsburg’s, however the penalty is that the common-mode voltage presented by the differential DAC is also monotonically decreasing throughout the SAR conversion, beginning at mid-rail ($\frac{V_{REFP}}{2}$) and ending at a half *LSB* above *GND* or V_{REFM} (specifically $V_{REFM} + \frac{V_{REFP}}{2^{N+1}}$). Because comparator gain and input-referred noise are dependent on the comparator’s common-mode input voltage, this switching scheme greatly complicates comparator design. This led to Liu reverting the first 3 MSB switches back to a differential method in his follow-up work, [27].

An even lower energy switching scheme was proposed in [36] and [28], which is often referred to as “merged capacitor switching,” or MCS. MCS is essentially the same as Ginsburg’s split-MSB scheme with the addition of a common-mode reference V_{CM} at $\frac{1}{2} \times V_{REF}$. After sampling the input against V_{CM} , every subsequent switching event is charged to half the

voltage difference needed in every other scheme, resulting in a 75% energy savings compared with charging to a full reference. Given a common-mode reference this switching method is advantageous, however voltage references, discussed later in Section [refsec:adc:references](#), require significant design effort as ADC resolution increases.

The most recent switching scheme cited in this work is the “detect-and-skip” method in [29], whose implementation in an ADC achieved a record-breaking Walden FoM of 0.85fJ/conv.-step. The detect-and-skip algorithm uses two separate DACs. The first DAC is a very small “coarse” DAC that is connected to a very low power comparator. This first DAC is only used to determine the first five MSBs, and due to the DACs small capacitor size the DAC voltage swings consume very little reference energy. Because the MSB trials present large DAC voltages to the comparator, the comparator used with the coarse DAC can be very low power as a large input referred noise can be tolerated. After determining the first five MSBs, the larger, second DAC is switched in one operation to its correct value. This results in a skipping of several MSB trials with large capacitors and a significant energy savings. The final 5 LSB decisions are performed with the larger DAC which is connected to a higher-power comparator with better input-referred noise performance. One bit of redundancy handles errors between the DAC and comparator switch.

Tunable Comparator Performance

Similar to traditional operational amplifiers, a comparator exhibits an input referred noise such that as its differential input magnitude approaches zero, the output is increasingly random. In an ideal comparator with zero offset, noise is assumed to be a zero-mean gaussian distribution about $V_{IN,DIFF} = 0V$. A comparator with DC offset will exhibit the same distribution, although it will be centered around the input offset voltage. In [41], Nuzzo uses the CDF of the comparator’s bit error rate along with an inverse error function and defines the $+1\sigma$ point on the CDF as the input-referred RMS noise voltage of a comparator, $V_{\sigma,COMP}$.

In typical SAR ADCs, comparators are designed so that this input-referred noise voltage is below $\frac{LSB}{2}$ in magnitude and preferably as small as possible given the desired power consumption. During a binary-weighted differential SAR conversion, the differential input

to the comparator will only be within $\frac{LSB}{2}$ once. Consider the case of an ideal 6-bit ADC where the input is only $0.1 \times LSB$ larger than the input common-mode, such that the output code should ideally be just over the MSB code transition and result in an output of “100000”. In this case the first (*MSB*) decision is the only decision where the comparator will have an input voltage is less than $\frac{LSB}{2}$ in magnitude.

In [30], Harpe presents a tunable comparator, where the internal capacitance and tail current can be varied to selectively enhance or degrade the input-referred noise for different SAR decisions. Harpe achieves a 50% reduction in comparator power by tolerating reduced noise performance for the first $N - 2$ trials, and increases noise performance for the final 3 comparisons which include an extra error-checking decision. Despite the total number of comparator decisions increasing by 1, this method still results in a net power savings.

1.4.2 Increased Speed and Conversion Bandwidth

Asynchronous Operation

The earliest paper cited describing the SAR timing as “asynchronous” is Chen’s 2006 ISSCC paper [31]. This term can be ambiguous but in most cases (and in this work) it means that the comparator initiates each successive SAR step immediately after regenerating, and does not necessarily have a fixed amount of time in which to make a decision. Some additional SAR background is required.

SAR ADCs were described by their *system clock* frequency, not their *sampling clock* frequency. A *system* clock is used to initiate each successive step of the SAR cycle, and would be a multiple of the effective *sampling* clock. For example, if a SAR ADC requires N *system* clock periods to complete a conversion, then the sampling rate F_S is $\frac{F_{system}}{N}$. Having each SAR step use the same amount of time means that the comparator has the same amount of time to make every decision. However, a regenerative comparator (used in all of the works cited in this section, and further discussed in Section refsec:adc:comparator), exhibits an exponential relationship between clock-to-output time and input voltage [42], which is not leveraged in system clock SARs.

Figure 1.8 shows an example of the exponential input voltage and clock-to-output rela-




Figure 1.8: Comparator output decision time vs. input voltage.

relationship, in which every $10\times$ reduction of differential input voltage results in a linear increase in clock-to-output time. This ratio of clock-to-output time increase vs. input characterizes the regenerative gain of the comparator [42]. During a SAR conversion, only one of the SAR trials will present an input to the comparator less than $\frac{LSB}{2}$ in magnitude, which means that most comparisons take a minimal amount of time while one or two comparisons will take much longer.

In [31], Chen noted that timing the entire SAR ADC with a fixed system clock around a worst-case comparator input voltage of $\frac{LSB}{2}$ results in wasted time that could instead be used for increased sampling rate and input bandwidth. His innovation to let the comparator “clock itself,” resulted in increased speed and has been replicated in almost every other SAR cited in this work, notably [32] and [26].

DAC Settling Time Reduction via Redundancy

In addition to the comparator decision, DAC switching is the other analog operation that must occur serially, in every SAR trial. When the SAR switches a group of DAC capacitors to or from the reference V_{REF} , the RC network between V_{REF} and the DAC requires a certain amount of time to settle that is dependent on the overall resolution of the DAC. Incomplete DAC settling results in code-dependent conversion errors, which diminish ADC linearity and greatly affect performance. DAC settling time can be calculated from the RC time constant

between the on-resistance R_{ON} of the reference-connecting switch on the DAC capacitors' bottom plates, and the equivalent series capacitance seen by the reference of the new DAC connections being charged to V_{REF} . This relationship is show below in equation 1.2.

$$t_{settle} = R_{ON}C_{EQ,DAC} \ln(2)(N + 1)$$

where

$$t_{settle} = \text{Required settling time}$$

$$R_{ON} = \text{On-resistance of reference-connecting switches on DAC bottom plates}$$

$$C_{EQ,DAC} = \text{Equivalent capacitance seen by combination of } V_{REF} \text{ and } V_{SS}\text{-connected DAC capacitors}$$

$$(N + 1) = \text{Required resolution of settling}$$

(1.2)

DAC settling time is accounted for in primarily two ways. In synchronous SARs, because the system clock frequency determines the period of each SAR trial, DAC settling time is predetermined as the time leftover when subtracting the comparator aperture time from each clock period ($t_{settle} = t_{clk} - t_{comp}$). However in asynchronous SARs, the DAC settling time is often accounted for by a fixed delay in the asynchronous timing loop such that after each DAC switching event, the next comparator decision is delayed by the amount of time required by equation 1.2. Note that both of these approaches have a fixed amount of DAC settling time for every SAR trial; similar to the inefficiency of system clocking vs. asynchronous clocking, the worst-case MSB decision (due to the largest capacitor group being switched) mandates excessive DAC settling time for all other trials. In addition, because the RC settling time is exponential in nature, as DAC resolution increases, an increasing amount of SAR conversion time is spent in DAC settling.

This relationship between conversion speed and DAC settling time is greatly improved in Kuttner's 2002 ISSCC non-binary radix SAR [33]. As discussed in the SAR architecture overview (Section refsec:intro:sararchitectureintro), most SAR ADCs use groupings of capacitors with a radix of 2 so the groupings are binary-weighted, which in turn results in the search space reducing by half for each SAR trial. In non-binary radix SARs such as

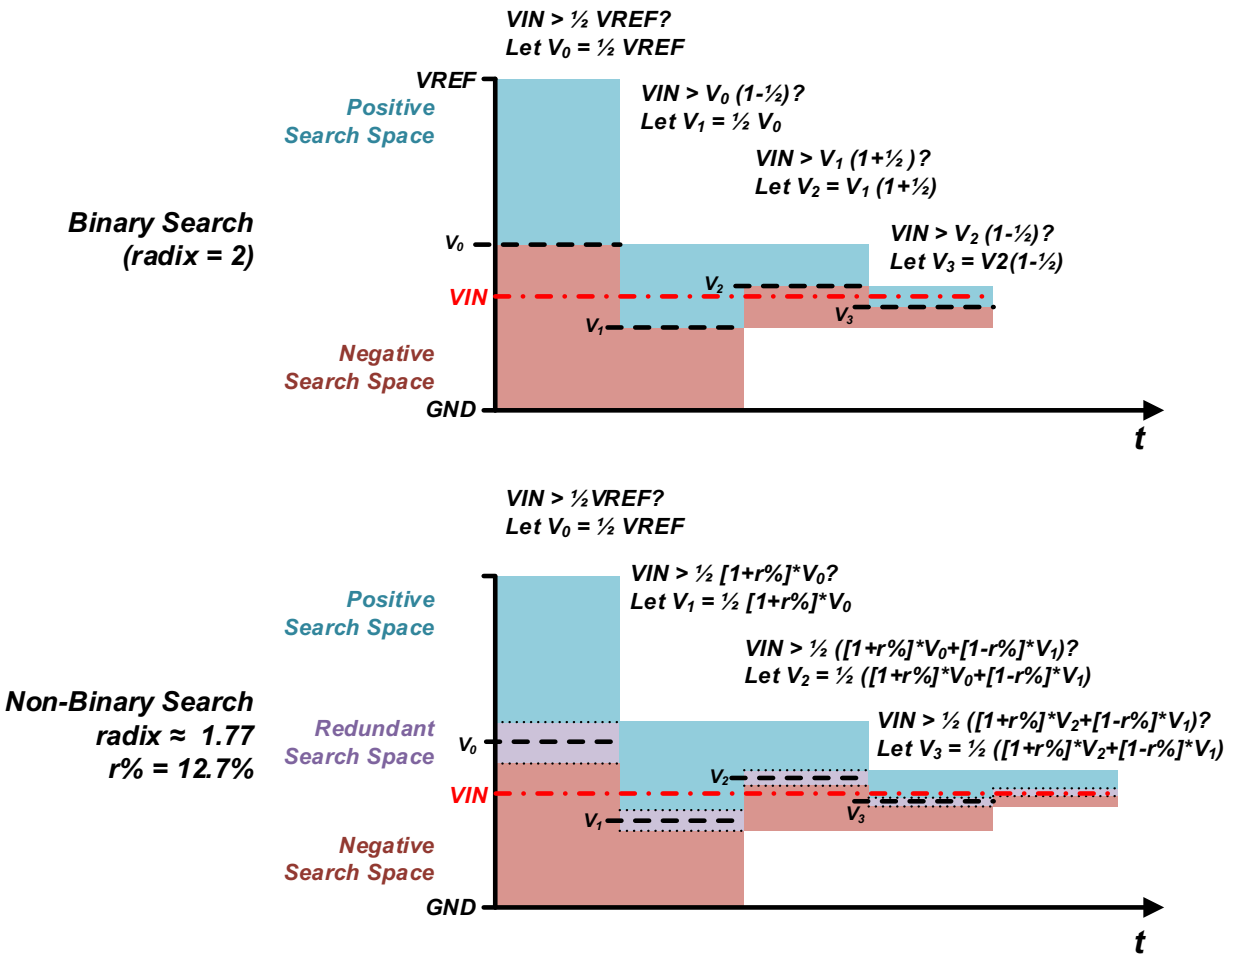


Figure 1.9: Comparison of binary vs. non-binary search highlighting redundant search area, from [33].

Kuttner's [33], and also Kang's [34], a radix of less than 2 is used, such that the search space is reduced by less than half for each SAR trial. Reducing the search space by less than half means that some portion of the previous trial's search space will be included in the next decision.

Figure 1.9 shows an example of binary vs. non-binary switching. In the binary switching example, each successive SAR trial reduces the search space by half as expected. In the non-binary switching example, while more trials are needed to resolve the input with the same precision (5 vs. 4 in this figure), more of the search space is carried over between SAR trials. The section of search space that is repeated between two successive trials is referred to as *redundancy*.

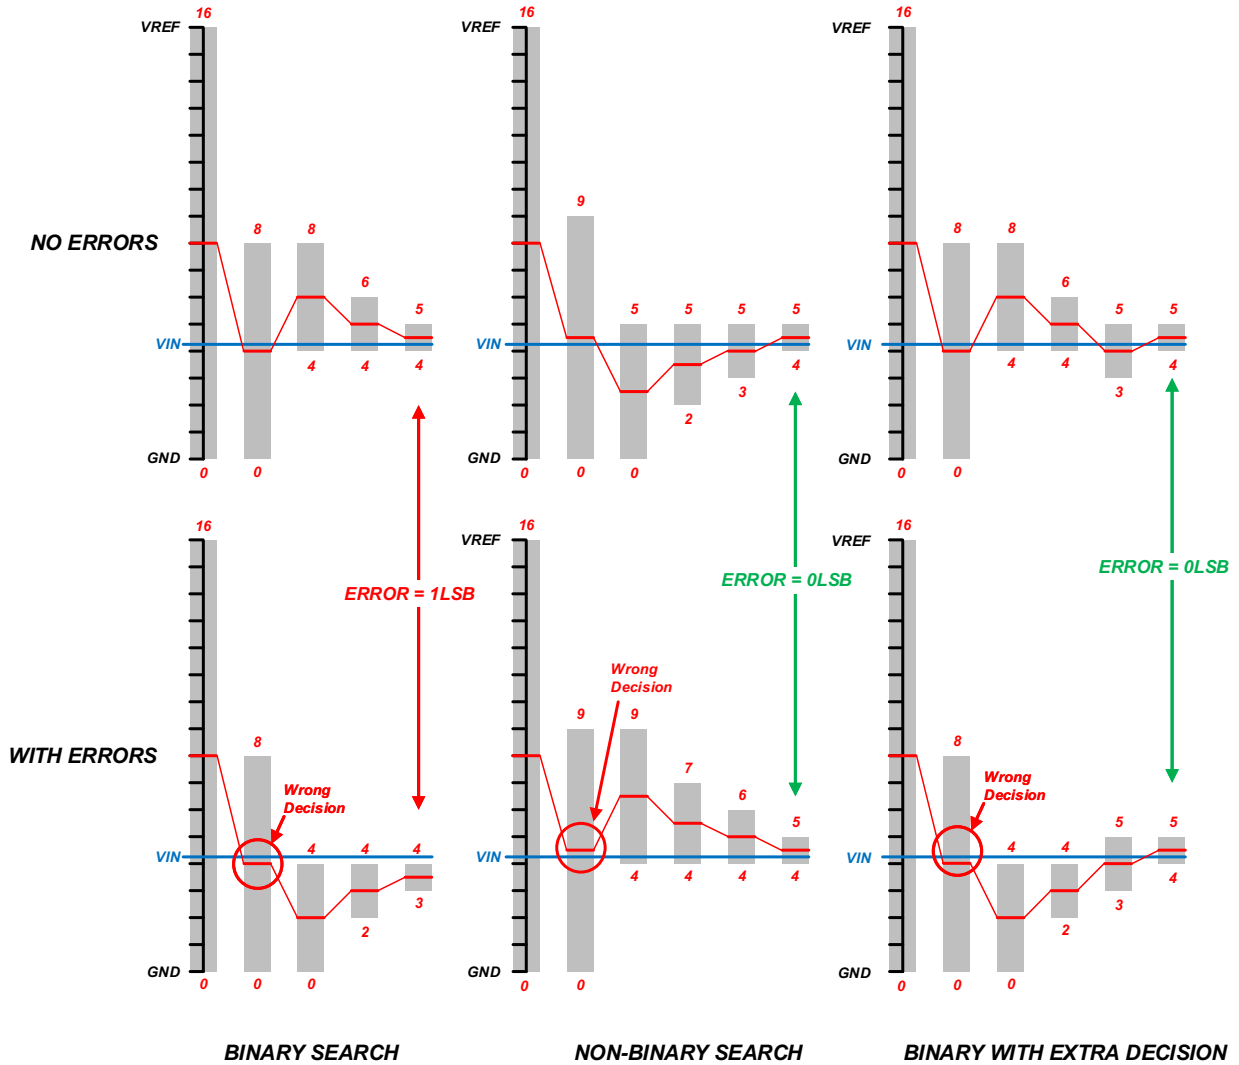


Figure 1.10: Example of ideal vs. error tolerance in binary and two types of redundant switching schemes, from [32].

