A Low-Power Pipeline ADC with Front-End Capacitor-Sharing

by

Guangzhao Zhang

A thesis submitted in conformity with the requirements
for the degree of Masters of Applied Science
Graduate Department of Electrical and Computer Engineering
University of Toronto

A Low-Power Pipeline ADC with Front-End Capacitor-Sharing

Guangzhao Zhang

Master of Applied Science, 2012

Graduate Department of Electrical and Computer Engineering

University of Toronto

Abstract

This thesis presents the design and experimental results of a low-power pipeline ADC that applies front-end capacitor-sharing. The ADC operates at 20 MS/s, resolves 1.5 bits/stage, and is implemented in IBM 0.13μm technology. The purpose of the technique is to reduce power consumption in the front-end S/H. This work is a proof-of-concept and it concentrates on the front-end design. A comparison is conducted between a capacitor-sharing ADC and a regular ADC and as a result, the technique reduces the power consumption in the front-end S/H by 39%. At an input frequency of 9.53 MHz and a sampling rate of 20 MS/s, the fabricated capacitor-sharing ADC consumes 4.7 mW at 1.2 V, and it achieves an ENOB of 8.5 bits and a FOM of 0.68 pJ/step. It has an ENOB as high as 8.67 bits at 0.4 MS/s and a FOM as low as 0.6 pJ/step when sub-sampling at 20 MS/s.

Acknowledgements

There were many people who played a crucial role in contributing to the success of this work.

Firstly, this thesis would not have been possible without the continual guidance and support of my supervisor, Professor David Johns.

I am also very grateful to Karim Abdelhalim, Mike Bichan, Yunzhi (Rocky) Dong, Safeen Huda, Sadegh Jalali, Hamed Mazhab-Jafari, Bert Leesti, Mario Milicevic, Alireza Nilchi, Alain Rousson, Amer Samarah, Ioannis Sarkas, Shayan Shahramian, Ravi Shivnaraine, Colin Tse, Kentaro Yamamoto, Hemesh Yasotharan, and Meysam Zargham for their advice and technical support that helped me get past those seemingly insurmountable obstacles. A special thanks to Rocky, the BA5000 circuits guru, for the countless invaluable discussions we had.

To all the students of BA5000 and BA5158, I thank you for making my graduate experience a fun and enjoyable adventure that I will forever remember.

I would also like to thank Professor Glenn Gulak, Professor Wai Tung Ng, and Professor Sean Victor Hum for being on my defense committee.

Finally, I owe my deepest gratitude to Mandy and my family for their understanding and unquestioning support when I needed it the most.

Contents

1 Introduction **1**

1.1 Thesis Objectives 1

1.2 Thesis Outline 2

2 Background **3**

2.1 Pipeline ADC 3

2.1.1 Figure of Merit 3

2.1.2 One-Stage pipeline 4

2.1.3 1-bit/Stage Pipeline 6

2.1.4 Digital Error Correction 7

2.1.5 1.5-bit/Stage Pipeline 11

2.1.6 Sample and Hold 12

2.1.7 Thermal Noise and Scaling of Pipeline Stages 13

2.2 Survey of Recently Published 10-bit Pipeline ADCs 14

2.3 Design Application and Target Specifications 14

3 Pipeline ADC Building Blocks and Design Methodology **17**

3.1 Opamp Design 17

3.1.1 Gain 17

3.1.2 Bandwidth 18

3.1.3 Output Voltage Swing 19

3.2 Building Blocks 20
3.2.1 Sub-ADC 21
3.2.2 Sample and Hold 22
3.2.3 Multiplying DAC 24
3.2.3.1 Flip-Around MDAC 24
3.2.3.2 Integrator MDAC 25
3.3 Design Procedure 26

4 Capacitor-Sharing Pipeline Design and Simulation 30
4.1 Front-End Capacitor-Sharing 30
4.2 Power Comparison 34
4.2.1 Regular versus Capshare ADC 35
4.2.2 Analysis 36
4.2.2.1 Initial Assumptions 36
4.2.2.2 Choosing the Sampling Capacitance 36
4.2.2.3 Two-Stage Opamp Design 37
4.2.2.4 Conclusions 39
4.3 Circuit Implementation 40
4.3.1 General Description 40
4.3.2 Two-Stage Opamp 41
4.4 Schematic Simulations 44

5 Experimental Results 48
5.1 Test Setup 48
5.2 Experimental Data 49
5.2.1 Measurement Results 49
5.2.2 Differential and Integral Non-Linearity 54
5.2.3 Conclusions 56

5.3 Dynamic Range Degradation 56

6 Conclusions and Future Work **60**

6.1 Conclusions 60

6.2 Future Work 61

References **62**

List of Tables

2.1 Previous 10-bit pipeline ADCs 15

2.2 Design specifications targeted for this work 16

4.1 Power of capshare pipeline ADC 38

4.2 Power of regular pipeline ADC 39

4.3 Transistor sizes in W/L $[\mu m/\mu m]$ within capshare ADC opamps 42

4.4 Transistor sizes in W/L $[\mu m/\mu m]$ within regular ADC opamps 43

4.5 $A\beta$, phase margin, and f_{3dB} of opamp closed-loop circuits in capshare ADC 43

4.6 $A\beta$, phase margin, and f_{3dB} of opamp closed-loop circuits in regular ADC 43

4.7 SNDR and SNR at $fin = 61/128 \cdot f_S$ 44

4.8 Power consumption of capshare versus regular ADC at $f_S = 2\ MHz$ . . . 46

4.9 Power consumption of capshare versus regular ADC at $f_S = 20\ MHz$ . . 47

5.1 List of input/output pins for the DUT 51

5.2 Experimental results of capshare ADC measured at different f_{in} and f_S . 51

5.3 Comparing this work to other 1.5-bit/stage 10-bit pipeline ADCs 56

5.4 SNDR and SNR at $f_{in} = 61/128 \cdot f_S$ after post-layout extraction 57

List of Figures

2.1 One-stage pipeline ADC 5

2.2 Ten-stage pipeline ADC that resolves 1-bit/stage 7

2.3 Transfer function of a 1-bit MDAC 8

2.4 Transfer function of a 1-bit MDAC with comparator offset 9

2.5 Transfer function of a 1.5-bit MDAC 10

2.6 Transfer function of a 1.5-bit MDAC with comparator offset 11

2.7 A 10-bit 1.5-bit/stage pipeline ADC with digital error correction 12

2.8 Digital error correction in a 1.5-bit/stage 10-bit pipeline ADC 12

2.9 General block diagram for a wireless radio receiver 15

3.1 Single-stage opamp with gain-boosting 19

3.2 Two-stage opamp with a folded-cascode and a common source stage . . . 20

3.3 Three-level sub-ADC 21

3.4 Charge distribution comparator 21

3.5 A commonly used S/H 22

3.6 Clock waveform with phases Φ_1 and Φ_2 labeled 23

3.7 Flip-around MDAC 25

3.8 Integrator MDAC 26

4.1 Front-end S/H and first stage flip-around MDAC 31

4.2 Front-end S/H and first stage integrator MDAC 32

4.3 Front-end S/H sharing Cs with first stage MDAC 32
4.4 Capacitor-sharing front-end with ping-pong scheme 33
4.5 Clock waveform with clock cycle phases and ping-pong phases labeled . . 34
4.6 Regular pipeline ADC with a flip-around MDAC in the first stage 35
4.7 Pipeline ADC with capacitor-sharing front-end 35
4.8 Two-stage opamp with with CMFB on second stage only 42
4.9 Spectrum of Simulations 4 and 8 in Table 4.7 46

5.1 Test setup for capshare pipeline ADC 49
5.2 PCB used to test capshare pipeline ADC 50
5.3 Die micrograph of pipeline ADC with front-end capacitor-sharing 50
5.4 Comparison of simulated and measured SNDR and SNR 52
5.5 Spectrum and SNDR vs Input Level at $f_{in} = 9.53\ MHz$ and $f_S = 20\ MHz$ 53
5.6 Spectrum at $f_{in} = 2.03\ MHz$ and $f_S = 20\ MHz$ 54
5.7 Spectrums at $f_{in} = 9.53\ MHz$ and $f_S = 20\ MHz$ 55
5.8 DNL and INL using histogram test at $fin = 2.03\ MHz$ and $f_S = 20\ MHz$ 55
5.9 MDAC gain error caused by parasitics across sampling capacitor 57
5.10 Spectrum of Simulations 5 and 6 in Table 5.4 58

List of Acronyms

ADC Analog to Digital Converter

BPF Band-Pass Filter

capshare capacitor-sharing

CM Common-Mode

CMFB Common-Mode Feedback

DAC Digital to Analog Converter

DEC Digital Error Correction

DNL Differential Non-Linearity

DSP Digital Signal Processing

DUT Device Under Test

ENOB Effective Number of Bits

FFT Fast Fourier Transform

FOM Figure of Merit

GBW Gain Bandwidth

INL Integral Non-Linearity

LNA Low-Noise Amplifier

LPF Low-Pass Filter

LSB Least Significant Bit

MDAC Multiplying Digital to Analog Converter

MIM Metal Insulator Metal

MSB Most Significant Bit

MS/s Mega Samples per Second

NMOS N-type Metal Oxide Semiconductor

PCB Printed Circuit Board

PMOS P-type Metal Oxide Semiconductor

PSD Power Spectral Density

RF Radio Frequency

S/H Sample and Hold

SR Slew Rate

SNDR Signal to Noise and Distortion Ratio

SNR Signal to Noise Ratio

T/H Track and Hold

VGA Variable Gain Amplifier

WLAN Wireless Local Area Network

Chapter 1

Introduction

1.1 Thesis Objectives

The front-end Sample and Hold (S/H) in a pipeline Analog to Digital Converter (ADC) typically makes up a large portion of total power consumption. This has motivated research into reducing the power consumption of this power-hungry block. For instance, [1] embeds the S/H within the first stage of the pipeline ADC. This thesis presents a novel front-end capacitor-sharing (capshare) technique that significantly reduces the power consumption in the front-end S/H. The technique is demonstrated in a 10-bit pipeline ADC that resolves 1.5 bits/stage. The goal of this thesis is a proof-of-concept of the technique and hence, it concentrates on the front-end design and not on attaining the best raw performance. The objectives of this thesis are as follows:

- Provide a background on pipeline ADCs.
- Introduce a novel front-end capshare technique that saves power in the front-end S/H.
- Show the theoretically power savings of the technique through a design comparison.

- Demonstrate via simulations and experimental results that the technique achieves the expected performance.

1.2 Thesis Outline

The next chapters in this thesis are organized as follows:

- Chapter 2 provides a background on pipeline ADCs.
- Chapter 3 presents the building blocks and the design methodology for pipeline ADCs. Then, an example design is done using the principles discussed in the chapter.
- Chapter 4 conducts the design of a pipeline ADC with front-end capshare and a regular pipeline ADC. Their theoretical and simulated performance is compared.
- Chapter 5 shows the experimental results of the capshare ADC fabricated in IBM 0.13 μm technology and analyzes a dynamic range issue through post-layout simulations.
- Chapter 6 summarizes the main conclusions and discusses the potential areas for future work.

Chapter 2

Background

This chapter introduces the pipeline ADC. Section 2.1 presents background material on pipeline ADCs, Section 2.2 provides a brief survey of previous work, and Section 2.3 describes the application and design specifications of this work.

2.1 Pipeline ADC

This section presents background material on pipeline ADCs.

2.1.1 Figure of Merit

An ADC quantizes an analog input signal into a digital output at a specific conversion resolution and accuracy. The resolution is equal to the number of bits, N, that are resolved, while the accuracy refers to how precise the output bits represent the input. The accuracy is typically measured in terms of Signal to Noise and Distortion Ratio (SNDR), which is the ratio of signal power to noise and distortion power, or alternatively in Effective Number of Bits (ENOB):

$$ENOB = \frac{SNDR - 1.76}{6.02} \quad [bits] \tag{2.1}$$

The noise power in the system comes from quantization noise, the conversion limit set by the resolution, and random noise (e.g. thermal noise). Inaccuracies in the ADC's physical components will appear in the output as distortion power. If an ADC is designed with perfectly accurate components, the distortion power is zero and the conversion accuracy is limited by the Signal to Noise Ratio (SNR) or the dynamic range. Dynamic range is the ratio of the maximum signal power to noise power. If all sources of random noise is below the N-bit level, the dynamic range is limited by the quantization limit set by the resolution. In general, ADCs are quantitatively compared using a Figure of Merit (FOM) defined by:

$$FOM = \frac{P_{total}}{2^{ENOB} \cdot 2f_{in}} \quad [pJ/step] \tag{2.2}$$

P_{total} is the total power consumed by the system and f_{in} is the frequency of the input signal. There are many types of ADCs, each suitable for different accuracies and conversion rates as discussed in [2]; however, this thesis will concentrate on the pipeline ADC.

2.1.2 One-Stage pipeline

A pipeline ADC is a type of switched-capacitor circuit that divides the quantization of an input signal into multiple steps. It does this by distributing the conversion over multiple stages so that each stage converts only a subset of the total number of bits, N. Pipeline ADCs are more power efficient than ADCs that quantize in only one step, like a flash ADC. Typically, the pipeline stages function in a two-step or two-phase cycle. To demonstrate this, Figure 2.1 shows a basic one-stage N-bit pipeline ADC. The flash-ADC inside the pipeline stage is called a sub-ADC as not to get mixed-up with the final flash-ADC. In the first half cycle (T/2), the ADC is in phase 1, Φ_1, and Track and Hold (T/H) T/H_1 and T/H_2 track the analog input. At the end of the half cycle, T/H_1 and T/H_2 sample the input. In the second half cycle, the ADC is in phase 2, Φ_2, and T/H_1 and T/H_2 output the voltages they sampled for the sub-ADC and summer block. The sub-ADC takes the input and quantizes it create the upper N/2 bits, which is the

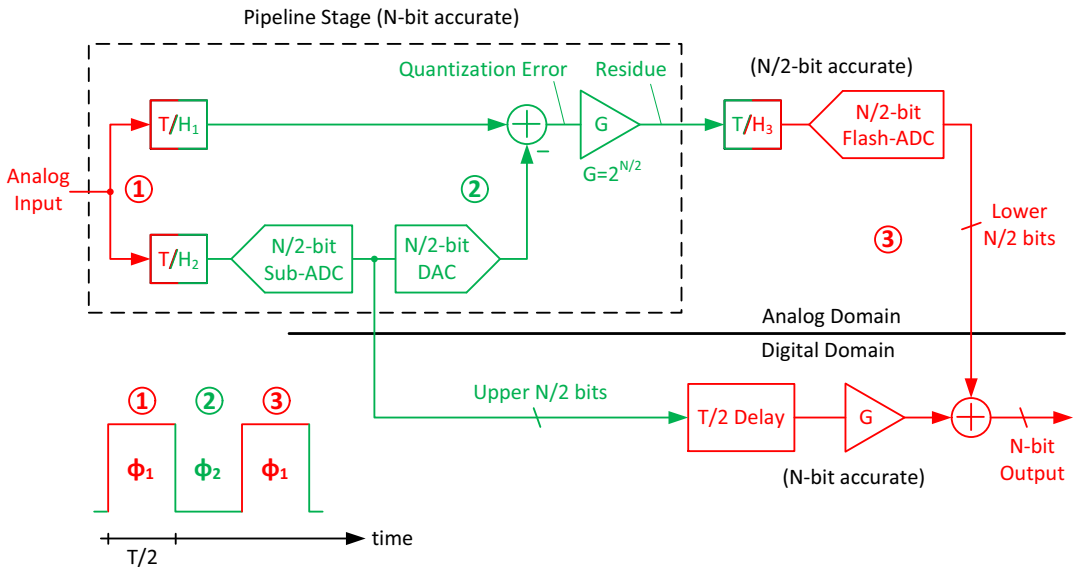

Figure 2.1: One-stage pipeline ADC

digital output of the pipeline stage. The upper N/2 bits are referred to as the Most Significant Bit (MSB)s. The MSBs are immediately converted back into an analog signal via the Digital to Analog Converter (DAC), which gets subtracted from the original analog input in the summer block to produce the quantization error for the conversion. The quantization error is then amplified by the stage gain, $G = 2^{N/2}$, to bring the voltage swing back to the input range. The output of the stage gain is called the residue output and is the analog output of the pipeline stage. At the end of the second half cycle, T/H_3 samples the residue output. In the third and final half cycle, the ADC is back in Φ_1. T/H_3 outputs the residue signal and the final flash-ADC quantizes it to create the lower N/2 bits. The lower N/2 bits are referred to as the Least Significant Bit (LSB)s. The MSBs, which are held in the digital domain so that they are available during this half cycle, are digitally scaled by the gain, 'G', to bring them to the correct magnitude, and combined with the LSBs to create a N-bit digital output. The input signal sampled in the first half cycle has now been converted. While this is happening at the end of the pipeline ADC, T/H_1 and T/H_2 track the next analog input and sample it at the end of the third half cycle. The process then repeats itself.

Ideally, the N-bit digital output should be be accurate to N bits (i.e. the ADC has an ENOB of 10-bits) as that would ensure the ADC is perfectly linear and not missing any conversion codes [3]. The analog and digital components in Figure 2.1 must be accurate to at least a certain number of bits to ensure N-bit accuracy. In the analog domain, all the components in the pipeline stage must be N-bit accurate to ensure the MSBs are N-bit accurate and the residue output settles to a value accurate to N bits. T/H_3 and the final flash-ADC must be N/2-bit accurate to ensure the LSBs are accurate to N/2 bits. In the digital domain, the T/2 delay, gain 'G', and summer should be N-bit accurate as they process and generate data accurate to N bits. The main design challenge is meeting the accuracy requirements in the analog domain. If the accuracy requirements are not met, there will be distortion in the digital output that limits the ADC's ENOB to below N bits.

In general, a pipeline stage outputs a sub-set of the total number of output bits and a residue signal, which are its digital and analog outputs respectively.

2.1.3 1-bit/Stage Pipeline

The procedure in Section 2.1.2 can be extended to multiple pipeline stages. Figure 2.2 shows a 10-bit pipeline ADC that resolves 1-bit per stage. Within the pipeline stage, T/H_1 and T/H_2 are integrated into the summer and sub-ADC respectively. The process in the analog domain is very similar to that of the one-stage pipeline. Starting from the first stage, the input signal is sampled by the sub-ADC and the summer. Next, the sub-ADC generates a 1-bit output. It immediately gets converted back into an analog voltage so it can be subtracted from the input sample to generate the quantization error. The quantization error is then amplified by the stage gain ($G_i = 2$) to bring it back to the input range, which produces the residue output for the next stage. The next stage and every stage afterwards repeat the cycle until all 10-bits are generated. The accuracy requirement in the first pipeline stage is the same as in the one-stage pipeline; however,

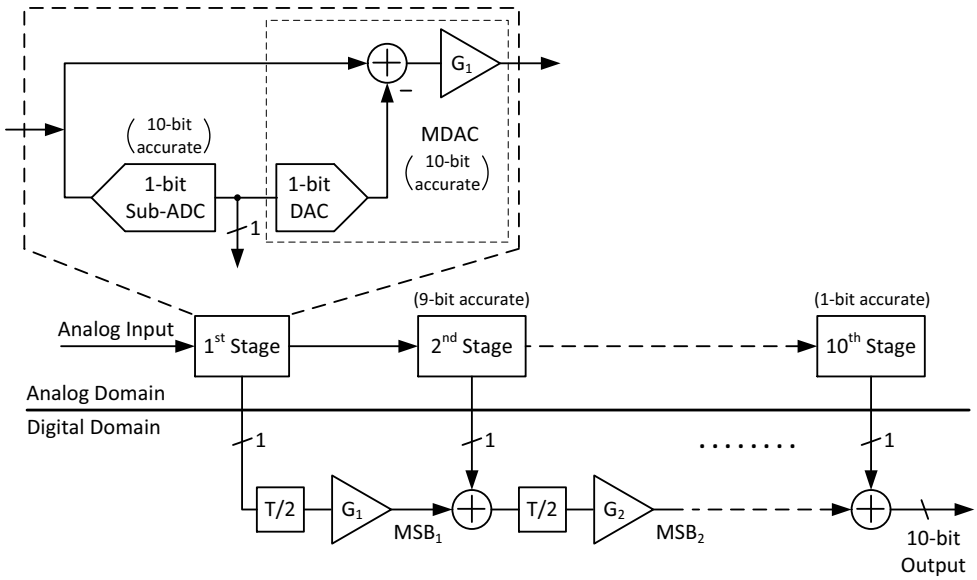

Figure 2.2: Ten-stage pipeline ADC that resolves 1-bit/stage

the accuracy requirement in every subsequent stage decreases by one bit. This is further discussed in Section 2.1.7.

