PIPELINED ADC ENHANCEMENT TECHNIQUES

By

IMRAN AHMED

A THESIS SUBMITTED IN CONFORMITY WITH THE REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSPOHY

EDWARD S. ROGERS Sr. DEPARTMENT OF
ELECTRICAL AND COMPUTER ENGINEERING
UNIVERSITY OF TORONTO

PIPELINED ADC ENHANCEMENT TECHNIQUES

Ph.D., 2008

IMRAN AHMED

Edward S. Rogers Sr. Department of Electrical and Computer Engineering
University of Toronto

ABSTRACT

In this work three techniques to improve pipelined ADC performance with respect to linearity and power consumption are presented. The first technique enables rapid background digital correction of both DAC and gain errors in the multi-bit first stage of an 11-bit pipelined ADC. The proposed rapid calibration scheme enables significantly shorter automated test times in an industrial production environment, hence facilitates larger production throughput and thus increased cost efficiency. Measured results from a prototype fabricated in 1.8V 0.18µm CMOS show the technique achieves an improvement in linearity by more than 20dB within only 10^4 clock cycles in an 11-bit 45MS/s pipelined ADC - more than two orders of magnitude faster than previously published reports. The second technique develops a new MDAC topology which enables a pipelined ADC to be designed without a front-end sample-and-hold, and thus allows for significant power reduction. Unlike previous reports, the proposed topology does not require a carefully matched and/or time consuming layout to work properly. Measured results from a prototype fabricated in 1.8V 0.18µm CMOS show that better than 51dB SNDR can be achieved using the proposed approach for input frequencies higher than 267MHz for a 10-bit pipelined ADC with a maximum sampling rate of 50MS/s, while consuming 20% less power than a similar chip which

required a front-end sample-and-hold. The third technique facilitates large ADC power reduction by replacing all opamps in a 10-bit 50 MS/s pipelined ADC with source followers, and uses a novel fully-differential passive gain technique to obtain an MDAC gain near 2x. Measured results from a prototype in 1.8V 0.18μm CMOS show the 50MS/s ADC to achieve a peak SNDR/SFDR of 58.2dB/66dB while only consuming 9.9mW for a figure of merit of 0.3pJ/step.

ACKNOWLEDGEMENTS

Researching a thesis is a unique proposition. One is forced to look into the depths of the unknown and find an answer to a question that does not necessarily have an answer. In some cases your answer fits the question – in some cases your answer fits the question like a square peg in a round hole. Regardless of the madness, the journey of developing a thesis from abstract ideas to ultimately a functional prototype is truly a unique and completely enriching experience - an experience that I for one am tremendously thankful for and very fortunate to have undergone. Acknowledging specific people in the development of an abstract piece of art as a thesis is somewhat partial, as undoubtedly every person one interacts with during the course of a thesis in some shape or form impacts the work. There are few however who have helped this piece of abstract art take form. Firstly I must thank my supervisor, Professor David Johns. Without his aid in developing the focus of this work, and his invaluable suggestions and advice throughout the duration of this degree, this work would not have been possible. I am also thankful to Professor Ken Martin whose rigor and bold ideas significantly helped this work take shape. Next I am tremendously indebted to the aid and friendship of the everyone in BA5000 who in addition to helping me develop and refine my skills as a mixed-signal designer, have made my tenure as a PhD student at U of T, truly enriching and thoroughly enjoyable. I also thank the support of the team of excellent designers at Broadcom Netherlands, especially Jan Mulder and Klaas Bult, who in addition to providing a wealth of knowledge, have inspired me to be excited about the future in mixed signal circuit design. Of course one cannot accomplish anything in life without the unquestioned pillar of support one's family offers. To my family I dedicate this work, for whom without I would not be where I am today.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ADC	Analog to Digital Converter
CBSC	Comparator Based Switched Capacitor
CM	Common Mode
CMFB	Current Mode Feed Back
CMPS	Current Modulated Power Scaling
CMR	Common Mode Rejection
DAC	Digital to Analog Converter
dBFS	dB relative to Full Scale
DLL	Delay Locked Loop
DNL	Differential Non-Linearity
DNW	Deep N-Well
ENOB	Effective Number of Bits
FOM	Figure of Merit
IC	Integrated Circuit
INL	Integral Non-Linearity
KCL	Kirchhoff's Current Law
MDAC	Multiplying Digital to Analog converter
MIM	Metal-Insulator-Metal
PGA	Programmable Gain Amplifier
S/H	Sample-and-hold
SAR	Successive Approximation Register
SFDR	Spurious Free Dynamic Range
SNR	Signal to Noise Ratio
SNDR	Signal to Noise plus Distortion Ratio

CHAPTER 1: INTRODUCTION

1.1: OVERVIEW

The pipelined topology is a popular option for ADCs which require resolutions on the order of 8 to 14 bits and sampling rates between a few MS/s to hundreds of MS/s. The popularity of the topology can be attributed to its relatively simple and repetitive core structure, as well as a significant reduction in the number of comparators required to achieve a fixed resolution when compared to other Nyquist-rate data converters such as Flash, folding + interpolating, etc.. Pipelined ADCs are used in a variety of applications such as: mobile systems, CCD imaging, ultrasonic medical imaging, digital receivers, base stations, digital video (e.g. HDTV), xDSL, cable modems, and fast Ethernet [1]. With the use of pipelined ADCs in many consumer products, research in improving the performance of pipelined ADCs has attracted much attention over the past decade, where the most popular areas of research have been: linearity enhancement, and power reduction. Linearity enhancement has been an active area of research as with deeper sub-micron technology low intrinsic gain, low supply voltages, and device mismatch have made achieving very linear data converters (i.e. >10-bit linear) challenging using conventional pipelined ADC design techniques. Low power consumption in pipelined ADCs is motivated by the fact that many mobile systems use pipelined ADCs, where low power consumption enables increased battery life and thus increased user productivity. In wired systems where many ADCs can be integrated on-chip in parallel, large net power consumption can generate high amounts of heat requiring expensive packaging for heat dissipation. Thus techniques to reduce power consumption in pipelined ADCs enable more cost effective integrated circuits for both mobile and wired systems.

In this dissertation three novel techniques which advance the state of the art in linearity enhancement and power reduction in pipelined ADCs will be discussed. The following paragraphs briefly outline the contribution of the advancements.

The first contribution [2], [3] demonstrates a new topology to rapidly correct for both DAC and gain errors in the multi-bit first stage of an 11-bit pipelined ADC. Using a dual-ADC based approach the digital background scheme is validated with a proof-of-concept prototype fabricated in a 1.8V 0.18μm CMOS process, where the calibration scheme improves the peak INL of a 45MS/s ADC from 6.4 LSB to 1.1 LSB after calibration. The SNDR/SFDR is improved from 46.9dB/48.9dB to 60.1dB/70dB after calibration. Calibration is achieved in approximately 10^4 clock cycles.

The second contribution [4], [5] demonstrates a sub-sampled pipelined ADC architecture which has a power consumption which scales with down sampled bandwidth. The ADC uses a novel technique to save power by eliminating the front-end sample-and-hold without relying on a carefully matched layout. The technique allows for a power savings of >20% compared to a previous design. A method to improve the settling behavior of Rapid Power-on Opamps is also developed. Measured results in a 1.8V 0.18μm CMOS process verify the removal of the front-end sample-and-hold does not cause gross MSB errors for input frequencies higher than 267MHz. With f_s=50MS/s, for f_{in}=79MHz the SNDR is 51.5dB, and with f_s=4.55MS/s for f_{in}=267MHz the SNDR is 52.2dB.

The third contribution demonstrates very low power consumption in a 1.5-bit/stage 10-bit pipelined ADC by replacing power hungry opamps with simple source followers and a fully-differential passive gain configuration. Foreground calibration is used to compensate for errors that arise from stage gains that are not exactly two. Measured results from a prototype operating at 50MS/s fabricated in a 1.8V 0.18μm CMOS process show the ADC to achieve a peak SNDR/SFDR of 58.2dB/66dB at 50MS/s while only consuming 9.9 mW, for a figure of merit of 0.3 pJ/step.

1.2: THESIS OUTLINE

- Chapter two provides an overview of ADC converters with an emphasis on Nyquist-rate data converters. The various tradeoffs between different ADC topologies are outlined in the chapter. A detailed discussion of circuit design issues and considerations in pipelined ADCs is given.
- Chapter three surveys state of the art pipelined ADC enhancement techniques, where topics most relevant to this dissertation, namely linearity enhancement, and low power techniques are detailed.
- Chapter four discusses a novel technique to rapidly measure and correct both DAC and gain errors in an 11-bit 45MS/s ADC. In the chapter the architecture, circuit design, and measured results of a prototype fabricated in 0.18µm CMOS are presented.
- Chapter five discusses a new technique to eliminate the front-end sample-and-hold, and thereby enable significant power reduction in a pipelined ADC. The architecture, circuit design, and measured results of a prototype fabricated in 0.18µm CMOS are presented in the chapter.
- Chapter six deals with a novel technique to significantly reduce pipelined ADC power consumption by replacing opamps with source followers and a passive gain technique to achieve an MDAC gain near 2x. In the chapter, the architecture, circuit design, and measured results of a prototype fabricated in 0.18µm CMOS are presented.
- Chapter seven summarizes the dissertation and provides a brief discussion of future work.

CHAPTER TWO: ADC ARCHITECTURES

2.1: OVERVIEW

In this chapter an overview of factors which commonly determine ADC resolution is outlined. A brief discussion of popular Nyquist-rate ADC topologies is given, where the topologies most relevant to the focus of this work (Flash ADC, SAR, and pipelined ADC) are discussed with the associated tradeoffs of each topology noted. A detailed discussion of pipelined ADC design issues at the circuit level is also given.

2.2: FACTORS WHICH DETERMINE ADC RESOLUTION AND LINEARITY

Digital transmission is the most common form of data communication due to its superior signal integrity in the presence of noise. Digital signals can for example be transmitted at baseband as shown in Fig. 2-1 or modulated by a carrier signal, LO, to a higher frequency as shown in Fig. 2-2.

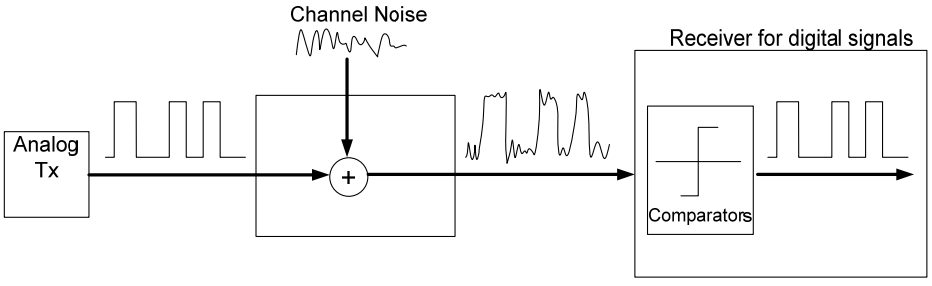

Fig. 2-1: Baseband binary digital transmission

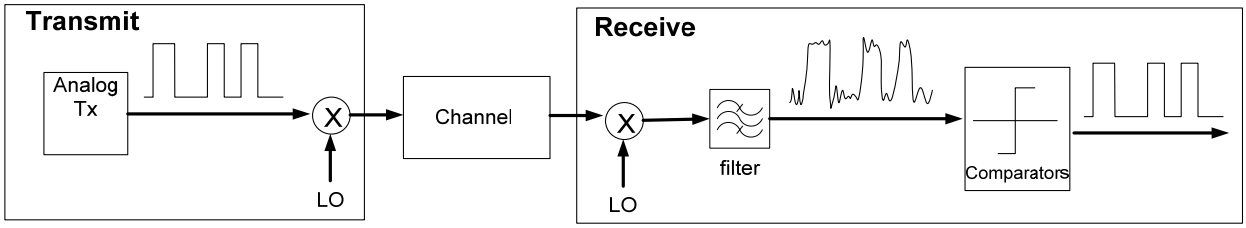

Fig. 2-2: Digital AM transmission

Different applications have different resolution requirements for the ADC in the receiver. For example, in many wireless receiver systems a receiver is required to resolve a received input with a minimum Signal to Noise and Distortion Ratio (SNDR), which is the ratio in dB of signal power to the power of all harmonics and total noise. The Effective Number Of Bits (ENOB) resolved by an ADC can be determined according to the following formula [7]:

$$ENOB = \frac{SNDR - 1.76}{6.02} \qquad \textit{(eqn. 2-1)}$$

In many applications the power of the received signal can vary significantly due to e.g. the distance changing between transmitter and receiver in mobile applications - yet the ADC is still required to linearly digitize the full range of analog inputs with a minimum resolution. In many communication standards, the receive signal is required to maintain a minimum SNDR, thus the total dynamic range of the input (i.e. desired SNDR of the ADC) becomes the sum of the minimum SNDR and variation in input signal strength as shown in Fig. 2-3.

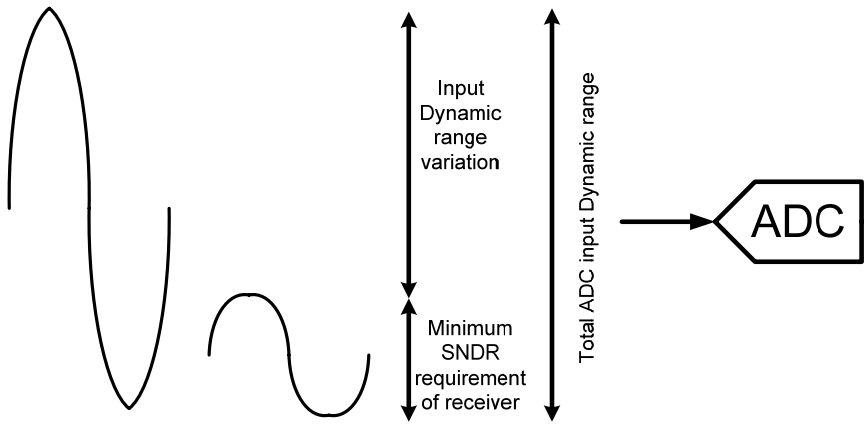

Fig. 2-3: Requirements of ADC input dynamic range

In some applications it may be possible to include an Automatic Gain Control (AGC) circuit before the ADC to compensate for the effect of input dynamic range variation, thus reduce the dynamic range hence resolution of the ADC as shown in Fig. 2-4.

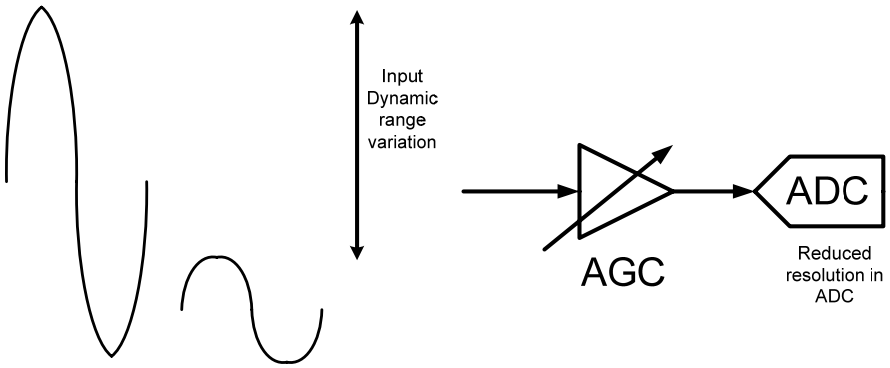

Fig. 2-4: AGC before ADC input to relax ADC input dynamic range requirements

ADC resolution can also be set by the fact that in some applications the signal of interest which is to be digitized is adjacent to another signal in the frequency domain which is orders of magnitude more powerful, as shown in Fig. 2-5.