While Kuttner [33] and Kang [34] implement redundancy via a non-binary radix, Liu used extra comparisons of binary weight to achieve the same end, as shown in [32], and later in [43], where additional binary-weighted decisions are used to error-check previous search spaces. Figure 1.10 shows how non-binary and binary redundancy [32] tolerate incomplete DAC settling error - despite an error during a SAR trial, the search space still converges on the correct code. Redundancy greatly reduces the required DAC settling time because any DAC settling error within the redundant search space is re-tried in the successive SAR step. Because a settling error within the redundant search space may now be tolerated, this decouples the DAC RC settling time equation from the overall DAC resolution (“ $N + 1$ ”

term in equation 1.2). While redundancy requires additional SAR trials, the net savings in DAC settling time ends up increasing the possible conversion speed of the ADC.

1.4.3 Comparator Noise Mitigation

An additional benefit of redundancy is that in addition to DAC settling errors, comparator errors can also be tolerated. As discussed above, any error within the redundant search space will have a second opportunity to be evaluated, including comparator error. In the case of Figure 1.10, instead of an ideal comparator with incomplete DAC settling, consider the identical outcome of the case with an ideal DAC settling but an incorrect comparator decision; due to redundancy the correct code is still converged upon. Kuttner [33], Kang [34] and Liu [32], [43], all discuss redundancy’s effect on comparator noise mitigation.

Comparator noise can also be mitigated through the use of extra comparator decisions. In [37], Giannini et. al. use an extra SAR decision to detect whether the final *LSB* decision is correct, in addition to separate low and high noise-performance comparators as in the later work by Harpe, [30].

An additional noise mitigation technique is to utilize time domain information of the comparison event. In [35], Agnes creates a “time domain comparator” which transforms the voltage input to a time difference, and the comparison is performed in the time domain. The goal of Agnes in [35] is to enable entirely subthreshold SAR circuitry by decoupling the comparator from supply voltage and headroom issues, and the result is a comparator with $190\mu\text{V}$ input referred noise and state-of-the-art power consumption at the time of publication.

Guerber et. al in [36] harness the exponential relationship between input voltage to comparator output time to provide extra information about the input during SAR conversion. Guerber et. al compare the length of comparator regeneration time with a known (and adjustable) time reference, which if exceeded indicates the input voltage to the comparator is small. This is additional information that is similar to “the 1.5 bit-per-stage [redundancy] seen in pipelined converters.” In addition to reducing DAC switching events and tolerating settling errors, this also provides “residue shaping” which “results in an effective extra 6dB of signal-to-quantization-noise ratio (SQNR) or an extra bit.” Guerber et. al. detail this

technique in a separate publication, [44].

1.4.4 Mismatch Mitigation

Mismatch, in this context, is the random variation in DAC elements created by process variation. The higher the degree of mismatch between DAC elements, the worse the linearity of the DAC and thus the overall SAR ADC. Mismatch mitigation is the key contribution of this thesis, and as such mismatch in general is more thoroughly analyzed in chapter 2. This section presents some recent approaches that, while effective within their context, illustrate the need for more scalable work in this area.

Assuming a nearly rail-to-rail dynamic range, comparator noise and DAC element mismatch constraints are critical but not insurmountable at 10 bits or less in sub-100nm processes; indeed most of the state-of-the-art SAR ADCs discussed in this chapter are also 10 bits or less. To achieve medium resolution (11-15 bits, depending on dynamic range), DAC element mismatch must be a central focus of the design.

A recent example of mismatch mitigation is Harpe et. al in their 2014 ISSCC paper [39], where a 12-to-14-bit SAR ADC is presented. A chopping technique is used to eliminate even-order distortions, while an innovative dithering technique utilizing the DAC itself helps reduce even-order distortions. A 14-bit mode ENOB of 12.99 bits is achieved, however the chopping and dithering techniques require oversampling.

In a successive attempt at a medium resolution SAR without relying on hybrid techniques such as oversampling, [38], Ding and Harpe et. al. add a tunable group of capacitors to correct for mismatch. While this correction is performed in the background at minimal power expense, a maximum ENOB of 10.5 out of 13 attempted bits is achieved. A similar post-fabrication capacitor tunability was also employed by Jung and Kim in [45].

1.5 Contributions of This Work

This thesis presents a novel approach to overcoming DAC mismatch design constraints, thereby enabling medium resolution SAR ADCs to benefit from process scaling without the use of hybrid ADC architectures. Currently the medium resolution (11-15 bit) ADC

publication space is dominated by hybrid ADC architectures [4], [19], [39], as some form of oversampling, noise-shaping, or other circuit technique is needed to provide additional resolution while keeping DAC capacitance at acceptable levels.

Chapter 2 discusses the theory and design of a completely reconfigurable capacitor DAC that eliminates the need to design for mismatch in medium resolution SAR ADCs, thereby enabling non-hybrid SAR ADCs to continue their current trajectory of improvement and increased coverage of ADC bandwidth and resolution. It will be shown that DAC reconfiguration severs the relationship between DAC capacitor size and linearity.

Chapter 3 discusses the theory and design of a SAR ADC utilizing the reconfigurable capacitor DAC. Measurements confirm the ADC linearity is enhanced via reconfiguration and thus immune from mismatch effects inherent in non-reconfigurable DACs.

CHAPTER 2

Completely Reconfigurable Capacitor DAC

2.1 Extending Scaling Benefits to Medium and Higher Resolution ADCs

2.1.1 Mismatch in Traditional DACs

In traditional capacitor DACs, the assignment of individual capacitors to SAR trial groupings is made during layout. The unit capacitors are hard-wired into groups of unit capacitors whose relative sizes determine the radix of the SAR algorithm. The radix refers to the base of each comparator trial in the SAR cycle. Traditionally most SARs have used a binary radix where each SAR trial represents a power of two. Some more recent SARs, beginning with [33] use non-binary radices such as 1.83. Figure 2.1 shows a sea of 8 unit elements whose bottom-plate connections create binary-weighted groups. The equivalent schematic is also shown. Thus in traditional DACs, whose element arrangements are made during design and layout, the overall linearity is solely determined by the mismatch characteristics of the fabrication process.

Element mismatch from fabrication is modeled by adding a random variable component to the nominal value of each DAC element. In the case of a capacitive charge redistribution DAC, one would draw an array of element capacitors with some value C_{UNIT} . The mismatch random variable is added with the term ΔC_{UNIT} . Element mismatch is characterized by the ratio of mismatch to the nominal value $\frac{\Delta C_{UNIT}}{C_{UNIT}}$, most often in units of percent. For example

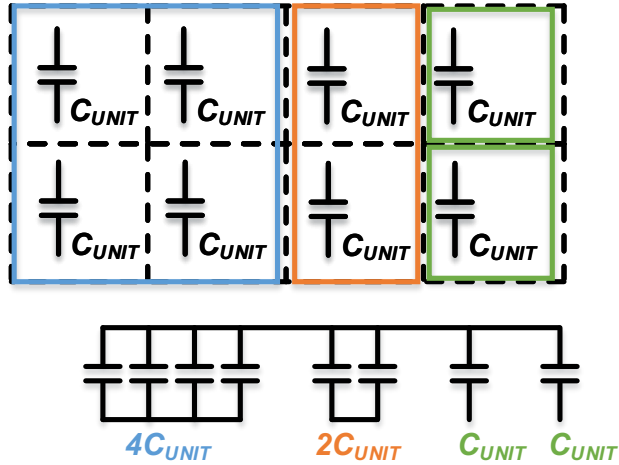


Figure 2.1: Example of 3-bit capacitor DAC assembled from unit element C_{UNIT}

a unit capacitor whose post-fabrication value is $0.98 \times C_{UNIT}$ would be said to have (negative) 2% mismatch.

The earliest work on monolithic charge-redistribution SARs [46],[47],[48] also included analysis of the MOS capacitor DAC elements [49]. Singh and Bhattacharyya later published more comprehensive analysis [50]. Perhaps the most-cited paper on mismatch in semiconductor fabrication is Pelgrom’s [51], which summarized “the variance of [the randomly varying part of a device] parameter ΔP between two rectangular devices...” is characterized by the following proportionality, where A_P is an area proportionality constant that is empirically specific to the fabrication process, and W, L are the width and length (area) of the device:

$$\sigma^2(\Delta P) \propto \frac{A_P^2}{WL} \quad (2.1)$$

More recent analysis [52], [53] of similar Metal-Oxide-Metal (MOM) capacitors, as used in this thesis, confirm that Pelgrom’s proportionality holds for these capacitor DAC elements. The end result of this analysis is that DAC element mismatch is inversely proportional to the area of each DAC element. As ADC system-level tolerance for element mismatch decreases, as is the case with increasing DAC or ADC resolution, the area of each DAC element must increase.

2.1.2 Linearity Metrics in ADCs and DACs

Linearity in ADCs and DACs is characterized using two metrics; Differential-Non-Linearity (DNL) and Integral-Non-Linearity (INL). As defined by Maloberti in [8], the DNL of a data converter is “the deviation of the [output code] step size of a real data converter from the ideal width of the bins, Δ ”, where “bins” refers to the quantized output code. Using notation from [8], let us define X_k as the transition point between successive codes $k-1$ and k , and the width of the k^{th} bin or input difference resulting in that output code, as $\Delta r(k) = (X_{k+1} - X_k)$; then the equation for DNL is:

$$DNL(k) = \frac{\Delta r(k) - \Delta}{\Delta} \quad (2.2)$$

Integral Non-Linearity (INL) describes the overall deviation between the real input-output transfer function of the converter from an ideal straight line between the minimum and maximum outputs. The INL for an output code k is simply the integral of all DNL values up to and including the DNL for code k (also using notation from [8]):

$$INL(k) = \sum_{i=1}^k DNL(i) \quad (2.3)$$

Linearity can also be measured via the frequency spectrum of a converter’s output. A converter’s output spectrum is measured by sending an input frequency F_{IN} through the converter that is mutually prime in relation to the sampling rate F_s (i.e. F_{in} and F_s have no common denominator), and within the converter’s bandwidth (e.g. below $\frac{F_s}{2}$ for a Nyquist ADC). Using a logic analyzer or other digital data acquisition device, a large number of samples are collected and processed using the Discrete Fourier Transform (DFT). The number of samples (M) must be large enough such that the frequency width of each DFT bin, $\frac{F_s}{M}$ [54], is sufficient to correctly represent the input frequency and its distortion products in other bins, and the processing gain [55], $10 \log_{10}(\frac{M}{2})$ lower in magnitude than ADC quantization noise, is large enough to distinguish between the quantization noise of the ADC and the distortion products of the input signal.

From the output spectrum, several metrics are defined. SNR (Signal-to-Noise Ratio)

neglects nonlinearities (or harmonics), and is defined as the ratio of the power in the input frequency bin and the average of the power in all non-harmonic bins. *THD* (Total Harmonic Distortion) neglects noise and is the ratio of the input bin power and the sum of power in all harmonic bins. *ENOB*, (Effective Number Of Bits) is the most conservative specification of converter linearity and is defined by:

$$ENOB = \frac{SNDR(dB) - 1.76}{6.02} \quad (2.4)$$

where *SNDR* is the Signal-to-Noise and Distortion Ratio (also referred to as *SINAD*), which is a measurement of a converter's output frequency spectrum when applying a single sine wave at its input that is mutually prime (coherent) with the converter's sampling frequency (f_s). *SNDR* is defined as:

$$SNDR = \frac{P_{fund}}{\sum_{i=DC+1}^{i=fund-1} P_i + \sum_{i=fund+1}^{i=\frac{F_s}{2}} P_i}$$

where

$$P_{fund} = \text{Power in fundamental DFT bin} \quad (2.5)$$

$$\sum_{i=DC+1}^{i=fund-1} P_i = \text{Sum of power in all DFT bins above DC up to fundamental bin}$$

$$\sum_{i=fund+1}^{i=\frac{F_s}{2}} P_i = \text{Sum of power in all DFT bins above fundamental bin}$$

ENOB and *SNDR* are interchangeable as ADC metrics. For example an ideal 10-bit ADC with zero non-linearity and noise other than ideal quantization error has an *ENOB* of 10, and an *SNDR* of 61.96 dB.

2.1.3 How Element Mismatch Affects Linearity

This thesis utilizes analysis from the work Fredenburg, who developed the first closed-form solution for *ENOB* and yield as a function of mismatch in [56]. As described in [56], for an N -bit capacitor DAC arranged in N binary-weighted switching groups, the DNL of

all 2^N codes can be determined with just N unique DNL measurements of the major code transitions, e.g. 0100 \rightarrow 1000 for the case of a 4-bit ADC.

As defined in 2.1.1, the mismatch ratio of each unit element is defined as $\frac{\Delta C_{UNIT}}{C_{UNIT}}$. Just as the unit capacitors themselves are grouped into binary weighted groups where $C_i = \sum_{k=1}^{2^i} C_{UNIT}$, so can their mismatch elements into what [56] refers as “fractional mismatch” γ_i which is defined as: $\frac{\Delta C_i}{C_{i,nominal}}$ where $C_{i,nominal}$ is the desired size of the i th group and ΔC_i is the error that group exhibits due to mismatch. Fredenburg ultimately derives the *ENOB* of a differential DAC as a function of these fractional mismatch parameters as:

$$ENOB = N - \log_4 \left[1 + 3\gamma_0^2 + 3 \sum_{i=1}^{N-1} (2^{i-1} \gamma_i)^2 \right] \quad (2.6)$$

The consequence of equation 2.6 is that for a differential DAC used in a SAR ADC comprised of binary-weighted groups of unit capacitances C_{UNIT} , the resulting *ENOB* statistics can be determined entirely by the element mismatch statistics (ΔC_{UNIT}) which in turn determine the fractional mismatch statistics γ_i .

This and the preceding sections describe the fundamental tradeoff in DAC design between mismatch and capacitor size, and how they affect linearity. In traditional DACs where groups are hard-wired during layout, the only way to minimize ΔC_{UNIT} and thus γ_i is by drawing larger and larger capacitors as the desired *ENOB* increases.

2.2 Overcoming Mismatch Constraints Through Reconfigurability

2.2.1 Canceling Mismatch via Reconfigurability

The work in this thesis severs the previously described relationship between *ENOB* and element mismatch ΔC_{UNIT} , by removing the hard-wired constraint. The DAC in this work is completely reconfigurable because after fabrication, any unit capacitor (C_{UNIT}) can be assigned to any trial grouping (C_i) in the SAR cycle, independent of layout.

This novel reconfigurability provides the opportunity to *re-group* capacitors such that

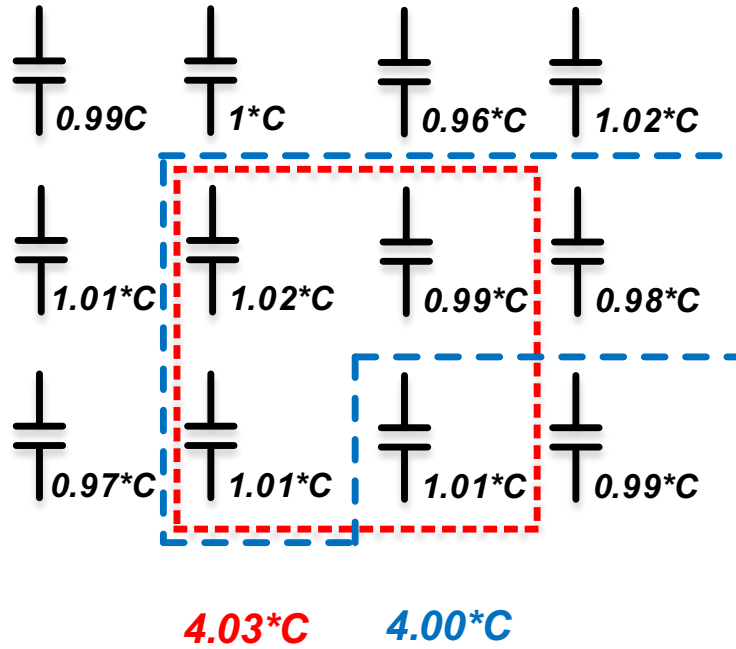


Figure 2.2: Initial mismatched grouping (red) vs. reconfigured, corrected grouping (blue).

their mismatches cancel, thereby overcoming the constraints described above in Section 2.1.3. For this DAC like any other there is a certain amount of C_{unit} mismatch from fabrication; however one is no longer forced to tolerate it, or over-size DAC elements to minimize it. Figure 2.2 shows an example where a sea of 12 unit capacitor elements (C_{UNIT}) are laid out in a 3×4 array. In this example, a binary weighted group of $4 \times C_{UNIT}$ is desired. The initial (red) grouping exhibits detrimental mismatch; its four elements sum to $4.03 \times C_{UNIT}$, which is greater than the intended $4 \times C_{UNIT}$. The reconfigured (blue) grouping replaces one mismatched cell with another, however the four elements now sum to the intended $4 \times C_{UNIT}$.