Digital circuitry in the digital domain processes the digital bits as it is outputted from each pipeline stage. As a single bit is generated by each subsequent stage every T/2, it is added to MSB_i generated previously. The sum is delayed by T/2 and the delayed sum gets scaled by the stage gain of the corresponding pipeline stage to generate the next MSB_{i+1}. This process ensures that, in the end, each bit has the correct magnitude weighting and all 10 bits corresponding to a specific input sample are aligned in time.

Because the first input sample must propagate through the pipeline before the first 10-bit output is generated, each stage contributes T/2 of latency to the pipeline ADC. After the first digital output is generated, a new output is generated every clock cycle, T. Thus, the conversion rate or sampling speed of a pipeline ADC, $f_S = 1/T$, is limited by the delay through a single pipeline stage.

2.1.4 Digital Error Correction

This section presents a technique called Digital Error Correction (DEC) [4] that reduces the accuracy requirement in the sub-ADC. When building a pipeline stage, typically the

DAC, summer block, and stage gain are combined into a single block called a Multiplying Digital to Analog Converter (MDAC). In a two phase process, the MDAC samples the input in the first phase and uses the digital bit(s) from the sub-ADC to generate the residue signal in the second phase. Figure 2.3 shows the transfer function of a 1-bit MDAC, which is used in the 10-bit pipeline ADC in Figure 2.2. The analog input has a

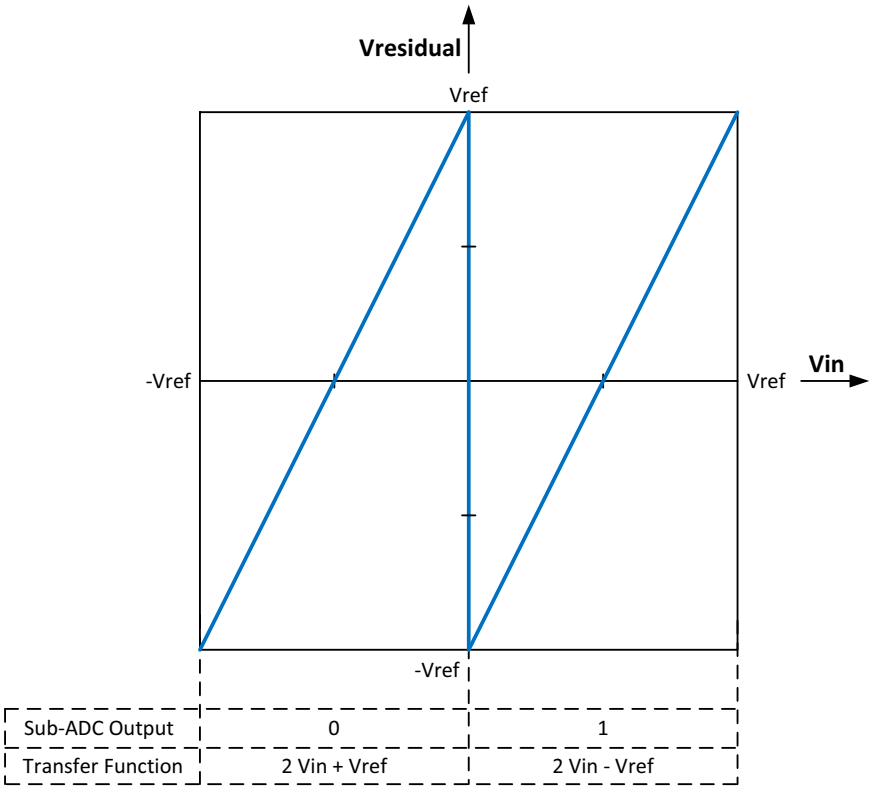

Figure 2.3: Transfer function of a 1-bit MDAC

maximum range defined from $-V_{ref}$ to V_{ref}, which is referred to as the full-scale range. Due to the stage gain, the residue signal will also span this full-scale range. V_{ref} is the reference voltage of the ADC. The 1-bit sub-ADC, which is simply a single comparator, generates the digital output for the pipeline stage and tells the MDAC whether the input is greater or less than zero. Using this information, the MDAC applies the corresponding residue transfer function. However, typically the comparator will have some threshold offset. Figure 2.4 shows the transfer function of a 1-bit MDAC when there is a threshold offset in the comparator. An offset results in the residue voltage exceeding the full-scale

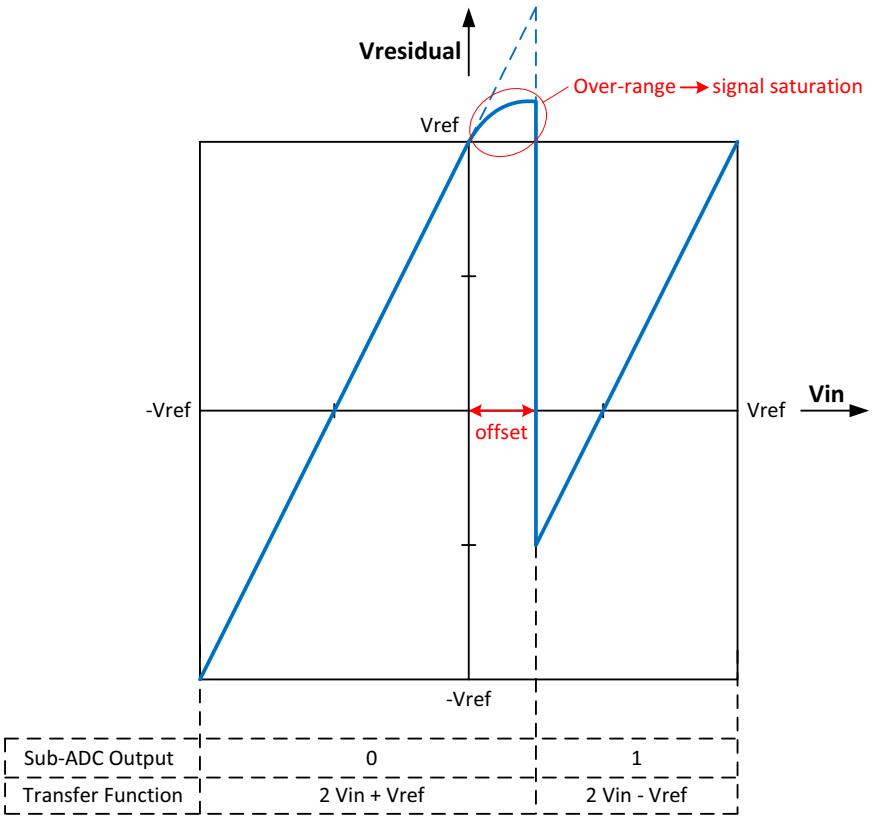

Figure 2.4: Transfer function of a 1-bit MDAC with comparator offset

range. Consequently, the residue output saturates and no longer follows the transfer function. This is because the stage gain is implemented by an opamp in a closed-loop circuit and this circuit has a maximum voltage swing equal to the full-scale range. Section 3.1.3 will discuss this further. Because the residue is not accurately passed onto the next stage, the pipeline has significantly reduced conversion accuracy. To maintain N-bit accuracy, the sub-ADC in the first pipeline stage must have an accuracy of N bits. That is to say, for a 10-bit pipeline with $V_{ref} = 0.8\ V$, the comparator must have a threshold accurate to less than 2 mV. This is very difficult to achieve. A technique called DEC greatly reduces the accuracy requirement of the sub-ADC. Figure 2.5 shows the transfer function of a 1.5-bit MDAC that applies DEC. A 1.5-bit sub-ADC quantizes the input

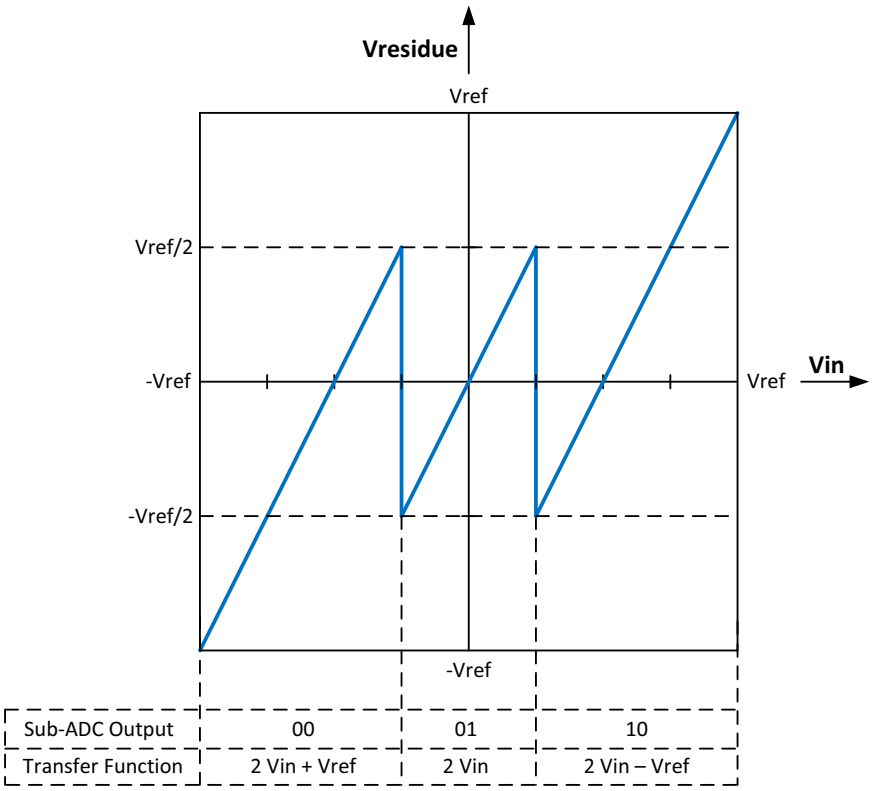

Figure 2.5: Transfer function of a 1.5-bit MDAC

to 3-levels and the residue is limited to half the full-scale range. If comparator threshold offsets cause the residue signal to fall outside it's normal range, the residue signal will be accurately passed onto the next stage provided the offset is within $\pm V_{ref}/4$. Figure

2.6 demonstrates this by showing the transfer function of a 1.5-bit MDAC when there are threshold offsets in the sub-ADC. The additional bit from each adjacent stage is

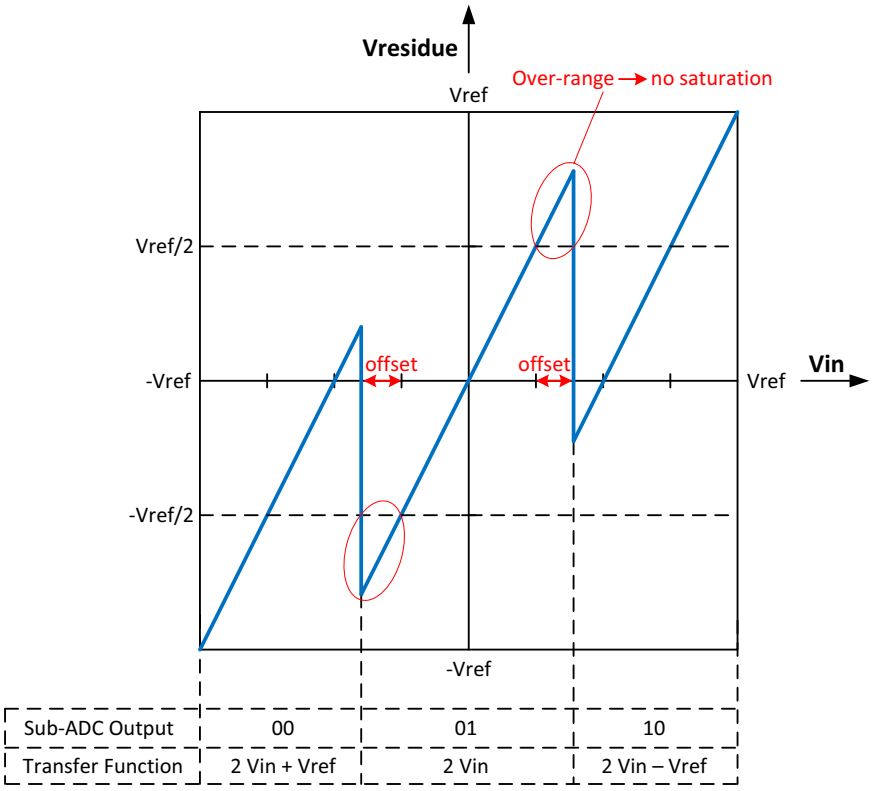

Figure 2.6: Transfer function of a 1.5-bit MDAC with comparator offset

overlapped to correct the over-range error. The fourth level is removed because the MDAC only needs to indicate whether the residue output is above or below Vref/2 and therefore, there is technically only an overlap of half a bit. As a result, the accuracy of the 1.5-bit sub-ADC is reduced from N to 2 bits with DEC. In general, with DEC, the accuracy of the sub-ADC can be reduced if fewer bits are resolved in each stage.

2.1.5 1.5-bit/Stage Pipeline

DEC is a very popular technique because it significantly reduces the accuracy requirements of the sub-ADC. Figure 2.7 shows a more practical version of the pipeline ADC in Figure 2.2. The pipeline ADC applies DEC and consists of a front-end S/H, 8 pipeline stages, and a 2-bit flash ADC at the end. Each pipeline stage resolves 1.5 bits, which is

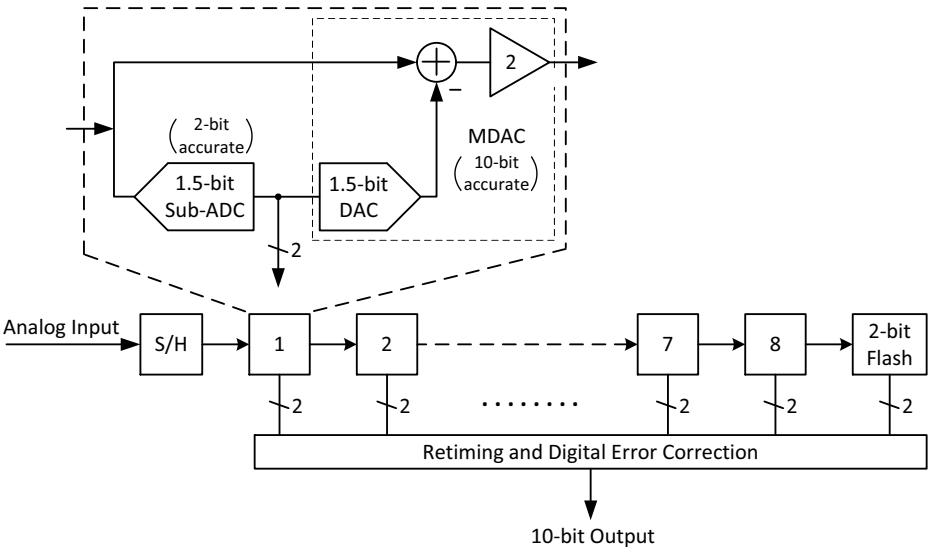

Figure 2.7: A 10-bit 1.5-bit/stage pipeline ADC with digital error correction

represented by a 2-bit output. In the end, a total of 18-bits are generated from a single input sample. To apply DEC, the bits from each adjacent pipeline stage overlap by half a bit and form the expected 10-bit output. This is demonstrated in Figure 2.8.

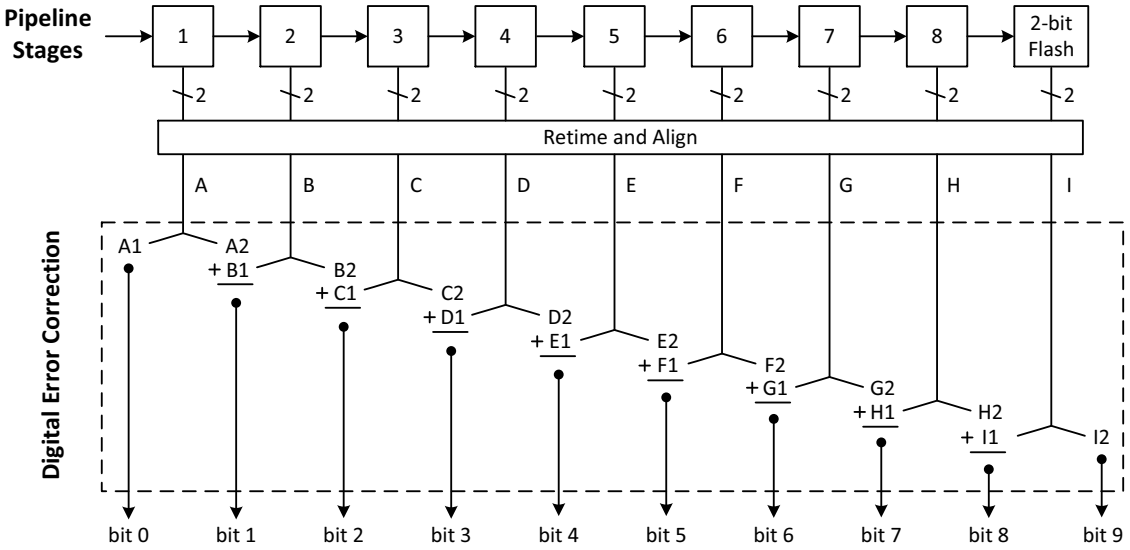

Figure 2.8: Digital error correction in a 1.5-bit/stage 10-bit pipeline ADC

2.1.6 Sample and Hold

As shown in Figure 2.7, a T/H is typically placed before the first stage. It is called a S/H to avoid confusion with the T/Hs discussed previously. The S/H needs to be

accurate to N bits so that it can provide the first stage with an input signal accurate to N bits. The function of the S/H is to ensure that, in the first stage, the sub-ADC and the MDAC sample the same input voltage. Typically the sampling capacitors in the MDAC are much larger than those in the sub-ADC. Therefore, if both were tracking and sampling a changing input signal, the MDAC would always lag the sub-ADC and thus sample a different voltage. For lower input signal frequencies, a pipeline ADC can function without a S/H since the input is changing slow enough for the sub-ADC and MDAC to sample the same value. However, a S/H is typically needed when the ADC is sampling high-frequency input signals, such as when the ADC is sub-sampling. Sub-sampling is where a pipeline ADC is sampling an input that has spectral content above it's Nyquist frequency. The signal frequency above the Nyquist rate gets aliased back into the in-band region, which is from DC to half the sampling frequency $f_S/2$. Once the signal is quantized, it spectral location is lost in the digital domain as there are many possible bands it could have originated from. However, if you limit the input frequencies to within a known region of half the sampling frequency, the in-band region will then correspond to only one possible band. Hence, the signal can be identified exactly in the digital domain.