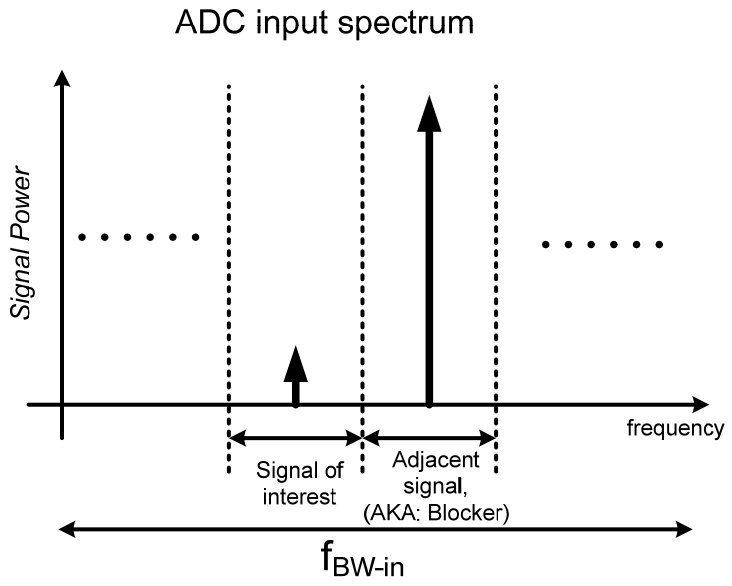

Fig. 2-5: ADC input spectrum illustrating case when desired signal is spectrally next to a much more powerful signal

In such applications it may not be feasible to economically or adequately suppress the adjacent signal (referred to as a blocker) using analog techniques before the ADC input. As a result in many applications the entire signal bandwidth $f_{BW\text{-}in}$ is digitized, and out of band signals efficiently eliminated using digital filters. In such applications to avoid introducing nonlinearity before digitally filtering out the blocker, the ADC must be able to linearly digitize both the large blocker and the signal of interest. Thus the dynamic range requirement of the ADC is set by the sum in dB of the ratio of largest to smallest signal within $f_{BW\text{-}in}$, and the minimum SNDR required to resolve the desired signal of interest.

The linearity [7] requirement of an ADC is set by the fact that in many applications a large number of signals over a large bandwidth are digitized by a single ADC, such as for example OFDM systems, an example spectrum of which is shown in Fig. 2-6.

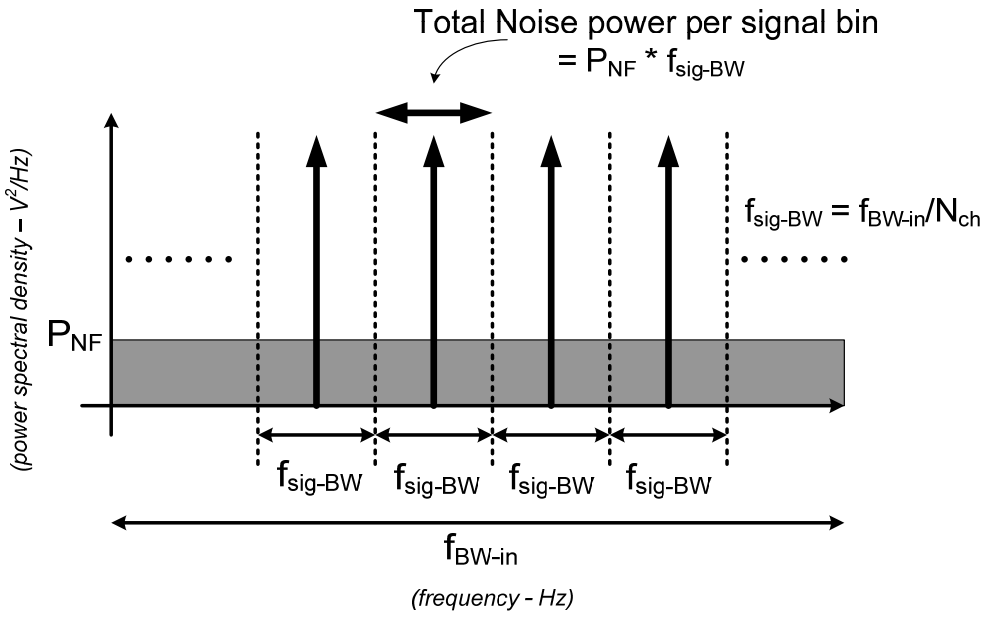

Fig. 2-6: Example OFDM spectrum

In an OFDM system the bandwidth of a single signal f_{sig-BW} is only a small fraction of the overall bandwidth f_{BW-in} digitized by the ADC. Thus the in-band thermal noise floor for a single signal is very small and given by $P_{NF}f_{BW-in}/N_{ch}$ (where N_{ch} is the total number of channels), and hence the ADC does not need to be designed with a very low input referred thermal noise floor. If the ADC has a nonlinear transfer characteristic however as shown in Fig. 2-7, each single tone input to the ADC produces an output with the same tone plus additional harmonic frequencies, reducing the linearity in the digital output spectrum. Linearity is commonly assessed in ADCs by taking the ratio in dB of the largest and smallest harmonic tone in-band – referred to as the Spurious Free Dynamic Range (SFDR).

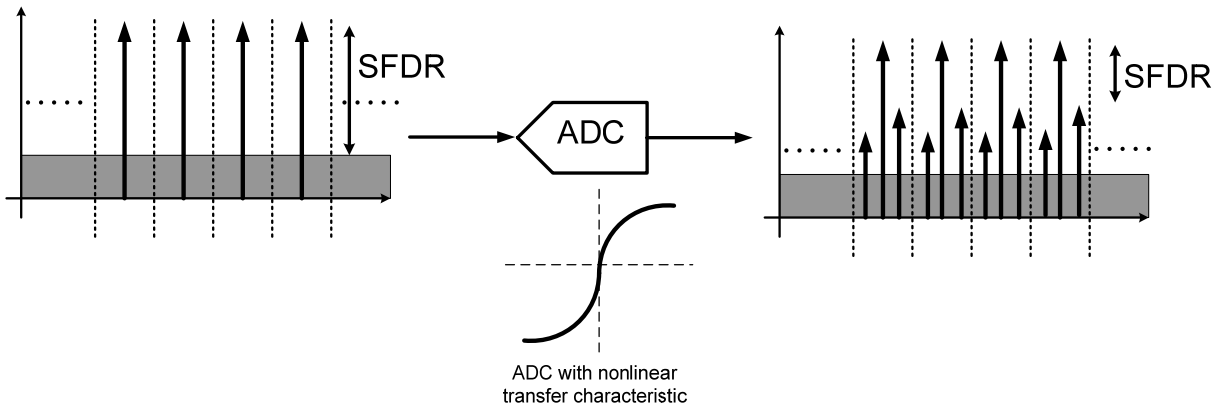

Fig. 2-7: Illustration of harmonic distortion in ADC output spectrum

Thus in a system in which the input consists of several frequency adjacent tones such as OFDM, harmonics generated by each unique OFDM signal can fall into the bandwidth of another OFDM tone, thereby reducing the minimum SNDR within each signal bandwidth f_{sig-BW}. As a result in many systems the linearity of an ADC is the key design parameter, rather than the SNDR. In fact in some publications the ADC ENOB is calculated using the SFDR rather than the SNDR.

From the discussion in this section it is clear that the required resolution and linearity of an ADC can be determined by a number of factors. Furthermore by modifying the receive path by using an AGC and/or a combination of analog filtering techniques, the requirements of an ADC can be relaxed. The optimal configuration for a receiver depends on the specific constraints of a system. It is noted that the general trend in industry is to try and perform as much as possible in the digital domain, thereby pushing the ADC closer to the receiver input, thus demanding ADCs with more resolution and bandwidth – making ADCs essential enablers for future technologies.

2.3: ADC ARCHITECTURES

As ADCs can consume a large percentage of power in a receiver, it is of vital interest to minimize ADC power consumption. Over the years several architectures have been developed which achieve optimal power consumption for different sampling rates, and resolutions as shown in Fig. 2-8.

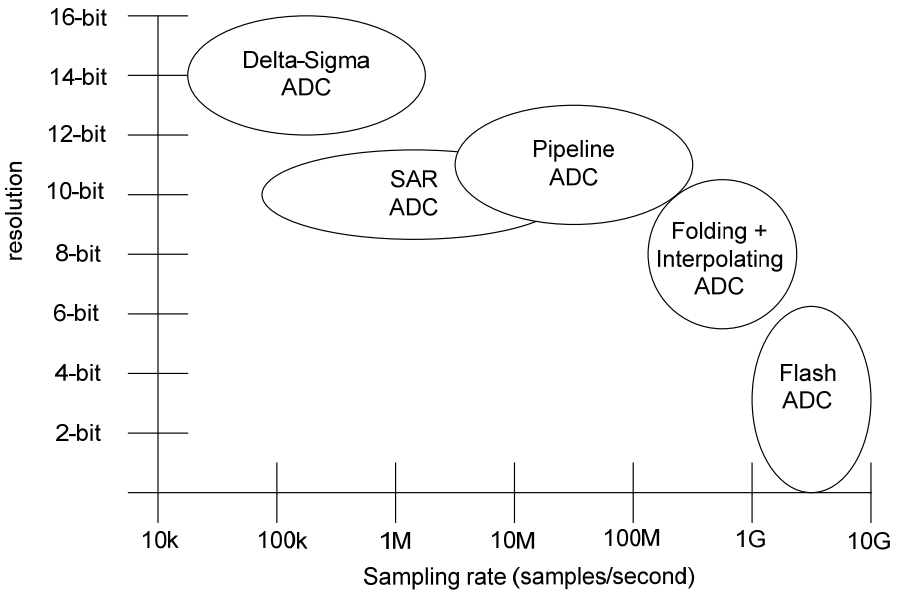

Fig. 2-8: ADC architecture comparison

In sections 2.5 - 2.7 the ADC topologies of Flash, SAR, and Pipelined are reviewed as they are essential to understand within the context of this work. A detailed discussion of topologies not discussed in this work (Delta Sigma, Folding + Interpolating, etc.) can be found in [7]. Table 2-1 summarizes the key tradeoffs of the different ADC topologies shown in Fig. 2-8.

Table 2-1: Comparison of ADC architectures

ARCHITECTURE	***LATENCY***	***SPEED***	***ACCURACY***	***AREA***
Flash	No	High	Low	High
SAR	No	Low-Medium	Medium-High	Low
Folding + Interpolating	No	Medium-High	Medium	High
Delta-Sigma	Yes	Low	High	Medium
Pipeline	Yes	Medium-High	Medium-High	Medium

2.4: ADC FIGURE-OF-MERIT

A popular Figure-of-Merit (FOM) used to compare different ADCs is

$$FOM = \frac{Power}{(2^{ENOB})(f_s)} \quad (pJ/step) \qquad \textit{(eqn. 2-2A)}$$

where f_s is the sampling rate in Nyquist-rate ADCs. This figure of merit is commonly used to compare published reports as the accuracy term is based on easily measured quantities, and calculates a value that has meaningful units (i.e. energy required per conversion step – thus lower FOM means a better ADC).

A slight variation of eqn. 2-2A is the following FOM:

$$FOM = \frac{Power}{(2^{2*ENOB})(f_s)} \quad (pJ/step) \qquad \textit{(eqn. 2-3B)}$$

In eqn. 2-2B the ENOB is multiplied by 2 to account for the fact that due to thermal noise limitations, to achieve twice the resolution 4x the power is required. In general similar FOMs can be achieved with different ADC topologies, however it is noted that ADCs with lower target resolutions tend to be able to achieve better FOMs using eqn. 2-2A.

2.5: FLASH ADC

Various ADC architectures have been developed over the years, each with different tradeoffs with respect to power, speed, and accuracy. Most ADC architectures however are in some form a variant of the Flash ADC or use a Flash ADC in their implementation.

Much like a ruler with a fixed resolution maps an infinite precision length to a finite accuracy (e.g. measure length in millimeters); Flash ADCs *measure* an analog signal into a digital signal by comparing an analog input to fixed reference values as shown in Fig. 2-9. The number of fixed references used determines the accuracy of the digital output. For example, 4-bit accuracy is obtained by comparing against 2^4-1=15 reference values, 10-bit accuracy by comparing against 2^{10}-1=1023 reference values. Determining which reference values the input is in-between forms a length 2^N bit (where N is the accuracy of the ADC) thermometer code representation of the analog input. Mapping the unique thermometer code to its binary equivalent forms a length N, binary representation of the analog input [7].

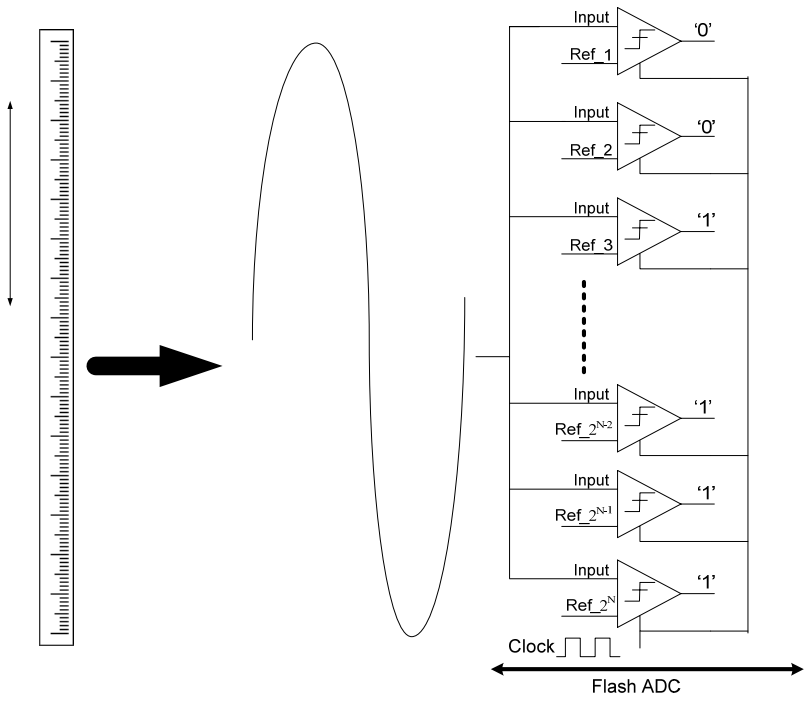

Fig. 2-9: Analogy between a ruler and a Flash ADC

In a Flash ADC, the number of comparators required is thus exponentially related to the desired resolution (in bits). It is noted that to achieve increased resolution in Flash ADCs, large devices are required to suppress process variation effects. Hence in the interest of minimizing area, Flash ADCs are most commonly used in applications where only low resolutions are required.