This process can be repeated for each trial group in a SAR cycle, thereby improving the overall linearity of the DAC transfer function. A 6-bit differential DAC example of layout-independent reconfiguration is shown in Figure 2.3. This compares an arbitrary arrangement DAC configuration with the DAC after it has been reconfigured.

2.2.2 Reconfiguration Techniques

With this new reconfigurability in hand, the next objective is to determine how best to reassign the unit cells and redistribute their random mismatch to maximize the resulting

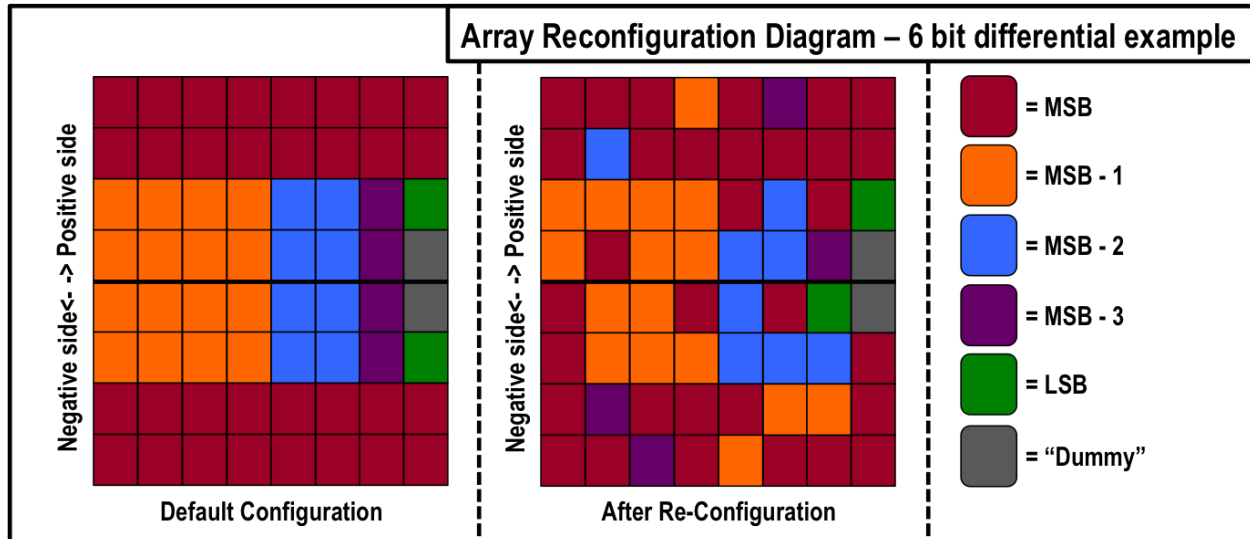


Figure 2.3: Before-and-after example of arbitrary layout-independent DAC reconfiguration.

$ENOB$ in equation 2.6 in Section 2.1.3. As equation 2.6 shows that $ENOB$ is negatively proportional to a summation of the weighted fractional mismatch $(2^{i-1}\gamma_i)^2$ across all capacitor groups (C_i), the maximum resulting $ENOB$ will occur when the **sum** is of minimum magnitude. This section will describe the reconfiguration techniques used and will confirm equation 2.6.

The number of possible configurations of unit elements to assemble a 10-bit binary weighted DAC (2^{10} , or 1024 unit elements) is staggering. In the case of a differential 10-bit DAC where each side (positive and negative) has 512 (or 2^9) elements, simply choosing the MSB_{-1} grouping of 256 capacitors follows the combinatorics formula:

$$C_{256}(512) = \binom{512}{256} = \frac{512!}{256!(512 - 256)!} \simeq 4.725 \times 10^{152} \quad (2.7)$$

While the DAC reconfiguration is now an optimization problem, the *optimal* algorithm to result in the most linear DAC is not the focus of this work. This work focuses on three intuitive approaches to reconfiguration:

MSB-first: Reconfigures the DAC beginning with the MSB group and ending with the LSB group.

LSB-first: Reconfigures the DAC beginning with the LSB group and ending with the MSB

group.

Random reconfiguration: Randomly reconfigure the entire DAC at once.

The MSB-first and LSB-first algorithms proceed in a similar fashion:

1. All capacitors in the DAC are made available for selection.
2. A SAR trial group is chosen to be reconfigured from the available unit capacitors (e.g. *MSB* in the case of MSB-first).
3. The desired number of unit capacitors for that SAR trial group is selected at random from the sea of available units. This would consist of $2^{(N-1)}$ units for the *MSB* grouping and 2^0 for the *LSB* grouping.
4. Linearity is measured using only the configured capacitors - for the first group active only, this would be a 2-bit ADC, then a 3-bit ADC for the second group, etc. When all groups have been reconfigured, then the linearity of the total ADC resolution will be measured. The linearity measurement is stored for future reference.
5. Steps 3 and 4 are repeated a user-defined number of times. All simulation results shown in this thesis using MSB-first or LSB-first reconfiguration algorithms repeat steps 3 and 4 100 times.
6. Store the configuration for the current SAR trial group being reconfigured that resulted in the best linearity measurement, and use this trial group's new configuration for all future iterations.
7. Proceed to the next SAR trial group to reconfigure and repeat steps 3-6.

Random reconfiguration is simply rearranging the entire array at once:

1. All capacitors in DAC are made available for selection.
2. All SAR trial groups are reassembled at one time without any iterative linearity measurement.

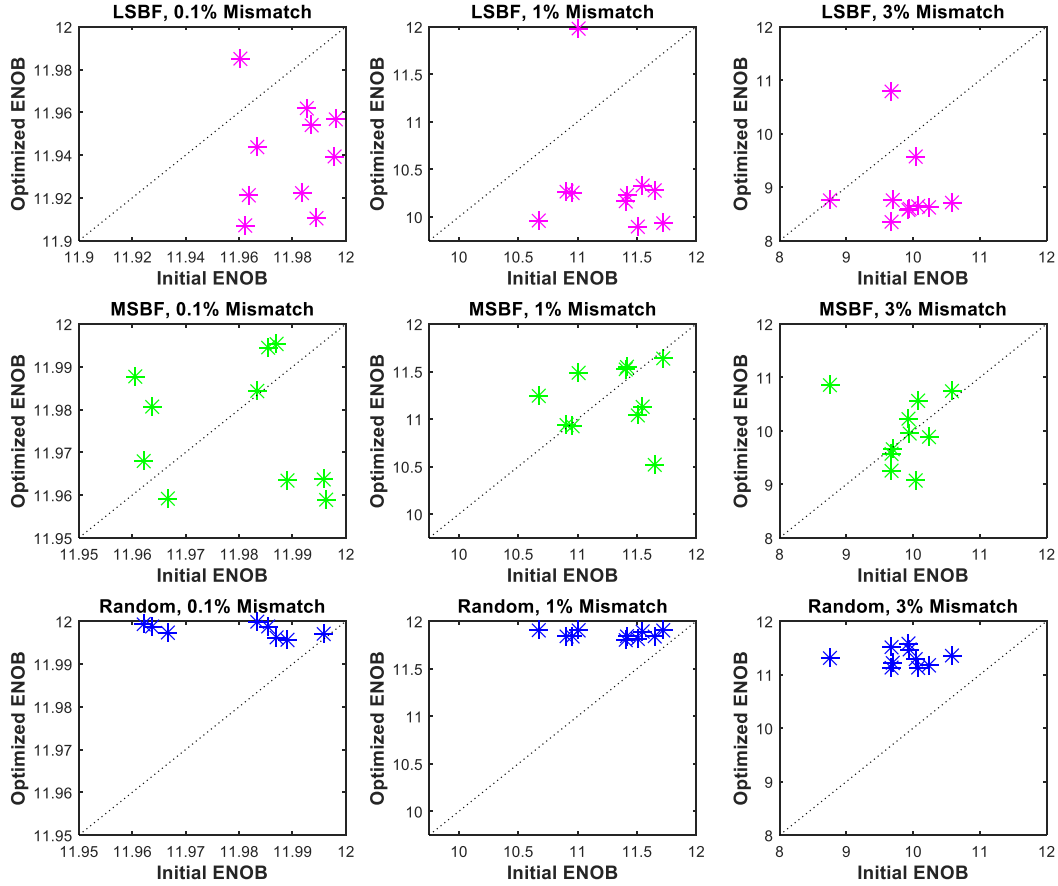


Figure 2.4: Comparison of the 3 described reconfiguration algorithms at 0.1%, 1% and 3% unit mismatch.

3. Linearity for the entire ADC is measured.
4. Repeat steps 2 and 3 a number of times. All simulation and measurement results shown in this thesis using random reconfiguration repeat steps 2 and 3 100 times.
5. The configuration with best linearity measurement is kept.

Figure 2.4 shows MATLAB simulation results comparing the three reconfiguration algorithms on capacitor arrays exhibiting three different amounts of unit element mismatch (0.1%, 1% and 3%). Each dot represents a separate capacitor DAC, where the x-axis value is that DAC's ENOB in its default configuration and the y-axis value is its ENOB after proceeding through the reconfiguration algorithm. The $y = x$ line is to clarify whether or

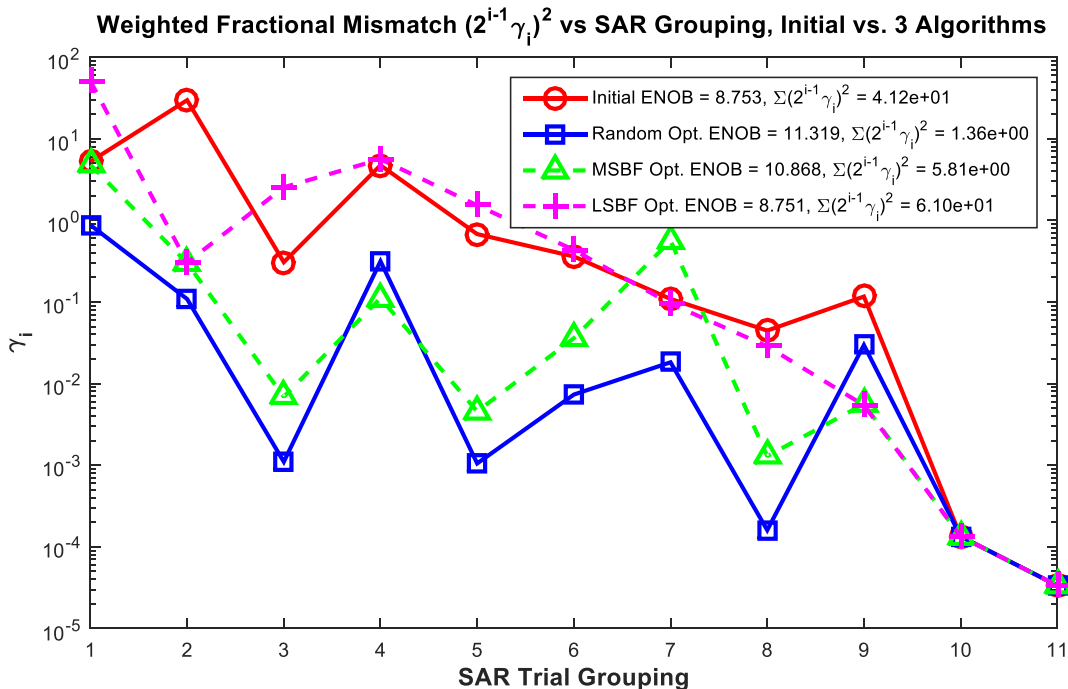


Figure 2.5: Weighted fractional mismatch $(2^{i-1}\gamma_i)^2$ of each SAR grouping of a 12-bit differential DAC with 3% unit element mismatch across 4 different DAC configurations. The resulting ENOB and sum of $(2^{i-1}\gamma_i)^2$ of each configuration are noted in the legend.

not reconfiguration has improved the DAC’s ENOB. The columns correspond to different amounts of unit element mismatch while the rows are for each reconfiguration algorithm. Figure 2.4 clearly shows that the MSB-first algorithm (row 2) is superior to LSB-first algorithm (row 1), although not every DAC is improved from its default state by optimization. This does not necessarily mean the algorithm would not have improved the DAC had it run longer than the 100 iterations described above; perhaps that DAC’s default configuration was already at a local optimum. With the number of iterations being only 100, random reconfiguration showed clear improvement in every DAC.

Figure 2.5 shows the weighted fractional mismatch $(2^{i-1}\gamma_i)^2$ from equation 2.6 of each SAR trial grouping for the initial (red) DAC configuration, and compares it with the outcome of the three reconfiguration algorithms discussed above. It uses the same 3% mismatch data set as Figure 2.4. The initial (red) DAC configuration from an array of unit elements with 3% element mismatch results in an ENOB of 8.753 and a weighted fractional mismatch sum of 41.2. The other three traces are weighted fractional mismatch after the DAC is

reconfigured using the three algorithms described in this section. The blue trace results from the random optimization algorithm and exhibits the best ENOB of 11.319 and lowest weighted fractional mismatch sum of 1.36. As shown in equation 2.6, the ENOB results directly from the weighted fractional mismatch, and the lower the sum of $(2^{i-1}\gamma_i)^2$ the higher and better the ENOB.

2.3 Completely Reconfigurable DAC Details

The remainder of this chapter describes the design details of the completely reconfigurable capacitor DAC, and how it achieves the reconfigurability detailed in Section 2.2.2. A top level description is given in Section 2.3.1 and the individual cells are described in Section 2.3.2.

2.3.1 Array Structure and Operation

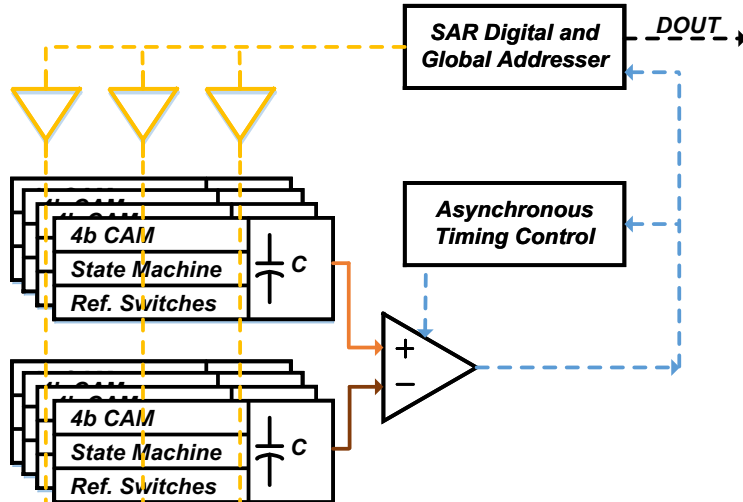


Figure 2.6: Conceptual diagram of DAC and ADC system.