2.1.7 Thermal Noise and Scaling of Pipeline Stages

The total input-referred noise referenced at the input of the front-end S/H is typically used to evaluate the thermal noise level of a pipeline ADC. Once formulated, it will be a combination of kT/C terms where k is the Boltzmann's constant, T is the operating temperature, and C is one of the capacitors used in the analog to digital conversion (such as the sampling capacitor). Therefore, increasing C will decrease the thermal noise level and increase the dynamic range of the ADC. However, the opamps are the devices that must drive this capacitance and thus increasing C also increases the power consumed by the opamps. Moreover, since opamps consume most of the system power, increasing

C significantly increases the power consumption of the entire system. Therefore, it is sub-optimal to increase C unnecessarily. Generally, setting the thermal noise level to just above the 10-bit level is done to maximize the dynamic range.

Input-referring the output noise of a gain block reduces the noise power by the square of the gain. Therefore, the effect of thermal noise originating further down the pipeline is reduced by each stage gain that is passed. For example, suppose each pipeline stage is identical and has a stage gain of two, then the contribution of noise at the output of the first stage is four times more than the contribution at the output of the second stage. Similarly, other sources of error that originate further down the pipeline have a reduced effect on total accuracy. This is why the accuracy requirement of the pipeline stages in Figure 2.2 can decrease along the pipeline. Similarly, in the 1.5-bit/stage pipeline ADC in Figure 2.7, the accuracy requirement of the MDAC decreases by one bit as you go down the pipeline. As a result, the power consumption and design complexity of the pipeline stage can be scaled. Once the pipeline has been scaled for optimal power consumption, the front-end dominates total power consumption. This is why the front-end S/H makes up a large portion of total power consumption. If the number of bits resolved per stage is increased, the subsequent stages can be relaxed even more and the front-end becomes even more dominant in power; however, the first stage becomes more difficult to design.

2.2 Survey of Recently Published 10-bit Pipeline ADCs

Table 2.1 presents some of the recently published 10-bit pipeline ADCs. The table is limited to conversion rates from 10 to 50 Mega Samples per Second (MS/s).

2.3 Design Application and Target Specifications

Figure 2.9 shows a general block diagram of a wireless radio receiver used in many popular receiver architectures today (e.g. WLAN, Bluetooth, etc). A wireless transmitter transmits the raw data packets across a physical channel, which is then received by the

Table 2.1: Previous 10-bit pipeline ADCs

Reference	Conversion rate [MS/s]	Supply [V]	Peak ENOB [bits]	Power [mW]	FOM [pJ/step]
[5]	10	1.0	8.9	8.1	1.7
[6]	12	1.2	8.4	3.3	0.80
[7]	20	1.2	9	5	0.68
[8]	20.5	1.5	9	19.5	0.19
[9]	30	1.0	8.8	4.7	0.43
[10]	40	1.2	9.5	18.3	0.63
[11]	50	1.8	9.4	9.9	0.30
[12]	50	1.0	8.0	1.9	0.15

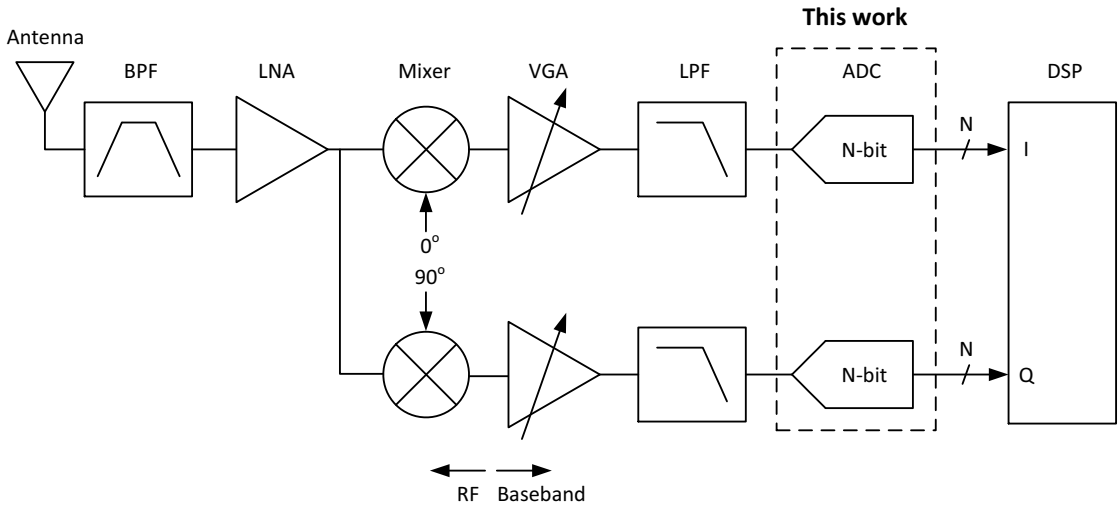

Figure 2.9: General block diagram for a wireless radio receiver

antenna. The signal content from a certain carrier frequency is selected by the Band-Pass Filter (BPF) and put through a Low-Noise Amplifier (LNA) to amplify the signal and suppress noise. The high-frequency Radio Frequency (RF) signal is then mixed down to the low-frequency baseband range. A Variable Gain Amplifier (VGA) scales the analog signal to a full-scale range before the spectral content above the baseband range is filtered by a Low-Pass Filter (LPF). The filtered signal is inputted to the ADC and converted into a digital signal. Once the signal is digitized, it gets processed in the Digital Signal Processing (DSP) block according to the receiver architecture. This thesis describes the design of a pipeline ADC with front-end capacitor-sharing for the ADC blocks within the radio receiver. In [13], a 10-bit pipeline ADC is designed in IBM 0.13 μm for a 802.11a/g Wireless Local Area Network (WLAN) receiver. It operates at 25 MS/s, resolves 1.5 bits/stage, and accepts an input signal up to a maximum range of 1.4 V_{PP} and a frequency of 12.5 MHz. This thesis uses these specifications as a guideline. Table 2.2 summarizes the target specifications for the pipeline ADC in this work. In [14], a capacitor-sharing technique similar to the one that was independently

Table 2.2: Design specifications targeted for this work

Design Parameter	Specification
Technology	0.13 μm
Resolution	10 bits
Sampling rate	20 MS/s
Maximum input frequency	10 MHz
Stage resolution	1.5 bits per stage
Reference voltage	0.8 V
Maximum input swing	1.6 V_{PP}
Supply	1.2 V
FOM	0.5 pJ/step

devised in this work is presented; however, the capacitor-sharing is performed between the pipeline stages and not between the front-end S/H and first stage.

Chapter 3

Pipeline ADC Building Blocks and Design Methodology

This section shows how to design a pipeline ADC for a specific operating speed and (thermal noise and settling) accuracy. In Section 3.1, opamp design considerations will be presented. Then in Section 3.2, three major pipeline building blocks are characterized. Finally, an example design of a N-bit pipeline ADC is described in Section 3.3.

3.1 Opamp Design

The opamp is used in the MDAC to implement the stage gain. It's configured in a closed-loop circuit where the closed-loop gain is equal to the stage gain. The three opamp design parameters that are discussed next are gain, bandwidth, and output swing.

3.1.1 Gain

As discussed in Section 2.1.2, the accuracy of the stage gain determines the settling accuracy of the residue signal. The settling accuracy refers to how close the residue output settles to it's intended value and therefore, it can limit the accuracy of the ADC. Assuming the opamp has sufficient time to settle to it's final value, the opamp's open-

loop gain, 'A', sets the settling accuracy. This is because a higher open-loop gain results in a more accurate stage gain and in turn, the residue output follows a more accurate transfer function. To ensure the residue settles to within Δ LSB, the loop gain of the closed-loop circuit, $A\beta$, must be:

$$A\beta > \frac{2^N}{\Delta} \tag{3.1}$$

where β is the feedback factor of the closed-loop circuit and 1 $LSB = 1/2^N$. The loop gain, as the name suggests, is the gain around the opamp closed-loop circuit. Considering there are other sources of error (e.g. thermal noise), a reasonable choice is $\Delta = 0.25$ LSB. For instance, the front-end S/H and first pipeline stage in a 10-bit pipeline ADC requires a loop gain of:

$$\frac{2^{10}}{0.25} = 4096 \; or \; 72dB \tag{3.2}$$

3.1.2 Bandwidth

Assuming the opamp has sufficient gain to accurately settle to it's final value, the opamp speed, which determines how fast the residue output settles to a final value, sets the settling accuracy. The bandwidth must be high enough for the opamp to settle to a sufficiently accurate value within the required time of half a sampling period, $0.5/f_S$. Equation 3.3 sets the bandwidth of the opamp closed-loop circuit, f_{3dB}, so that the residue settles to within 0.5 LSB in half a sampling period.

$$f_{3dB} = \frac{(N+1) \cdot ln2 \cdot f_S}{\pi} \tag{3.3}$$

Like with loop gain, the bandwidth of the opamp closed-loop circuit is the product of the opamp's Gain Bandwidth (GBW) product and β:

$$f_{3dB} = GBW\beta \tag{3.4}$$

3.1.3 Output Voltage Swing

Clearly, high opamp gain is needed at the front-end to achieve the desired settling accuracy. Previous publications ([6], [8], [1] have typically used single-stage opamps with gain-boosting (Figure 3.1) or two-stage Miller-compensated opamps (Figure 3.2) to achieve these gains. In Figures 3.1 and 3.2, V_{CMFB} is a signal generated by a Common-Mode Feedback (CMFB) circuit to control the Common-Mode (CM) output voltage.

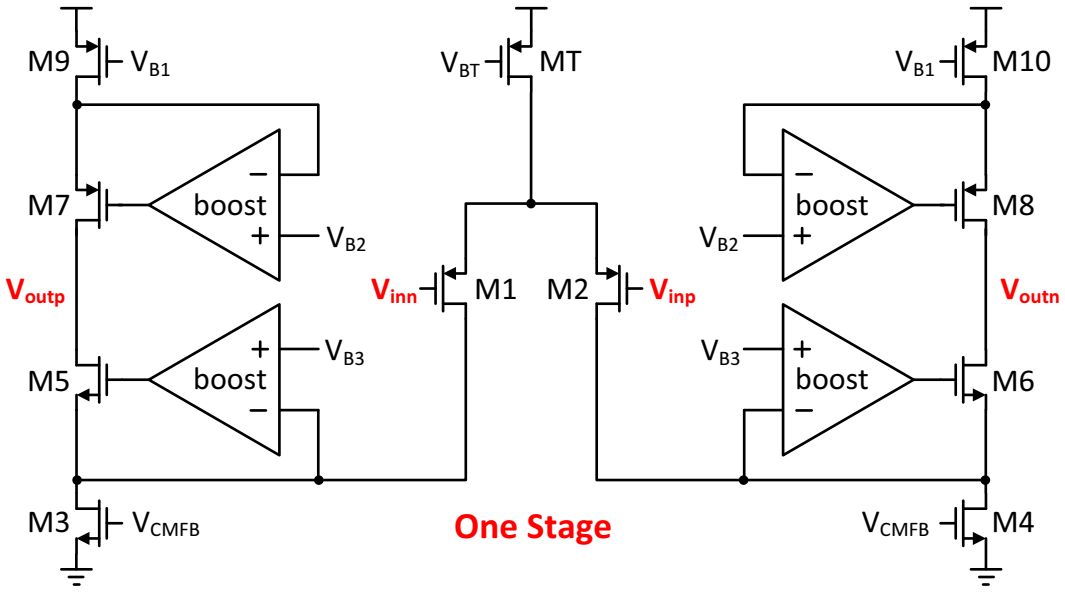

Figure 3.1: Single-stage opamp with gain-boosting

The two-stage opamp has more output swing than the single-stage opamp because the output transistors in the two-stage opamp are cascoded. Each transistors must have a drain-to-source voltage, V_{DS}, of at least one overdrive voltage, V_{eff}; otherwise, the transistors drop out of saturation and the gain dramatically decreases. A transistor's V_{eff} is the difference between its gate-to-source voltage, V_{GS}, and its threshold, V_t. At a supply voltage of 1.2 V and a V_{eff} of 150 mV, the differential output swing in Figure 3.1 is limited to:

$$2\ (1.2 - 4V_{eff}) = 1.2\ V_{pp} \tag{3.5}$$

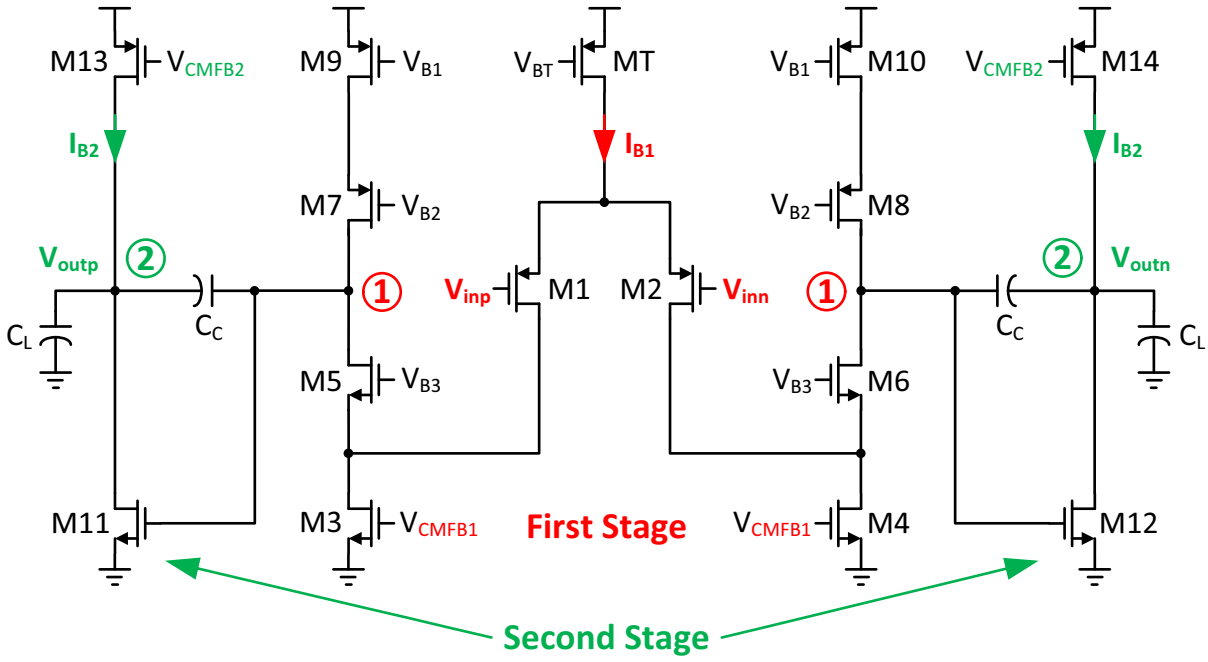

Figure 3.2: Two-stage opamp with a folded-cascode and a common source stage

For the two-stage case, the differential output swing in Figure 3.2 is limited to:

$$2\ (1.2 - 2V_{eff}) = 1.8\ V_{pp} \tag{3.6}$$

Equations 3.5 and 3.6 specify the absolute maximum swing; however, the gain drops even as the output transistors near the edge of saturation. Instead of designing a much higher gain to accommodate for the drop, a simpler method is to design for greater swing. For a full-scale range of 1.6 V_{pp}, this work adopts a two-stage opamp like the one in Figure 3.2. The first stage is a folded cascode, which will generate most of the gain, and the second stage is a simple common source, which supports an output swing of up to 1.8 V_{pp}.

3.2 Building Blocks

In this section, the 1.5-bit sub-ADC and choice of comparator are presented. The front-end S/H and two 1.5-bit MDAC blocks are then analyzed to show that their thermal noise level sets the dynamic range of the pipeline ADC.

3.2.1 Sub-ADC

As described in Section 2.1.3, the sub-ADC samples the input voltage in Φ_1 and then in Φ_2, digitizes it to a sub-resolution equal to the number of bits resolved per stage. The bits are used by the MDAC in Φ_2 to calculate the residue for the next stage. In this work, a 1.5-bit sub-ADC is used in each pipeline stage. A 1.5-bit sub-ADC is composed of two comparators and digital logic as shown in Figure 3.3. The comparator used in this

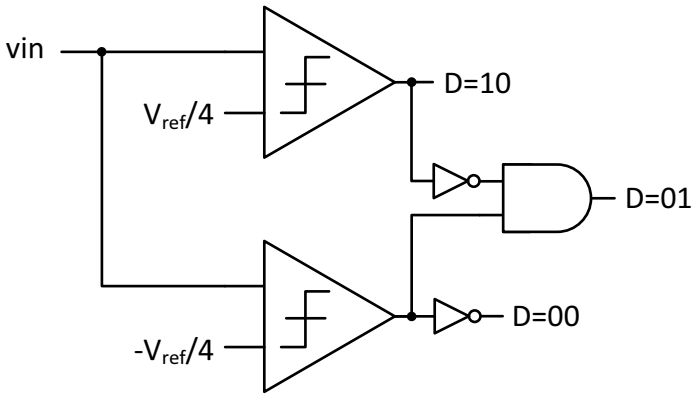

Figure 3.3: Three-level sub-ADC

work is the the Charge-Distribution comparator [15] shown in Figure 3.4. The threshold

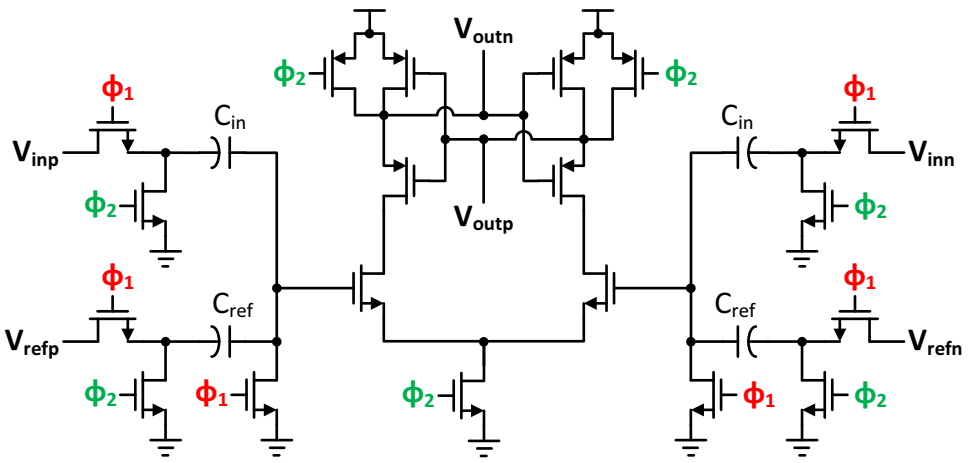

Figure 3.4: Charge distribution comparator

of this comparator is set by:

$$V_{thres} = \frac{C_{ref}}{C_{in}}(V_{refp} - V_{refn}) \tag{3.7}$$

The total input capacitance of the sub-ADC during Φ_1 is $C_{iT,subADC}$, which is equal to $2C_{in}$. It is desirable to make C_{in} small because $C_{iT,subADC}$ loads the opamp in the previous stage. Since the sub-ADC only needs to be accurate to 2 bits, the sampling capacitors C_{in} and C_{ref} can be made small without worrying about thermal accuracy. In fact, the smallest value C_{in} and C_{ref} can be is most likely limited by the fabrication technology.