One of the key advantages of the flash topology is that it has a potential latency of only one clock cycle – that is the digital output is available within one clock cycle of the input being sampled. In certain systems where an ADC is required in a feedback path (e.g.: quantizer in a Delta Sigma [7]), it is critical to implement the ADC with as low a latency as possible to maximize closed-loop stability. As low latency is an attractive feature in some systems, many techniques have been developed to enable increased resolution while reducing area consumption, using essentially the Flash topology (e.g.: folding, interpolating, and averaging).

2.6: SAR ADC

The algorithm which forms the basis of the Successive Approximation Register (SAR) ADC has been known since the 1500's, however it was first patented as an algorithm for use in ADCs in 1958 by Bernard M. Gordon [8]. SAR became a popular topology to implement ADCs in the 1970's with the availability of several logic ICs from companies such as AMD. The SAR ADC is relevant to this work as from Fig. 2-8 the SAR and pipelined ADC (the topology of focus of this work) have overlapping areas of use – roughly for low-medium speed applications with 8 to 10 bits of desired resolution. Thus a clear understanding of each topology and their associated tradeoffs is beneficial in understanding under which circumstances a designer would choose one or the other.

The algorithm used in Successive Approximation is based on a binary search algorithm, and thus is more component efficient than Flash ADCs which use a brute force approach to perform data conversion. Fig. 2-10 illustrates the topology of a standard SAR ADC.

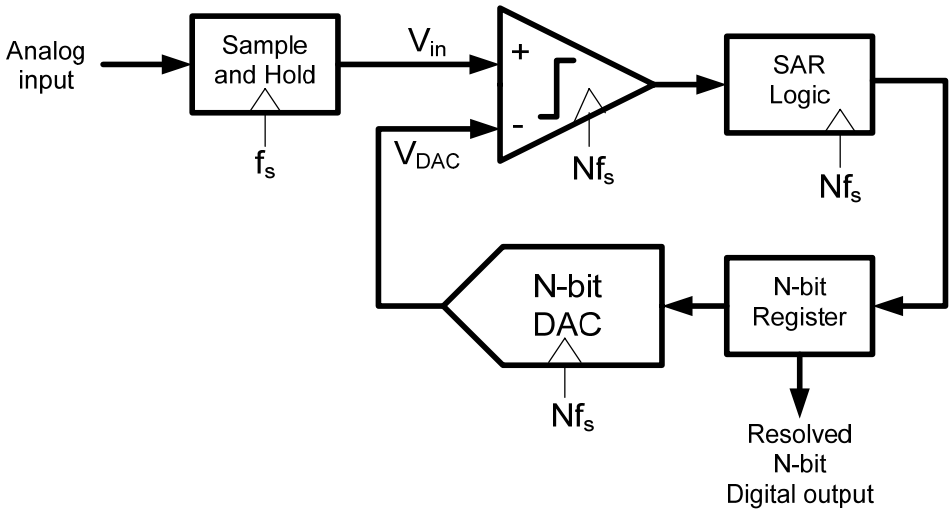

Fig. 2-10: SAR ADC topology

In a SAR ADC the analog input is sampled by a sample-and-hold circuit which operates at the effective Nyquist sampling rate of the ADC (f_s). A sequential binary search is performed on the sampled input by initializing the N-bit register to midscale, which forces the decision threshold of the comparator to be $V_{ref}/2$ (where V_{ref} is the full scale input voltage). As a result if the sampled input is greater than $V_{ref}/2$, the MSB of the N-bit register remains at '1', whereas if the comparator output is '0', the MSB of the N-bit register is changed to '0' [9]. By successively repeating the same algorithm and initializing the next bit in the N-bit register to '1', the digital representation of the analog input can be determined to N bits, where N is the resolution of the SAR ADC and the number of times the algorithm is repeated. Fig. 2-11 illustrates an example of 4-bit conversion.

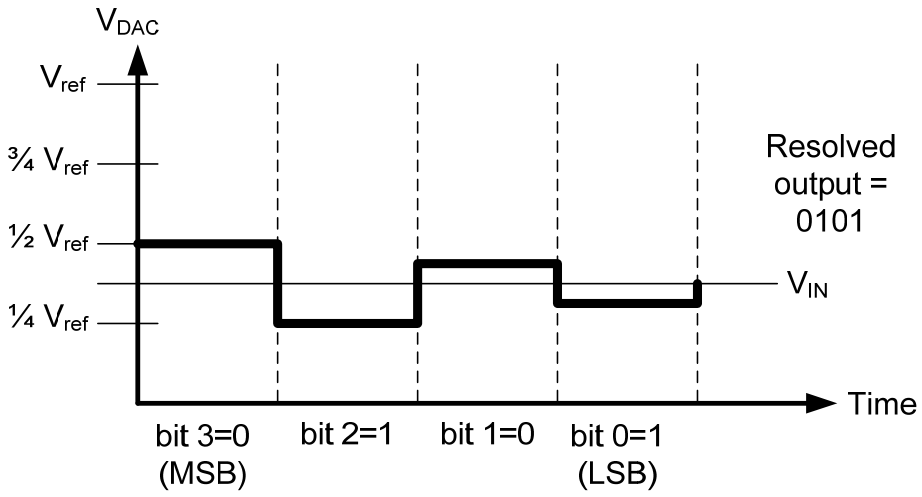

Fig. 2-11: Example SAR conversion - DAC voltage

The significant advantage of the SAR ADC is that it uses only a few analog components (notably only a single comparator) to implement N-bit data conversion, resulting in a compact area and simple design. Furthermore since the topology produces a new digital output every $1/f_s$, the latency of the ADC with respect to the sampling rate is only one clock cycle of the Nyquist-rate clock f_s. As a result the SAR topology can be useful in systems which require ADCs in feedback.

For the SAR ADC to operate with an effective sampling rate of f_s however, the comparator, DAC, and SAR logic are required to operate at Nf_s. For example, if the desired sampling rate and resolution are 100MHz and 10-bits respectively, the DAC, comparator, and SAR logic are required to operate at 1GHz. Thus although the SAR ADC allows for a significant reduction in the number of analog components it comes at the cost of restricting the maximum sampling rate to only a fraction of the maximum speed available by a given technology. From a system designer's point of view a SAR ADC may not be feasible in some systems where the available clock is only at the Nyquist sampling rate. As a result SAR ADCs have traditionally been restricted to low to medium speed, and medium to high accuracy applications.

From a design perspective, it is noted that while the SAR topology determines 1-bit of the final digital output every clock cycle, the DAC is required to settle to the full accuracy of the ADC every clock cycle. Also, while any static offset in the comparator appears as an input referred offset to the ADC, the comparator is required to be able to resolve inputs as small as the LSB of the ADC. As a result much effort is required to optimize the DAC and comparator blocks for high speed and high accuracy. It should be noted however that over the past two years published reports [10], [11] have begun to emerge which show that low power passive DACs based on charge sharing between capacitors can be used to drastically reduce power consumption and significantly increase the operating speed of SAR ADCs to the point that in 90nm operating speeds of 50MS/s can be achieved with resolutions on the order of 8 to 9 bits while only consuming 0.7mW [10].

2.7: PIPELINED ADC

Pipelined ADCs are capable of resolving medium to high resolutions like the SAR ADC topology, however unlike the SAR topology the pipelined approach is able to achieve very

high sampling rates as it does not require a large clock frequency to realize high resolution conversion.

In Fig. 2-12 a two-step ADC or single stage pipelined ADC topology is presented. During the first clock cycle the N/2 Most Significant Bits (MSBs) are resolved (where N is the number of bits in the final ADC output). During the second clock cycle the resolved N/2 MSBs are removed from the input, the residue amplified by 'A' to full scale to maintain the dynamic range, and reuse reference voltages, and subsequently the remaining N/2 bits are resolved. Of note, the gain 'A' forms the radix of the digital output, and is often referred to as the 'stage-gain'.

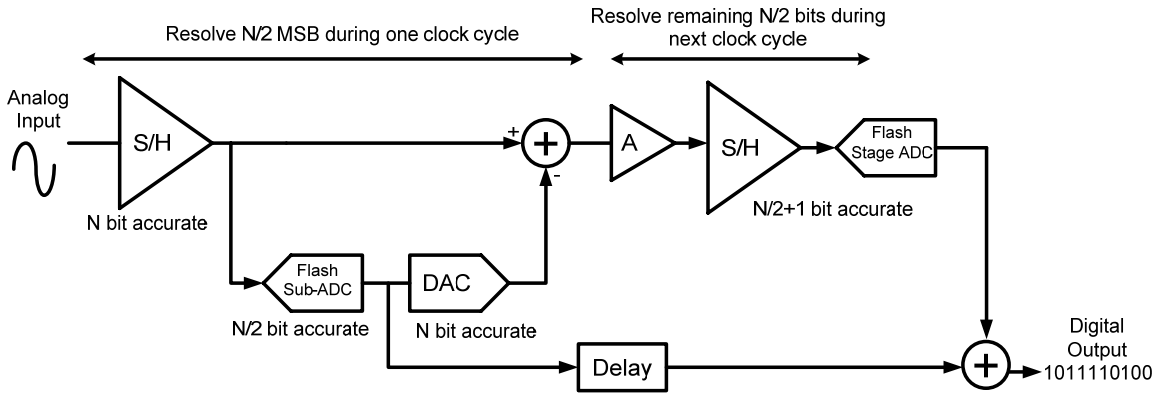

Fig. 2-12: Two step N-bit accurate pipelined ADC

Thus the number of comparators required in the two-stage approach is $2^{N/2+1}$, which is lower than the Flash ADC for N>2. As the pipelined approach implements a queue structure, the maximum speed of the topology is limited by the delay through only a single pipeline stage. However the tradeoff in implementing a queue structure is that the ADC has large conversion latency. That is, rather than the digital outputs being available one clock cycle after the analog input is first sampled (as in the Flash architecture), two clock cycles are required to generate the digital output in the two-step approach.

From a design perspective to ensure the output of the first pipelined stage is sufficiently linear when referred to the input, the *Digital to Analog Converter* (DAC), and subtraction blocks of the first stage must be precise to at least N-bits. Of note however, as the second sample-and-hold is divided by the stage gain, 'A' when referred to the input, the second sample-and-hold only requires N/2+1 bits settling accuracy. In [12] it is shown that by using redundant bits in the sub-ADC, a large offset in the sub-ADC comparators can be tolerated.

The divide and conquer approach used in the two step ADC can be extended further, such that several clock cycles are used, and only a few bits resolved per stage as illustrated in Fig. 2-13; this generalized approach forms the basis of a pipeline ADC [7].

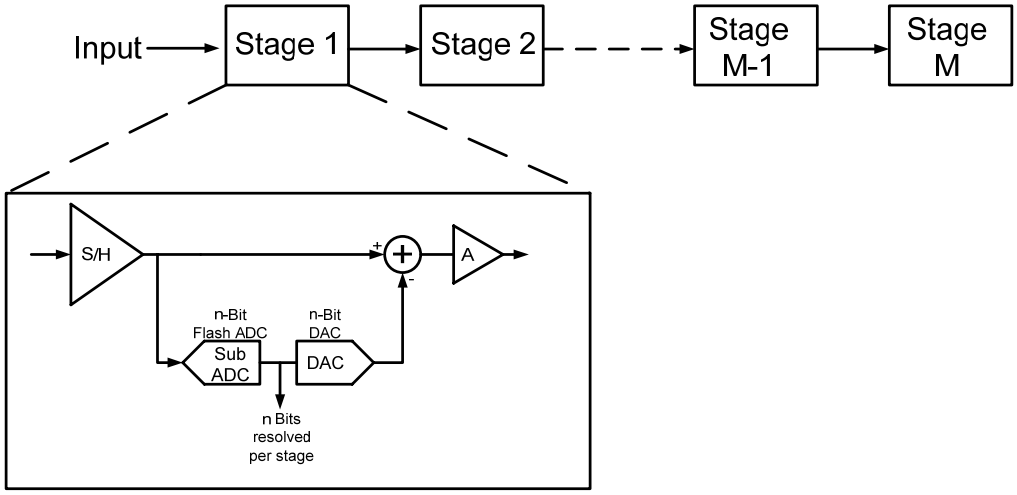

Fig. 2-13: General pipelined ADC architecture

Although several clock cycles are required for an analog value to be digitized in a pipelined ADC, a new digital output is available every clock cycle. Thus the throughput of the pipeline (i.e. speed) is limited only by the delay through a *single* pipeline stage [7] - unlike a SAR ADC where by virtue of the same stage being reused to resolve the entire digital code the SAR ADC's throughput is limited to f_s/N as described in section 2.6. Compared to a SAR

ADC however, pipelined ADCs require more area as instead of one stage being continuously reused to resolve a digital output, the pipelined approach requires N/n pipeline stages to be implemented on-chip (where 'n' bits are resolved per stage). Pipeline ADCs are useful in configurations where latency is not critical (e.g.) where the ADC is in an open-loop or feed-forward signal path.

The precision requirements of each pipeline stage decrease along the pipeline (i.e.) the first stage must be most precise, subsequent stages need only be as precise as the previous stage less the number of bits resolved previously. Thus analog design complexity can be reduced along the pipeline [13] as shown in Fig. 2-14 (less opamp gain and bandwidth for later stages – see section 2.8). As the backend pipeline stages have relaxed precision requirements, they can be designed with smaller area and lower power consumption [13]. Hence it is possible to significantly reduce total power consumption and area by having many stages, where each subsequent stage in the pipeline is sized smaller than the previous stage.

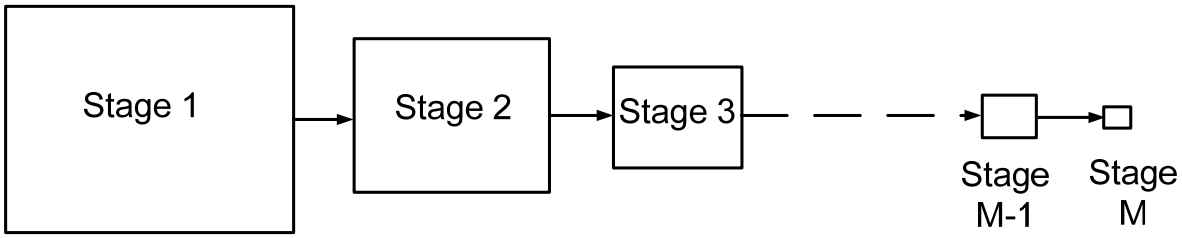

Fig. 2-14: Pipeline stage scaling – stages are sequentially smaller

2.8: PIPELINED STAGE TOPOLOGY – CIRCUIT DETAILS

In section 2.7, the topology of a pipelined ADC was described, where it was shown that a pipelined topology was derived from a cascade of pipelined ADC stages, each resolving a fraction of the overall output code. Within a pipelined stage the functions of sample-and-hold, subtraction, DAC, and gain are commonly combined into a single switched capacitor

circuit, referred to as a *Multiplying Digital-to-Analog Converter* (MDAC). The MDAC and the sub-ADC make up each pipelined stage in a pipelined ADC as shown in Fig. 2-15.