Array Structure

A conceptual diagram of our DAC and ADC system is shown in Figure 2.6, and a more detailed representation is shown in Figure 2.7. The DAC is a symmetric, differential pairing

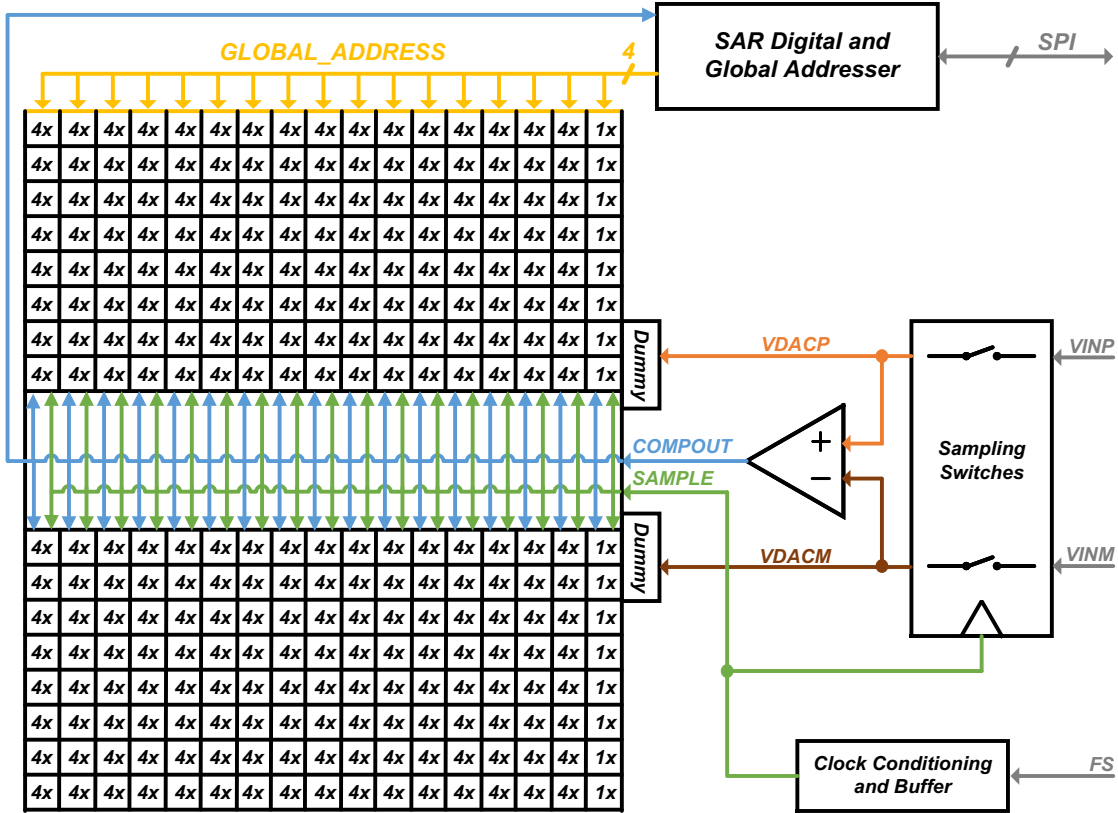


Figure 2.7: Detailed DAC and system diagram.

of two arrays, each containing 521 unit capacitors, corresponding to a 9-bit array with 9 additional unit capacitors. This extra capacitance can be used if desired (e.g. in a non-binary radix); when unused it slightly attenuates the dynamic range relative to a strictly binary weighted DAC. Connected to the dummy capacitor on each array is a 2-bit resistor ladder sub-DAC used for the final two SAR trials, yielding a total differential DAC resolution of 12 bits.

Shown in Figure 2.7, the reconfigurable DAC is constructed with 16 rows and 17 columns. Rows 1-8 are the positive side of the differential DAC while rows 9-16 are the negative side. All cells in columns 1-16 are the 4x cells, containing $4 \times C_{UNIT}$ capacitors grouped together by connecting their bottom plates within the cell. The cells in column 17 also contain C_{UNIT} capacitors but only one of these is connected to the top plate of the array and the switching circuitry below, thereby creating the 1x cell column with $1 \times C_{UNIT}$ cells. The cells are assigned to SAR trial groupings and activated via 4 bits of Content-Addressable Memory (CAM) within each cell. More detail is provided about the cells in Section 2.3.2,

and their assignment in Section 2.3.1.

While *random* variation (*i.e.* mismatch) is no longer of concern due to the novel reconfigurability, *systematic* variation still cannot be tolerated, because systematic variation does not necessarily adhere to the mismatch characteristics described in Section 2.1.1 upon which this work is based. To eliminate systematic variation, dummy capacitors (identical metal patterns without circuitry beneath) surround the exterior sides of the array, ensuring that the entire structure has consistent metal density well past the actual array’s area. As the total number of connected usable capacitors is $1040 \times C_{UNIT}$, corresponding to a binary-weighted 10-bit array with 16 extra capacitors, the two additional bits are achieved with the sub-DAC whose total capacitance is equal to $1 \times C_{UNIT}$ and is discussed in Section 2.3.2.

Array Operation

In a conventional SAR ADC, capacitors are connected to reference voltages depending on the comparator decision. Because the groupings of unit capacitors are predetermined during layout, the logic that decides whether to switch that group can be very simple; a shift-register or counter-based approach is common.

In the case of this work, the traditional SAR control logic is split into two sections: the logic within the miniature state machine in every cell across the DAC (discussed in Section 2.3.2), and the addresser that sends out the 4-bit words that activate each SAR trial grouping in the DAC via the cells’ Content Addressable Memory (CAM). The addresser system is synthesized along with the SPI interface, array programming controls, digital output registers, and all other non-ADC-operational logic. It is located separately from the array.

Complete reconfigurability is achieved by activating programmable groups of capacitors during the steps of SAR operation. The addresser sends out the 4-bit word corresponding to the current SAR trial to activate the cells whose CAM matches that word. When a cell is activated, the cell switches between references (V_{REFP}, V_{REFM}) according to the comparator result of the current SAR trial. A simplified schematic of the addresser is shown in Figure 2.8. The addresser is essentially a 4-bit-wide (to match the CAM memory in DAC cells) by 13-row (12 DAC states with 1 sampling state) long circular chain of flip-flops. Although




Figure 2.8: Simplified addresser schematic

the digital control block (via SPI) initializes the addresser to either default or custom user-programmable values, once the ADC is in operation it is in a separate clocking domain to enable its asynchronous ADC operation further discussed in Section 3.2.2.

The output from the addresser block is sent directly to the top of the array where it is buffered before becoming the complimentary bitlines seen by each cell’s CAM. The complimentary bitlines are driven by a row of buffers on the top of the array which receive the initial 4-bit address from the addresser unit almost simultaneously (the maximum sized output buffer is used by the synthesized digital for the addresser outputs). These complimentary bitline buffers above the array dominate the power consumption of the entire chip during ADC operation. The ADC timing is discussed further in Section 3.2.2, but here we note that the tradeoff of bitline propagation delay and buffer power consumption is the fundamental speed-power trade-off in this work. The bitlines are minimum-width metal to reduce the amount of capacitance as much as possible, and the matching detection circuits that load them have minimal capacitance as discussed in Section 2.3.2.

The comparator outputs *COMPPPOS* and *COMPNEG* (which correspond to a positive or negative decision) are sent to the center of the array between the positive and negative sides directly from the comparator block, and are simplified in Figure 2.7 as the signal *COMPOUT*. The comparator outputs are then buffered using similar buffers as the complimentary bitlines. The only difference between positive and negative sides of the DAC is that the comparator output polarity is flipped in routing for the negative side; the DAC cells themselves are identical on both sides of the DAC but the *COMP** buffers in the center swap *COMPPPOS* and *COMPNEG* to the negative side of the array.

Table 2.1: Default Binary-Weighted Array Programming

Comparator Decision	Num. of Cells $\times C_{UNIT}$	Reference During Sampling	Binary 4-bit Address
1*	32×4	V_{REFP}	0010
1*	32×4	V_{REFM}	0011
2*	16×4	V_{REFP}	0110
2*	16×4	V_{REFM}	0111
3*	8×4	V_{REFP}	1110
3*	8×4	V_{REFM}	1111
4*	4×4	V_{REFP}	1010
4*	4×4	V_{REFM}	1011
5	4×4	V_{REFP}	1000
6	2×4	V_{REFP}	1100
7	1×4	V_{REFP}	0100
8	2×1	V_{REFP}	0101
9	1×1	V_{REFP}	1101
10	$\frac{1}{2}$	V_{REFP}	$UNIT_1^*$
11	$\frac{1}{4}$	V_{REFP}	$UNIT_2^*$
12	N/A**	V_{REFP}	SAMPLE**

* Groups are switched differentially
 $UNIT_X^*$ Done by sub-DAC discussed in 2.3.2
SAMPLE** After this final decision initiate sampling

Cell Assignment

Beneath the capacitors in each cell are four bits of content-addressable memory (CAM), and a small state-machine that controls the switching of that cell between the two voltage references V_{REFP} and V_{REFM} . Four bits of memory and thus 16 possible programmings are needed to identify each of the 12 SAR trial groupings and the additional sampling state. These circuits (detailed further in Section 2.3.2) are identical for both 4x and 1x cells; as all cells contain 4 capacitors for layout uniformity. The only difference between 4x and 1x cells is the number of bottom plates tied to the local reference switches. During SAR operation, the addresser described in Section 2.3.1 cycles through a series of 4-bit words that correspond to each SAR trial, and this word is sent across the entire DAC via the complimentary bit lines that run vertically across the DAC. During each SAR trial, the switching of the capacitors between references depends on the relationship of the 4-bit local word stored in the content addressable memory in each cell, and the 4-bit global word being sent to all cells by the addresser.

Table 2.1 lists the default mapping of 4-bit words to SAR trials, although the addresser can be reprogrammed to have any order of 4-bit words. The first 4 MSB capacitor groups are switched differentially to keep the comparator input common-mode within $\frac{1}{16}$ of its initial value of $\frac{V_{REF}}{2}$, and all subsequent groups use set-and-down switching. This is almost identical to the combination of switching methods in [27] and is further discussed in Section 3.2.3. The content addressable memory beneath each cell determines if the 4-bit global word matches the 4-bit local word in that cell. If so, the state machine is enabled and monitors the comparator decision to switch or not switch. The comparator output and sampling clock edge are routed to the center of the DAC and buffered symmetrically up the positive and down the negative arrays. After the SAR algorithm proceeds through all the 4x cell and 1x cell groupings which contribute the first 10 bits, the dummy capacitor sub-DAC is used for 2 extra decisions resulting in a total resolution of 12 bits. After the final comparator decision, sampling the input for the next conversion occurs until the next falling edge of the sampling clock. A timing diagram for the entire ADC operation is shown in Section 3.2.2.

2.3.2 Cell Structure

The most significant challenge in creating this reconfigurable DAC is transferring the SAR operational logic from the top system level to an individual cell level. This is necessary for the cells to be able to self-determine their own actions during SAR operation without the addition of significant digital circuits. Due to the simple and minimal logic required in conventional SARs, from a top-level ADC perspective this potentially reduces efficiency, so great care was taken to minimize the penalties incurred by this novel approach.

Each capacitor cell has a 4-bit content-addressable-memory (CAM) and miniature state machine which are placed directly beneath the MoM capacitor structure. All cells in each SAR trial grouping share the same 4-bit memory programmed value. The memory is fully programmable, so that each capacitor cell can be assigned to any group for complete reconfigurability. The CAM and state machine fit directly under the MOM capacitors in each cell structure, so there is minimal area penalty for this new architecture. Figure 2.9 shows the 3-D structure of a 4x or 1x cell, all of which measure $10\mu\text{m} \times 10\mu\text{m}$.

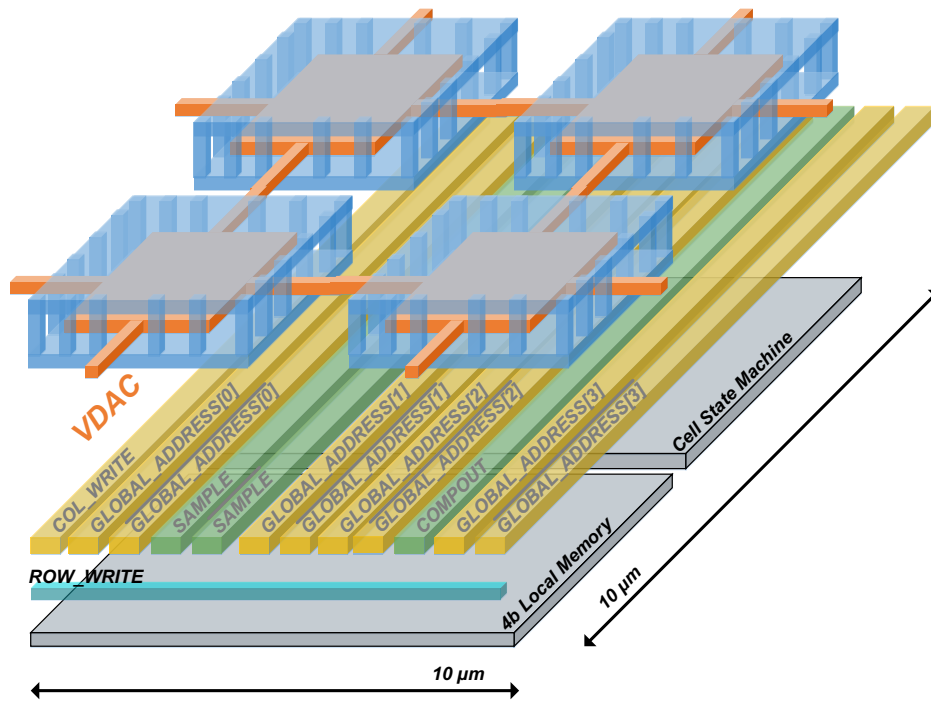


Figure 2.9: Simplified 3-D diagram of 4x or 1x cell structure.

Out of the 521 unit capacitors in each half array, 512 unit capacitors are grouped into 128 cells of 4 capacitors (“4x cells”), 8 unit capacitors are in their own individual cells (“1x cells”), and the dummy capacitor has its own unique cell that controls the sub-DAC switching. Figure 2.10 shows a simplified schematic of the 4x and 1x cells.

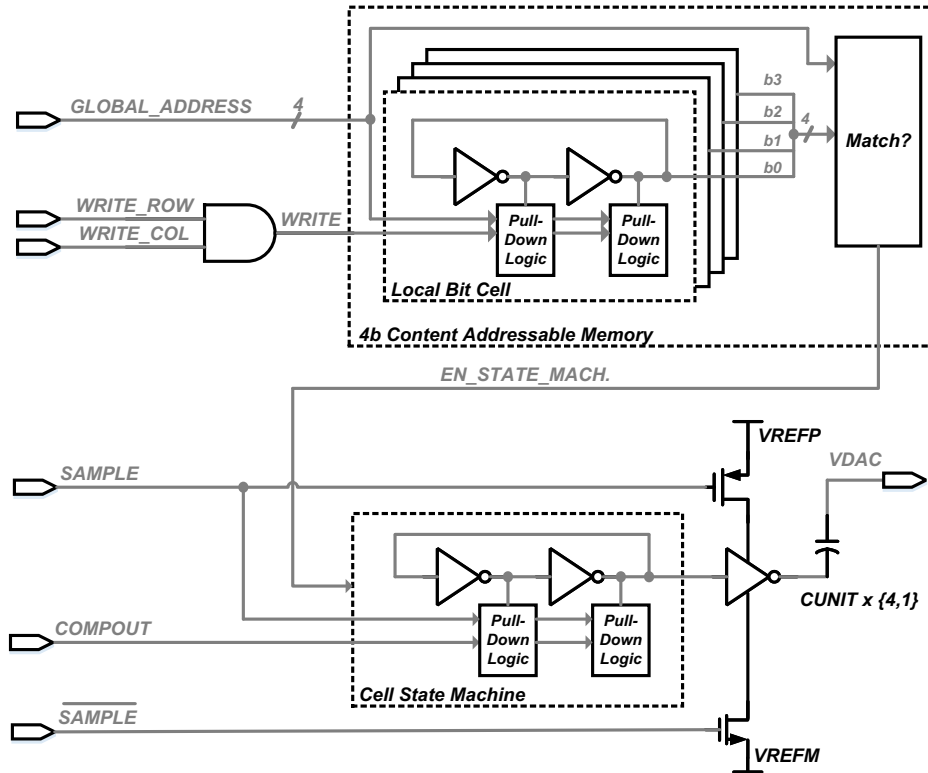


Figure 2.10: Simplified schematic of one DAC cell.

Content Addressable Memory Bitcell

The cell memories are programmed and addressed by four complimentary bitlines (e.g. $b_0, \bar{b}_0, b_1, \bar{b}_1, \dots, b_3, \bar{b}_3$) that span each column of the array, and each column is identical in that every cell in the entire array sees the same “address” simultaneously. Because this 4-bit address is seen by the entire array simultaneously, it is referred to as the “global address.” During the initial programming or reconfiguration, write signals are sent to each row and column so each cell can be written individually if desired. When both a row and column write signal are asserted, the cell’s SRAM stores the global address that is present on the bitlines at that time. Figure 2.11 shows the schematic of an individual memory bit cell.