The comparators in a 1.5-bit sub-ADC have thresholds at $V_{thres} = \pm V_{ref}/4$ as shown in the transfer function of the 1.5-bit MDAC in Figure 2.5. If only flip-around MDACs [16] are used and only V_{ref} is made available for the MDAC, then from Equation 3.7, $C_{in}/C_{ref} = 4$. Thus, $C_{in} = 4C_{ref}$ and $C_{iT,subADC} = 8C_{ref}$. However, if an integrator MDAC [16] is used and $V_{ref}/2$ is made available, then C_{in}/C_{ref} can be reduced to 2 and hence, $C_{in} = 2C_{ref}$ and $C_{iT,subADC} = 4C_{ref}$. Therefore, making a fraction of V_{ref} available on-chip reduces $C_{iT,subADC}$.

3.2.2 Sample and Hold

The S/H located at the front of the pipeline performs the first step of sampling the external input signal. Figure 3.5 shows a commonly used S/H. Phase Φ_1 and Φ_2 are

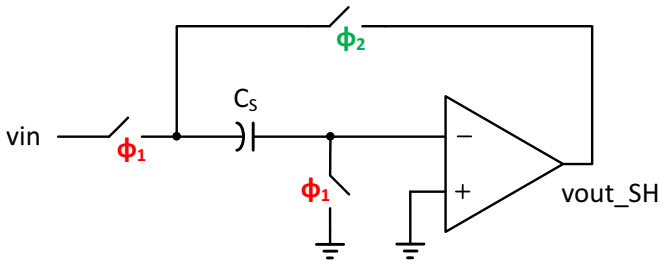

Figure 3.5: A commonly used S/H

indicated on the clock waveform in Figure 3.6. The total input sampling capacitance the external source must drive in Φ_1 is $C_{iT,S/H} = C_S$. In Φ_2, the S/H opamp drives a load capacitance to hold the sampled voltage for the first pipeline stage. The capacitor C_S does not load the opamp because it's top plate has no path to ground. As a result, the

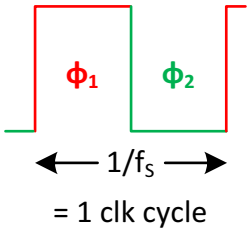

Figure 3.6: Clock waveform with phases Φ_1 and Φ_2 labeled

feedback factor of the S/H is:

$$\beta_{S/H} = \frac{C_S}{C_S + 0} = 1 \tag{3.8}$$

However, the total input sampling capacitance of the sub-ADC, $C_{iT,subADC}$, and of the MDAC, $C_{iT,MDAC}$, in the next stage does load the S/H opamp. The opamp is also loaded by an additional capacitance, C_{n2}, which is the total parasitic capacitance connected to node 2 in Figure 3.2. It consists of the drain capacitance of M11 and M13, and the wire capacitance of the metal used to connect voutp/voutn. Therefore, the total output load on the S/H opamp is:

$$C_{L,S/H} = C_{iT,subADC} + C_{iT,nextstageMDAC} + C_{n2} \tag{3.9}$$

From [17], the differential input-referred thermal noise power of a S/H or an MDAC is:

$$\overline{v_{i/p}}^2 = \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}(\beta)(1 + n_f)\right) \tag{3.10}$$

where C_C is the compensation capacitor used to stabilize the opamp closed-loop circuit.

As you can see, the size of the sampling and compensation capacitors set the thermal noise level and hence, the dynamic range of the pipeline ADC. Specifically looking at the S/H, as the input signal is being sampled in Φ_1, thermal noise from the transistors is sampled across C_S. The opamp is not used and therefore, does not contribute any noise. The first term in Equation 3.10 is the noise contribution from transistors generated during Φ_1. Because a differential configuration is used, the term becomes $2kT/C_S$. In Φ_2, both

the opamp and transistors contribute thermal noise. However, the dominate source of noise is from the opamp's first stage in Figure 3.2. The noise from the second stage is negligible because, once input-referred, it is greatly reduced by the gain of the opamp. Therefore, the second term in Equation 3.10 is the noise contribution from the opamp's first stage generated during Φ_2. The noise fraction, n_f, is defined in Equation 3.11 and it's presence in Equation 3.10 accounts for the noise contributed by transistors M3/M4 and M9/M10 in Figure 3.2. The formula is based on a similar calculation performed on a simple opamp in [18].

$$nf = \frac{g_{m3} + g_{m9}}{g_{m1}} \tag{3.11}$$

Since typically $n_f = 1$ [17] and $\beta_{S/H} = 1$, the input-referred thermal noise power of the S/H block is:

$$\overline{v_{i/p,S/H}}^2 = \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}(1)(2)\right) \tag{3.12}$$

3.2.3 Multiplying DAC

In this work, a 1.5-bit MDAC is used in each pipeline stage. This section analyzes two well-known 1.5-bit MDACs.

3.2.3.1 Flip-Around MDAC

Figure 3.7 shows the very popular flip-around MDAC [16]. As can be seen from the transfer function of a 1.5-bit MDAC in Figure 2.5, the value of V_{DAC} is either $+V_{ref}$, 0 V, or $-V_{ref}$ depending on the output of the sub-ADC. The total input sampling capacitance in Φ_1 is $C_{iT,MDAC} = C_s$. In Φ_2, the two $C_S/2$ capacitors apply a load of $C_S/4$ on the MDAC opamp. Taking into consideration the loading effects of the next stage and of the parasitic capacitance C_{n2}, as discussed in Section 3.2.2, the total output load on the flip-around MDAC opamp is:

$$C_{L,flipMDAC} = C_S/4 + C_{iT,subADC} + C_{iT,nextstageMDAC} + C_{n2} \tag{3.13}$$

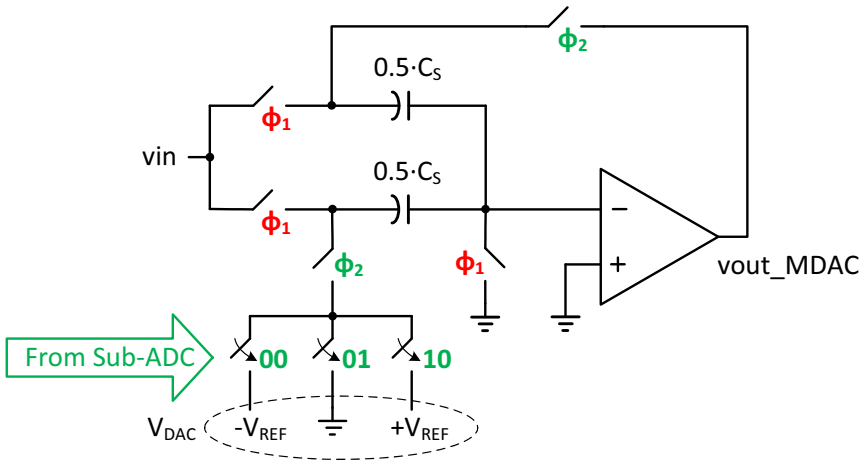

Figure 3.7: Flip-around MDAC

The feedback factor of the flip-around MDAC is:

$$\beta_{flipMDAC} = \frac{C_S/2}{C_S/2 + C_S/2} = \frac{1}{2} \tag{3.14}$$

Using Equation 3.10, the differential input-referred noise of the flip-around MDAC is:

$$\overline{v_{i/p,flipMDAC}}^2 = \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}\left(\frac{1}{2}\right)(2)\right) \tag{3.15}$$

As you can see, increasing C_S increases C_L and decreases the thermal noise level, and vice versa. Therefore, there is a trade-off between C_L and the thermal noise level. Since C_L determines the opamp power consumption and the thermal noise level limits the accuracy of the ADC, the trade-off is actually, as expected, between power and accuracy.

3.2.3.2 Integrator MDAC

Figure 3.8 shows the integrator MDAC [16]. Due to the inherent gain between the sampling capacitors, the magnitude of V_{DAC} and hence V_{ref} in the integrator MDAC must be half of that in the flip-around MDAC. Therefore, $V_{ref}/2$ is required if integrator MDACs are used. Again, the total input sampling capacitance in Φ_1 is $C_{iT,MDAC} = C_S$. In Φ_2, C_S and $C_S/2$ apply a load of $C_S/3$ on the MDAC opamp. Taking into consideration

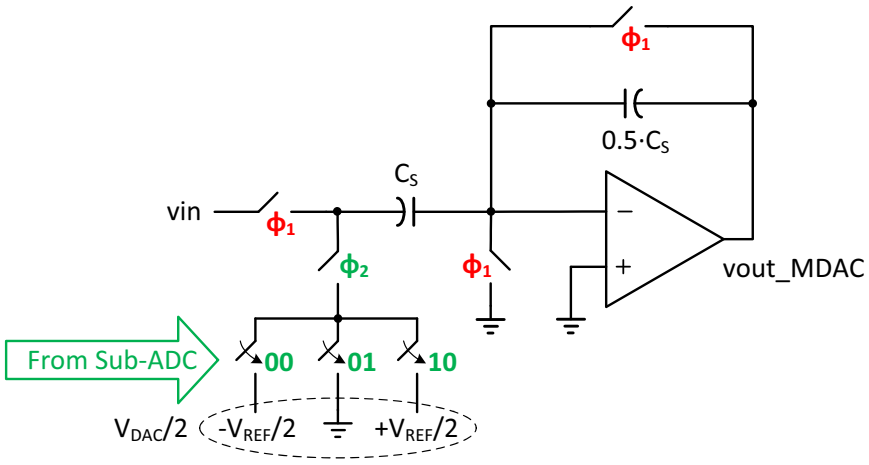

Figure 3.8: Integrator MDAC

the loading effects of the next stage and of the parasitic capacitance C_{n2}, as discussed in Section 3.2.2, the total output load on the integrator MDAC opamp is:

$$C_{L,intMDAC} = C_S/3 + C_{iT,subADC} + C_{iT,nextstageMDAC} + C_{n2} \tag{3.16}$$

The feedback factor of the integrator MDAC is:

$$\beta_{intMDAC} = \frac{C_S/2}{C_S + C_S/2} = \frac{1}{3} \tag{3.17}$$

From Equations 3.10, the differential input-referred noise of the integrator MDAC in Figure 3.8 is:

$$\overline{v_{i/p,intMDAC}}^2 = \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}\left(\frac{1}{3}\right)(2)\right) \tag{3.18}$$

3.3 Design Procedure

This section goes through an example design of an N-bit pipeline ADC. First the size of C_S is chosen and then the two-stage opamp for a specific block is designed. The theory and equations are taken from [19].

Step 1

Using the thermal noise equations in Equations 3.12, 3.15, and 3.18, the total input-referred noise power, $\overline{v_{i/p,total}}^2$, referenced at the input of the front-end S/H is formulated in terms of C_S and C_C. Setting $C_C = C_S$ is a good starting point, but it is checked afterwards in Step 9.

Step 2

Add up the output load capacitance from each stage to create $C_{L,total}$. Since the majority of power is consumed in the opamps and opamp power consumption is proportional to the load it drives, the total power consumed by the ADC is proportional to $C_{L,total}$.

Step 3

According to Section 2.1.7, the accuracy requirement decreases as you go down the pipeline. Therefore, scale C_S along the pipeline for an optimal trade-off between $\overline{v_{i/p,total}}^2$ and $C_{L,total}$, which, as mentioned in Section 3.2.3.1, are inversely proportional to each other.

Step 4

Once a choice in scaling has been made, apply Equation 3.19 to set the size of C_S and C_C so that the input-referred noise level of the ADC is at 0.25 LSB:

$$\overline{v_{i/p,total}}^2 = \left(1.6 \cdot \frac{0.25}{2^N}\right)^2 \tag{3.19}$$

Step 5

Now that the values of C_S and C_C have been chosen, the two-stage opamp in Figure 3.2 can be designed. Use Equation 3.3 to calculate f_{3dB} and then apply Equation 3.4 to calculate GBW.

Step 6

The following equation shows that GBW is dependent on the transconductance of the input pair transistors M1/M2, g_{m1}, and on C_C:

$$GBW = \frac{g_{m1}}{2\pi C_C} \tag{3.20}$$

Apply Equation 3.20 to calculate g_{m1}.

Step 7

The following equation specifies the opamp's non-dominant pole, f_{nd}:

$$f_{nd} = \frac{g_{m12}}{2\pi C_L} \cdot \frac{C_C}{C_C + C_{n1}} \tag{3.21}$$

where g_{m12} is the transconductance of the input transistors M11/M12 in the second stage. The total capacitance at node 1 in Figure 3.2 is the parasitic capacitance C_{n1}. It mainly consists of the drain capacitance of M5/M7 and gate capacitance of M11/M12. To achieve sufficient phase margin, a reasonable choice is $f_{nd} = 3f_{3dB}$. Equation 3.21 can then be rearranged as:

$$g_{m12} = (3f_{3dB})2\pi C_L \cdot \frac{C_C + C_{n1}}{C_C} \tag{3.22}$$

With an estimate of C_{n1}, apply Equation 3.22 to calculate g_{m12}.

Step 8

The Slew Rate (SR) of the two-stage opamp in Figure 3.2 is expressed as:

$$SR = \frac{I_{B1}}{C_C} \tag{3.23}$$

The relationship between the current I, transconductance g_m, and V_{eff} of a MOSFET in saturation is:

$$g_m = \frac{2I}{V_{eff}} \tag{3.24}$$

Combining Equations 3.24, 3.23, and 3.20 and rearranging to solve for the ratio between SR and GBW gives:

$$\frac{SR}{GBW} = 2\pi V_{eff} \tag{3.25}$$

A sufficient SR is needed to prevent the opamp from slewing; however, the value required depends on the bandwidth of the opamp. Equation 3.25 implies that, if the V_{eff} of M1/M2 is large, the SR is also large for a given GBW. Therefore, a reasonable V_{eff} for M1/M2 should be chosen to prevent the opamp from slewing. After choosing a reasonable V_{eff} for M1/M2, let that be the V_{eff} of M11/M12 as well. Use Equation 3.24 to calculate I_{B1} and I_{B2}.

Step 9

The following equation specifies a recommended range for C_C:

$$2C_{n1} < C_C < C_L/2 \tag{3.26}$$

Use Equation 3.26 to check that the value of C_C set in Step 1, $C_C = C_S$, is within the recommended range.

Step 10

Use Equation 3.1 to specify the gain and size the transistors accordingly.

Chapter 4

Capacitor-Sharing Pipeline Design and Simulation

This chapter demonstrates the power savings of the the proposed front-end capshare technique. In Section 4.1, the front-end capshare technique is first explained. In Section 4.2, a power comparison between a regular ADC and capshare ADC is presented. The circuit implementation of the fabricated capshare ADC is then described in Section 4.3. Finally, the simulation results of the capshare design are presented in Section 4.4.

4.1 Front-End Capacitor-Sharing

In this section, the theory and power benefits of front-end capacitor-sharing are explained. The most common configuration for the front-end S/H and the first-stage MDAC is shown in Figure 4.1. In Step 1, the S/H samples the input signal onto C_S. In Step 2, the S/H holds it's value so the MDAC and sub-ADC in the first-stage can resample it. In Step 3, the MDAC uses the bits produced by the sub-ADC to generate the residue output for the next pipeline stage to sample. From Equations 3.12 and 3.15, the input-referred

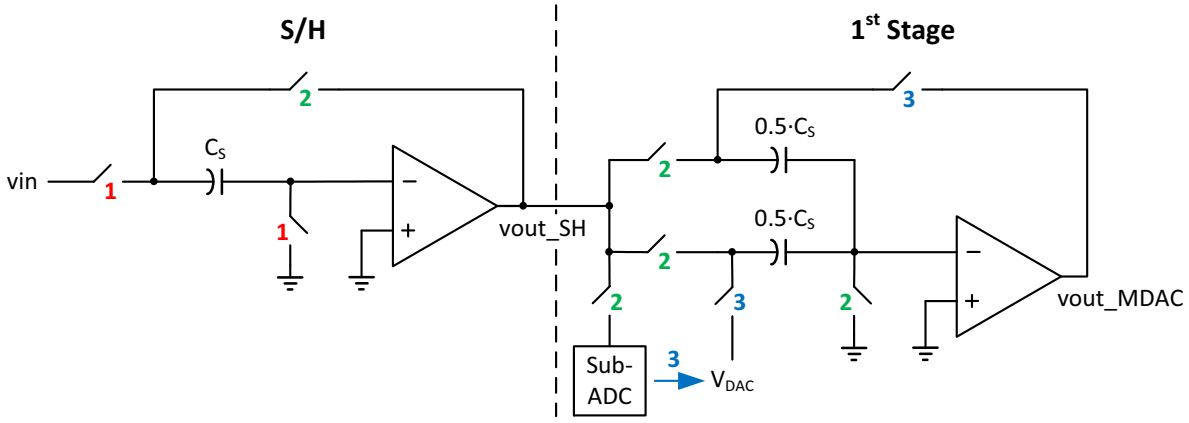

Figure 4.1: Front-end S/H and first stage flip-around MDAC

noise of Figure 4.1 is:

$$\begin{aligned} \overline{v_{i/p}}^2 &= \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}(1)(2)\right) + \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}(0.5)(2)\right) \\ &\approx \left(4.67\frac{kT}{C_S}\right) + \left(3.33\frac{kT}{C_S}\right) = \left(8\frac{kT}{C_S}\right) \end{aligned} \tag{4.1}$$

The approximation initially assumes $C_C = C_S$ to simplify the equation. According to Equation 3.9, the S/H experiences a load of $C_S + C_{iT,subADC}$. From Equation 3.13, the flip-around MDAC experiences a load of $C_S/4$ since, for the time being, the stages that come after the first stage are ignored. For simplicity, the analysis also ignores the parasitic capacitance C_{n2}.

An integrator MDAC can be used instead of using a flip-around MDAC as shown in Figure 4.2. Steps 1 to 3 are the same as the steps in Figure 4.1. From Equations 3.12 and 3.18, the input-referred noise of Figure 4.2 is:

$$\begin{aligned} \overline{v_{i/p}}^2 &= \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}(1)(2)\right) + \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}(0.33)(2)\right) \\ &\approx \left(4.67\frac{kT}{C_S}\right) + \left(2.89\frac{kT}{C_S}\right) = \left(7.56\frac{kT}{C_S}\right) \end{aligned} \tag{4.2}$$

Again, the S/H experiences a load of $C_S + C_{iT,subADC}$ and according to Equation 3.16, the integrator MDAC experiences a load of $C_S/3$ in this configuration.