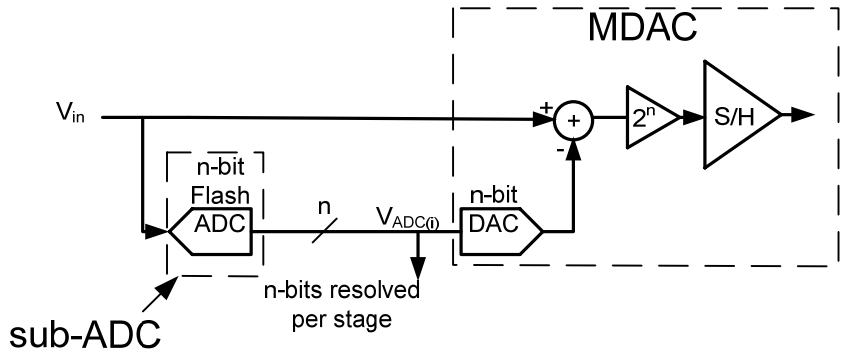

Fig. 2-15: Pipeline stage functionality

Fig. 2-16 illustrates a generic switched capacitor circuit which implements the functionality of an n-bit MDAC. Opamp based switched capacitor circuits are by far the most commonly used topology to implement the MDAC in both academia and industry. Opamp based approaches enable MDACs which have closed-loop gains that are primarily a function of capacitor mismatch (which can be better than 10-bits) and DC opamp gain.

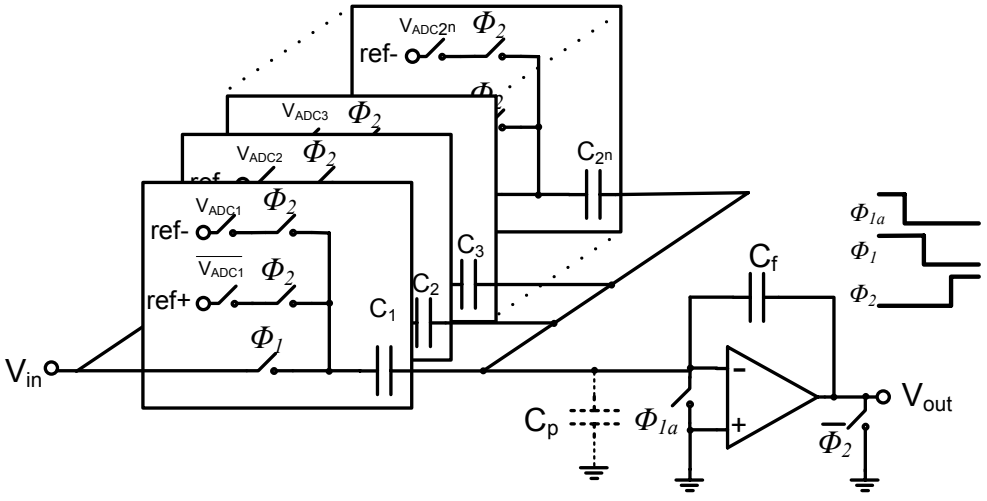

Fig. 2-16: Generic MDAC circuit (shown single-ended but can also be implemented fully-differentially)

From Fig. 2-16, during Φ_{1a} the input is bottom plate sampled on capacitors C_0-C_2^n [7]. Bottom plate sampling significantly reduces the signal dependent charge injection on the sampling capacitors. During Φ_2, the outputs of the sub-ADC are latched and the MDAC configured to multiply the sampled input by: $\sum_1^{2^n} C_i/C_f$, and perform a DAC operation by connecting one end of each sampling capacitor to either $+V_{ref}$ or $-V_{ref}$ as a function of the sub-ADC output. Assuming the opamp has fully settled, at the end of Φ_2, V_{out} can be expressed as:

$$V_{out} = \frac{\sum_1^{2^n} C_i}{C_f} V_{in} - \left[\frac{\sum_1^{k} C_i}{\sum_1^{2^n} C_i} V_{ref} - \frac{\sum_{k+1}^{2^n} C_i}{\sum_1^{2^n} C_i} V_{ref}\right], \qquad \text{(eqn. 2-4)}$$

where *k* is set by the output of the sub-ADC in the pipeline stage.

2.8.2: OPAMP DC GAIN REQUIREMENTS

To understand the gain requirements of the opamp in an MDAC circuit, consider the hypothetical two stage pipelined ADC topology of Fig. 2-17 which shows a pipelined stage that has its MDAC implemented by switched capacitor circuits, an ideal n-bit sub-ADC, and an ideal backend N-1 bit Flash ADC.

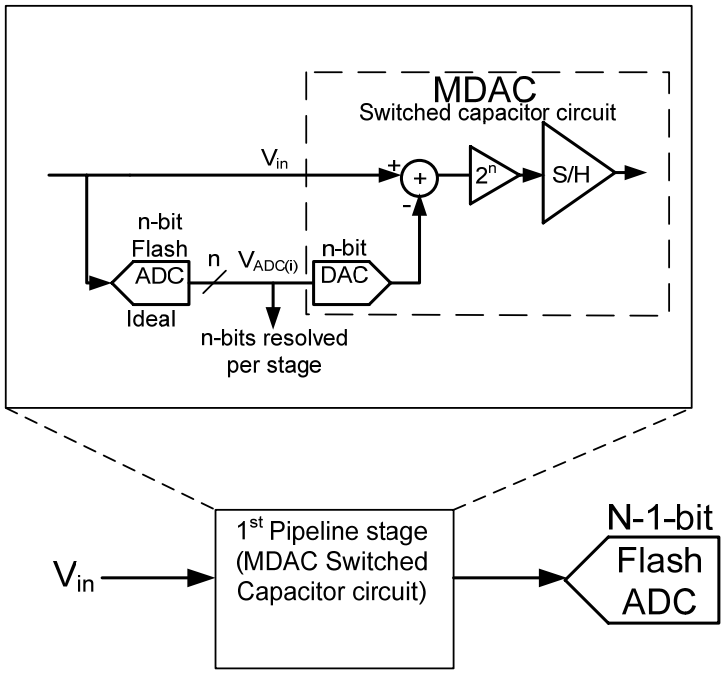

Fig. 2-17: Hypothetical pipeline ADC for illustration purposes

Fig. 2-18 illustrates the residue transfer characteristic (i.e. input/output plot for pipeline stage) of the 1st stage MDAC when the DC gain of the opamp is infinite, the opamp bandwidth infinite, and no capacitor mismatch. Also shown is the total ADC output which is generated by a summation of the bits generated by the sub-ADC from the first pipeline stage and the backend ideal N-1 bit Flash ADC.

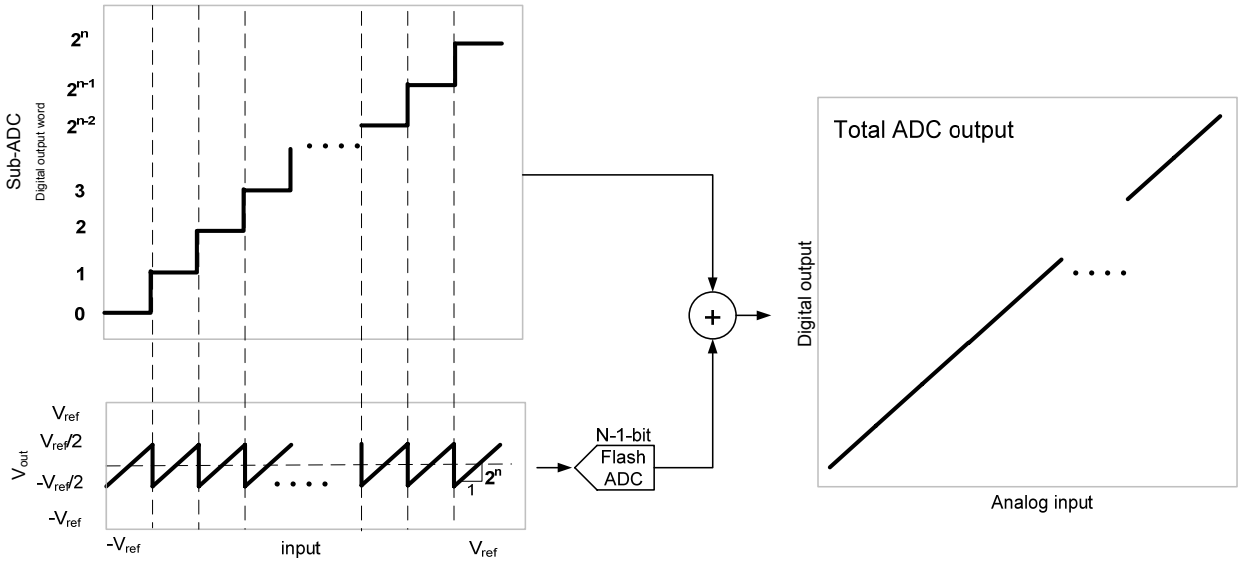

Fig. 2-18: MDAC residue transfer curve, and total ADC output when opamp gain is infinite, infinite opamp bandwidth, and capacitor mismatch ignored

If the opamp of Fig. 2-16 has a DC gain of A, the error in the fully settled output, Δ, due to finite DC opamp gain can be found to be [17]:

$$\Delta \approx \frac{1}{A\beta} \qquad \textit{(eqn. 2-5)}$$

where β is the feedback factor, and is given by:

$$\beta = \frac{C_f}{C_f + C_p + \sum_1^{2^n} C_i} \quad , \qquad \textit{(eqn. 2-6)}$$

and C_p is the parasitic capacitance at the input of the opamp. Clearly for a finite A, the gain error is non-zero. Fig. 2-19 illustrates the impact on the residue of the first pipeline stage when the effect of finite DC opamp gain is included while ignoring the effect of capacitor mismatch and assuming the opamp output has fully settled.

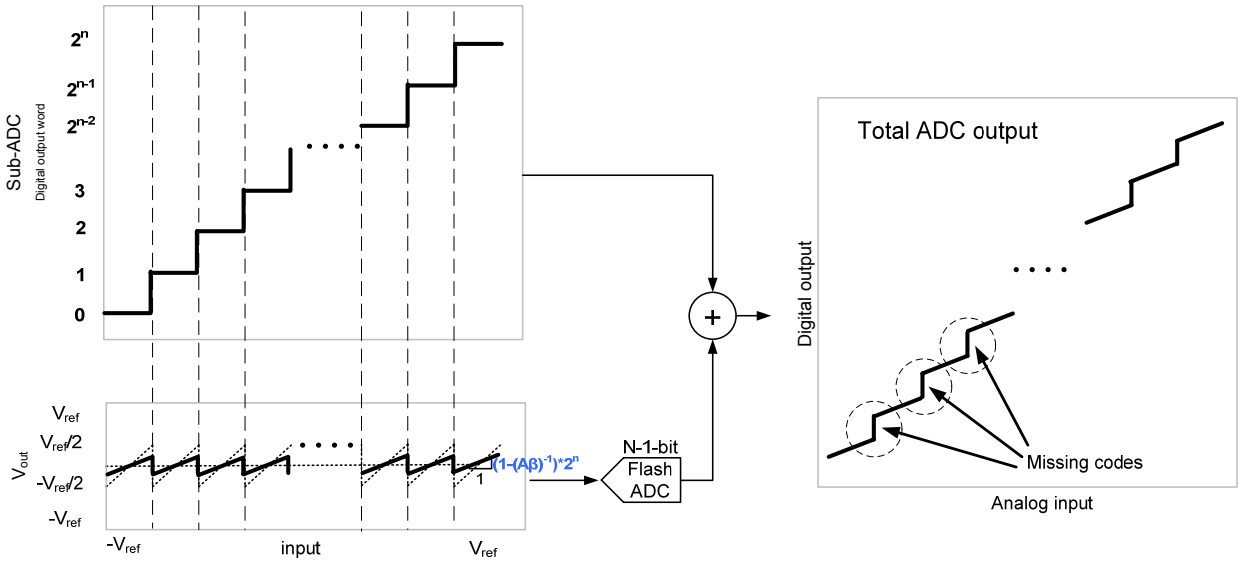

Fig. 2-19: Residue transfer curve of pipeline stage when opamp gain error is included. (Original i.e. error free residue curve shown in dashed lines)

From Fig. 2-19 it is clear that opamp gain error alters the jump in the residue transfer curve whenever the sub-ADC output changes, from the ideal value of V_{ref}, to a scalar of V_{ref}. As a result missing codes are generated in the overall pipelined ADC output whenever the sub-ADC output changes. Missing codes in the ADC output result in harmonic distortion which limits ADC accuracy and linearity [7]. As a result to achieve a sufficiently linear pipelined ADC missing codes need to be eliminated. The minimum gain required to achieve no missing codes can be derived by referring the gain error Δ to the input and forcing the total

gain error to be below the quantization noise floor of the ADC. For example, for the case of Fig. 2-17, the gain error referred to the input is:

$$\Delta_{input} \approx \frac{1}{A\beta 2^n} \qquad \textit{(eqn. 2-7)}$$

For an N-bit pipelined ADC, the total sum of errors is required to be less than $1/2^N$. Ignoring all other non-idealities, thus:

$$\frac{1}{2^N} > \frac{1}{A\beta 2^n} \qquad \textit{(eqn. 2-8)}$$

hence:

$$A > \frac{2^{N-n}}{\beta} \qquad \textit{(eqn. 2-9)}$$

For example: if N=11-bits, n=3-bits, and β=1/8, the minimum gain is ~ 66dB to suppress the effect of the finite gain of the first stage only. Of note, to a 1st order $\beta \approx 2^{-n}$, thus eqn. 2-8 could be also expressed as $A>2^N$. Clearly a large DC gain is required when opamp based switched capacitor circuits are used to implement the MDAC in medium to high resolution pipelined ADCs. In practice the backend ADC is not ideal and introduces additional non-ideality, forcing the gain of the first stage to be even larger. The net effect of all non-idealities can be found by referring each non-ideality to the input and forcing the sum of all errors to be below the quantization noise floor, i.e. less than 2^{-N}

With modern sub-micron technology nodes offering less and less intrinsic gain in the active region for CMOS transistors (due to short channel effects), attaining opamps with large DC gain can be difficult. Large DC gain is commonly achieved in CMOS using techniques such as gain boosting [18], multi-stage opamps [19], or using long channel lengths for key transistors [7]. However such gain enhancement techniques increase the power of the opamp, and/or increase the design complexity of the opamp. Rather than using analog techniques to suppress the effect of finite DC gain, digital techniques can be used where the error due to finite gain can be compensated in the digital output of the ADC. Such techniques are referred to as calibration and are covered in more detail in section 3.2.

2.8.3: OPAMP BANDWIDTH REQUIREMENTS

In section 2.8.2 the DC gain requirements of an opamp based MDAC were derived under the assumption that the opamp had sufficiently settled such that the gain error was limited only by the finite gain error of the opamp. In order to have an opamp which is sufficiently settled within a given timeframe the opamp must have enough bandwidth. As will be seen in this section, an opamp which has a large bandwidth requires large power consumption - thus to minimize power it is critical to optimize opamp bandwidth.