Figure 2.11: Schematic of individual memory bit cell.

Global vs. Local Bit Match Indicator

During SAR operation, a custom match indicator cell determines whether or not the cell should be active for the present SAR trial, as shown in Figure 2.12. Each global bitline traverses the entire column of the DAC, so the load presented by one match indicator cell is replicated 16 times; one for each of the 16 rows of the DAC. Because the propagation delay of the bitlines is a critical path for determining the overall conversion speed of the ADC, the prohibitively high gate capacitance of a conventional inverter buffer is replaced with the much lower drain-bulk capacitance C_{DB} . While there is a local buffer (inverter in Figure 2.11) that buffers the global bit line during programming, that is the only time it is employed. When all four match indicators assert because the global address present on the bitlines match the locally programmed memory, the cell activates its state machine and monitors the comparator decision.




Figure 2.12: Schematic of global bit vs. local bit match indicator cell.

Miniature State Machine

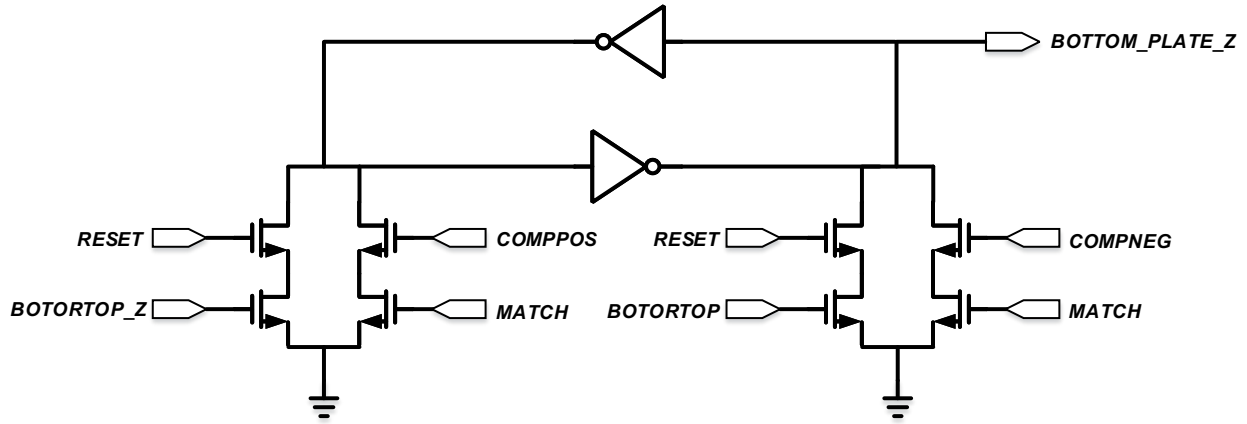


Figure 2.13: Schematic of miniature state machine in each DAC cell.

The miniature state machine in each cell controls the bottom-plate switches of the unit capacitors in each cell. Once activated by the matching indicator circuit, the state machine monitors the comparator output and depending on the memory contents (SAR trial assignment) of the cell, switches the bottom-plate to either $VREFP$ or $VREFM$. The state machine is similar to the CAM bitcell as it is simply a pair of inverters in feedback with pull-down logic that overpowers the appropriate inverter to change state. A schematic of the state machine is shown in Figure 2.13. The assignment of DAC cells to SAR trial groupings is discussed in Section 2.3.1, and the SAR reference switching technique is discussed in Section 3.2.3.

The *NMOS* and *PMOS* devices shown in Figure 2.10 that connect the cell's bottom

plate to the references $VREFP$ and $VREFM$ are sized for a DAC settling time of 500ps, as described by equation 2.8:

$$\begin{aligned}
 t_{settle} &= RC(N + 1) \ln(2) \\
 t_{settle} &= targetDACsettlingtime \\
 R &= cell\ switch\ resistance \\
 C &= cell\ capacitance \\
 N &= total\ ADC\ resolution\ (12)
 \end{aligned}
 \tag{2.8}$$

Capacitor Design

Due to all of the signaling going on underneath the capacitors, it is very important that the critical top capacitor plate connected to the comparator inputs is shielded as much as possible. Any parasitics are kept mostly to the bottom plate which is driven by references at all times and thus not sensitive to other coupling. A structure similar to the one used in [26] is chosen, and is shown in Figure 2.9.

Metal layers 4 through 6 form bottom plate “cages” (blue) around the top plate (orange), which is in the center on layer 5. The top plate is connected across cells on layer 5, and the bottom plate connections within each cell (creating the 4x cells) are directly below the top plate route on layer 4 to assist with shielding the top plate metal.

The parasitic extraction tools estimate the top-plate to bottom-plate coupling at nearly 99%, with a value of 2.43 fF. In addition to shielding the top plate from signal noise below, the high top-bottom plate coupling leaves little room for any parasitic top plate capacitance, so dynamic range attenuation is greatly reduced.

Unit Capacitor Sub-DAC

As mentioned in sections 2.3.1, and 2.3.2, the final two steps of the SAR cycle use a sub-DAC to resolve the final two bits of the search space. A schematic of the unitcap sub-DAC is shown in 2.14. The unitcap sub-DAC cell is the same width ($10\mu m$) as the other array cells but is twice as tall ($20\mu m$) to accommodate the resistor ladder and additional switching logic.




Figure 2.14: Schematic of unitcap sub-DAC cell for final two bits of DAC resolution.

The sub-DAC has a capacitance of $1 \times C_{UNIT}$ although it has a three-bit resistor ladder underneath it that connects the bottom plate of its capacitor to multiples of $1/8 \times (V_{REFP} - V_{REFM})$. This is equivalent to switching the same fraction of C_{UNIT} , and an example of switching from full reference ($V_{REFP} - V_{REFM}$) to $3/4 \times (V_{REFP} - V_{REFM})$ is shown in equation 2.9.

$$\begin{aligned}
 C_{UNIT} \times (V_{DAC1} - (V_{REFP} - V_{REFM})) &= Q_1, \text{ before switching} \\
 C_{UNIT} \times (V_{DAC2} - \frac{3}{4}(V_{REFP} - V_{REFM})) &= Q_2, \text{ after switching} \\
 Q_2 &= Q_1, \text{ cancel } C_{UNIT} & (2.9) \\
 V_{DAC2} - \frac{3}{4}(V_{REFP} - V_{REFM}) &= V_{DAC1} - (V_{REFP} - V_{REFM}) \\
 V_{DAC2} &= V_{DAC1} - \frac{1}{4}(V_{REFP} - V_{REFM})
 \end{aligned}$$

However this sub-DAC method only works if all of the C_{UNIT} capacitors match within the final (12-bit) ADC resolution tolerance, because the change in V_{DAC} affected by the sub-ADC is still relative to $\frac{C_{UNIT}}{C_{DAC,TOTAL}}$. For this case of a 12-bit ADC using a 10-bit array of C_{UNIT} , if C_{UNIT} is not within $\frac{1}{2} \times \frac{C_{DAC,TOTAL}}{2^{12}}$ then the DAC voltages created by switching the sub-DAC will not be scaled correctly to provide the desired final two bits of search space resolution. In other words, creating a sub-DAC with a resistor ladder beneath a unit

capacitor does not free the DAC from the overall capacitor matching constraints. One can not make a DAC with 8 bits of linearity due to mismatch into a 12-bit linear DAC nearly for free, simply by adding a sub-DAC beneath the unit capacitor. Furthermore, the use of a sub-DAC in this work does not diminish the challenges of overcoming DAC element mismatch.

The resistors beneath the specialized unitcap cell are unsilicided polysilicon resistors arranged in 16 segments of 1.7 k Ω each, creating a total resistance between references of nearly 27.5 k Ω . According to equation 2.8 this results in a worst-case DAC settling of 0.847 ns.

The switches connecting $\frac{4}{4}$, and $\frac{3}{4} \times (V_{REFP} - V_{REFM})$ to the bottom plate of C_{UNIT} are PMOS devices because they function as pull-up devices to higher reference voltages, while the switches connecting $\frac{0}{4}$, and $\frac{1}{4} \times (V_{REFP} - V_{REFM})$ are NMOS devices because they function as pull-down devices to lower reference voltages. The connection to $\frac{1}{2} \times (V_{REFP} - V_{REFM})$ is a transmission gate node of complimentary NMOS and PMOS devices to ensure low on-resistance despite $|V_{GS}|$ being less than 2 threshold voltages for either device.

Because they are used during normal ADC operation, the $\frac{8}{8}, \frac{3}{4}, \frac{1}{2}$, and $\frac{1}{4}$ switches require their own state machine logic similar to the other cells.

2.3.3 Brief Note Regarding Novelty

While there are published examples of post-silicon capacitor adjustments via additional capacitors [45],[38], to the authors knowledge there is not a previously published DAC that is completely reconfigurable, with the possible exception of the DAC used in the seminal work [33]. The author of [33] mentions the local decoder is small enough to fit under the capacitors, however as shown in figure 10.6.2 of [33] this only refers to reference switching and not SAR trial assignment; this implies SAR trial assignment is done completely by row and column activation. Furthermore the diagram of the DAC in [33] implies significant digital overhead if it were to achieve complete reconfigurability as opposed to a finite set of switching patterns. In addition, no explicit mention of reconfigurability is made, nor is there any mismatch mitigation.

2.3.4 Other Contributors to This Work

Andres Tamez

This work continues previous research by Andres Tamez. Tamez’s work included more analysis of the optimization of reconfiguration techniques, however a functional prototype ADC was not realized. Tamez’s SAR ADC architecture was almost entirely different from this work, however the CAM bitcells (Section 2.3.2) and match indicator circuit (Section 2.3.2) remain in the ADC discussed in this thesis.

Lu Jie

Lu Jie provided significant modeling and simulation assistance with the algorithm comparison in Section 2.2.2. Jie was also a significant contributor during ADC testing, discussed in Section 3.4.

CHAPTER 3

A Mismatch-Immune 12-bit SAR ADC

3.1 ADC Design Considerations

As discussed and defined in Section 2.1.2, the most conservative performance metric of any ADC is its Signal-to-Noise-and-Distortion Ratio (SNDR), which is the ratio of the output power in the fundamental bin, divided by the sum of power in all other bins except DC (equation 2.5). Here the equation is rewritten with the denominator separated into terms representing known SAR ADC design components such as sampling noise, comparator error, ideal or uniform-step quantization error, sampling error from jitter (as mentioned in equation 3.2) and increased quantization error from non-uniform quantization or DNL/INL errors.

$$\begin{aligned} SNDR &= \frac{P_{FUND}}{\frac{kT}{C_{sampling}} + V_{\sigma,COMP}^2 + \varepsilon_q + V_{RMS,jitter}^2 + P_{DNL/INL}} \\ P_{FUND} &= \text{Power in fundamental bin} \\ \frac{kT}{C_{sampling}} &= \text{Sampling thermal noise power} \\ V_{\sigma,COMP}^2 &= \text{Noise power of comparator error} \\ \varepsilon_q &= \frac{LSB^2}{12} = \text{Spectrum/noise of quantization error (ideal)} \\ V_{RMS,jitter}^2 &= \text{Noise power of sampling jitter} \\ P_{DNL/INL} &= \text{Noise power of DNL/INL error} \end{aligned} \tag{3.1}$$

Because this design should eliminate mismatch, $P_{DNL/INL}$ is neglected for the purposes of calculating requirements of the ADC.

3.2 Implementation of Recent Innovations

To demonstrate the reconfigurable DAC in a state-of-the art ADC, many of the innovations mentioned in Section 1.4 are used. Below the context and implementation details of the following techniques is discussed:

- Dynamic comparator
- Asynchronous comparator timing
- Set-and-down DAC switching
- Flexible transfer function

3.2.1 Comparator Design

The comparator used is based on Miyahara’s work [57]. A schematic of the modified version is shown in Figure 3.1.

The comparator has two stages; a dynamic input stage followed by a regenerative cross-coupled latch. The input stage, while bearing resemblance to a classic differential pair with active load, behaves differently. During comparison, the input devices differentially discharge two capacitive nets ($MIDP$, $MIDM$), which are precharged to V_{DD} during the reset phase. These internal capacitive nets are the inputs to the cross-coupled regenerative latch. This work is modified from the original in two ways: there is no offset compensation, and device sizes have been increased for higher speed operation and better noise performance, whereas in the original work all devices were uniformly minimum-sized.

A timing diagram for the comparator is shown in Figure 3.2. During the **reset phase**, the PMOS devices $M4$, $M5$ that would normally serve as an active load are instead active in opposite polarity from the input devices, and only reset the internal nodes of the first stage to V_{DD} . $M4$ and $M5$ must reset the internal nodes sufficiently that any remaining

roughly $\frac{3}{5}LSB$. This is not sufficient for the noise performance required by 12-bit operation with V_{REF} of 1V, however as will be discussed in Section 3.3.2 due to the chosen sampling method larger input capacitance from the gates of input devices $M2$, $M3$ cannot be tolerated.

3.2.2 Asynchronous Timing

As discussed in Section 1.4.2, idea of an asynchronous SAR ADC is to let the comparator “clock itself,” such that for large input voltages where $V_{IN} \gg V_{LSB}$, which are likely to occur for all but two of the decisions in the SAR cycle, the comparator makes a quick decision and the SAR cycle can immediately proceed, thus decreasing overall conversion time.

A figure detailing the sampling, addressing, and comparator decision switching the cells during ADC operation is shown in 3.3. After each comparator decision propagates through the DAC (causing the state machines in each activated cell to switch reference connection appropriately), the addresser discussed in Section 2.3.1 is clocked and sends out a new word across the array, such that the next SAR trial group is activated. This is illustrated by the “comp out” and “DAC Address” fields in Figure 3.3. This forms the critical timing loop that determines the overall conversion speed of the ADC.

The speed of this asynchronous loop is limited by the propagation delays across the array, which are longer than the minimum comparator input decision time. There is an added programmable delay cell that can delay the comparator activation, which ensures that there is a requisite amount of DAC settling time and that the cells have all been activated as needed. The programmable delay cell is largely a safety measure; measurements in Section 3.4 show that ADC performance increases for the first 5 delay settings and further added delay yields negligible improvement.

3.2.3 DAC Switching Method

The DAC switching method used is very similar to C.C. Liu’s in [27]. Liu first introduced what he named (and is now commonly referred to in literature) the “set-and-down switching” method in [26] which was discussed in Section 1.4.1.

In the first set-and-down switching paper [26], the input is sampled onto the top plates




Figure 3.3: Asynchronous timing diagram of ADC.

of the DAC, with V_{DD} on the bottom plates of all capacitors. Then after each comparator decision only one side of the DAC switches a capacitor group from V_{DD} to V_{SS} .

The advantages of this technique are reduced power consumption and circuit complexity:

- Only one side of DAC is switching
- Always switching in one direction (V_{REFP} to V_{SS})
- Only need an inverter as a switch beneath DAC bottom plates

while the disadvantages of this technique are:

- Non-symmetric switching reduces even-order harmonic cancellation properties of differential DAC switching
- Severe common-mode variation (from $\frac{V_{REF}}{2}$ to $\frac{LSB}{2}$)
- Requires top-plate sampling, or extra input polarity-check comparison if bottom-plate sampling¹

It should be noted that both of Liu's set-and-down papers [26] and [27] are 10-bit designs so the disadvantages were easily overcome, although the common-mode variation is the biggest performance limitation of [26].

In [27], Liu addresses the common-mode variation issue by changing the first three capacitor switches to be symmetric; instead of only switching the higher side down, Liu also switches the lower side up. By keeping the DAC common-mode at $\frac{V_{REF}}{2}$ for the first three decisions, the common-mode variation from $\frac{V_{REF}}{2}$ to $\frac{V_{REF}}{16}$ has been greatly reduced, which

¹If a common-mode shift and non-binary switching method less than 2 is used as in [23], then a polarity check is not needed however the common-mode shift will still consume extra time.