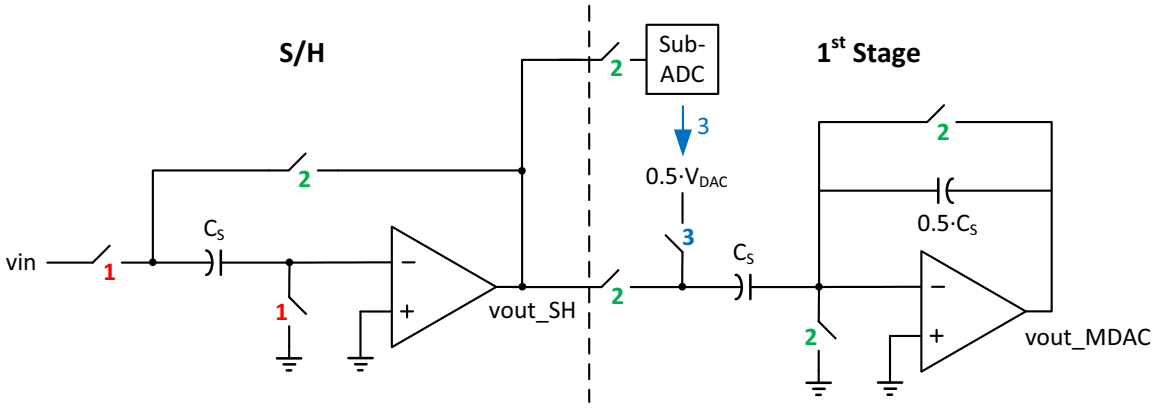

Figure 4.2: Front-end S/H and first stage integrator MDAC

In Step 2 of Figures 4.1 and 4.2, $C_{iT,MDAC}$ is charged to the same voltage that was sampled by the S/H in Step 1. If the signal across C_S in Step 1 could be reused in Step 2, the second charge could be avoided. Figure 4.3 illustrates how this is done with the capacitor-sharing technique. As you can see, Figure 4.2 is modified so that C_S is shared

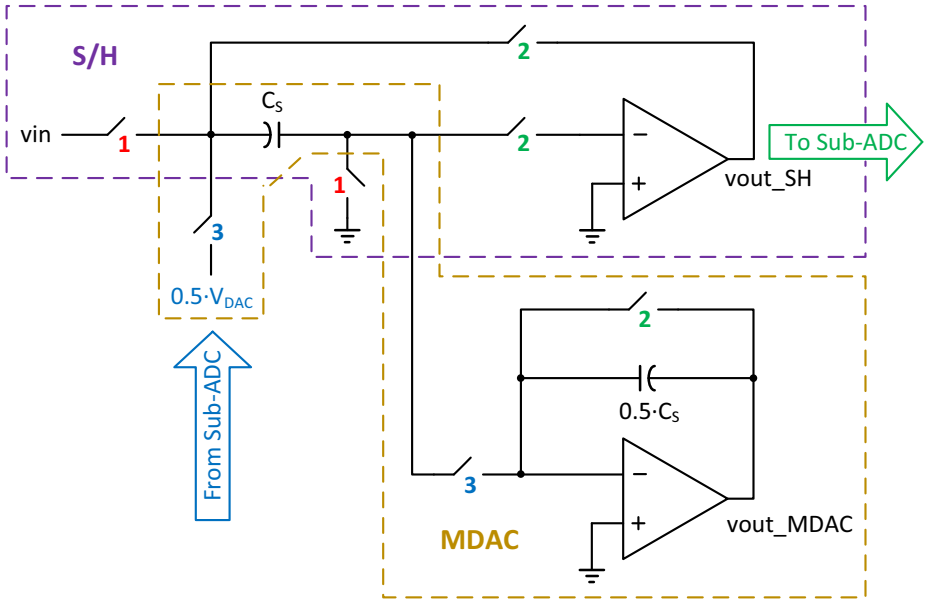

Figure 4.3: Front-end S/H sharing Cs with first stage MDAC

between the S/H and the MDAC in the first stage. As a result, $C_{iT,MDAC} = C_S$ no longer loads the S/H and the thermal noise generated by the second charge is removed. Like in Figure 4.1 and 4.2, the input signal is sampled onto C_S in Step 1. However, in Step 2, the S/H only needs to charge $C_{iT,subADC}$ in first stage. In Step 3, C_S, which still has the input

signal across it, is reused as the sampling capacitor for the first stage and the residue is generated for the next stage. Since there are only two phases, Step 1 and 3 map to Φ_1 and Step 2 maps to Φ_2. The problem is that in Φ_1, both Step 1 and 3 require the use of C_S. This is solved by implementing a ping-pong scheme where two parallel paths can operate in the same clock cycle but in different ping-pong phases. Figure 4.4 shows how Figure 4.3 is modified to apply the ping-pong scheme. The ping-pong clock waveforms are shown in Figure 4.5. As you can see, when the S/H is sampling the input in phase A (B), the residue $vout_{MDAC}$ is being generated based on the previously sampled input in phase B (A). If this scheme was not applied, the S/H could only sample the input every second clock cycle.

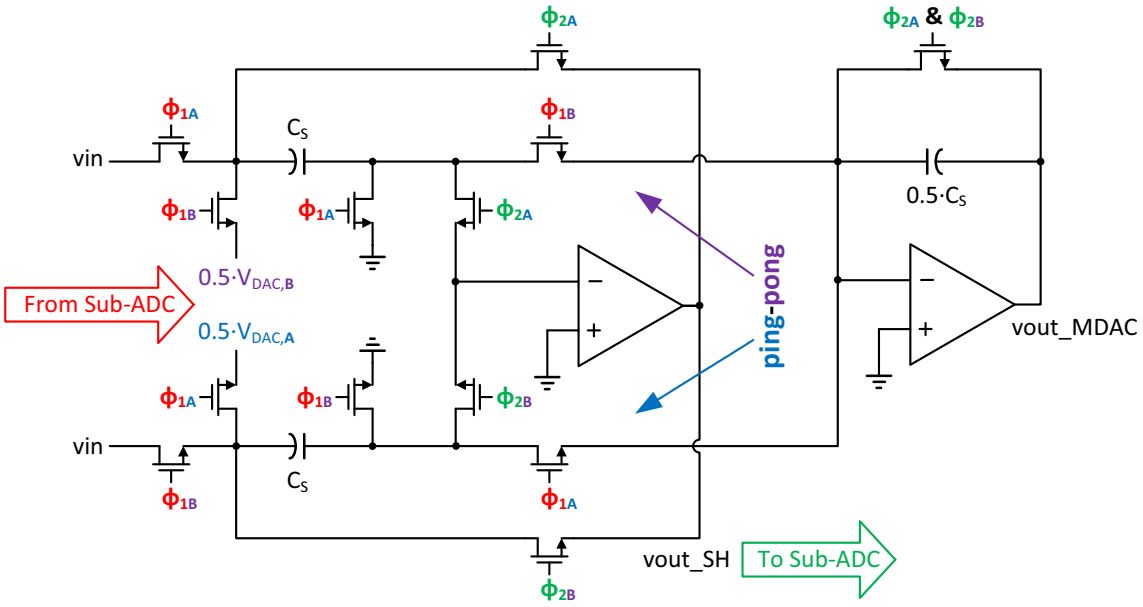

Figure 4.4: Capacitor-sharing front-end with ping-pong scheme

In Equation 4.2, the $2kT/C_S$ term is the thermal noise term generated by the second charge. Since the second charge has effectively been removed, the new input-referred

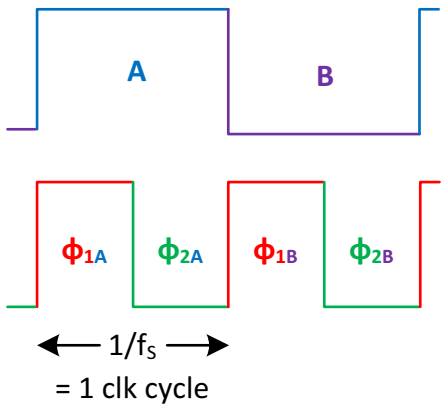

Figure 4.5: Clock waveform with clock cycle phases and ping-pong phases labeled

noise is simply Equation 4.2 minus $2kT/C_S$:

$$\begin{aligned} \overline{v_{i/p}}^2 &= \left(2\frac{kT}{C_S} + \frac{4}{3}\frac{kT}{C_c}(1)(2)\right) + \left(0 + \frac{4}{3}\frac{kT}{C_c}(0.33)(2)\right) \\ &\approx \left(4.67\frac{kT}{C_S}\right) + \left(0.89\frac{kT}{C_S}\right) = \left(5.56\frac{kT}{C_S}\right) \end{aligned} \tag{4.3}$$

As a result of the capshare technique, the S/H is only loaded by $C_{iT,subADC}$. The flip-around MDAC still experiences a load of $C_S/3$ like it does in Figure 4.2. Comparing the configurations in Figures 4.1 and 4.2 with the capshare configuration in Figure 4.3, this analysis concludes that front-end capacitor-sharing of C_S between the S/H and first stage: (1) eliminates the loading effect of $C_{iT,MDAC}$ on the S/H, and (2) reduces the thermal noise in the system. This saves a significant amount of power in the front-end S/H as will be shown next.

4.2 Power Comparison

This section demonstrates the power savings of the front-end capacitor-sharing technique by comparing a pipeline ADC that applies the technique to a regular pipeline ADC that does not.

4.2.1 Regular versus Capshare ADC

To demonstrate the power savings, the thesis compares two 20 MS/s 10-bit 1.5-bit/stage pipeline ADC designs:

- A regular design that uses the popular flip-around MDAC in the first stage as shown in Figure 4.6
- A novel design that applies front-end capacitor-sharing between the S/H and first stage MDAC as shown in Figure 4.7

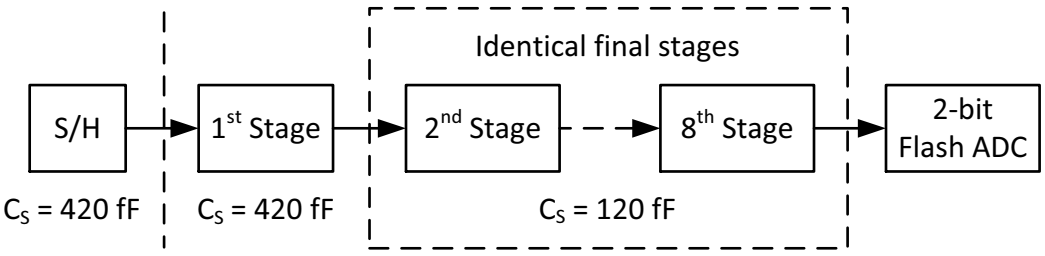

Figure 4.6: Regular pipeline ADC with a flip-around MDAC in the first stage

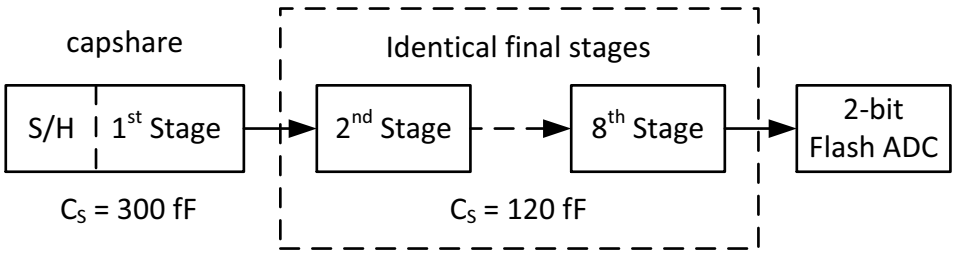

Figure 4.7: Pipeline ADC with capacitor-sharing front-end

This thesis will concentrate on the design of the front-end S/H and first stage. The final seven stages, stages 2 to 8, will be identical in each design to reduce the design complexity. They will use a flip-around MDAC with 60 fF sampling capacitors, which is the minimum size Metal Insulator Metal (MIM) capacitors in IBM 0.13 μm technology. Hence, the final pipeline stages have a $C_{iT,MDAC} = C_S = 120fF$.

4.2.2 Analysis

Next, an analysis is done to estimate and compare the power consumptions of the two designs. Since most of the power is consumed in the opamps, only the opamp power will be considered. The analysis follows the example design performed in Section 3.3 and uses the target specifications in Table 2.2.

4.2.2.1 Initial Assumptions

Before starting the analysis, several initial assumptions are presented:

- $V_{ref} = 0.8\ V$ meaning the full-scale range is 1.6 V_{PP}.
- Each pipeline stage use the same 1.5-bit sub-ADC. Referring to Section 3.2.1, the following analysis assumes $V_{ref}/2$ is available on-chip, allowing $C_{iT,subADC} = 4C_{ref}$. To minimize $C_{iT,subADC}$, set $C_{ref} = 60\ fF$, which is the smallest MIM capacitor allowed. Therefore, $C_{iT,subADC} = 240fF$.
- The opamp parasitic capacitance at node 1 and 2 in Figure 3.2 is approximately 150 fF (i.e. $C_{n1} = C_{n2} = 150\ fF$).
- As indicated by Equation 3.25, a reasonable V_{eff} for M1/M2 in Figure 3.2 is required to ensure that the opamp is not slew limited. Set the V_{eff} of M1/M2 and M11/M12 to be 150 mV.
- An over-design factor of 1.5x is used.
- Initially set $C_C = C_S$. Later, Equation 3.26 will be used to check that C_C is within the recommended range.

4.2.2.2 Choosing the Sampling Capacitance

In this section, the size of C_S and C_C for both the regular and capshare ADC are chosen so that both have the same dynamic range. Consider the regular pipeline ADC in Figure

4.6. The total input-referred noise taken at the input of the front-end S/H is calculated by adding the noise contribution of stages 2 to 8 to Equation 4.1:

$$\begin{aligned}\overline{v_{i/p,total}}^2 &= \left(2\frac{kT}{C_S}+\frac{4}{3}\frac{kT}{C_c}(1)(2)\right) + \left(2\frac{kT}{C_S}+\frac{4}{3}\frac{kT}{C_c}(0.5)(2)\right) \\ &\quad + \left(2\frac{kT}{(120\ fF)}+\frac{4}{3}\frac{kT}{C_c}(0.5)(2)\right)\times\sum_{n=1}^{7}\left(\frac{1}{2^{2n}}\right) \\ &\approx \left(8.44\frac{kT}{C_S}\right)+\left(0.67\frac{kT}{(120\ fF)}\right)\end{aligned} \tag{4.4}$$

Applying Equation 3.19 to set the thermal noise level to 0.25 LSB limits $C_S > 415\ fF$. Therefore, let $C_S = C_C = 420\ fF$ for the regular ADC.

The same analysis can be performed for the capshare pipeline ADC in Figure 4.7. The total input-referred noise taken at the input of the front-end S/H is calculated by adding the noise contribution of stages 2 to 8 to Equation 4.3:

$$\begin{aligned}\overline{v_{i/p,total}}^2 &= \left(2\frac{kT}{C_S}+\frac{4}{3}\frac{kT}{C_c}(1)(2)\right) + \left(0+\frac{4}{3}\frac{kT}{C_c}(0.33)(2)\right) \\ &\quad + \left(2\frac{kT}{(120\ fF)}+\frac{4}{3}\frac{kT}{C_c}(0.5)(2)\right)\times\sum_{n=1}^{7}\left(\frac{1}{2^{2n}}\right) \\ &\approx \left(6\frac{kT}{C_S}\right)+\left(0.67\frac{kT}{(120\ fF)}\right)\end{aligned} \tag{4.5}$$

Applying Equation 3.19 to set the thermal noise level at 0.25 LSB limits $C_S > 295\ fF$. Therefore, let $C_S = C_C = 300\ fF$ for the capshare ADC. The chosen values of C_S and C_C ensure that the capshare and regular ADC at are the same thermal noise level and hence, have the same dynamic range.

4.2.2.3 Two-Stage Opamp Design

This section designs the opamp within each pipeline stage. The corresponding opamps in the capshare and regular ADC should have the same bandwidth so that the power consumption between the capshare and regular ADC can be fairly compared. From

Equation 3.3, in order for the residue signal to settle within 0.5 LSB at a sampling frequency of 20 MS/s, the f_{3dB} of the closed-loop system should be:

$$\frac{(10+1)\ln 2(20\ MHz)}{\pi} = 48.5\ MHz \tag{4.6}$$

After applying an over-design factor of 1.5, each opamp targets $f_{3dB} = 72.8\ MHz$.

Next, the power for each stage is computed and compared. Table 4.1 applies the capacitor values calculated above, $C_S = C_C = 300\ fF$, to compute the opamp power in the front-end S/H, first stage, and final stages (i.e. stages 2 to 8) for the capshare ADC. Table 4.2 applies the capacitor values calculated above, $C_S = C_C = 420\ fF$, to compute

Table 4.1: Power of capshare pipeline ADC

Step	Parameter	S/H	1st stage	Stages 2-8	Reference
A	β	1	1/3	1/2	Sections 3.2.2, 3.2.3.2, & 3.2.3.1
B	C_L [fF]	390	610	540	Equations 3.9, 3.16, & 3.13
C	GBW [MHz]	72.8	218	146	Equation 3.4
D	g_{m1} [μA/V]	137	412	274	Equation 3.20
E	g_{m12} [mA/V]	1.11	1.56	1.42	Equation 3.22
F	I_{B1} [μA]	20.6	61.8	41.2	Equation 3.24
G	I_{B2} [μA]	83.4	117	107	Equation 3.24
H	I_{total1} [μA]	187	296	255	$I_{B1} + 2I_{B2}$

the opamp power in the front-end S/H and first stage for the regular ADC. The opamp power for stages 2 to 8 are the same as in the capshare ADC so only the steps for the S/H and first stage are presented. As you can see from comparing Step H in Tables 4.1 and 4.2, the capshare ADC has a 42.5% power savings in the front-end S/H.

Earlier in this section, the assumption $C_C = C_S$ was made. Since $2C_{n1} = 300\ fF$, Equation 3.26 can be rewritten as:

$$300\ fF < Cc < C_L/2 \tag{4.7}$$

By plugging the values of C_L in Tables 4.1 and 4.2 into Equation 4.7, it can be concluded

Table 4.2: Power of regular pipeline ADC

Step	Parameter	S/H	1st stage	Reference
A	β	1	1/2	Sections 3.2.2 & 3.2.3.1
B	C_L [fF]	810	615	Equations 3.9 & 3.13
C	GBW [MHz]	72.8	146	Equation 3.4
D	g_{m1} [μA/V]	192	384	Equation 3.20
E	g_{m12} [mA/V]	1.98	1.57	Equation 3.22
F	I_{B1} [μA]	28.8	57.6	Equation 3.24
G	I_{B2} [μA]	148	118	Equation 3.24
H	I_{total2} [μA]	325	294	$I_{B1} + 2I_{B2}$
I	$I_{total1} : I_{total2}$ [%]	57.5	101	I_{total1}/I_{total2}

that $C_C = 300fF$ for the capshare ADC and $C_C = 420fF$ for the regular ADC are acceptable choices.

4.2.2.4 Conclusions

The design comparison concludes the following:

- For the same dynamic range and settling performance, the capshare ADC consumes 42.5% less power in the front-end S/H.
- The power consumed in the first stage is approximately the same in each design.
- Stages 2 to 8 consume the majority of total power.

Even though there is significant power savings in the front-end S/H, the total power of the capshare ADC is only 5.7% less. This is because stages 2 to 8, which are the same in each design, consume the majority of total power. A higher ADC resolution and proper scaling of stages 2 to 8 would make the front-end power consumption more dominant; however, as mentioned above, this thesis concentrates on the front-end design.

Because $C_{n2} = 150\ fF$ and $C_{iT,subADC} = 240\ fF$, they make up a large portion of C_L in Tables 4.1 and 4.2. Therefore, if there was an opportunity to revise the design, the opamp power consumption in every stage would be decreased by reducing $C_{iT,subADC}$ and C_{n2}. This would be particularly effective in reducing the power consumption in the

front-end S/H as it is loaded by only $C_{iT,subADC}$ and C_{n2}. $C_{iT,subADC}$ can be made smaller by making C_{ref} in Figure 3.4 smaller. C_{ref} was set to be 60 fF, which as mentioned, is the minimum size MIM capacitor allowed. A smaller effective capacitance could have been created by putting multiple 60 fF MIM capacitors in series; however, this was not realized until after the design.