Assuming a first order response of $A(s) = A/(1 + \frac{s}{\omega_p})$ for the opamp of Fig. 2-16 (where ω_p is the dominant pole frequency of the opamp), near the unity gain frequency ω_{ta}, $A(s) \approx \omega_{ta}/s$. Thus the closed-loop transfer function, H(s), of the MDAC topology of Fig. 2-16 during Φ_2 is approximately given by:

$$H(s) = \frac{A(s)}{1+A(s)\beta} \approx \frac{1}{\beta}\left(\frac{1}{1+\frac{s}{\beta\omega_{ta}}}\right) \qquad \textit{(eqn. 2-10)}$$

Of note, the unity gain frequency of the closed-loop, ω_t, is only **a fraction** of the unity gain frequency of the open-loop, i.e.:

$$\omega_t = \frac{\omega_{ta}}{1/\beta} \qquad \textit{(eqn. 2-11)}$$

The step response in the time domain of eqn. 2-10 is given by:

$$h(t) = \frac{1}{\beta}(1 - e^{-t/\tau}) \qquad \textit{(eqn. 2-12)}$$

where $\tau=(\beta\omega_{ta})^{-1}$. The relative error in the output of the first pipeline stage due to finite bandwidth (Δ_{BW}) after T/2 seconds (where T is the period of clock Φ_2 from Fig. 2-16) is thus given by:

$$\Delta_{BW} = e^{-T/2\tau} \qquad \textit{(eqn. 2-13)}$$

Referring the settling error of the first pipeline stage to the input of the ADC, and noting that the total error must be less than the quantization noise (i.e. $<2^{-N}$), the required unity gain frequency of the opamp f_{ta} to achieve N-bit settling is thus given by:

$$f_{ta} > \frac{(N-n)log2}{\beta\pi} f_s \qquad \textit{(eqn. 2-14)}$$

where f_s = 1/T is the sampling rate of the pipelined ADC. From eqn. 2-13 if parasitic capacitance C_p at the input of the opamp is ignored, every time n increases by 1-bit, β ideally decreases by a factor of 2. Using this relationship Fig. 2-20 plots f_{ta}/f_s vs. n for different N, where it is seen that for high resolution ADCs as the number of bits per stage n increases the minimum required open-loop unity gain frequency of the opamp to ensure sufficient settling also increases.

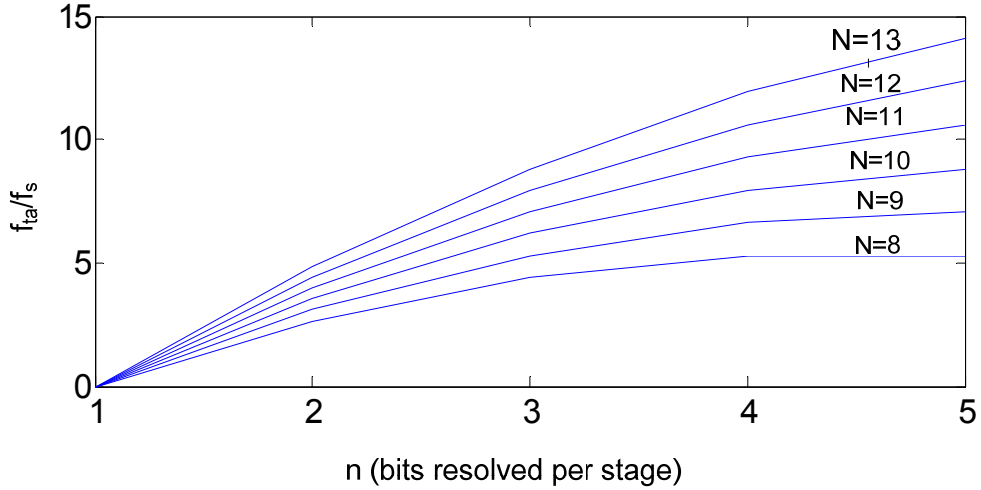

Fig. 2-20: Variation of required unity gain frequency relative to sampling rate with number of bits resolved in the first pipeline stage

If parasitic capacitances are included, β is made smaller, thus even more bandwidth from the opamp is required to achieve the desired settling.

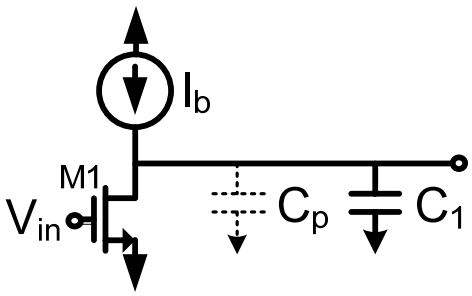

Fig. 2-21: Modeling an opamp by a single transistor

If an opamp is modeled as a single transistor in the active region as shown in Fig. 2-21, it is shown in [7] that the unity gain frequency is given by:

$$\omega_{ta} = \frac{g_m}{C_{load}} \qquad \textit{(eqn. 2-15)}$$

where g_m is the transconductance of M1 and $C_{load}=C_1+C_p$. From [7], the total parasitic capacitance is given by the sum of drain to gate (C_{dg}) and drain to bulk (C_{db}) capacitances so that:

$$C_p = C_{dg} + C_{db} = WL_{OV}C_{OX} + A_dC_{jd} + P_dC_{j-sw} \qquad \textit{(eqn. 2-16)}$$

Where W, L are the width and length of M1, A_d the area of the drain, P_d the perimeter of the drain, and L_{ov}, C_{ox}, C_{jd}, and $C_{j\text{-}sw}$ process dependent parameters. Since:

$$g_m = \sqrt{2\mu C_{ox}\frac{W}{L}I_D} \qquad \textit{(eqn. 2-17)}$$

it can thus be seen that:

$$\omega_{ta} = \frac{\sqrt{2\mu C_{ox}\frac{W}{L}I_D}}{WL_{OV}C_{OX}+A_dC_{jd}+P_dC_{j-sw}+C_1} \qquad \textit{(eqn. 2-18).}$$

It can be seen from eqn. 2-17 that in order to double the unity gain frequency, if the width of the device is increased by a factor of two (i.e. two devices in parallel), the drain current is required to more than double. This is because as the area of the transistor is increased to increase g_m, the effect of the parasitic capacitor also increases (while C_1 remains fixed), and thus reduces the efficiency of the transistor. For large unity gain frequencies - which could be required from an opamp if fast settling is required with a large closed-loop gain (i.e. small β), the parasitic capacitance C_p can be on the order of the load capacitance C_1. Thus in such situations much power is wasted by the transistor in charging both the load capacitance C_1 and the transistor's own parasitics. Clearly to avoid such a situation it is beneficial to achieve the maximum speed possible with the lowest unity gain frequency possible – which occurs when the gain, 1/β, is smallest, i.e. when the fewest number of bits are resolved per stage.

In [7] it is also shown that $g_m=2I/V_{eff}$, where I is the bias current of the transistor, and V_{eff} the overdrive voltage ($V_{eff}=V_{gs}-V_t$). Hence it can also be shown that: $\omega_t=2\beta I/V_{eff}$. Thus if the gain of the MDAC is increased (i.e. β decreased), to maintain the same closed-loop settling time, the current 'I' needs to be *increased* by the same amount (assuming V_{eff} is made constant). Hence there exists a clear tradeoff between closed-loop gain and power in MDACs.

2.8.4: THERMAL NOISE REQUIREMENTS

In section 2.8.3 it was shown that MDAC settling accuracy was a direct function of the opamp's unity gain frequency. Since the unity gain frequency of an opamp is given approximately by g_m/C_L [7], the question thus arises what determines the size of the load capacitance? One of the main issues which determines capacitor size is thermal noise. This section introduces the concept of sampled noise and discusses how capacitor size affects sampled noise.

Although capacitors are ideally noiseless elements, in a sampled system, capacitors which hold discrete time values capture noise generated by noisy elements such as resistors (from sampling switches) and opamps. Consider the following noise analysis of a capacitor sampling resistor noise as shown in Fig. 2-22:

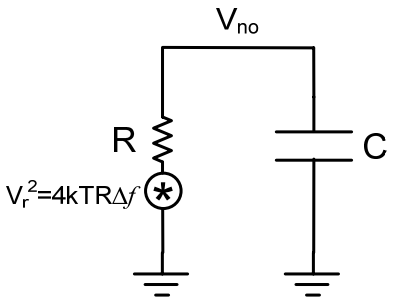

Fig. 2-22: RC noise model

from [7] it is shown the equivalent noise bandwidth is $\frac{\pi}{2} f_0$,

$$\therefore V_{no-RMS}^2 = \frac{\pi}{2} f_0 V_R^2(f) \qquad \textit{(eqn. 2-19)}$$

$$\because f_0 = \frac{1}{2\pi RC} \rightarrow V_{no-RMS}^2 = \frac{kT}{C} \qquad \textit{(eqn. 2-20)}$$

From the above example it is clear increasing the size of the sampling capacitor reduces the power of thermal noise. As thermal noise represents a dynamic noise source that reduces ADC SNR, a sufficiently large capacitance for the sampling capacitors of Fig. 2-16 must be used to suppress thermal noise of the MDAC to a level below the effects of ADC quantization. Thus thermal noise imposes a tradeoff between power and accuracy – the larger the capacitance (thus power consumption), the lower the thermal noise hence higher the accuracy. In high accuracy ADCs (≥10-bits) large capacitive loads result in large, power hungry opamps. In general the opamp is the largest consumer of power in pipelined ADCs, thus most power saving techniques in pipelined ADCs deal with some variation of the opamp.

Thermal noise from opamps also contributes to reducing the thermal noise floor in pipelined ADCs. An analysis of the power spectral density of various opamp topologies is given in [20].

2.8.5: MDAC DESIGN CONSIDERATIONS - CAPACITOR MATCHING/LINEARITY

Recalling that the residue of a generic pipeline stage (from Fig. 2-16) is given by:

$$V_{out} = \frac{\sum_1^{2^n} C_i}{C_f} V_{in} - \left[\frac{\sum_1^k C_i}{\sum_1^{2^n} C_i} V_{ref} - \frac{\sum_{k+1}^{2^n} C_i}{\sum_1^{2^n} C_i} V_{ref} \right],$$

if there is capacitor mismatch between the sum of C_i and C_f there will be an error in the desired gain of the pipeline stage, producing similar errors to those illustrated in Fig. 2-19.

Capacitor mismatch *between* sampling capacitors also affects the linearity of the DAC in Fig. 2-16. For example, if n=4 and C_1=100fF, C_2=95fF, C_3=98fF, and C_4=103fF (i.e. mismatch between sampling capacitors C_i), for k=1,2,3,4, the amount of V_{ref} subtracted is: -0.4949V_{ref}, -0.015V_{ref}, .4797V_{ref}, and 1V_{ref}, respectively. For the DAC operation to be error free, every time the sub-ADC output increases by 1, the DAC voltage subtracted should increase by 0.5V_{ref}, however as illustrated with the example, when capacitor mismatch is included this is not the case. The result of these DAC errors is a **unique** jump in the residue transfer function every time the output of the sub-ADC changes as illustrated in Fig. 2-23, resulting in missing codes and substantial harmonic distortion.

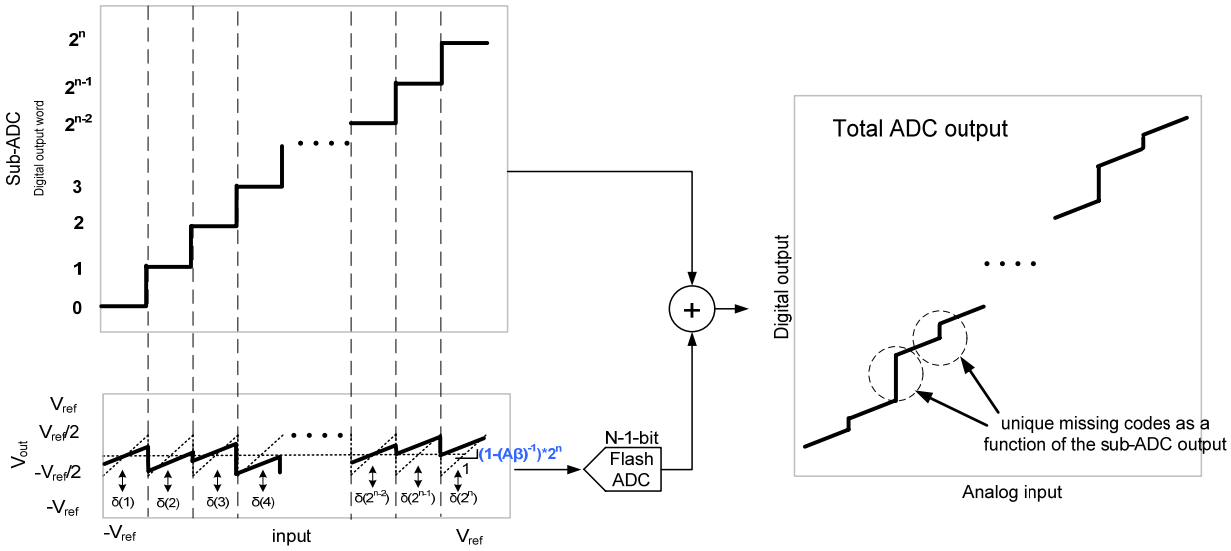

Fig. 2-23: Illustration of DAC and gain errors in pipelined ADC output – ideal residue transfer curve shown by dashed lines

Relative capacitor mismatch is determined by the area of a capacitor where:

$$capacitor\ mismatch\ \propto \frac{1}{\sqrt{capacitor\ area}} \qquad \textit{(eqn. 2-21)}$$

Thus large capacitors are required to achieve a high degree of matching. Hence capacitor size is a function of thermal noise *and* capacitor mismatch. As discussed in section 2.2 in some applications thermal noise requirements are low, however the linearity requirements are high, thus forcing a pipelined ADC designer in such instances to use large capacitors to suppress DAC errors, even though low a low thermal noise floor is not necessary.

Improved matching in layout can be achieved by using arrays of common unit size capacitors in a highly symmetric configuration and or dummy capacitors [20]. Capacitor mismatch can also be improved by using Metal Insulator Metal (MiM) capacitors. However, as MiM capacitors require an extra layer during fabrication, they are not always available in processes used to make commercial products, as commercial products tend to be optimized in cost for digital circuits.

Due to process variation however, layout techniques and additional design layers can only improve capacitor mismatch so far. Thus an arbitrarily high level of matching cannot be achieved using good layout techniques alone, hence limiting pipelined ADC resolution to the medium-high range as noted in Fig. 2-8. One of the strategies used to overcome limitations due to process variation is to use digital calibration techniques. Calibration techniques can measure and compensate the effect of non-idealities such as finite DC opamp gain, and capacitor mismatch. More details on calibration are provided in section 3.2.

2.8.6: SUB-ADC DESIGN - COMPARATOR

From section 2.5, it is clear that the number of comparators required to implement the Flash sub-ADC increase exponentially with each additional bit resolved per stage. From section 2.4 it was shown that if a very large resolution is demanded from a Flash ADC, the Flash ADC can become very large and power hungry to adequately suppress mismatch effects. As a result from the standpoint of minimizing complexity in the sub-ADC it is desirable to minimize the number of bits resolved per stage in the sub-ADC.

As discussed in section 2.7, due to redundancy in the pipeline stage, comparators with large offsets can be used in the Flash sub-ADC. Typically dynamic comparators are used for the sub-ADC as they have low power consumption and complexity, but high offset. In [21] different dynamic comparator topologies for use in pipelined ADCs are analyzed.

2.9: FRONT-END SAMPLE-AND-HOLD

In a pipelined ADC the analog input is sampled by both the MDAC and sub-ADC with different sampling circuits. Due to mismatches in the signal paths as well as threshold mismatches in the sampling switches, the analog inputs sampled by the sub-ADC and MDAC in the first pipeline stage (when driven by an analog input source) are different, as illustrated in Fig. 2-24.