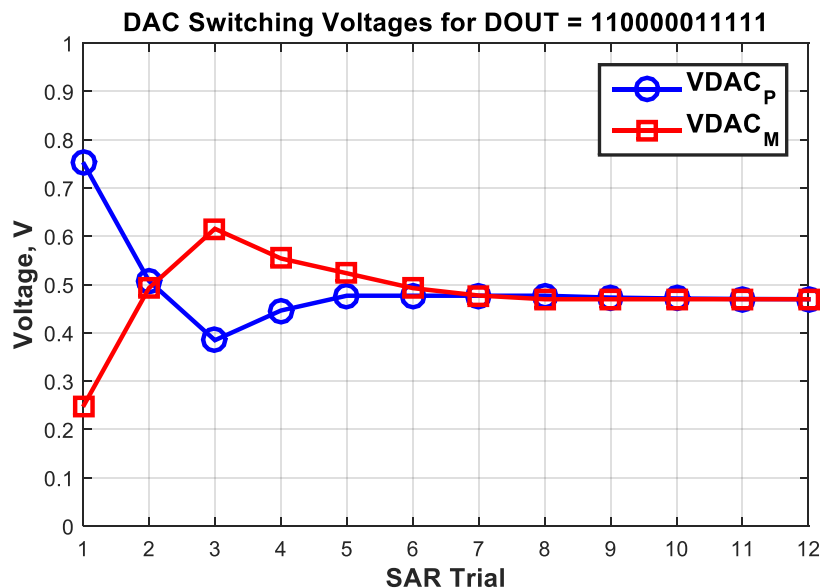


Figure 3.4: DAC voltages during SAR cycle for output code 3103. Note the first 4 switching events are symmetric and differential, while the final 8 are set-and-down as in [27].

significantly relaxes comparator design requirements and keeps the input devices in a lower-overdrive, higher g_m biasing.

Figure 3.4 shows an example of the DAC voltages while converting an input that results in the ADC output code 3103. In the default programming scheme, the DAC is switched symmetrically for the first four comparator decisions, which in the case of V_{REFP} being $1.0V$ results in a common-mode DAC voltage of $\frac{V_{REFP}}{2} = 500mV$ for the first four decisions. The final 8 decisions (beginning with SAR trial 5 in Figure 3.4) the set-and-down switching method can be seen, which in the event of every set-and-down switching being the same decision (so 7 “1s” or 7 “0s”) would result in a worst-case common mode decrease to $468.75mV$ for the final decision.

In [27] Liu also incorporates a “windowing” scheme to save DAC switching power by eliminating unnecessary capacitor switches via extra comparators. This technique is not implemented in this work, although for a future larger and higher resolution DAC where power consumption due to reference switching becomes more significant, windowing could be implemented to good effect.

3.2.4 Flexible Transfer Function

In a conventional DAC, the decision of what transfer function the DAC should exhibit, e.g. binary or non-binary such as [33], must be made in layout prior to fabrication as explained in Section 2.2.1. However in this work where the capacitors are truly reconfigurable, then the end user can modify the DAC transfer function at will. Because input sampling occurs onto the entire array, if one chooses to utilize fewer than the maximum available number of capacitors, the input dynamic range must be reduced by the same proportion of unused capacitors to total capacitors (this is a known penalty of non-binary radix conversion, also discussed in [43]). As will be discussed in Section 3.4.2, for the peak performance of this work a non-binary weighed MSB (smaller than 2^{N-1} codes) is used to reduce the conversion space from a full-scale input to roughly $\frac{3}{4}$ due to input dependent nonlinearities in the sampling circuitry. For any system where the input dynamic range is less than V_{REF} this flexible resolution ability is advantageous.

3.3 Other Design Considerations

3.3.1 Voltage References

On-chip voltage reference buffers are required in many medium to high performance ADCs, or most ADCs with higher-frequency sampling rates, because of the difficulty of getting clean reference voltages through packaging parasitics. The most detrimental of these parasitics is bondwire inductance. Even if off-chip references are used (as is often the case in academic and prototype designs), they require additional input pads to reduce bondwire inductance via parallelization. Separating high-accuracy (reference and precision analog) power domains on-chip is another inductance mitigation technique which further increases routing complexity and pad count. As pad count increases, so does packaging size, and thus also the length of the bondwires; *i.e.* the more pads on chip, the more pins in package, the bigger the package must be, and the longer the bondwires must be.

“Referenceless” operation (where the power supply is used as a reference) is not practical for most higher-resolution designs. In higher resolution ADCs with larger DACs, the

amount of power required to overcome the increased thermal noise requirements result in too much circuitry and ripple on V_{DD} for it to be used as a DAC reference. The choice of a referenceless design is also a function of the amount of decoupling capacitance placed on chip, and quality of power supply off-chip. In this case, great care was taken in timing design of the asynchronous loop operation to ensure that no other events were occurring on the power supply V_{DD} during DAC switching and comparator decisions. Early prototypes referenceless operation was attempted proved unsuccessful, and ultimately separate reference supplies were used.

All unused area on this chip is filled with as much decoupling capacitance as allowed by the conservative set of design rules; 2.5V thick-oxide MOSCAP devices are used for low Q, and low gate leakage, as well as a seven layer stack of opposing metal fingers between V_{DD} and V_{SS} wherever possible. If this ADC was not pad-limited (perhaps by use of a serialized D_{OUT} instead of parallel I/O, and by 2-wire digital control instead of 5), more supply and reference pads could be employed to reduce bondwire inductance via parallelization and possibly result in increased performance. The primary concern for the DAC layout was robust, low resistance routing across the array, utilizing the top two (thickest) metal layers, so the DAC would get as clean a reference as possible and the logic and CAM cells beneath the DAC would have a stable supply and not lose programming due to supply droop from series resistance. All of the dummy capacitors surrounding the DAC are used as decoupling capacitors.

3.3.2 Sampling Technique

To reduce the routing complexity of the DAC, the choice was made to sample the input on the top plates as in [26], and [27]. However unlike those top plate sampling examples, the bottom plates are disconnected before the top plates. This method will be referred to as a “hybrid” sampling technique in table 3.1. This differs from the conventional, or “full” bottom plate sampling method such as shown in the schematic of Figure 3.5. A summary of the three techniques is shown in table 3.1.

The main advantage of bottom plate sampling this work benefits from is disconnecting the




Figure 3.5: Conventional SAR DAC example. A single-ended 3-bit DAC samples signal V_{IN} onto the capacitors' bottom plates, and during conversion will switch between V_{REF} and GND . From [5].

Table 3.1: Sampling Techniques Summary

Technique:	Complexity	Input-dependent non-linear charge injection from V_{SB}	Input-dependent non-linear charge sampling from C_{DB} and comparator input
Top Plate	Most Simple	Yes	Yes
Bottom Plate	Most Complex	No	No
Hybrid	Compromise	No	Yes

capacitor from the reference before it is disconnected from the input. A cell-level schematic and timing diagram of the hybrid sampling technique is shown in Figure 3.6. Disconnecting the bottom plates first prevents input-dependent charge injection from the variation in the source-body voltage (V_{SB}) of the input NMOS switch. This is because when the bottom plate switch is disconnected, current cannot flow through the input switch into a capacitor that is not connected on the other side as there is no longer a path to a lower voltage. This bottom-plate disconnection was accomplished by surrounding the reference-switching inverter beneath each DAC cell's capacitors (also shown in Figure 3.6) with series devices that isolate them from the references during the sampling instant. The primary cost of this

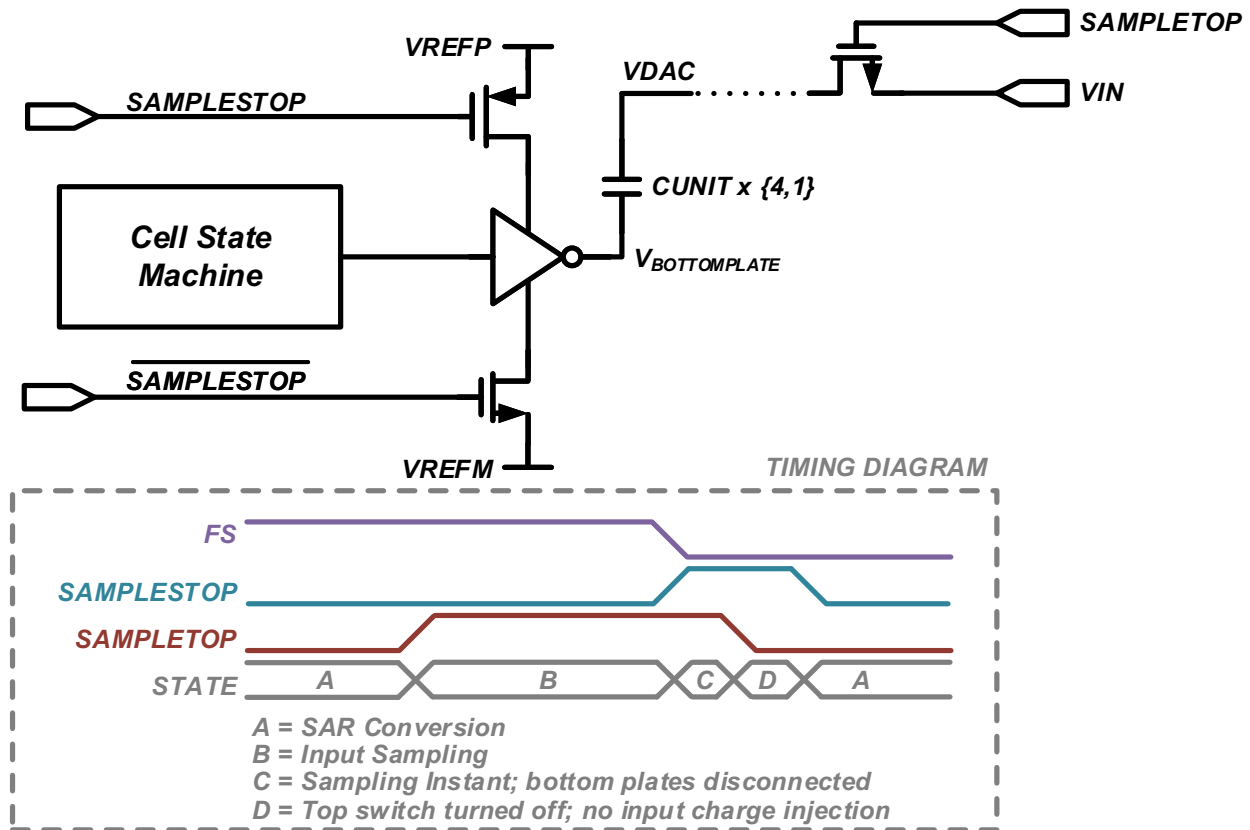


Figure 3.6: Cell-level schematic and timing diagram for hybrid sampling method.

feature was the routing of the complimentary $SAMPLE_{STOP}$ signals, which are sent from the same cells in the center of the array as the $COMP_X$ signals described in Section 2.3.2. This is a much easier solution than routing the input voltage across the array as discussed below.

The main disadvantage of the hybrid sampling method compared with the traditional bottom plate method is that the input is sampled onto the same node as the comparator inputs. This means any nonlinear capacitance at either the drain of the sampling switch ($C_{DB,switch}$) or the input gates of the comparator ($C_{GS,comparator}$) will affect each sample. As the input voltage and thus the voltage across $C_{DB,switch}$ and $C_{GS,comparator}$ varies rail-to-rail, their nonlinear voltage-dependent capacitance will vary in an input-dependent fashion. Because the DAC samples *charge* from the input, this results in an input-dependent charge variation that is as nonlinear as those capacitances' ratio with the total capacitance of the DAC.

The main reason the traditional bottom plate sampling technique was not implemented was to avoid increased routing complexity across the array, which would require the following additions:

- Low resistance interconnect from the input voltage across every row and column to connect to the bottom plates in each cell,
- Additional switch in each cell to connect the bottom plate to the input voltage,
- Low resistance interconnect from a boosted or “boot-strapped” input voltage to connect to the gate of that additional switch.

As stated in Section 3.3.1, the primary concern for the DAC layout was robust, low resistance power routing across the array, utilizing the top two (thickest) metal layers, so it was ensured that the CAM bit cells and internal state machines retained data and did not suffer from supply droop. While this routing prioritization was a prudent decision, in future work if attempting a higher-resolution version of this design the possibility of true bottom-plate sampling should certainly be explored.

The $SAMPLE_{STOP}$ signal is generated by a clockbuffer cell on the negative edge of the sampling clock, F_S . The clockbuffer has two sections, a 4-stage inverter chain sized to not add more than $388fs$ of jitter to the incoming clock, according to equation 3.2:

$$\begin{aligned}
 V_{FS} 2\pi f_{in,max} t_{jitter} &= V_{jitter\ error} \\
 &or \\
 V_{FS} 2\pi f_{in,max} t_{jitter,RMS} &= \frac{V_{jitter\ error}}{\sqrt{2}} = V_{jitter\ error,RMS} \\
 &where \\
 &V_{FS} = Fullscale\ input\ voltage \\
 &f_{in,max} = Maximum\ input\ frequency \\
 &t_{jitter} = Maximum\ tolerable\ sampling\ jitter \\
 &V_{jitter\ error} = Maximum\ tolerable\ voltage\ error = V_{FS}/(2^{N+1})
 \end{aligned} \tag{3.2}$$

The second section of the clockbuffer uses a NOR gate and programmable delay cell to generate the $SAMPLE_{STOP}$ signal with a width between 40ps to 1.2ns in 40ps steps. Following the $SAMPLE_{STOP}$ signal generation both the F_S and $SAMPLE_{STOP}$ signals are buffered and sent to the digital control block and DAC array, respectively.

3.4 Measurements

3.4.1 Die Photo




Figure 3.7: Die photograph of ADC.

Figure 3.7 is a photograph of the ADC, fabricated in 65nm CMOS. Power routing and metal fill prevent smaller features from being visible in the die photo, so the digital control block, DAC, comparator, and input sampling circuits are outlined. The die is $1.04\text{mm} \times 1.04\text{mm}$, however the DAC area is $228\mu\text{m} \times 228\mu\text{m}$, and the entire area of all circuits except for I/O routing and buffers but including reference decoupling is 0.112 mm^2 .

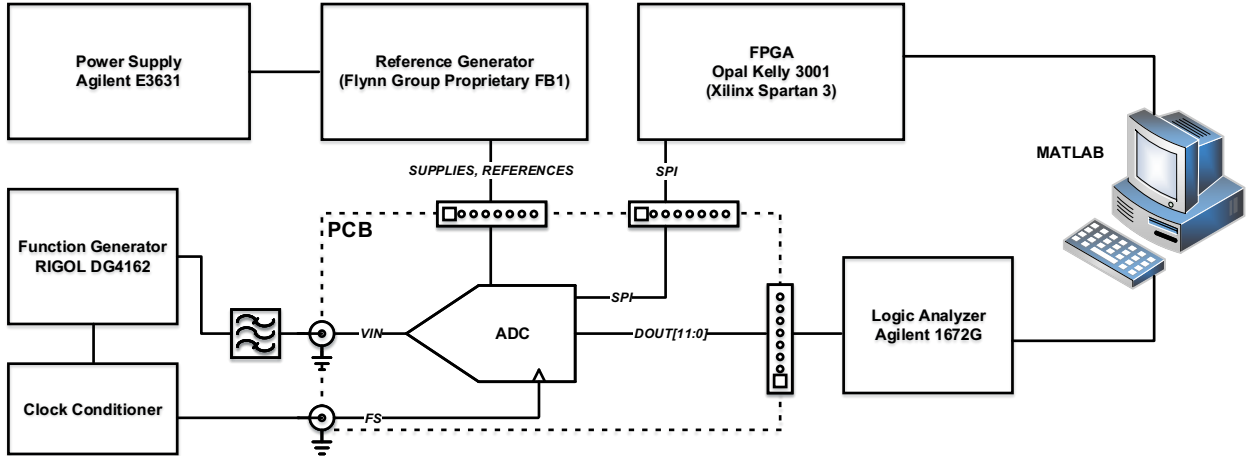


Figure 3.8: Block diagram of test setup.

3.4.2 Measurement Method

Test Setup

Figure 3.8 shows a block diagram of the testing setup. All voltage supplies and references are generated by a proprietary reference generation board². This reference generation board exhibits extremely low noise due to two stages of regulation and heavy filtering.

The RIGOL DG4162 function generator achieves 12.6 bits of linearity when tested with a 16-bit Analog Devices ADC and an inline bandpass filter (MiniCircuits) as shown in Figure 3.8. The RIGOL has an internal PLL that ensures the ADC sampling clock F_S and the ADC input frequency VIN are phase-locked and remain coherent. Because the RIGOL is generating sine waves, the sampling clock is passed through a clock conditioner to become a 1V square wave, compatible with the ADC's 1V CMOS logic level. The clock conditioner is an HP unit however no additional features such as duty cycle manipulation are used.