4.3 Circuit Implementation

This section describes the implementation of the main circuit blocks within the regular ADC (Figure 4.6) and the capshare ADC (Figure 4.7). The two designs are implemented in IBM 0.13 μm technology.

4.3.1 General Description

The ADCs are designed to operate at a sampling frequency of 20 MS/s and a resolution of 10 bits. As discussed in Section 4.2.1, each pipeline ADC consists of a front-end S/H, 8 pipeline stages, and a final 2-bit flash-ADC. Each pipeline stage contains a 1.5-bit sub-ADC that produces 2 output bits and, as mentioned in Section 3.2.1, charge-distribution comparators are used in the sub-ADC. The final flash-ADC generates the final 2 bits. A total of 18 bits are generated for each input sample. The bits are re-timed on-chip using flip-flops so that all the bits corresponding to a specific input sample are aligned in time. DEC is performed off-chip to process the 18 bits into the expected 10-bit output.

A wide-swing cascode current mirror [20] is used as the bias-generation circuit to generate vb2/vb3/vb4 in each of the three opamps as shown in Figure 3.2. An off-chip potentiometer is connected to an on-chip diode-connected P-type Metal Oxide Semiconductor (PMOS) transistor to set the bias current. Three node voltages from each opamp are multiplexed off-chip to monitor each opamp's operating point. The bias point is tuned via the potentiometer until the node voltages are near the simulated values. Two additional diode-connected PMOS transistors are connected to potentiometers

to directly generate vb5 and the bias voltage for the opamps' second-stage CMFB.

A non-overlapping clock generation circuit [21] is used to generate non-overlapping and advanced clocks. Advanced clocks are used in the MDAC to prevent charge-injection errors.

Three power domains (analog, digital, and buffer) are used to reduce power-supply noise for the sensitive analog blocks. The analog supply, vdd_a, powers the MDACs, the opamps, the bias-generation circuit, and the comparators. The digital supply, vdd_d, powers the digital logic in the sub-ADC and the clock-generation circuitry. The buffer supply, vdd_{buf}, powers the chip-level input/output buffers, ESD diodes, and the digital retiming circuitry. The circuits powered by vdd_{buf} are excluded from the FOM calculations. The fabricated chip uses a supply voltage of 1.2 V.

Two differential reference signals, $V_{ref} = 0.8\ V$ and $V_{ref,half} = 0.4\ V$, and a single-ended CM voltage, $V_{cm} = 0.6\ V$, are provided from off-chip and used in the MDAC and sub-ADC. The capshare ADC operates at a full-scale range of $2V_{ref} = 1.6\ V_{PP}$.

4.3.2 Two-Stage Opamp

This section describes the opamps used in the regular and capshare ADC. In each design, three different opamps are needed for each of the three blocks (i.e. the S/H, the first pipeline stage, and stages 2 to 8). Figure 4.8 from [22] shows the two-stage opamp that is used. As shown in Figure 3.2, two CMFB signals are needed to control the CM output of the first and second stage. In the two-stage opamp in Figure 4.8, transistors M3 and M4 in the first stage are split and connected in a cross-coupled configuration. The cross-coupled connections of M3A/M3B and M4A/M4B reduce the CM gain while maintaining high differential gain in the first stage so that the CM output in the first stage can be set without CMFB. The switched-capacitor CMFB circuit in [23] is used to control the CM output of the second stage.

Table 4.3 summarizes the width, W, and length, L, of the transistors in each opamp

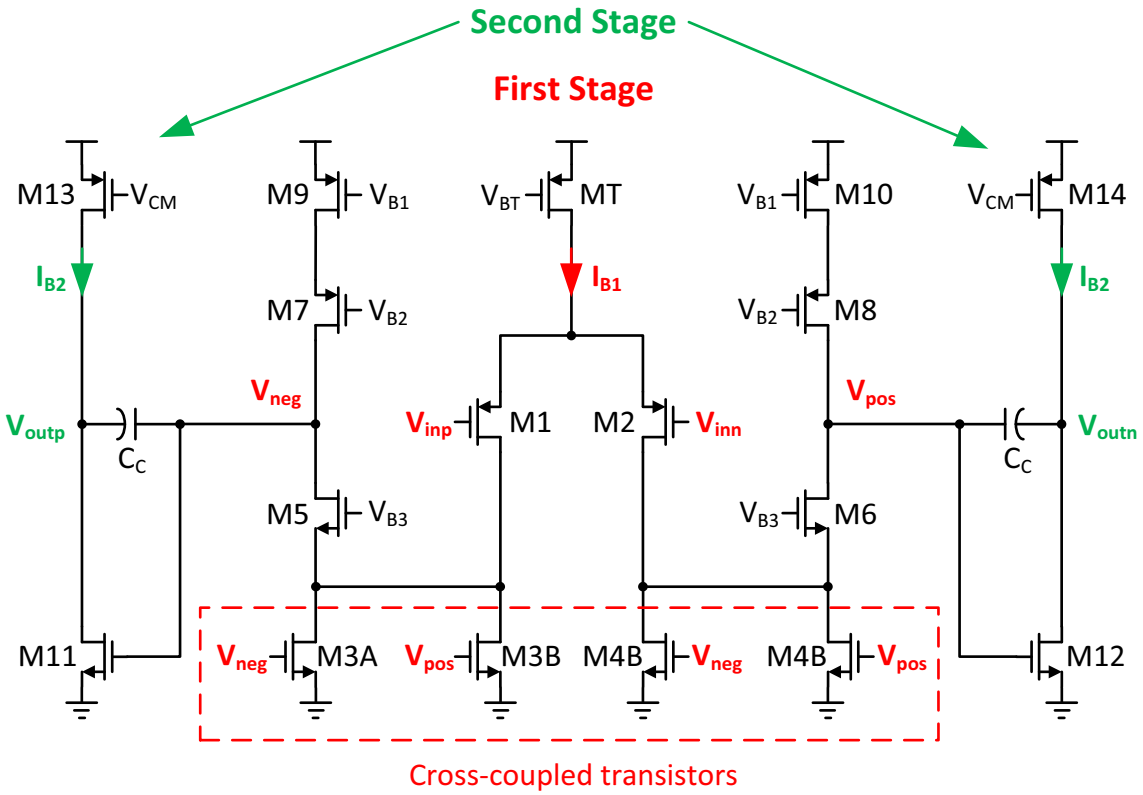

Figure 4.8: Two-stage opamp with with CMFB on second stage only

for the capshare ADC. Table 4.4 summarizes the width and length of the transistors in

Table 4.3: Transistor sizes in W/L $[\mu m/\mu m]$ within capshare ADC opamps

Transistor	capshare S/H	capshare 1st stage	Stage 2 to 8
MT	20 / 0.4	52 / 0.4	36 / 0.4
M1, M2	6 / 0.5	15 / 0.25	9 / 0.25
M3A/B, M4A/B	6 / 0.5	10 / 0.5	12 / 0.5
M5, M6	10 /0.5	15 / 0.5	15 / 0.5
M7, M8	15 / 0.5	25 / 0.5	25 / 0.5
M9, M10	12 / 0.5	15 / 0.5	15 / 0.5
M11, M12	36 / 0.6	42 / 0.6	24 / 0.4
M13, M14	120 / 0.6	156 / 0.6	114 / 0.6

each opamp for the regular ADC. Table 4.5 lists the loop gain, phase margin, and f_{3dB} of the opamp closed-loop circuit in the capshare ADC. Table 4.6 lists the loop gain, phase margin, and f_{3dB} of each opamp closed-loop circuit in the regular ADC. As expected, the f_{3dB} for both designs are near 70 MHz.

Table 4.4: Transistor sizes in W/L [$\mu m/\mu m$] within regular ADC opamps

Transistor	S/H	1st stage	Stage 2 to 8
MT	40 / 0.4	48 / 0.4	36 / 0.4
M1, M2	8 / 0.5	12 / 0.25	9 / 0.25
M3A/B, M4A/B	12 / 0.5	15 / 0.5	12 / 0.5
M5, M6	15 /0.5	20 / 0.5	15 / 0.5
M7, M8	20 / 0.5	30 / 0.5	25 / 0.5
M9, M10	15 / 0.5	20 / 0.5	15 / 0.5
M11, M12	90 / 0.6	66 / 0.6	24 / 0.4
M13, M14	252 / 0.6	186 / 0.6	114 / 0.6

Table 4.5: $A\beta$, phase margin, and f_{3dB} of opamp closed-loop circuits in capshare ADC

	S/H	1st stage	2nd to 8th stage
$A\beta$ [dB]	77.0	72.6	70.1
Phase Margin [°]	81.0	74.1	77.3
f_{3dB} [MHz]	70.0	70.1	72.4

Table 4.6: $A\beta$, phase margin, and f_{3dB} of opamp closed-loop circuits in regular ADC

	S/H	1st stage	2nd to 8th stage
$A\beta$ [dB]	79.0	74.8	70.1
Phase Margin [°]	79.6	79.2	77.3
f_{3dB} [MHz]	70.1	68.7	72.4

4.4 Schematic Simulations

In this section, transient simulations are performed on the capshare and regular ADC implemented in Section 4.3 to demonstrate the benefits of the capshare technique and to compare their relative performance. Neither design has been layed out at this point; however, layout parasitics are modeled in the schematic. The post-layout simulations of the capshare ADC are described in Chapter 5. A 1.6 V_{PP} full-scale input sine wave is used to test the ADCs up to the rated specification in Table 2.2. The SNDR/SNR values and spectrum plots are generated using a modified version of the delta-sigma tool box in [18]. Table 4.7 summarizes the SNDR and SNR measurements taken from the transient simulations. They are performed at a sampling frequency of 2 MHz and 20 MHz and at an input frequency of $61/128 \cdot f_S$. A ratio of 61/128 sets the input frequency near the Nyquist rate and maximizes the number of quantization levels used in the conversion. A total of 512 sample points are put through a Fast Fourier Transform (FFT) to produce a single SNDR/SNR pair.

Table 4.7: SNDR and SNR at $fin = 61/128 \cdot f_S$

Sim	Pipeline Type	f_S [MHz]	Thermal Noise	SNDR [dB]	SNR [dB]	ENOB [bits]
1	Regular	2	No	62.1	62.5	10.0
2	Capshare	2	No	60.9	61.4	9.82
3	Regular	2	Yes	59.0	59.4	9.51
4	Capshare	2	Yes	59.2	59.3	9.54
5	Regular	20	No	52.2	61.4	8.37
6	Capshare	20	No	54.1	62.6	8.70
7	Regular	20	Yes	51.2	58.2	8.21
8	Capshare	20	Yes	52.7	58.7	8.47

The following explains the simulations in Table 4.7:

- Simulations 1 and 2 are performed at a slower sampling rate so that each stage has more than enough time to settle. Hence, the opamp bandwidth should not affect SNDR. The simulation has thermal noise deactivated so thermal noise should not

limit SNR. The opamp's finite open-loop gain is the remaining source of error, which reduces the the settling accuracy and hence reduces SNDR. The regular pipeline has a higher SNDR because the regular ADC's front-end S/H and first stage have a higher loop gain as shown in Table 4.6.

- Simulations 3 and 4 introduce thermal noise. The noise causes the SNR to drop below the 10-bit level and hence, the dynamic range of the ADCs are thermally noise limited. Since both ADCs have the same SNR, the capshare ADC attains the same thermal noise performance but with a smaller C_S (300 fF as opposed to 420 fF). The thermal noise level is below the 10-bit level because the noise fraction, n_f, in the input-referred noise equation, Equation 3.10, is actually greater than 1, which was the initial assumption in Section 4.2.2.2. The noise fractions are actually measured to be between 3 and 5 in the two designs, which suggests that the opamp contributes significantly more noise than originally assumed.

- Simulations 5 and 6 test if the opamp bandwidth is sufficient to operate at $f_S = 20\ MHz$. The SNR is approximately at the 10-bit level because there is no thermal noise. However, the capshare ADC has a greater SNDR. This shows that the front-end S/H and first stage opamps in the regular ADC are not over-designed. Doing so would be unfair and give the appearance the capshare ADC has more power savings than it actually has.

- Simulations 7 and 8 are performed at $f_S = 20\ MHz$ with thermal noise activated. Except for the effects of post-layout parasitics and mismatch, all sources of error are included. The resulting SNDRs show that the capshare ADC performs slightly better than the regular ADC.

The spectral plots for Simulations 4 and 8 are shown in Figures 4.9(a) and 4.9(b), respectively. The dots on the spectrum mark the input signal's harmonic distortion terms that have been aliased back into the in-band frequency region. Their contribution to noise

is removed when calculating SNR, since SNR excludes distortion power. Next, Table

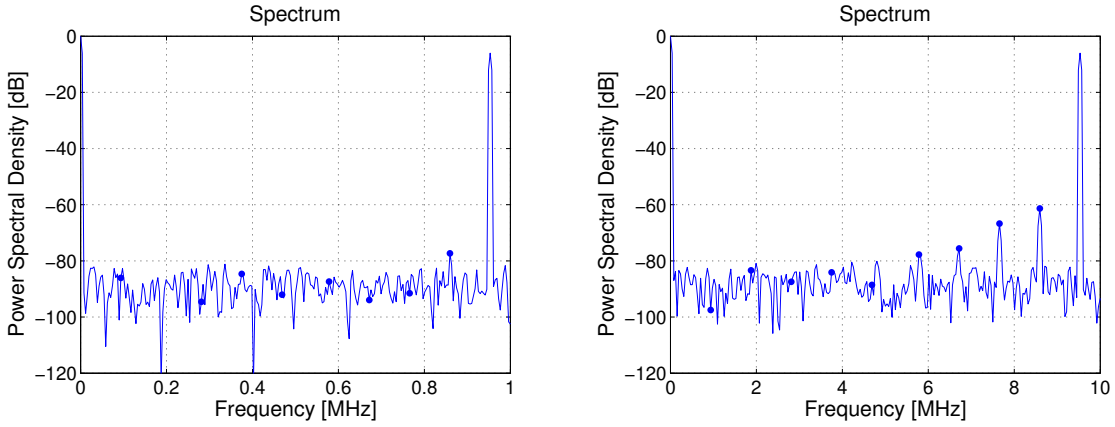

(a) Spectrum at $f_S = 2\ MHz$ in Sim 4 (b) Spectrum at $f_S = 20\ MHz$ in Sim 8

Figure 4.9: Spectrum of Simulations 4 and 8 in Table 4.7

4.8 and 4.9 present the power consumption of the front-end S/H, first stage, and final stages (i.e. stages 2 to 8). The digital power is consumed by the digital logic in the sub-ADC and the clock-generation circuit. The analog power is consumed by mainly the opamps. Table 4.8 presents the power measurements from Simulations 3 and 4, and Table 4.9 presents the power measurements from Simulations 7 and 8. As you can see, at

Table 4.8: Power consumption of capshare versus regular ADC at $f_S = 2\ MHz$

Pipeline	Stage(s)	I_{analog} [μA]	P_{analog} [μW]	$I_{digital}$ [μA]	$P_{digital}$ [μW]
Regular	S/H	520.4	624.5	0	0
	1st	474.8	569.8	0.5	0.6
	2 to 8	323.6	388.3	0.5	0.6
	All	3812	4574	31.3	37.6
Capshare	S/H	313.8	376.6	0	0
	1st	440.3	528.4	0.6	0.7
	2 to 8	323.3	388.0	0.5	0.6
	All	3571	4285	32.7	39.2

$f_S = 20\ MHz$, the S/H in the capshare ADC consumes 39% less power than the S/H in the regular ADC, which is close to the 42.5% estimated in Section 4.2.2.3. The capshare ADC's first stage consumes approximately the same power as the regular ADC's first stage, which again is what was estimated in Section 4.2.2.3. Therefore, Tables 4.7 to 4.9

Table 4.9: Power consumption of capshare versus regular ADC at $f_S = 20\ MHz$

Pipeline	Stage(s)	I_{analog} [μA]	P_{analog} [μW]	$I_{digital}$ [μA]	$P_{digital}$ [μW]
Regular	S/H	514.4	617.3	0	0
	1st	478.1	573.7	5.0	6.0
	2 to 8	329.2	395.0	4.6	5.5
	All	3854	4625	310	372
Capshare	S/H	314.9	377.9	0	0
	1st	449.8	539.8	5.8	7.0
	2 to 8	326.4	391.7	4.7	5.6
	All	3615	4338	323	388

show that the capshare ADC can attain the same dynamic range as a regular ADC but with a smaller C_S and with less power. At the targetted specifications of $f_{in} = 9.53\ MHz$ and $f_S = 20\ MHz$, the capshare ADC acheives a FOM of 0.7 pJ/step while the regular ADC has a FOM of 0.88 pJ/step. As an aside, from comparing Tables 4.8 and 4.9, it is evident that digital power scales with f_S, while analog power, the dominant source of power consumption, does not.

Chapter 5

Experimental Results

This chapter will presented the experimental results of the capshare pipeline ADC. In Section 5.1, the experimental setup and fabricated design is shown. The measurement results are then presented in Section 5.2. Finally, a 6 dB degradation in dynamic range is debugged with post-layout simulations in Section 5.3.

5.1 Test Setup

Figure 5.1 shows the setup used to test the capshare ADC. The Hewlett Packard (HP) 8064A function generator generates a single-ended sine wave, which after going through two Minicircuits LPFs to remove harmonic distortion, provides the single-ended input signal, *vin*. On the Printed Circuit Board (PCB), a single-ended to differential converter converts *vin* into a differential signal with a common mode voltage of 0.6 V, which then goes to the Device Under Test (DUT). The Stanford Research Systems DS360 signal generator was used to generate *vin* for input frequencies below 200 kHz. The clock signal for the DUT is provided by the Rohde & Schwarz SMT03 function generator. The power to the PCB is provided by the HP E3631A and the supply voltages for the DUT are generated by voltage regulators implemented on the PCB. The National Instruments PCI-6534 captures the 18 digital bits that are outputted from the DUT and as they are

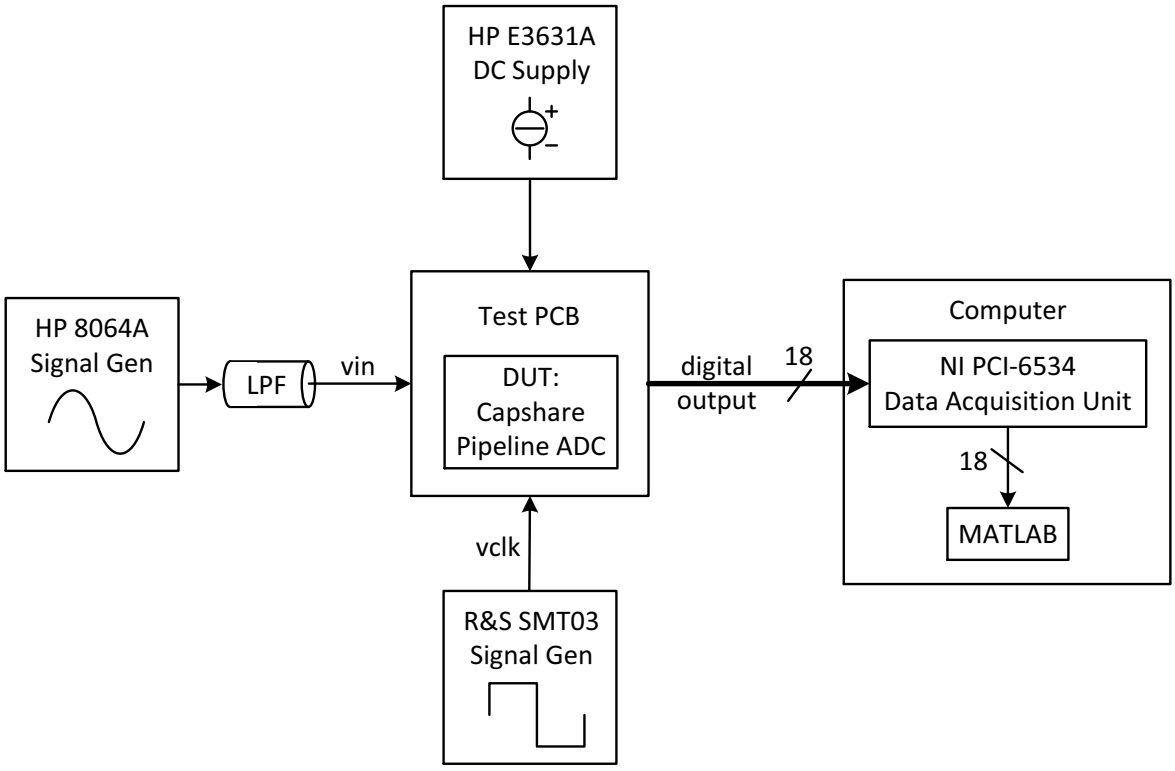

Figure 5.1: Test setup for capshare pipeline ADC

being captured on a computer, C++ code applies digital error correction to combine the 18 bits into the expected 10 bits. MATLAB code, using the delta-sigma toolbox in [18], processes the bits to compute the experimental results. A total of 2^{14} capture points are used in a single FFT and an average of 100 FFTs produce a single SNDR/SNR data point. Figure 5.2 shows the PCB used to test the capshare ADC. Figure 5.3 shows the die micrograph of the capshare ADC under test and Table 5.1 summarizes the function of each pin.