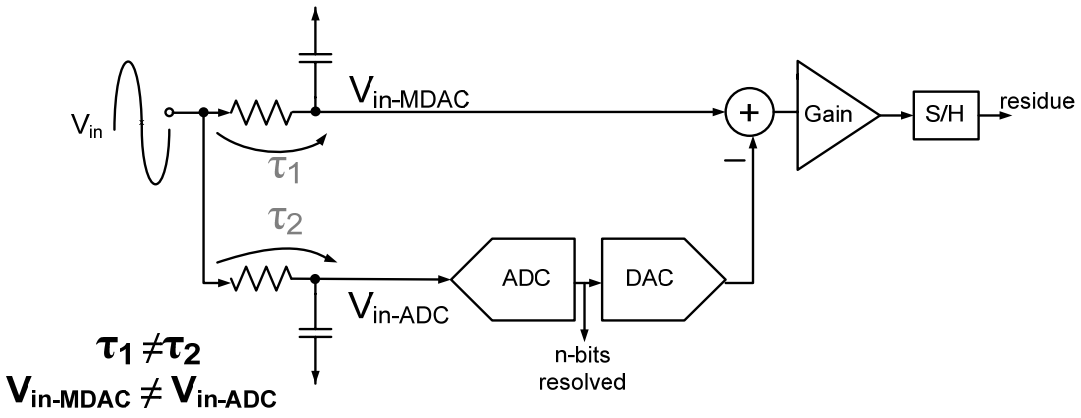

Fig. 2-24: Timing mismatch between sub-ADC and MDAC when first pipeline stage is connected directly to analog input

For a difference in effective sampling time between the sub-ADC and MDAC of Δ_{skew}, and an input sinusoid with frequency f_{in}, and peak voltage V_{peak}, the maximum difference in input voltage sampled by the MDAC and sub-ADC is given by

$$V_{skew-\max} = 2\pi f_{in}\Delta_{skew}V_{peak} \qquad \textit{(eqn. 2-22)}$$

For very large input frequencies the MDAC and sub-ADC can sample vastly different inputs, resulting in massive harmonic distortion in the ADC output. For example, if a full scale sinusoid of 270MHz is applied to a pipelined ADC which has a sampling skew of 140ps between MDAC and sub-ADC, the difference between inputs sampled by the MDAC and sub-ADC can be as high as a quarter of the full scale voltage.

To ensure the sub-ADC and MDAC see the same input, a front-end Sample-and-Hold (S/H) is commonly used before the first pipelined stage so as to make the input to the first pipeline stage discrete time, thus independent of input frequency given a sufficient settling time. Since a front-end S/H has a gain of one, its inherent thermal noise contributes directly to reduce the dynamic range of the ADC. To thus maintain input dynamic range, the S/H is required to have a noise floor and distortion lower than that of the pipelined ADC following it. As a result the power of the ADC is significantly increased by using a front-end S/H. It is not unusual for the front-end S/H to be the largest power consumer in an ADC. Fig. 2-25 illustrates a common topology used to implement a front-end S/H [22].

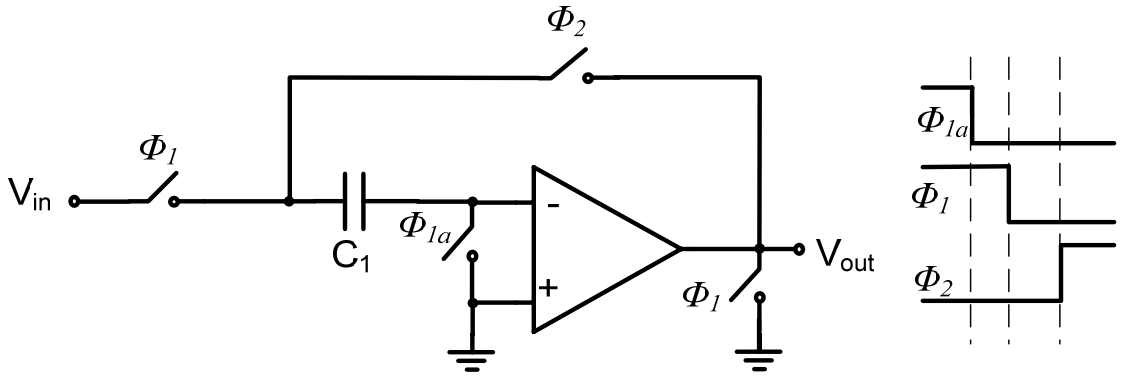

Fig. 2-25: Commonly used front end S/H topology

2.10: SUB-SAMPLING

The majority of Nyquist-rate ADCs use a sampling rate which is twice the highest frequency component of the input - this guarantees that the sampled input can be perfectly reconstructed (ignoring the ADC's resolution). In some applications however it is desirable to sample at a fraction of the input frequency – this type of sampling is referred to as sub-sampling. Sub-sampling is commonly used to alias a high frequency input down to a lower frequency. To understand how this happens consider a sinusoidal input *y(t)* of frequency (f_0+f_{LO}), which is sampled by an ADC at a sampling rate of f_s. The discrete time values of the input digitized by the ADC are given by:

$$y[m] = sin\left(2\pi \frac{f_0+f_{LO}}{f_s} m\right) \qquad \textit{(eqn. 2-23)}$$

If f_0+f_{LO} is greater than f_s, the periodicity of the sine function results in y[m] appearing as if it were sampled at a much lower frequency: modulo(f_0+f_{LO},f_s). Thus rather than use a mixer to translate an input bandwidth to a lower frequency [1] as shown in Fig. 2-26, sub-sampling can be used to efficiently perform frequency translation as shown in Fig. 2-27. The tradeoff in using a sub-sampled ADC is the maximum input frequency of the ADC is significantly increased.

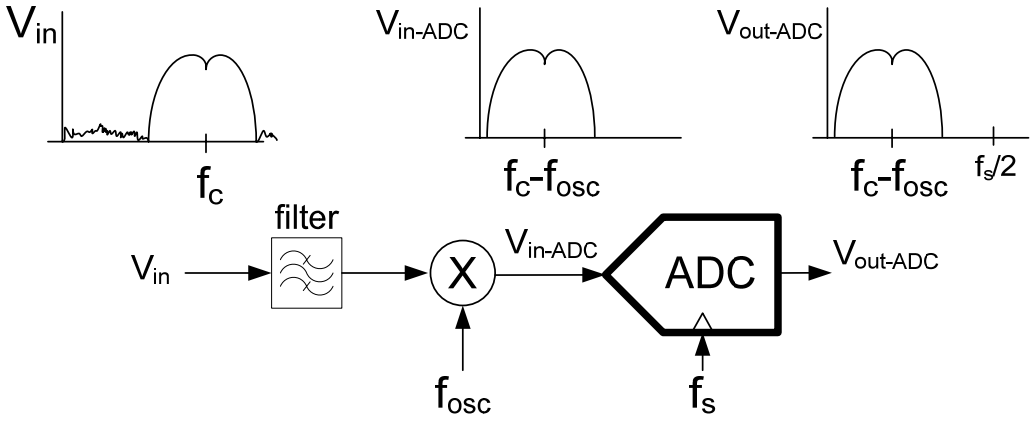

Fig. 2-26: Frequency translation using a front-end mixer

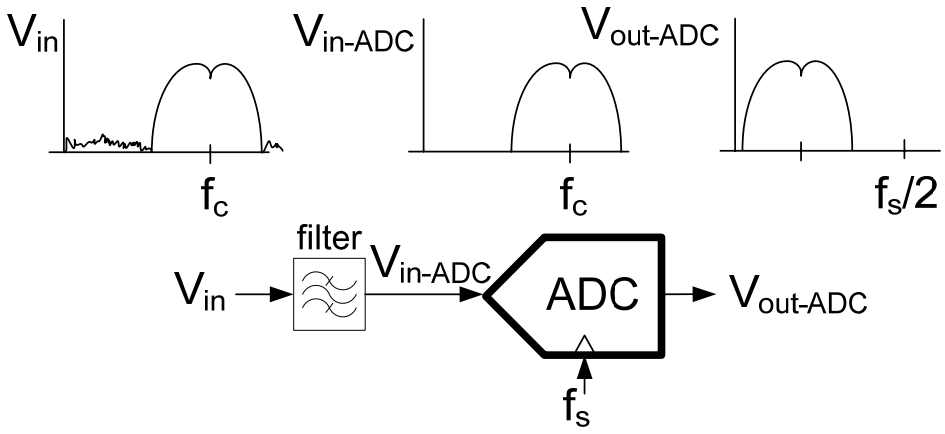

Fig. 2-27: Frequency translation using sub-sampling

2.11: SUMMARY

This chapter, Nyquist-rate ADC architectures were introduced, where the key design issues related to the design of pipelined ADCs were discussed in detail. Circuit issues and tradeoffs discussed include: opamp DC gain, opamp bandwidth, thermal noise, capacitor matching, and sub-ADC comparator design. The role of a front-end sample-and-hold was also discussed, where it was shown that a front-end sample-and-hold ensures functionality for very high input frequencies. Sub-sampling was also discussed as a way to eliminate mixers in receivers.

CHAPTER THREE: STATE OF THE ART PIPELINE ADC ENHANCEMENT TECHNIQUES

3.1: OVERVIEW

In this chapter state of the art research in pipelined ADC enhancement techniques which address some of the circuit issues and tradeoffs introduced in Chapter three will be discussed. In section 3.2 the concept of calibration in pipelined ADCs is introduced as a way of reducing the impact of non-idealities such as low DC opamp gain, and capacitor mismatch. Examples of pipelined ADCs which use calibration to improve their performance will be discussed, where the tradeoffs and limitation in each approach will be discussed. The issue most relevant to the work presented in this dissertation, namely rapid calibration, will be discussed in most detail. In section 3.3 techniques to reduce power consumption in ADCs will be reviewed. The discussion will focus primarily on areas of research most applicable to this dissertation, namely: techniques to eliminate the front-end sample-and-hold, and non-opamp based amplifier topologies. The chapter concludes with a summary of the discussion.

3.2: CALIBRATION IN PIPELINED ADCS

As discussed in sections 2.8.2 and 2.8.5, low DC opamp gain and capacitor mismatch are major sources of missing codes which degrade the linearity of a pipelined ADC. To enable higher linearity in processes which offer only low opamp DC gain and/or insufficiently matched capacitors, digital calibration techniques which measure and compensate the effect of missing codes can be used. As newer deep sub-micron technologies offer digital circuits with lower power consumption and area compared to older technologies, digital calibration

has become increasingly attractive to realize low power highly-linear ADCs. Sections 3.2.3 and 3.2.4 review how gain and DAC errors are corrected in a pipeline stage. Sections 3.2.5 and 3.2.6 compare and contrast foreground and background calibration techniques.

3.2.2: REVIEW OF ERROR SOURCES

Fig. 3-1 illustrates a 3+1-bit pipelined ADC stage which has its gain altered from the ideal value by (1-γ), and DAC output in error by δ(MSB).

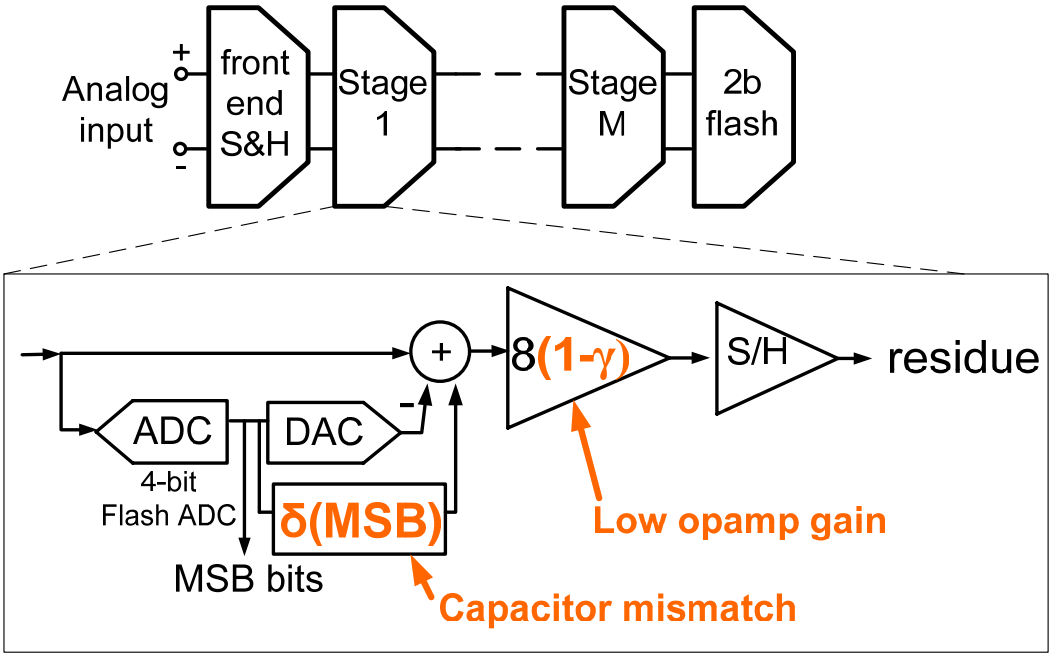

Fig. 3-1: Illustration of errors sources in pipelined ADC

Fig. 3-2, Fig. 3-3, and Fig. 3-4 illustrate the residue transfer curve of Fig. 3-1 when the pipelined stage is: 1.) ideal, 2.) has a non-ideal gain, and 3.) has a non-ideal gain *and* DAC errors respectively.

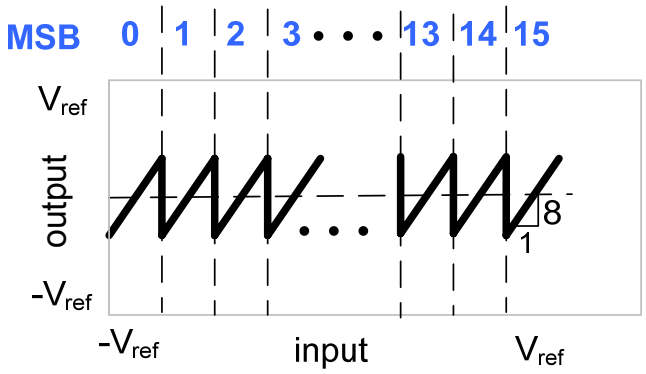

Fig. 3-2: Ideal residue transfer curve of a 3+1-bit pipelined stage

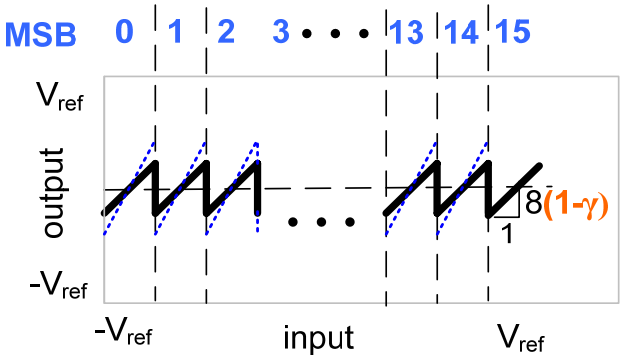

Fig. 3-3: Residue transfer curve of a 3+1-bit pipelined stage with gain errors included

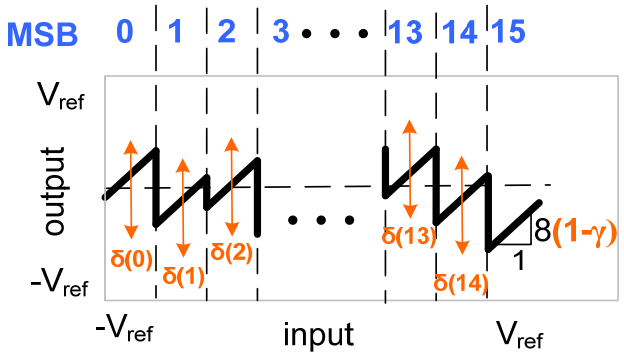

Fig. 3-4: Residue transfer curve of a 3+1-bit pipelined stage with gain and DAC errors included

3.2.3: GAIN ERROR CORRECTION

Gain errors can be corrected by simply scaling the backend digital output by the inverse of the gain error factor (1-γ). Fig. 3-5 illustrates an example of an architecture which compensates for the effect of the non-ideality γ in the first pipeline stage, assuming the value of γ is already known. Gain errors in subsequent pipeline stages can be corrected by inductively using the same error correction technique.