An Opal Kelly 3001 FPGA board (with a Xilinx Spartan 3 FPGA) is used to generate the SPI commands sent to the ADC that control DAC reconfiguration and basic ADC operation. The FPGA is controlled from a MATLAB testing environment on the lab PC. The MATLAB testing scripts generate data packets that are sent to the FPGA, where they are relayed to the ADC using SPI protocol at a 2MHz clock frequency.

²The reference generation board is designed by Fred Buhler, and shared amongst the lab.

Printed Circuit Board (PCB) Design

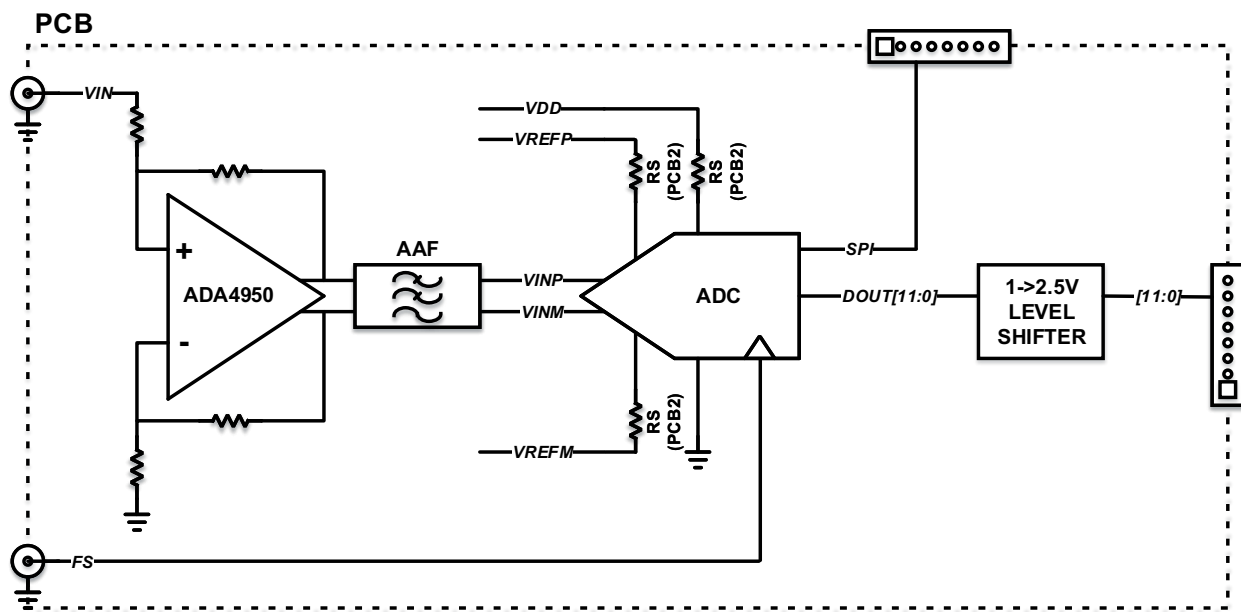


Figure 3.9: Simplified schematic of PCB used for ADC testing.

A simplified schematic of the PCB used for testing the ADC is shown in Figure 3.9. The Analog Devices ADA4950 is a high linearity ADC driver exhibiting 98dB SFDR at 20 MHz (16 bit linearity) and performs single-ended to differential conversion. This differential signal is then lowpass filtered by a RC network anti-aliasing filter (“AAF” as shown) before connecting to the ADC input pads.

The data used in this thesis was taken with two versions of the test PCB: PCB2 has series resistors in the supply and reference paths to the chip (“RS” as shown) and PCB1 does not have series resistors. A series resistance of 5Ω was used for RS to dampen supply ringing on PCB2.

Finally a 1-to-2.5V level shifter converts the ADC’s CMOS logic output to 2.5V levels for the Agilent 1672G logic analyzer.

A photo of the PCB is shown in Figure 3.10. The PCB is a 4-layer board and is shown connected to the voltage reference generation board. Differential input traces are routed symmetrically, while digital and analog grounds were only tied underneath the chip for short return paths (the chip has only one ground net, V_{SS}). Several 100nF NP0 capacitors are placed in parallel at every supply and reference pin, as close as possible to the ADC. Return




Figure 3.10: Photo of PCB used for ADC testing.

paths for all pins are as short as possible. Voltage supplies have hundreds of μF of decoupling at the board input and originate from the reference board as previously described.

Supply Ringing

As described in Section 3.2.2, there is a programmable delay cell that allows tuning of the DAC settling time. This is the time between a DAC switching event and the next comparator decision. This programmable delay ensures that DAC settling time is not a degrading factor in ADC performance.

On PCB1 (without series resistance in supply and reference traces), supply ringing can be seen in the output of the ADC during a sweep of the programmable delay cell values,




Figure 3.11: ENOB vs. DAC settling time settings on PCB *without* series resistance.




Figure 3.12: ENOB vs. DAC settling time settings on PCB *with* series resistance.

shown in Figure 3.11. Supply ringing occurs because of the LC tank created by the bondwire inductance and internal capacitance on the supply and reference nets. The peaks and valleys shown in Figure 3.11 are due to the mixing of the reference supply and the equivalent frequency of comparator decisions.

Figure 3.12 shows an identical sweep on the second PCB *with* series resistance on supply and reference nets. This behavior matches what one would expect from an ideal supply, which is increasing ENOB as the amount of DAC settling time increases.

Noise Averaging

As previously discussed in Section 3.2.1, noise performance beyond 10.5 bits was not expected out of the ADC due to comparator design constraints. However, DAC linearity of 11 bits or more was expected. To prevent the comparator noise (and other thermal or white noise sources) from veiling the DAC linearity, $32\times$ averaging of ADC output samples is performed, however the binary precision of this average remained at 12 bits - this prevents the quantization error from decreasing below the level of a 12-bit ADC. If the resulting average were not truncated to 12 bits, it would be “cheating” because the quantization noise would be artificially reduced beneath 12-bit precision.

Figure 3.13 confirms that the ADC measurements are linearity-limited and not noise-limited, so DAC linearity is measured up to the limit determined by the large-input-distortion level. Figure 3.13 shows the measured second and third order distortion levels (colored traces)




Figure 3.13: Noise floor, 2nd, and 3rd harmonics before and after noise averaging. The X-axis varies the total resolution of the ADC from an 11-bit ADC with a half-scale input to a 12-bit ADC at full-scale input.

and the average of the noise floor (black traces) across all non-harmonic frequency bins. The X-axis is the total resolution of the ADC from an 11-bit ADC with a half-scale input to a 12-bit ADC at full-scale input; from left to right, increasing amounts of MSB capacitance and input amplitude are being added. Distortion products are still apparent at equal levels both with and without noise averaging, while the noise floor drops approximately 8-10 dB. As input signal amplitude increases from $\frac{1}{2}$ full-scale (at 11 bits) to full-scale (at 12 bits), the 3rd order distortion steadily increases. This indicates there is a large-signal distortion present, most likely a detrimental consequence of the chosen sampling method discussed in Section 3.3.2.




Figure 3.14: Random scrambling algorithm in progress: ENOB vs. reconfiguration iteration.

3.4.3 Linearity Enhancement Results Via Reconfiguration

Random Reconfiguration in Progress

As discussed in Section 2.2.2, random reconfiguration predicted linearity improvement for every DAC in simulation. Figure 3.14 shows the measured results of 100 random reconfiguration - this is the real-time ENOB output of a random reconfiguration algorithm. For this ADC, an ENOB variation of 0.5 bits is observed over 100 separate configurations.

Harnessing a Flexible Transfer Function

Figure 3.15 shows the measured ENOB resulting from sweeping the input amplitude from 50% full-scale to 100% full-scale, and also the amount of MSB capacitors used as a fractional radix from 11 bits (no MSB capacitors) to 12 bits (all MSB capacitors in a default binary weighted grouping). Using a fractional amount of MSB capacitors allows the input dynamic range to be scaled between ADC resolutions without sacrificing conversion space. For example, if a 50% full-scale input is converted with an 11-bit ADC (no MSB capacitors used), then with a fractional MSB, a 75% full-scale input can be converted with an 11.5-bit ADC. This scaling is done to determine whether the input dynamic range is contributing

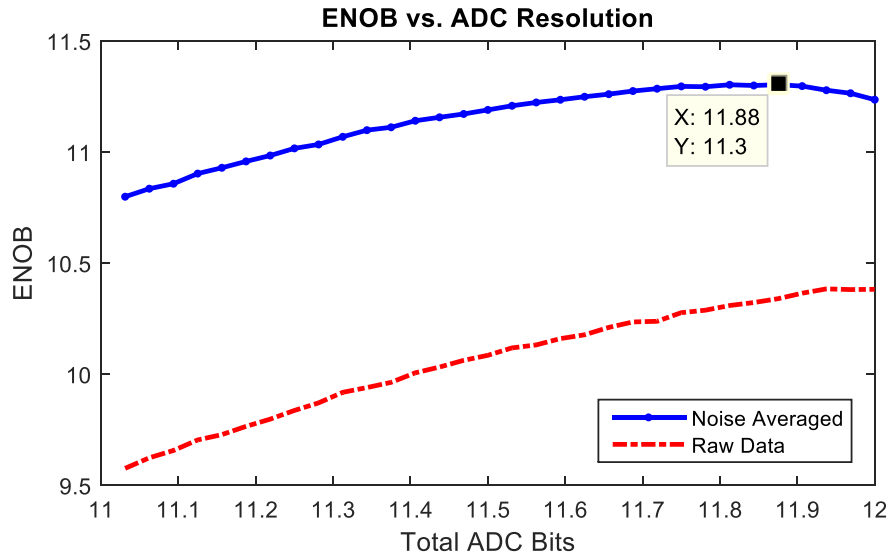


Figure 3.15: ENOB vs. MSB used as non-binary group).

to linearity degradation due to non-linear capacitances. At 50% full-scale input, the non-linear capacitances during sampling (Section 3.3.2) are experiencing a lesser degree of voltage variation than at the rail-to-rail 100% full-scale input. This figure confirms that the ADC is limited by non-linearities at large input voltages. As can be seen in this example, peak ENOB is achieved slightly below the full-scale input, and ENOB begins degrading after the peak value of 11.3. If the peak ENOB was not determined by non-linearities, the ENOB on this plot would keep rising towards the upper-right corner. This plot uses the same data set as the noise and harmonic analysis in Figure 3.13.

Peak Performance

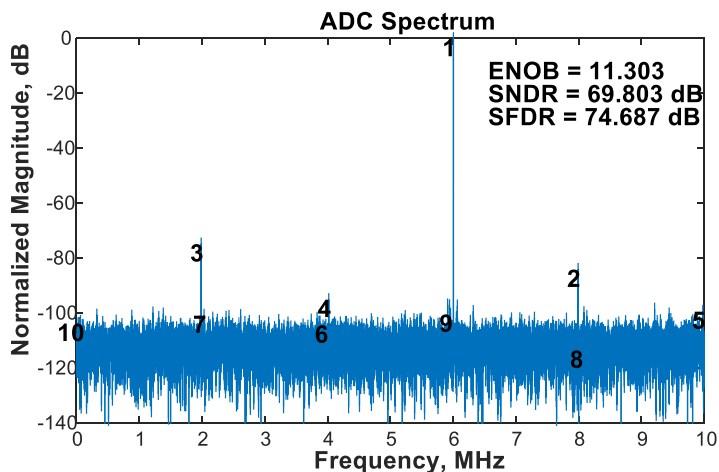


Figure 3.16: Noise-averaged ADC output spectrum (best measured): $f_{in} = 6\text{MHz}$, $f_s = 20\text{MHz}$.

The best spectrum measured is shown in Figure 3.16. ENOB is 11.3 bits, SNDR is 69.8 dB, and SFDR is 74.69 dB. The second and third harmonics are the dominant non-idealities in the spectrum, which is from the same dataset as figures 3.13 and 3.15. This spectrum is the result of 100 random scrambling iterations with an input frequency of 6MHz, sampling rate of 20MS/s, and a V_{REF} and V_{DD} of 1.1V.

Figure 3.17 shows minimum and maximum ENOB vs. a sweep of both input and sampling frequencies. For each one of the points on this plot, random scrambling was run 100 times, and the minimum and maximum ENOBs are plotted. A peak ENOB of over 11 bits is observed in all cases.

Maximum ENOB does not degrade either with input or sampling frequency. This confirms that the observed nonlinearities limiting the ADC from higher ENOB are not frequency-dependent, at least within the tested range of 1-10 MHz. Successful operation at sampling frequencies of 35 MHz nearly achieves the targeted design frequency of 40 MHz, and as discussed in Section 3.2.2 is limited by the propagation delay across the DAC. The ADC also achieves a peak ENOB of 10.5 bits at a sampling rate of 38 MHz, however ceases to function at higher sampling frequencies. This plot confirms successful ADC design overall with the exception of the choice of sampling technique (Section 3.3.2).

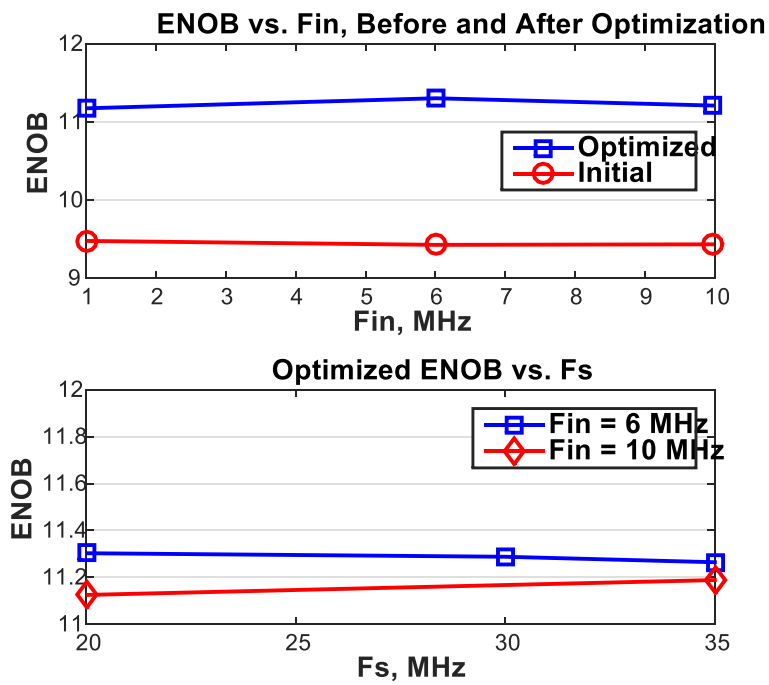


Figure 3.17: ENOB vs. Input Frequency and Sampling Frequency.




Figure 3.18: Measured ADC ENOB, optimized vs. default configuration, separated by test PCB number.

The following measurements separate ADC data by PCB to illustrate ADC performance across 2 chips. A “golden” chip was not used; as the ADCs were soldered down onto the boards for robustness over lengthy testing sessions, all data presented is simply from the first 2 ADCs that were tested.

Figure 3.18 shows the optimized vs. default ENOB for all measured ADCs in a similar format to the algorithm comparison in Figure 2.4. This data is taken across input and sampling frequencies from the same set as Figure 3.17. Data is also separated by PCBs used for testing. PCB1 (red points) is the board without series resistance in supply and reference traces, and PCB2 (blue points) is the board with series resistance. As every final ENOB is above the $y = x$ line, this confirms ENOB was improved for every chip, under every testing condition. This confirms the reconfiguration simulations (Section 2.2.2) and mismatch-immunity of this design.

Figure 3.19 displays the previous data from Figure 3.18 in histogram form, to better display the amount and consistency of ENOB improvement. Each count (y-axis) signifies a separate reconfiguration run with 100 iterations of randomization. Figure 3.19 shows a most consistent improvement amount of 1-1.3 bits. Data is again separated by PCB.




Figure 3.19: ENOB improvement histograms for PCB 1 and PCB 2.




Figure 3.20: ADC core power consumption (mW)

A detailed breakdown of the core (supply and reference)³power consumption is shown in Figure 3.20. These numbers are taken with the ADC operating at $F_S = 10\text{MS/s}$, $F_{IN} = 1\text{ MHz}$, with V_{DD} and V_{REF} both at 1.1V. These V_{DD} and V_{REF} values are the same for all of the results in this thesis. The total for the primary supply V_{DD} , reference V_{REF} , and clock buffer supply $V_{DD,CLK}$ is 1.456mW.