5.2 Experimental Data

The section presents the measurement results.

5.2.1 Measurement Results

Table 5.2 lists the experimental results at different f_{in} and f_S. The input sine-wave is at a full-scale range of 1.6 V_{PP}. Because the HP 8064A function generator produces harmonic

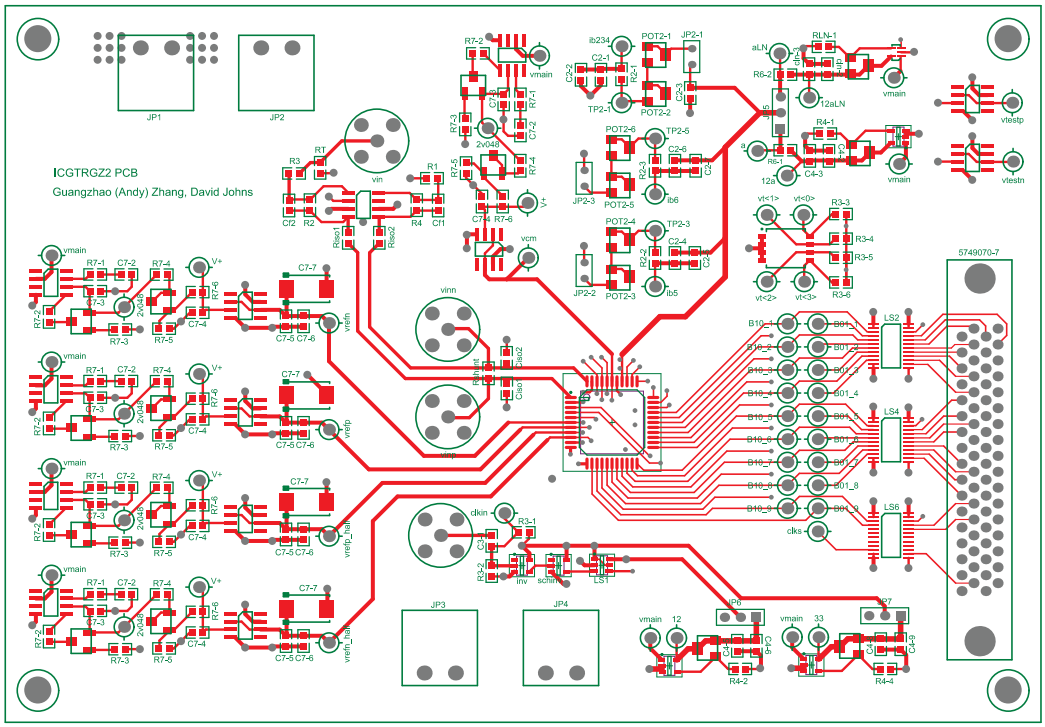

Figure 5.2: PCB used to test capshare pipeline ADC

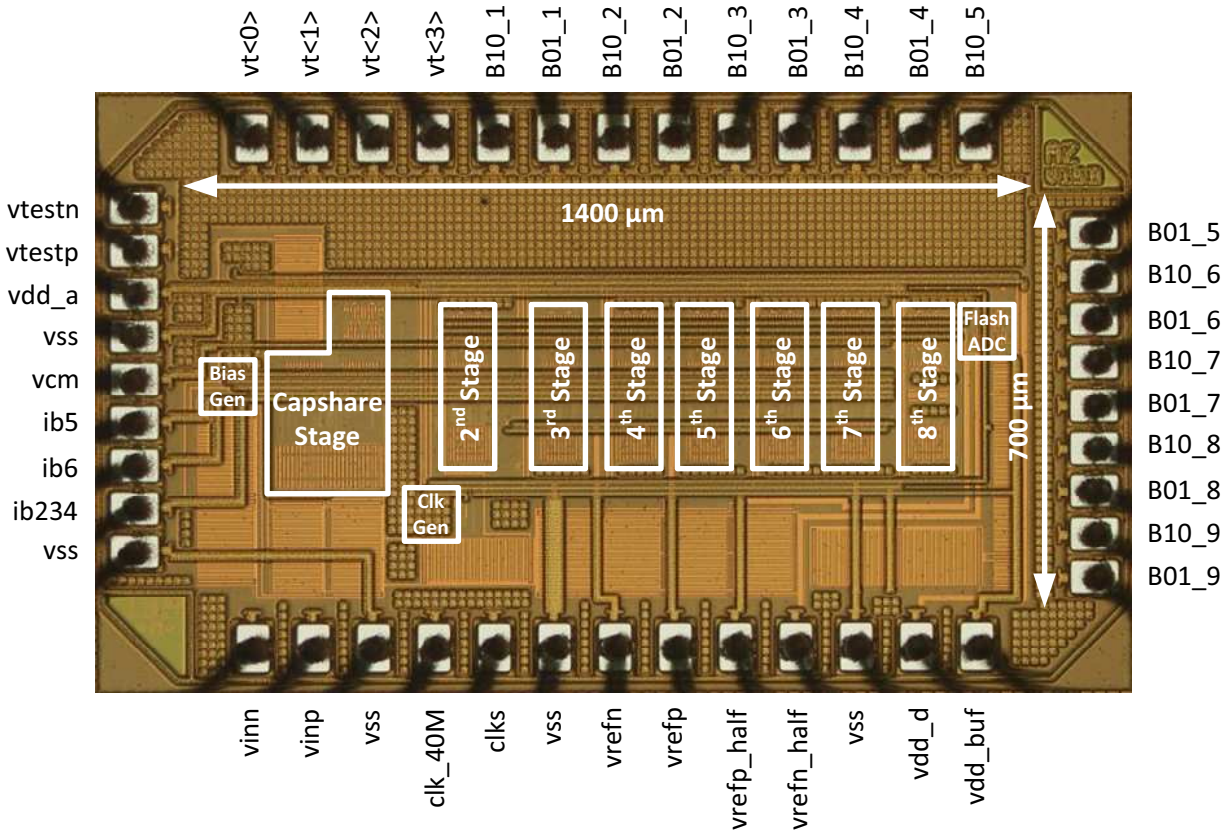

Figure 5.3: Die micrograph of pipeline ADC with front-end capacitor-sharing

Table 5.1: List of input/output pins for the DUT

Pin Name	Description
vtest(p/n)	differential test points to measure operating point within opamps
vt<0:3>	select lines to pick which test points to output on vtest
vin(p/n)	differential input signal
B10_1:9, B01_1:9	18 output bits aligned in time and ready for DEC
clks	output clock at frequency f_S, synchronous to output bits
clk_40M	input clock at frequency $2f_S$
vref(p/n)	differential reference (1 V, 0.2 V)
vref(p/n)_half	half of differential reference (0.8 V, 0.4 V)
vcm	single-ended common-mode reference (0.6 V)
ib234, ib5, ib6	input bias current nodes that go to off-chip potentiometers
vdd_(a/d/buf)	three supply domains (all 1.2 V)
vss	ground reference (0 V)

distortion terms that need to be removed by the Minicircuits LPFs, the choice of which f_{in} to take measurements at was limited by the LPFs that were available. When comparing

Table 5.2: Experimental results of capshare ADC measured at different f_{in} and f_S

Measurement	f_S [MHz]	f_{in} [MHz]	SNDR [dB]	SNR [dB]	ENOB [bits]	P_{analog} [mW]	$P_{digital}$ [μW]	FOM [pJ/step]
1	0.4	0.19	53.9	54.2	**8.67**	4.06	8.64	26.2
2	5	2.38	53.8	54.3	8.65	4.21	106	2.25
3	10	4.77	53.8	54.4	8.65	4.25	212	1.17
4	14	6.67	53.1	53.7	8.52	4.28	297	0.93
5	20	2.03	53.0	53.9	8.51	4.26	422	3.15
6	20	4.84	53.1	54.4	8.53	4.28	422	1.32
7	20	9.53	53.0	54.1	8.51	4.30	423	0.68
8	20	12.0	51.4	53.5	8.25	4.32	422	0.65
9	20	13.6	51.0	52.6	8.18	4.31	423	**0.60**
10	24	11.4	51.9	53.6	8.33	4.34	507	0.66
@ $v_{dd} = 1.1\ V$	20	9.53	51.4	54.2	8.25	3.49	356	0.67
@ $v_{dd} = 1.3\ V$	20	9.53	53.5	54.2	8.60	5.30	503	0.80

the SNR values in Tables 5.2 to the thermally noise limited SNR values in Table 4.7, it is evident that, at least experimentally, the dynamic range is being limited by something other than thermal noise. To see if the dynamic range drop is caused by an experimental artifact, the experimental results are compared to simulation results from simulating the capacitor-only post-layout capshare ADC with thermal noise activated. Figures 5.4(a)

and 5.4(b) show plots of SNDR vs f_S and SNR vs f_S, respectively, at $fin = 61/128 \cdot f_S$. Figure 5.4(c) shows the SNDR vs f_{in} at $f_S = 20\ MHz$. One evident difference is that the

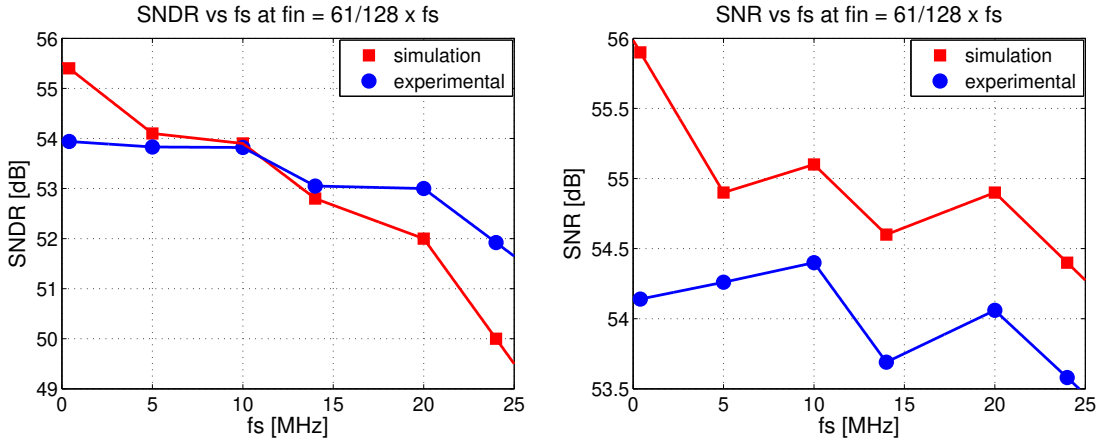

(a) SNDR vs f_S at $fin = 61/128 \cdot f_S$

(b) SNR vs f_S at $fin = 61/128 \cdot f_S$

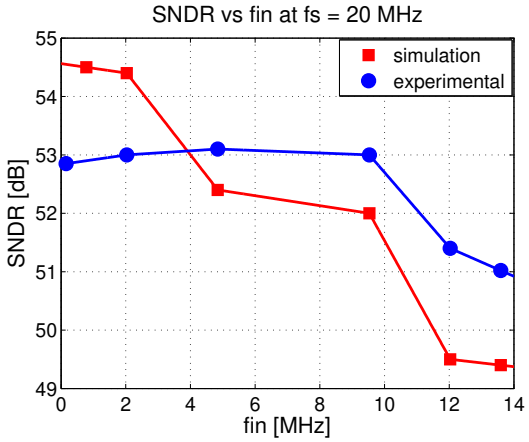

(c) SNDR vs f_{in} at $f_S = 20\ MHz$

Figure 5.4: Comparison of simulated and measured SNDR and SNR

simulated SNR is higher than the measured SNR. This is mainly caused by digital noise from the output coupling back into the sensitive inputs of the DUT. The voltages level of the output bits coming from the DUT have to be level-shifted to TTL-compatible levels before the NI PCI-6534 can capture them. A total of 18 level-shifters are used and they produces a considerable amount of digital noise. The measured SNR decreases as the voltage level that the level-shifter shifts to is increased, suggesting digital noise is coupling back into the input. The digital noise can be reduced by lowering the level-

shifter voltage; although, it needs to be at least 2.4 V to provide TTL-compatible signals. However, even if all the sources of error introduced by the experimentation is ignored, the dynamic range in the post-layout simulations with thermal noise is still approximately 4 dB less than in the schematic simulations with thermal noise in Table 4.7. This dynamic range degradation is revisited in Section 5.3.

Figure 5.5(a) plots the spectrum at $f_S = 20\ MHz$ and $f_{in} = 9.53\ MHz$. At this operating point, it achieves a FOM of 0.68 pJ/step as shown in Measurement 7 in Table 5.2. This is very close to the simulated result of 0.7 pJ/step from the schematic simulations in Section 4.4. Again, the dots on the spectrum mark the input signal's harmonic distortion terms that have aliased back into the in-band region. Figure 5.5(b) shows how the SNDR increases with the input amplitude, reaching it's maximum of 53 dB at the full-scale input. Next, the FOM in measurement 7 is compared to the FOM in other measurements.

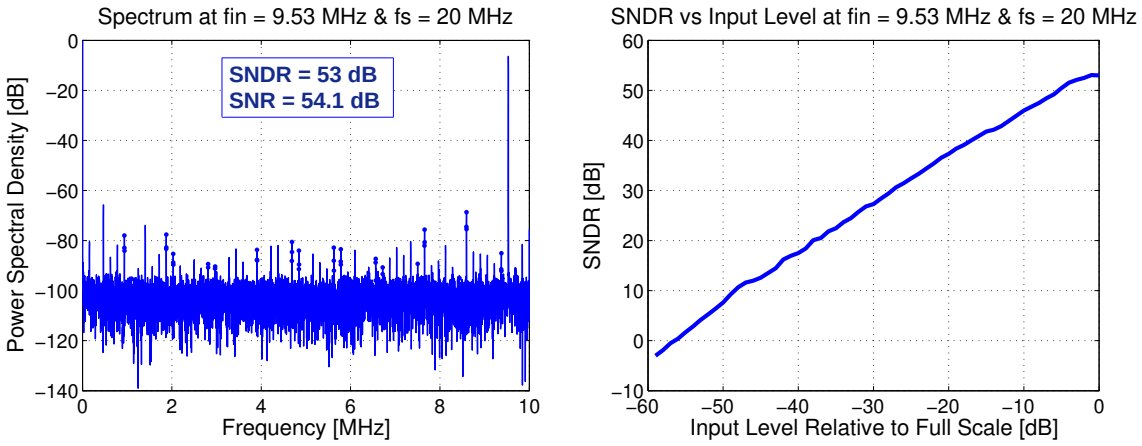

(a) Power Spectral Density

(b) SNDR vs Input Level

Figure 5.5: Spectrum and SNDR vs Input Level at $f_{in} = 9.53\ MHz$ and $f_S = 20\ MHz$

When the capshare ADC is sub-sampling, in measurements 8 and 9, it achieves a FOM as low as 0.6 pJ/step. This is because the increase in input frequency is greater than the decrease in SNDR. This likely would not be possible if a front-end S/H was not used.

In measurement 5, the input frequency is well below the Nyquist rate. Because the power consumption does not scale with the input frequency, the FOM is significantly higher than in measurement 7. Figures 5.6 plots the spectrum for Measurement 5.

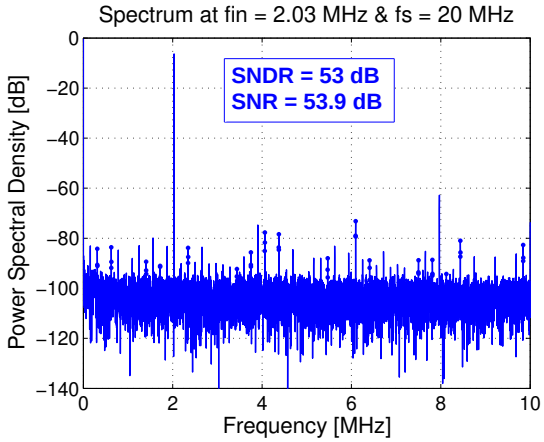

Figure 5.6: Spectrum at $f_{in} = 2.03\ MHz$ and $f_S = 20\ MHz$

The final two measurement in Table 5.2 consider how changing the supply voltage effects performance. The reference voltages are scaled accordingly as the supply voltage is changed. When the supply is raised to $v_{dd} = 1.3\ V$, SNDR increases to 53.5 dB; however, because the power consumption increases even more, the FOM increases to 0.8 pJ/step. When the supply is reduced to $v_{dd} = 1.1\ V$, the drop in SNDR balances the decrease in power consumption and hence, the FOM is approximately the same as in measurement 7. Figure 5.7(a) and 5.7(b) plots the spectrum for $v_{dd} = 1.1\ V$ and $v_{dd} = 1.3\ V$, respectively.

5.2.2 Differential and Integral Non-Linearity

Because the DC approach to testing static Differential Non-Linearity (DNL) and Integral Non-Linearity (INL) is difficult to perform, the histogram-based approach is adopted in this work. MATLAB code from [24] is used to generate the DNL and INL plots in Figure 5.8(a) and 5.8(b), respectively. The test was conducted at $f_{in} = 2.03\ MHz$ and $f_S = 20\ MHz$, and a total of 2^{21} data points are used to generate each plot. The DNL and INL have a peak-to-peak swing of 1.26 LSB_{PP} and 1.73 LSB_{PP}, respectively.