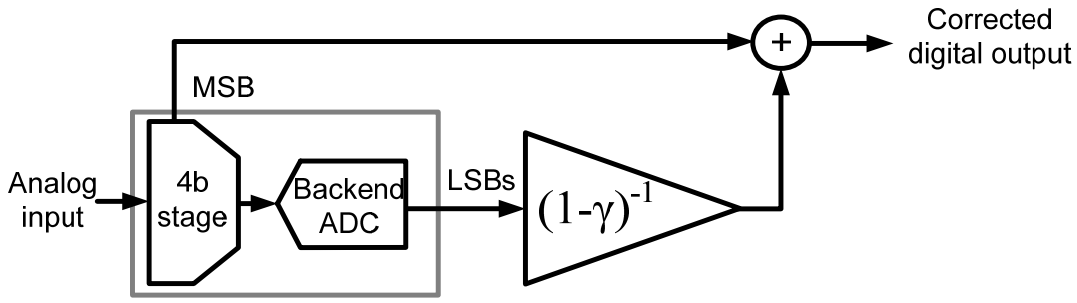

Fig. 3-5: Gain error correction of 1st pipeline stage

Most stages which resolve 1.5-bits/stage use the circuit topology shown in Fig. 3-6 for the MDAC.

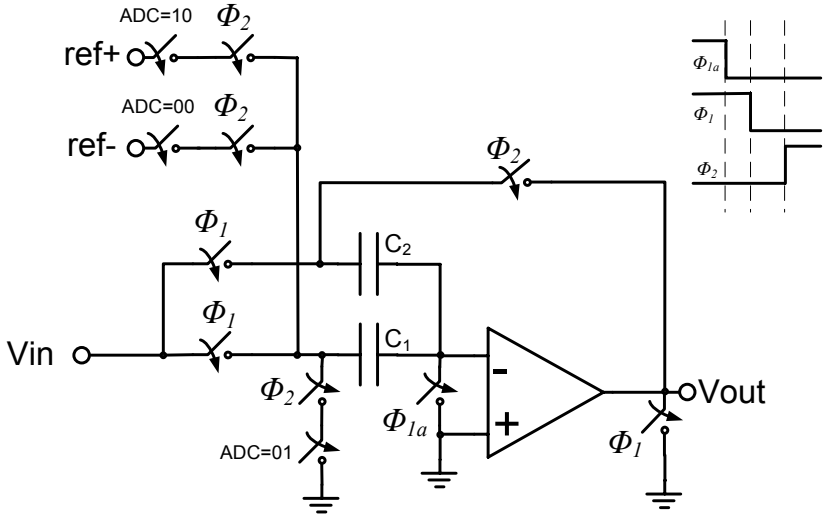

Fig. 3-6: Commonly used circuit topology for 1.5-bit based MDAC

From the topology it is noted that the DAC is implemented with only one capacitor (C_1), thus there is no DAC error in such topologies. As the majority of pipelined ADCs in literature use 1.5-bit/stage topologies, most calibration techniques in literature deal only with the correction of γ. For stages which resolve more than 1.5-bit/stage however, the DAC error can be significant.

3.2.4: DAC ERROR CORRECTION

From section 2.8.5 it was shown that capacitor mismatch in the DAC results in unique missing codes at every MSB transition, yielding substantial harmonic distortion. Since unique errors are produced at each MSB transition a gain calibration scheme (which effectively only corrects for the average number of missing codes at each transition) cannot be used to digitally correct the ADC output. To correct DAC errors (e.g.: [23], [24], [25]) a separate corrective term for *each* MSB transition is required, significantly increasing the complexity of the correction scheme over gain-only correction techniques. For example, with a 3+1-bit pipeline stage, 15 correction parameters for 16 unique DAC outputs are required to be estimated, whereas a gain-only correction scheme has only one parameter to estimate.

Comparing Fig. 3-3 and Fig. 3-4, it is noted that missing codes produced by gain errors look the same as missing codes produced by DAC errors where the DAC error is constant at every MSB transition. Thus in a DAC calibration scheme (where the missing codes are corrected as a function of *each* MSB), the gain errors are also corrected in addition to DAC errors. Fig. 3-7 illustrates an example of a system which corrects for both DAC and gain errors in the first stage of a pipelined ADC, when the non-idealities δ(i) are already known.

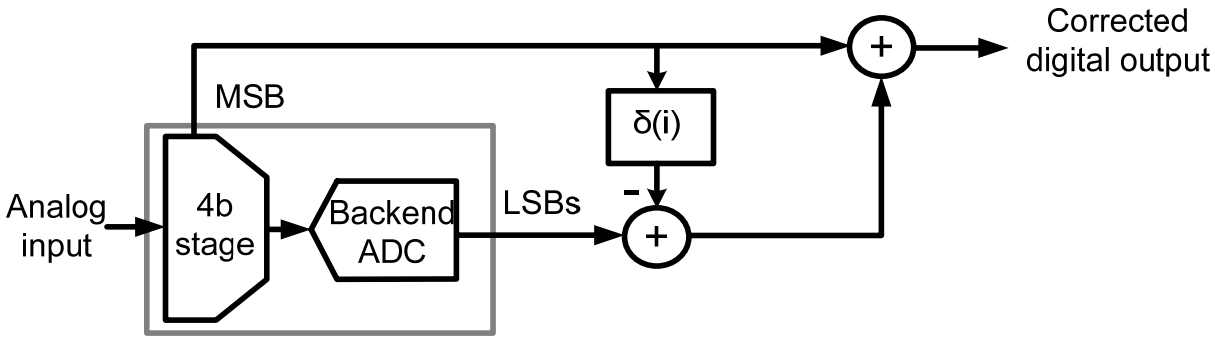

Fig. 3-7: Correction of gain and DAC errors in 1st pipeline stage

3.2.5: FOREGROUND CALIBRATION

Sections 3.2.3 and 3.2.4 discussed how gain and DAC errors can be corrected when the amount of error is already known in advance. In reality however the error is unknown to the designer before fabrication; furthermore the magnitude of each error source varies from chip to chip due to process variation. Thus a scheme to adaptively measure the unknown and unique error sources in an ADC needs to be implemented. In Fig. 3-8, a foreground calibration scheme is shown.

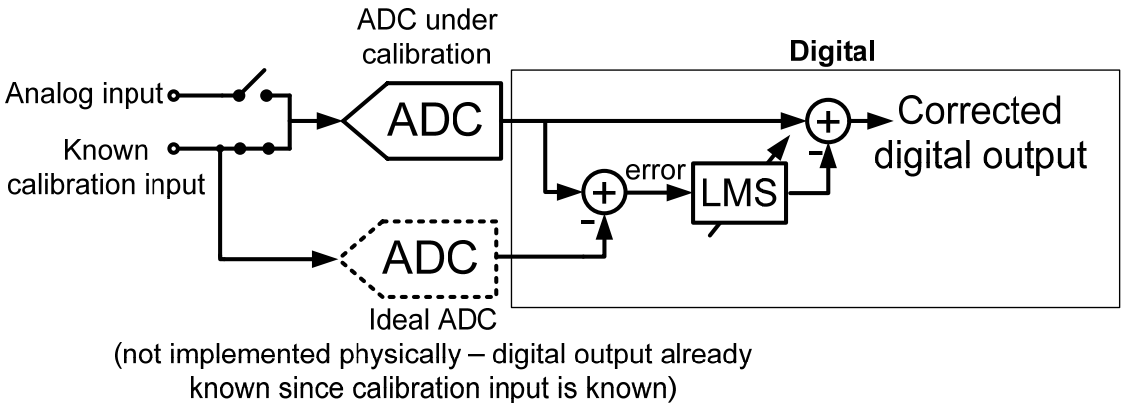

Fig. 3-8: Principle of foreground calibration

Foreground calibration estimates the unknown errors sources by interrupting normal ADC operation and applying a known input sequence to the ADC. By comparing the output of the ADC to the expected ADC output under ideal conditions (i.e. no non-idealities) the impact of each error source can be measured and corrected. Examples of foreground calibration in publications can be found in [26] and [27].

The advantage of a foreground scheme is that calibration can be achieved within a small number of clock cycles, since the error signal labeled in Fig. 3-8 is highly correlated with the error sources causing the missing codes. The disadvantage of foreground calibration is that the ADC is required to be taken offline every time calibration is performed, which in some applications may not be possible.

3.2.6: BACKGROUND CALIBRATION

Background calibration continuously measures and corrects the effect of non-idealities in a pipeline stage, thus has the significant advantage that the ADC is not required to be taken offline to perform calibration. As such the vast majority of calibration based publications are focused on background techniques.

Several topologies have been proposed recently to implement background calibration, where the vast majority of the schemes use a statistics based approach to realize calibration. In a statistical scheme, the input of the pipeline stage under calibration is combined with a known pseudo random sequence, where by correlating the digital output of the ADC with the known pseudo random sequence, the impact of missing codes can be determined. To avoid significantly altering the ADC output spectrum, the pseudo-noise sequence is typically made very long to avoid correlations with the analog input, as well as small in amplitude so that the injected pseudo random sequence which appears as an additional white noise source at the output only consumes a small portion of the dynamic range. Fig. 3-9 shows the basic principle of background calibration.

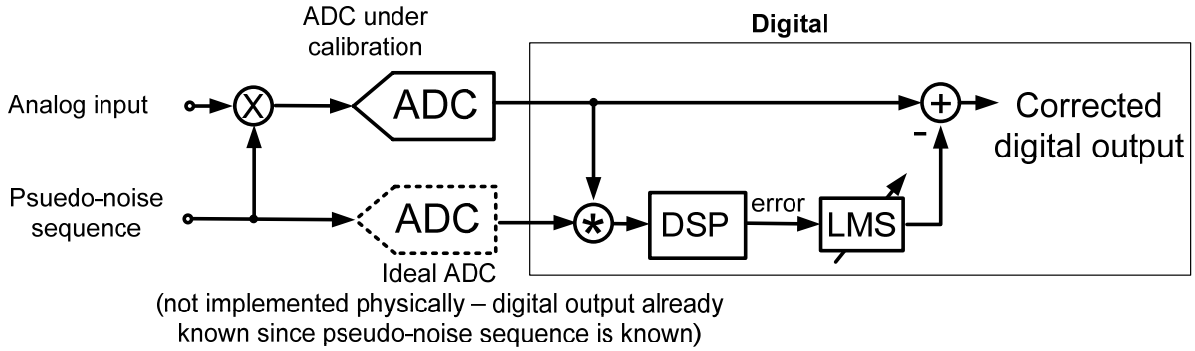

Fig. 3-9: Principle of background calibration

With statistics based background calibration schemes however, since the digital output of the ADC is highly correlated with the analog input and weakly correlated with the pseudo random sequence, a large number of clock cycles are required to accurately extract the pseudo random sequence from the digitized analog input in the ADC output. For example, in [25] ~10^7 cycles were required to achieve 13-b linearity, and in [24] ~10^8 clock cycles were required to achieve >14b linearity. In [28] it was shown empirically that statistical techniques required on the order of 2^{2N} clock cycles to calibrate gain errors only. For 11-bit linearity approximately 4 million clock cycles are required to only correct gain errors using statistics-based background calibration. Thus while background schemes are popular as they enable

continuous ADC operation, the calibration time of background approaches is very lengthy. Example publications which use background calibration are: [29] to [39]

It is noted that in DAC calibration schemes, all MSB bits are required to be exercised to measure each unique DAC error (as the DAC error is different for each MSB). Since an MSB bit can change at most only once every clock cycle, DAC calibration schemes can require M-1 times more cycles to converge than gain calibration schemes (where M is the number of unique MSB outputs).

3.2.7: RAPID CALIBRATION OF ADC ERRORS

In an industrial environment where ICs are mass produced, ICs are tested for functionality by automated testers. In ADCs which use background-statistical techniques to achieve calibration, long calibration times can lead to excessive test times thus limiting IC production throughput and hence revenue. For example, with 4 million calibration cycles, even with a reasonably high sampling rate of 40MS/s, $1/10^{th}$ of a second would be required at minimum to test each ADC. For higher resolution and/or lower speed ADCs the test time can be much higher [28]. In the interest of larger production throughput it is highly desirable to reduce calibration time.

Reducing calibration time has become an active area of research over the past few years. One topology which has proven to be highly effective in reducing calibration times in background schemes is the 'dual-ADC' or 'split-ADC' approach [33], [28], [37]. As shown in Fig. 3-10, the split-ADC takes a single ADC and splits it into two almost identical ADCs where each ADC has half the area, and half the thermal noise floor (thus half the power) of the overall ADC. The final ADC output is derived by taking the average of each ADC output, hence power and area of the split-ADC topology to a first order are not increased over a conventional ADC [28].

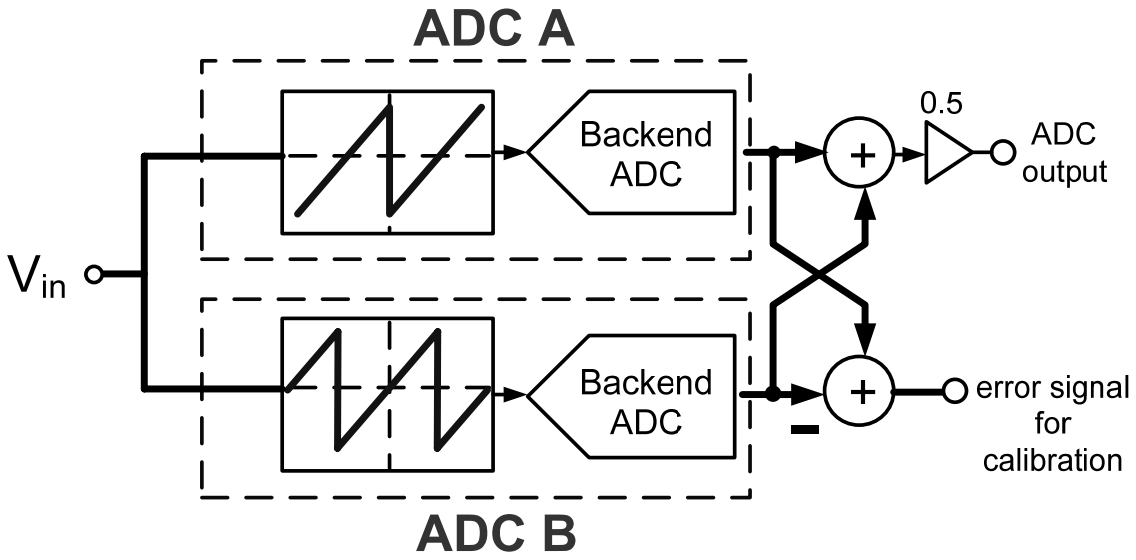

Fig. 3-10: Split-ADC topology

Each ADC is identical, except the residue transfer curve of the stage under calibration in one ADC is designed differently than the other. As a result when the ADCs are free of errors both ADCs produce the same output, however when errors are included each ADC produces different outputs. Since the analog input effectively appears as common mode to the split-ADCs, the error signal which is formed by the difference of the two ADCs is very weakly correlated to the analog input. However the error sources are very highly correlated with the difference in ADC outputs (i.e. error signal) due to each residue transfer curve being designed slightly differently than the other signal [33]. Thus error sources can be estimated very quickly in the background by only looking at a small number of clock cycles of the error signal.