³Typically in ADC papers, only supply and reference power consumption are used for FoM calculation.

CHAPTER 4

Conclusion

4.1 Summary

- A novel, mismatch-immune, completely reconfigurable DAC has been realized.
- The goal of decoupling mismatch from linearity was successfully validated in a 65nm CMOS prototype.
- ADC performance is as intended and is consistent across both input and sampling frequency ranges.
- Every reconfiguration procedure resulted in improved linearity from the default fabricated configuration, regardless of test conditions.
- All measurements indicate performance was limited not by the novel DAC but by nonlinear input capacitance and/or insufficient reference conditioning.

This work is inserted below on the 4 major outputs of Murmann's ADC Performance Survey [4]; energy consumption vs. resolution in figure 4.1, bandwidth vs. resolution in figure 4.2, Walden FoM vs. speed in figure 4.3, and Schreier FoM vs. speed in figure 4.4. These plots use the raw, not noise averaged, ENOB from the ADC, and this accounts for a large penalty in performance vs. other ADCs. Had this work utilized a full bottom-plate sampling method allowing for improved comparator performance, these figures of merit would improve by nearly a factor of two.




Figure 4.1: Energy vs. Resolution comparison from [4].




Figure 4.2: Bandwidth vs. Resolution comparison from [4].




Figure 4.3: Walden FoM vs. speed comparison from [4].

4.2 Future Work

As discussed in section 3.3.2, the addition of true bottom-plate sampling would remove any input-dependent variation in sampling capacitance, and more importantly would decouple the comparator design from any linearity concerns. The current cell layout leaves space for the additional routing required to implement true bottom-plate sampling in a future revision. Furthermore, true bottom-plate sampling would enable higher resolution than 12-bits, at which point reference quality would become the limiting factor for the design.

The comparator design was only revised once during this research, and after increased experience with similar comparator architectures a comparator with 11.5 - 12 bit noise performance could be achieved for a minimal power increase, if any. In addition one or several of the noise mitigation innovations discussed in section 1.4 would help achieve noise performance equal to the targeted resolution.

V_{DD} accounts for nearly 88% of core power consumption, which is dominated by the column buffers driving the CAM bitlines that traverse the entire DAC array (as discussed in




Figure 4.4: Schreier FoM vs. speed comparison from [4].

2.3.1). Designing a memory cell that does not require complimentary bitlines would reduce core power consumption by a factor of nearly 1.5.

With a combination of the above improvements, a 14-15 bit version of this design would be feasible and competitive among recent publications.

BIBLIOGRAPHY

- [1] Robert L. Howald. QAM Bulks Up Once Again: Modulation to the Power of Ten. *Motorola Broadband Communications [whitepaper]*, <http://educyclopedia.karadimov.info/library/QAM.pdf>.
- [2] Analog Devices Inc. AD9877 [Datasheet] - Mixed-Signal Front End, Set-Top Box, Cable Modem. <http://www.analog.com/media/en/technical-documentation/datasheets/AD9877.pdf>.
- [3] Analog Devices Inc. AD7152/7153 [Datasheet] - 12-bit Capacitance-to-Digital Converter. http://www.analog.com/media/en/technical-documentation/datasheets/AD7152_7153.pdf.
- [4] Boris Murmann. ADC Performance Survey 1997-2016. <http://web.stanford.edu/~murmman/adcsurvey.html>, 2017.
- [5] Analog Devices Inc. *The Data Conversion Handbook*. Newnes, Burlington, MA, USA, 2005.
- [6] Claude Shannon. Communication in the Presence of Noise. In *Proc. Institute of Radio Engineers*, volume 37, pages 10–21, 1949.
- [7] Conor Donovan and Michael P. Flynn. A Digital 6-bit ADC in 0.25 μ m CMOS. *IEEE Journal of Solid-State Circuits*, 37(3):432–437, 2002.
- [8] Franco Maloberti. *Data Converters*. Springer, Dordrecht, The Netherlands, 2007.
- [9] Richard Schreier and Gabor C. Temes. *Understanding Delta-Sigma Data Converters*. John Wiley and Sons, Inc., Hoboken, NJ, 2005.
- [10] Matthias Steiner and Nigel Greer. A 22.3b 1kHz 12.7mW Switched-Capacitor $\Delta\Sigma$ Modulator with Stacked Split-Steering Amplifiers. In *Proc. International Solid-State Circuits Conference*, pages 284–285, 2016.
- [11] Gordon E. Moore. Cramming More Components onto Integrated Circuits. *Electronics*, 38(8), 1965.
- [12] Don Clark. Intel Rechisels the Tablet on Moore’s Law. *Wall Street Journal - Digits Tech Blog [Online]*, <http://blogs.wsj.com/digits/2015/07/16/intel-rechisels-the-tablet-on-moores-law/>, 2015.
- [13] Rick Merritt. TSMC, GF-Samsung Battle at 7nm. *EETimes [Online]*, http://www.eetimes.com/document.asp?doc_id=1330657, 2016.
- [14] Neil Weste and David Harris. *CMOS VLSI Design: A Circuits and Systems Perspective*. Addison-Wesley Publishing Company, USA, 4th edition, 2010.
- [15] David A. Fick. *Large Scale System Design in Advanced VLSI Process Technologies [Dissertation]*. University of Michigan, Ann Arbor, MI, 2011.

- [16] Ming-Horn Tsai and Tso-Ping Ma. The Impact of Device Scaling on the Current Fluctuations in MOSFET's, journal=IEEE Transactions on Electron Devices. 41(11):2061–2068, 1994.
- [17] Jorge Pernillo and Michael P. Flynn. A 1.5-GS/s Flash ADC With 57.7-dB SFDR and 6.4-Bit ENOB in 90 nm Digital CMOS. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 58(12):837–841, 2011.
- [18] Skyler Weaver, Benjamin Hershberg, Daniel Knierim, and Un-Ku Moon. A 6b Stochastic Flash Analog-to-Digital Converter Without Calibration or Reference Ladder. In *Proc. IEEE Asian Solid-State Circuits Conference*, pages 373–376, 2008.
- [19] Yong Lim and Michael P. Flynn. A 100MS/s 10.5b 2.46mW Comparator-less Pipeline ADC Using Self-Biased Ring Amplifiers. In *Proc. International Solid-State Circuits Conference*, pages 202–203, 2014.
- [20] Benjamin Hershberg, Skyler Weaver, Kazuki Sobue, Seiji Takeuchi, Koichi Hamashita, and Un-Ku Moon. A 61.5dB SNDR Pipelined ADC Using Simple Highly-Scalable Ring Amplifiers. In *IEEE Symposium on VLSI Circuits*, pages 32–33, 2012.
- [21] Dong-Young Chang, Carlos Munoz, Denis Daly, Soon-Kyun Shin, Kevin Guay, Thomas Thurston, Hae-Seung Lee, Kush Gulati, and Matthew Straayer. A 21mW 15b 48MS/s Zero-Crossing Pipeline ADC in 0.13 μ m CMOS with 74dB SNDR. In *Proc. International Solid-State Circuits Conference*, pages 204–205, 2014.
- [22] Hyungil Chae, Jaehun Jeong, Gabriele Manganaro, and Michael Flynn. A 12mW Low-Power Continuous-Time Bandpass $\Delta\Sigma$ Modulator with 58dB SNDR and 24MHz Bandwidth at 200MHz IF. In *Proc. International Solid-State Circuits Conference*, pages 148–149, 2012.
- [23] Hung-Yen Tai, Hung-Wei Chen, and Hsin-Shu Chen. A 3.2fJ/c.-s. 0.35V 10b 100KS/s SAR ADC in 90nm CMOS. In *IEEE Symposium on VLSI Circuits*, pages 92–93, 2012.
- [24] Chang-Yuan Liou and Chih-Cheng Hsieh. A 2.4-to-5.2fJ/conversion-step 10b 0.5-to-4MS/s SAR ADC with Charge-Average Switching DAC in 90nm CMOS. In *Proc. International Solid-State Circuits Conference*, pages 280–281, 2013.
- [25] Brian P. Ginsburg and Anantha P. Chandrakasan. An Energy-Efficient Charge Recycling Approach for a SAR Converter With Capacitive DAC. In *Proc. ISCAS*, pages 184–187, 2005.
- [26] Chun-Cheng Liu, Soon-Jyh Chang, Guan-Ying Huang, and Yin-Zu Lin. A 0.92mW 10-bit 50-MS/s SAR ADC in 0.13 μ m CMOS Process. In *IEEE Symposium on VLSI Circuits*, pages 236 – 237, 2009.
- [27] Chun-Cheng Liu, Soon-Jyh Chang, Guan-Ying Huang, Ying-Zu Lin, and Chung-Ming Huang. A 1V 11fJ/Conversion-Step 10bit 10MS/s Asynchronous SAR ADC in 0.18 μ m CMOS . In *IEEE Symposium on VLSI Circuits*, pages 243 – 244, 2010.

- [28] Yan Zhu, Chi-Hang Chan, U-Fat Chio, Sai-Weng Sin, Seng-Pan U, Rui Paulo Martins, and Franco Maloberti. A 10-bit 100-MS/s Reference-Free SAR ADC in 90 nm CMOS. *IEEE Journal of Solid-State Circuits*, 45(6):1111–1121, 2010.
- [29] Hung-Yen Tai, Yao-Sheng Hu, Hung-Wei Chen, and Hsin-Shu Chen. A 0.85fJ/conversion-step 10b 200kS/s Subranging SAR ADC in 40nm CMOS. In *Proc. International Solid-State Circuits Conference*, pages 196–197, 2014.
- [30] Pieter Harpe, Yan Zhang, Guido Dolmans, Kathleen Philips, and Harmke De Groot. A 7-to-10b 0-to-4MS/s Flexible SAR ADC with 6.5-to-16fJ/conversion-step. In *Proc. International Solid-State Circuits Conference*, pages 472–473, 2012.
- [31] Shuo-Wei Mike Chen and Robert W. Brodersen. A 6b 600MS/s 5.3mW Asynchronous ADC in 0.13 μ m CMOS. In *Proc. International Solid-State Circuits Conference*, pages 2350–2351, 2006.
- [32] Chun-Cheng Liu, Soon-Jyh Chang, Guan-Ying Huang, Ying-Zu Lin, Chung-Ming Huang, Chih-Hao Huang, and Chih-Chung Tsai. A 10b 100MS/s 1.13mW SAR ADC with Binary-Scaled Error Compensation. In *Proc. International Solid-State Circuits Conference*, pages 250–251, 2010.
- [33] Franz Kuttner. A 1.2V 10b 20MSample/s Non-Binary Successive Approximation ADC in 0.13 μ m CMOS. In *Proc. International Solid-State Circuits Conference*, volume 1, pages 176–177, 2002.
- [34] Joshua J. Kang and Michael P. Flynn. A 12b 11MS/s Successive Approximation ADC with two comparators in 0.13 μ m CMOS. In *IEEE Symposium on VLSI Circuits*, pages 240–241, 2009.
- [35] Andrea Agnes, Edoardo Bonizzoni, Piero Malcovati, and Franco Maloberti. A 9.4-ENOB 1V 3.8 μ W 100kS/s SAR ADC with Time-Domain Comparator, 2008.
- [36] Jon Guerber, Hariprasath Venkatram, Manideep Gande, Allen Waters, and Un-Ku Moon. A 10-b Ternary SAR ADC With Quantization Time Information Utilization. *IEEE Journal of Solid-State Circuits*, 47(11):2604–2613, 2012.
- [37] Vito Giannini, Pierluigi Nuzzo, Vincenzo Chironi, Andrea Baschirotto, Geert Van der Plas, and Jan Cranickx. An 820 μ W 9b 40MS/s Noise-Tolerant Dynamic SAR ADC in 90nm Digital CMOS. In *Proc. International Solid-State Circuits Conference*, pages 238–239, 2008.
- [38] Ming Ding, Pieter Harpe, Yao-Hong Liu, Benjamin Busze, Kathleen Philips, and Harmke de Groot. A 5.5fJ/conv-step 6.4MS/s 13b SAR ADC Utilizing a Redundancy-Facilitated Background Error-Detection-and-Correction Scheme. In *Proc. International Solid-State Circuits Conference*, pages 460–461, 2015.
- [39] Pieter Harpe, Eugenio Cantatore, and Arthur van Roermund. An Oversampled 12/14b SAR ADC with Noise Reduction and Linearity Enhancements Achieving up to 79.1dB SNDR. In *Proc. International Solid-State Circuits Conference*, pages 194–195, 2014.

- [40] R. H. Walden. Analog-to-digital converter survey and analysis. *IEEE J. Select. Areas Commun.*, 17:539–550, 1999.
- [41] P. Nuzzo, F. De Bernardinis, P. Terreni, and G. Van der Plas. Noise Analysis of Regenerative Comparators for Reconfigurable ADC Architectures. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 55(6):1441–1454, 2008.
- [42] Stacy Ho and Michael P. Flynn. EECS 511 Lecture Notes. Ch. 4:28, 2009.
- [43] Chun-Cheng Liu. A 10-bit 320-MS/s Low-Cost SAR ADC for IEEE 802.11ac Applications in 20-nm CMOS. In *Proc. IEEE Asian Solid-State Circuits Conference*, pages 77–80, 2014.
- [44] Jon Guerber, Manideep Gande, and Un-Ku Moon. The Analysis and Application of Redundant Multistage ADC Resolution Improvements Through PDF Residue Shaping. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 59(8):1733–1742, 2012.
- [45] In-Seok Jung and Yong-Bin Kim. A 12-bit 32MS/s SAR ADC Using Built-In Self Calibration Technique To Minimize Capacitor Mismatch. In *International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pages 275–279, 2014.
- [46] Ricardo E. Suarez, Paul R. Gray, and David A. Hodges. An All-MOS Charge-Redistribution A/D Conversion Technique. In *Proc. International Solid-State Circuits Conference*, pages 194–195, 1974.
- [47] James L. McCreary and Paul R. Gray. All-MOS Charge Redistribution Analog-to-Digital Conversion Techniques - Part I. *IEEE Journal of Solid-State Circuits*, SC-10(6):371–379, 1975.
- [48] Ricardo E. Suarez, Paul R. Gray, and David A. Hodges. All-MOS Charge Redistribution Analog-to-Digital Conversion Techniques - Part II. *IEEE Journal of Solid-State Circuits*, SC-10(6):379–385, 1975.
- [49] James L. McCreary. Matching Properties, and Voltage and Temperature Dependence of MOS Capacitors. *IEEE Journal of Solid-State Circuits*, SC-16(6):608–616, 1981.
- [50] Rajinder Singh and A. B. Bhattacharyya. Matching Properties of Linear MOS Capacitors. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 36(3):465–467, 1989.
- [51] Marcel J. M. Pelgrom, Aad C. J. Duinmaijer, and Anton P. G. Welbers. Matching Properties of MOS Transistors. *IEEE Journal of Solid-State Circuits*, 24(5):1433–1440, 1989.
- [52] Vaibhav Tripathi and Boris Murmann. Mismatch Characterization of Small Metal Fringe Capacitors. *IEEE Transactions on Circuits and Systems I: Regular Papers*, pages 2236–2242, 2014.

- [53] Hesham Omran, Hamzah Alahmadi, and Khaled N. Salama. Matching Properties of Femtofarad and Sub-Femtofarad MOM Capacitors. *IEEE Transactions on Circuits and Systems I: Regular Papers*, pages 763–772, 2016.
- [54] Kent H. Lundberg. Analog-to-Digital Converter Testing. <http://www.mit.edu/~klund/A2Dtesting.pdf>, 2017.
- [55] Walt Kester. MT-001: Taking the Mystery out of the Infamous Formula “ $\text{SNR} = 6.02N + 1.76\text{dB}$,” and Why You Should Care. <http://www.analog.com/media/en/training-seminars/tutorials/MT-001.pdf>, 2017.
- [56] Jeffrey A. Fredenburg and Michael P. Flynn. Statistical Analysis of ENOB and Yield in Binary Weighted ADCs and DACs With Random Element Mismatch. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 59(7):1396–1408, 2012.
- [57] Masaya Miyahara and Akira Matsuzawa. A Low-Offset Latched Comparator Using Zero-Static Power Dynamic Offset Cancellation Technique. In *Proc. IEEE Asian Solid-State Circuits Conference*, pages 233–236, Nov. 2009.