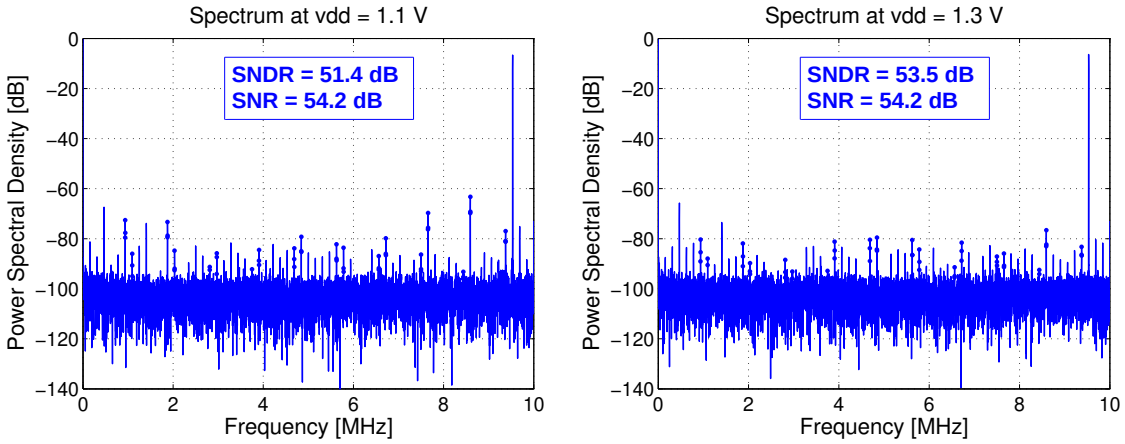

(a) Spectrum at $v_{dd} = 1.1\ V$ (b) Spectrum at $v_{dd} = 1.3\ V$

Figure 5.7: Spectrums at $f_{in} = 9.53\ MHz$ and $f_S = 20\ MHz$

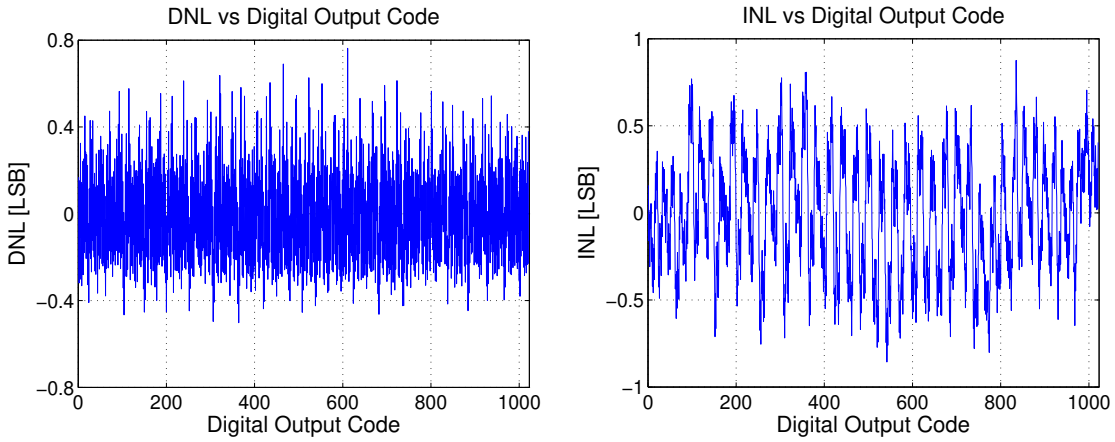

(a) DNL (b) INL

Figure 5.8: DNL and INL using histogram test at $fin = 2.03\ MHz$ and $f_S = 20\ MHz$

5.2.3 Conclusions

In conclusion, the capshare ADC has experimentally shown that it can function at the targeted specifications in Section 4.3. At $f_{in} = 9.53\ MHz$ and $f_S = 20\ MHz$, the DUT consumes 4.7 mW of power and achieves a ENOB of 8.5 bits and a FOM of 0.68 pJ/step. The DUT can achieve an ENOB as high as 8.67 bits at $f_S = 0.4\ MHz$ and a FOM as low as 0.6 pJ/step when sub-sampling at $f_{in} = 13.6\ MHz$ and $f_S = 20\ MHz$. The FOM is slightly higher than the targeted 0.5 pJ/step; however, because stages 2 to 8 were not optimized, the back-end consumes significantly more power than it should. If the back-end was scaled for optimum power consumption, a lower FOM could be achieved. Table 5.3 compares the capacitor-sharing front-end pipeline ADC in this work to the regular pipeline ADC that was designed in Section 4.2 and to two other comparable designs. Since the pipeline ADC in [1] is adaptable over a range of sampling frequencies, only the performance at a single f_S is presented in the table.

Table 5.3: Comparing this work to other 1.5-bit/stage 10-bit pipeline ADCs

	Capshare ADC	Regular ADC	[25], [6]	[1]
Technology	0.13μm	0.13μm	90nm	0.18μm
Supply [V]	1.2	1.2	1.2	1.8
Active area [mm^2]	0.98	N/A	0.3	1.21
Input swing [V_{PP}]	1.6	1.6	1.1	1.6
f_S [MHz]	20	20	12	@24
f_{in} [MHz]	9.53	9.53	4.87	@10.3
ENOB [bits]	8.5	8.2	8.4	9.2
S/H Power [μW]	**378**	617	750	0
Total Power [mW]	4.7	5.0	3.3	12.6
FOM [pJ/step]	0.68	0.88	0.97	1.02

5.3 Dynamic Range Degradation

This section debugs the capacitor-only post-layout capshare ADC to discover why the dynamic range decreases when capacitor layout parasitics are introduced. Table 5.4 shows the SNDR/SNR when portions of the post-layout design are extracted and simulated with

thermal noise deactivated. As you can see from comparing Simulation 1 to 2, the SNR drops from 61.4 dB to 55.5 dB once the entire design is extracted and the post-layout parasitic capacitors are introduced. Because the SNR is below the thermal noise level of 59.3 dB, as shown in Table 4.7, the capshare ADC is no longer thermally noise limited.

Table 5.4: SNDR and SNR at $f_{in} = 61/128 \cdot f_S$ after post-layout extraction

Sim	Extraction	Description	f_{in} [MHz]	SNDR [dB]	SNR [dB]
1	None	schematic	2	60.9	61.4
2	Complete chip	post-layout	2	54.7	55.5
3	Complete chip	calibrate final stages	2	56.4	56.6
4	Capshare stage	post-layout	2	55.9	56.5
5	Capshare stage	post-layout	20	52.7	56.5
6	Capshare stage	ping-pong off	20	56.3	60.6

When there is parasitic capacitance across the sampling capacitors, as shown in Figure 5.9, it produces a constant gain error in the MDAC's stage gain. As a result, the ADC's linearity suffers and the dynamic range decreases. If the gain constants (G_1 to G_9) in the

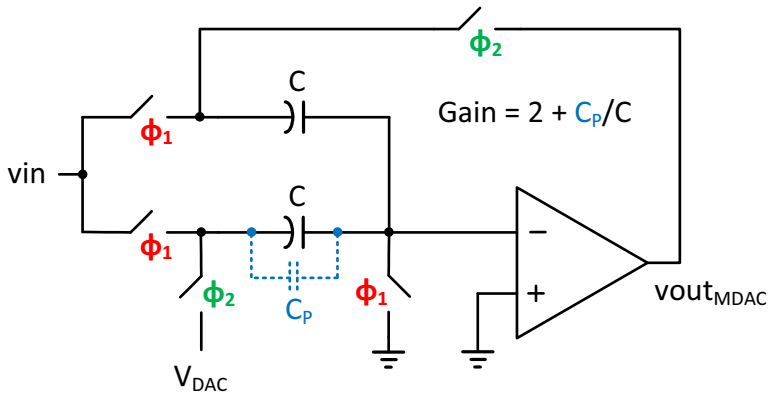

Figure 5.9: MDAC gain error caused by parasitics across sampling capacitor

digital domain, as shown in Figure 2.1, are corrected to reflect the new MDAC stage gains, the drop in SNDR and SNR can be calibrated out. Simulation 3 successfully calibrates out the gain error in stages 2 to 8. In Simulation 4 and 5, only the capacitor-sharing stage (Figure 4.4) is extracted. Since the SNR remains the same as in Simulation 3, the dynamic range problem is narrowed down to within the capshare stage. Calibration does

not improve the dynamic range in Simulations 4 and 5, which suggests that the problem within the capshare stage is not due to a constant gain error.

Because the layout of the capshare stage is very compact, there is parasitic capacitance between the two ping-pong paths. In Φ_1, phase A (B) is sampling the next input at the same time phase B (A) is generating the residue, $vout_{MDAC}$, based on the previously sampled input. Consequently, charge from the input sampling is coupling into the capacitors used to generate the residue, creating an error in $vout_{MDAC}$. This error cannot be calibrated out because the amount of charge coupling varies based on the input voltage level sampled. The coupling should stop if one of the ping-pong phases (A or B) is shut down as that would cause the sampling of the input and generation of $vout_{MDAC}$ to occur on opposite clock cycles. Simulation 6 shows that the dynamic range improves to near the 10-bit level when the ping-ponging is turned off, indicating that the coupling within the capshare stage was the source of the problem. The spectral plots in Figures 5.10(a) and 5.10(b) show the spectrum when ping-ponging is turned on and when it's turned off, respectively.

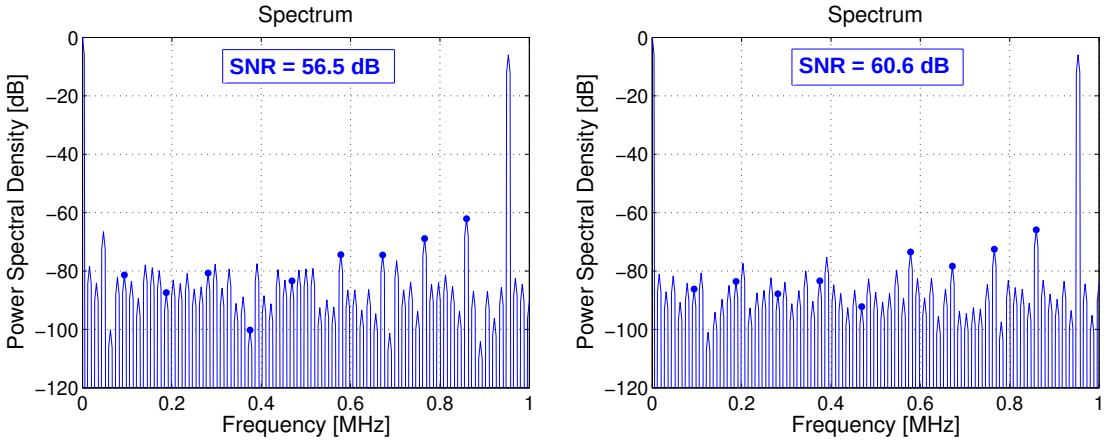

(a) Spectrum when ping-ponging is on

(b) Spectrum when ping-ponging is turned off

Figure 5.10: Spectrum of Simulations 5 and 6 in Table 5.4

The transistors and sampling capacitors in each of the two ping-pong paths were layed out very compactly to achieve good matching. Consequently, it created cross coupling

between the two paths, which limited the dynamic range to below the thermal noise level. If the capshare design could be redesigned, the layout of ping-pong phase A and phase B would be separated. This effectively divides the layout of the S/H and the first stage into two different blocks, as is typically done in a regular pipeline ADC.

Chapter 6

Conclusions and Future Work

This chapter summarizes the main conclusions in this thesis and discusses the areas for future work.

6.1 Conclusions

This thesis presents a novel front-end capacitor-sharing technique that significantly reduces the power consumption in the front-end S/H. The goal of this thesis was a proof-of-concept for the technique in pipeline ADCs and not to attain the lowest FOM. The application is for wireless radio receivers. A comparison was conducted between a capshare pipeline ADC and a regular pipeline ADC to show the benefits of the technique. As a result, the capshare technique achieved a power savings of 39% in the front-end S/H. The capshare ADC design was fabricated in IBM 0.13 μm technology and then tested. At a input frequency of 9.53 MHz and a sampling rate of 20 MS/s, which are the target specifications, the DUT consumes 4.7 mW of power and achieves a ENOB of 8.5 bits and a FOM of 0.68 pJ/step. Furthermore, it was tested to have an ENOB as high as 8.67 bits, at low sampling speeds, and a FOM as low as 0.6 pJ/step when sub-sampling at 20 MS/s. Finally, post-layout simulations discovered that cross coupling within the ping-pong circuit in the front-end capshare stage caused the dynamic range

to drop below the thermal noise level.

6.2 Future Work

If the capshare ADC in this work could be redesigned, the following changes would be made to improve the front-end performance:

- Because the two ping-pong paths within the capshare stage were layed out compactly, there was cross coupling between the two paths, which limited the dynamic range to below the thermal noise level. If the design could be revised, the layout of each parallel ping-pong phase would be separated so that the layout of the front-end S/H and the first stage are divided into two separate blocks.

- Because the sampling capacitance of the sub-ADC makes up a large portion of the load capacitance for each opamp, the opamp power consumption in every stage could be decreased by making the sub-ADC sampling capacitors C_{in} smaller than 60 fF. This would effectively reduce the power consumption in the front-end S/H as it is loaded by only the sub-ADC and it's opamp parasitic capacitance. If the design could be revised, the sampling capacitors in the sub-ADC would be implemented by multiple 60 fF minimum-sized MIM capacitors placed in series to create a smaller effective capacitor.

The goal of this work was a proof-of-concept for front-end capacitor-sharing. Hence, it concentrated on the front-end design and did not aim for the lowest FOM. However, to realize the full potential of the power-saving technique, the front-end must dominate total power consumption. This can be achieved by:

- Optimizing and scaling the back-end sampling capacitors in stages 2 to 8 for low-power consumption.

- Increasing the ADC's resolution so the sampling capacitors in the system are larger and hence, there is more room for scaling.

- Resolve more bits per stage, which would relax the power requirements at the back-end and increase the design complexity and power requirements of the front-end.

A potential future project would be to design an 11-bit pipeline ADC that resolves 2.5 bits/stage, operates at 25 MS/s, and applies front-end capacitor-sharing between the front-end S/H and the 2.5-bit first stage.

References

[1] I. Ahmed and D.A. Johns. A high bandwidth power scalable sub-sampling 10-bit pipelined adc with embedded sample and hold. *Solid-State Circuits, IEEE Journal of*, 43(7):1638 –1647, july 2008.

[2] D. Johns and K.W. Martin. *Analog integrated circuit design*. John Wiley & Sons, 1997.

[3] S.I. Ahmed. *Pipelined ADC Design and Enhancement Techniques*. Analog Circuits and Signal Processing. Springer, 2010.

[4] S.H. Lewis and P.R. Gray. A pipelined 5-msample/s 9-bit analog-to-digital converter. *Solid-State Circuits, IEEE Journal of*, 22(6):954 – 961, dec 1987.

[5] H. Balasubramaniam and K. Hofmann. A 10 bit, 1.5b/stage pipeline adc using a fully differential current conveyor with foreground calibration. In *Signals, Circuits and Systems (ISSCS), 2011 10th International Symposium on*, pages 1 –4, 30 2011-july 1 2011.

[6] R. Wang, K. Martin, D. Johns, and G. Burra. A 3.3 mw 12 ms/s 10b pipelined adc in 90 nm digital cmos. In *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, pages 278 –279 Vol. 1, feb 2005.

[7] Young-Deuk Jeon, Seung-Chul Lee, Kwi-Dong Kim, Jong-Kee Kwon, Jongdae Kim, and Dongsoo Park. A 5-mw 0.26-mm2 10-bit 20-ms/s pipelined cmos adc with

multi-stage amplifier sharing technique. In *Solid-State Circuits Conference, 2006. ESSCIRC 2006. Proceedings of the 32nd European*, pages 544 –547, sept 2006.

[8] O.A. Adeniran and A. Demosthenous. A 19.5mw 1.5v 10-bit pipeline adc for dvb-h systems in 0.35 /spl mu/m cmos. In *Circuits and Systems, 2006. ISCAS 2006. Proceedings. 2006 IEEE International Symposium on*, page 4 pp., may 2006.

[9] Young-Deuk Jeon, Seung-Chul Lee, Kwi-Dong Kim, Jong-Kee Kwon, and Jongdae Kim. A 4.7mw 0.32mm2 10b 30ms/s pipelined adc without a front-end s/h in 90nm cmos. In *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pages 456 –615, feb 2007.

[10] J. Pacheco, N. Paulino, M. Figueiredo, and J. Goes. Design of a 10-bit 40 ms/s pipelined adc using 1.5-bit with current-mode reference shifting. In *Mixed Design of Integrated Circuits and Systems (MIXDES), 2011 Proceedings of the 18th International Conference*, pages 299 –302, june 2011.

[11] I. Ahmed, J. Mulder, and D.A. Johns. A low-power capacitive charge pump based pipelined adc. *Solid-State Circuits, IEEE Journal of*, 45(5):1016 –1027, may 2010.

[12] P. Lajevardi, A.P. Chandrakasan, and Hae-Seung Lee. Zero-crossing detector based reconfigurable analog system. *Solid-State Circuits, IEEE Journal of*, 46(11):2478 –2487, nov 2011.

[13] Shenjie Wang, V. Maheshwari, and W.A. Serdijn. Instantaneously companding baseband sc low-pass filter and adc for 802.1 la/g wlan receiver. In *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, pages 2215 –2218, 30 2010-june 2 2010.

[14] S.Q. Malik and R.L. Geiger. Simultaneous capacitor sharing and scaling for reduced power in pipeline adcs. In *Circuits and Systems, 2005. 48th Midwest Symposium on*, pages 1015 –1018 Vol. 2, aug 2005.

[15] L. Sumanen, M. Waltari, V. Hakkarainen, and K. Halonen. Cmos dynamic comparators for pipeline a/d converters. In *Circuits and Systems, 2002. ISCAS 2002. IEEE International Symposium on*, volume 5, pages V–157 – V–160 vol.5, 2002.

[16] M. Gustavsson, J.J. Wikner, and N. Tan. *CMOS data converters for communications*. Kluwer international series in engineering and computer science. Kluwer Academic, 2000.

[17] Hao Meng, Jia Sun, and A. Paasio. A 1.2v 10bit 83msps pipeline adc in 130nm cmos. In *Microelectronics Electronics, 2009. PrimeAsia 2009. Asia Pacific Conference on Postgraduate Research in*, pages 177 –180, jan 2009.

[18] R. Schreier and G.C. Temes. *Understanding delta-sigma data converters*. IEEE Press, 2005.

[19] W.M.C. Sansen. *Analog design essentials*. Number v. 1 in Analog circuits and signal processsing series. Springer, 2006.

[20] J.N. Babanezhad and R. Gregorian. A programmable gain/loss circuit. *Solid-State Circuits, IEEE Journal of*, 22(6):1082 – 1090, dec 1987.

[21] Andrew M. Abo. *Design for Reliability of Low-Voltage, Switched-Capacity Circuits*. PhD thesis, EECS Department, University of California, Berkeley, 1999.

[22] M. Dessouky and A. Kaiser. Very low-voltage digital-audio delta; sigma; modulator with 88-db dynamic range using local switch bootstrapping. *Solid-State Circuits, IEEE Journal of*, 36(3):349 –355, mar 2001.

[23] R. Castello and P.R. Gray. A high-performance micropower switched-capacitor filter. *Solid-State Circuits, IEEE Journal of*, 20(6):1122 – 1132, dec 1985.

[24] Au and Hofner. Histogram testing determines dnl and inl errors, june 2003. Tutorial 2085.

[25] Robert Wang. A low-voltage low-power 10-bit pipeline adc in 90nm digital cmos technology. Master's thesis, University of Toronto, 2004.