It should be noted that the split-ADC technique while very effective at reducing calibration time, has thus far only been demonstrated to correct gain errors in 1.5b/stage based ADCs. In multi-bit pipelined stages (i.e. ≥2 bits resolved per stage), the effect of DAC errors also needs to be measured and corrected. To date all DAC background calibration techniques in

literature use statistical techniques to achieve calibration, and hence suffer from very lengthy calibration times. As noted in section 3.2.4, since DAC errors are a function of the MSB of a pipeline stage, DAC error measurement can take much longer than gain error calibration, hence in an industrial environment it is even more critical for calibration schemes which correct DAC errors to have short calibration times.

3.3: POWER REDUCTION TECHNIQUES IN PIPELINED ADCS

In many electronic systems achieving low power consumption is a key performance metric. In mobile systems low power consumption allows for increased battery life, and in wired systems low power consumption can in some cases enable low cost packaging - in both cases allowing for more cost effective and thus attractive system solutions. As a result a large percentage of papers published on pipelined ADCs discuss techniques to reduce power consumption. In this section the current state of the art in power reduction techniques in the areas most relevant to this dissertation will be discussed: the elimination of the front-end S/H, power scalability, and the use of non-opamp based amplifiers to reduce power.

3.3.2: FRONT-END S/H REMOVAL

As discussed in section 2.9, a front-end S/H is commonly used to ensure the MDAC and sub-ADC in the first pipeline stage see the same input. However as was noted earlier, this results in a significant increase in power consumption. As a result several techniques have been proposed which allow a pipelined ADC to work without a front-end S/H and thus reduce power.

In previous publications, the front-end S/H was eliminated by relying on the redundancy of the first pipelined stage (e.g. [40], [41], [42]). For example, in a 1.5b/stage architecture the comparator offset in the sub-ADC can be as large as $V_{ref}/4$ (where V_{ref} is the maximum peak voltage of the input). Thus so long as the difference in input operated on by the MDAC and

sub-ADC (as shown in Fig. 2-24) is less than $V_{ref}/4$, the effect of sampling skew appears as an input-referred offset on the sub-ADC comparator, and the effect of the offset is eliminated by the redundancy of the first pipeline stage. Hence the front-end S/H can be eliminated without any further modification to the ADC. From eqn. 2-21, assuming a sinusoidal input to the ADC with a maximum peak voltage, and assuming no inherent offset in the sub-ADC comparators, the maximum allowable skew time that can be corrected by the redundancy of a 1.5b stage is $(8\pi f_{in})^{-1}$. Thus high frequency inputs require low skew between sub-ADC and MDAC; for f_{in}=270MHz the maximum skew allowable is 140ps. The practical allowable skew between sub-ADC and MDAC however must also take into consideration the mismatch of the comparators in the sub-ADC (which in a 1.5b stage are typically made large) as well as a design safety margin – with these factors included the skew must be significantly lower than 140ps. In a sub-sampled ADC where the input frequencies are necessarily very large, clearly great care [41] and thus increased time (hence increased costs) must be taken in layout and simulation if the first pipeline stage's redundancy is used to eliminate the front-end sample-and-hold such that the system works over all process corner variations and temperatures.

3.3.3: POWER SCALABILITY WITH RESPECT TO SAMPLING RATE

The growing demand for mobile systems which can provide multi-standard compatibility in a single solution [43] has stimulated much research in systems which allow for multiple design specifications to be met with only one low power design (e.g.: [44]-[50]). A single reconfigurable solution significantly reduces design time (and thus time to market). Sections 3.3.3.2 and 3.3.3.3 discuss published techniques which enable ADC power to scale with different sampling rates.

3.3.3.2 : REVIEW OF PRIOR POWER-SCALING TECHNIQUES

Digital CMOS circuits have a power which explicitly scales with operating frequency according to $\frac{1}{2}fCV^2$. Analog circuits however have a power which by in large does not scale explicitly with sampling rate as static bias currents are used to place analog transistors in the

active region. Since ADCs contain primarily analog circuits, the power of an ADC remains approximately fixed for different sampling rates. A common method to achieve a scalable analog power is to adaptively scale transistor bias currents with the sampling rate (e.g.: [44], [51], [52], [53]). Since more settling time is made available for lower sampling rates, transistor bandwidths can be reduced hence bias currents correspondingly reduced. Although bias current scaling can be effective for a small variation of sampling rates, for a very wide variation of sampling rates the bias currents could be forced to vary by orders of magnitude to maintain a reasonable power consumption, thereby driving MOS transistors to be deep in the weak inversion region. As MOS transistors in the weak inversion region are more susceptible to mismatch effects than devices in strong inversion [54], in the interest of yield and reduced sensitivity to external noise sources it is highly desirable to avoid operating deep in the weak inversion region. In [45] a Current Modulated Power Scaling (CMPS) technique was presented which allowed for very wide variations of power with sampling rate without commensurate variations in bias currents.

3.3.3.3 : REVIEW OF CMPS

As illustrated in Fig. 3-11, CMPS achieves power scaling by digitizing the analog input within a fixed time interval t_{ON}, then powering off the ADC for t_{OFF} until the next input sample is required to be sampled.

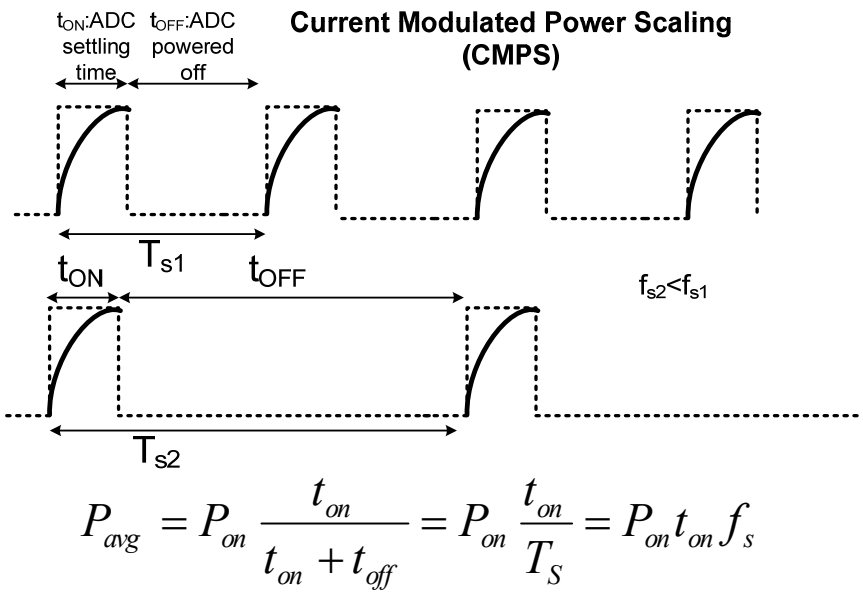

$$P_{avg} = P_{on}\frac{t_{on}}{t_{on}+t_{off}} = P_{on}\frac{t_{on}}{T_S} = P_{on}t_{on}f_s$$

Fig. 3-11: Illustration of Current Modulated Power Scaling (CMPS)

Thus with CMPS, by changing how long the ADC is powered off, different sampling rates with linearly scaled power consumption (due to time averaging) are achieved. The settling times for each pipeline stage (t_{ON}) remain constant for different sampling rates, hence bias currents are thus constant for different sampling rates. Hence linearly scaled power for different sampling rates can be achieved using CMPS without commensurate variations in bias current, thus operating in the weak inversion region can be avoided using CMPS.

From [45], it is noted that when the pipelined ADC completely powers off its stages are reset, thus when the next analog input is sampled, the input is required to be processed by each pipeline stage before a valid digital output can be generated. As a result, when using CMPS in a pipelined ADC, the ADC effectively operates as an algorithmic ADC, limiting the maximum sampling rate to $1/t_{lat}$ (where t_{lat} is the total latency of the pipelined ADC). From [45] it is noted that the total latency (t_{lat}) includes the delay of the pipeline stages as well as additional clock cycles required to initialize various voltages (e.g.: bias, clock generator) before the ADC can process the next sampled input. Current scaling is used to achieve power scalability for sampling rates not covered with CMPS [45]. Fig. 3-12 illustrates how CMPS

can be used in combination with current scaling to achieve a very wide variation in power with sampling rate.

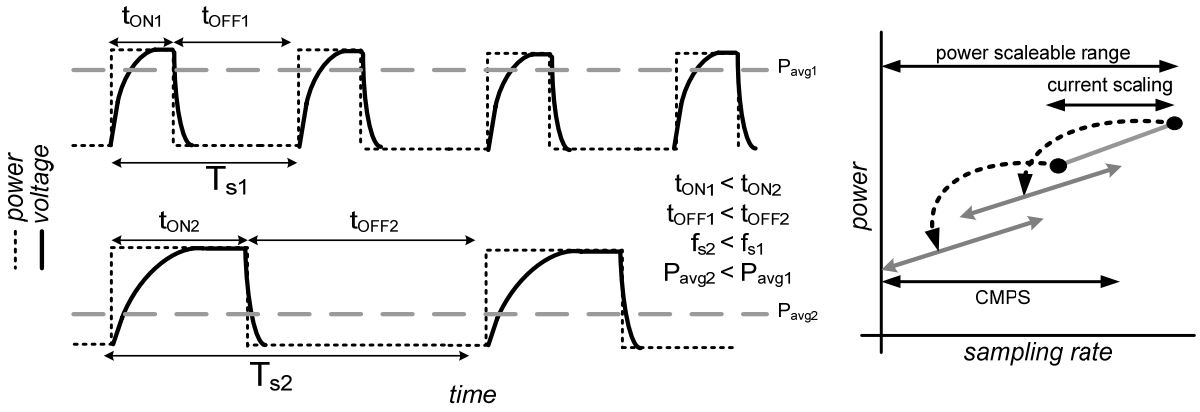

Fig. 3-12: Illustration of CMPS with current scaling

Thus by using CMPS in combination with current scaling a very wide power scalable range can be realized without a commensurate variation in bias currents. This is highly desirable in many systems as the system could be configured for different standards with widely different specifications, while having the power scale with bandwidth.

From [45] it is noted that to achieve high sampling rates, the ADC is required to rapidly power-on so as to digitize the next analog input. As a result analog circuits which are capable of rapidly powering-on and settling to the full accuracy demanded by the ADC are required. As opamps form the main analog circuit blocks in the pipeline stages of the ADC, rapid power-on opamps are required to enable CMPS at high sampling rates [45]. To enable rapid power-on in [45], a novel rapid power-on opamp was proposed and is illustrated in Fig. 3-13.

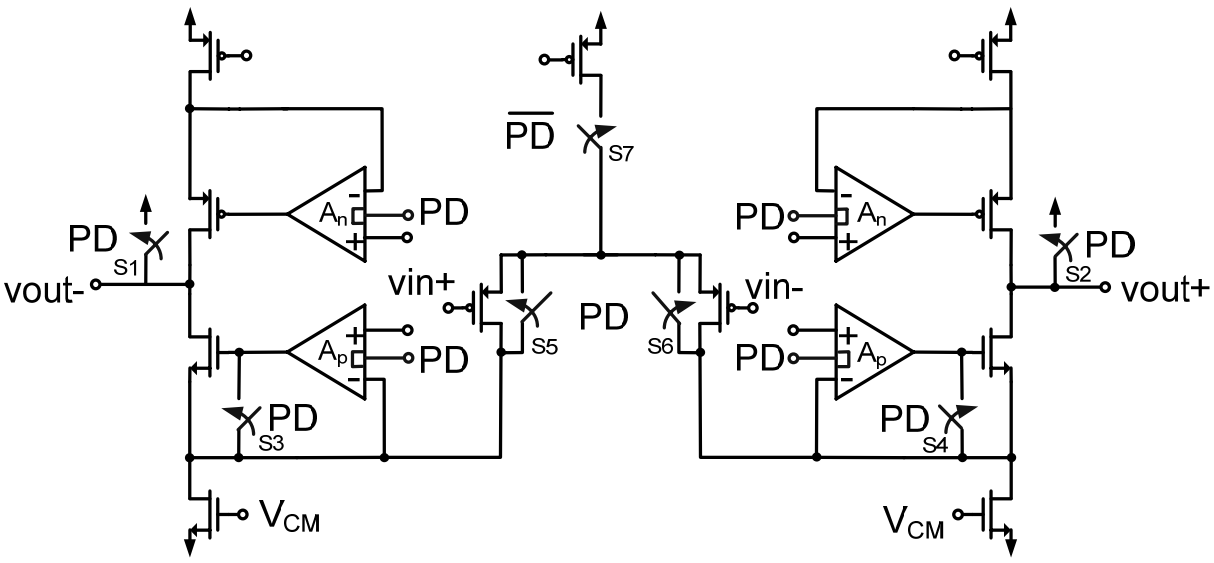

Fig. 3-13: Rapid Power-On Opamp

Rapid power-on is achieved by modulating the opamp's power by switching nodes which are loaded with a small capacitance and driven by a sufficiently large current [45]. This is contrasted with bias-node modulation, where the power of the opamp is modulated by switching bias voltages [55] (e.g.: [56]-[59]). In bias voltage switching the opamp power on time is excessive as bias nodes are typically loaded with large decoupling capacitors and are set by bias circuits which are driven with small currents.

3.3.4: NON-OPAMP BASED PIPELINE STAGES:

The opamp is a key, yet power consuming and poorly adapted to technology scaling block in the construction of a pipeline ADC stage. In sections 3.3.4.4 and 3.3.4.5, alternative non-opamp based pipeline stage topologies which are better adapted to scale with technology and consume lower power than opamp based approaches are reviewed.

3.3.4.4 : OPEN-LOOP AMPLIFIER APPROACH

In section 2.8.3 it was shown that while feedback around an opamp enables a highly linear and precise gain, it comes at the cost of bandwidth reduction in the closed-loop. Namely by closing the feedback loop, the unity gain frequency of the closed system is that of the open-loop divided by 1/β. As a result pipeline stages with large closed-loop gains come at the cost of decreased maximum speed. To overcome this limitation, some researchers have investigated using open-loop amplifiers to reduce power consumption.

In [60] an open-loop technique is used for a 4-bit first pipeline stage in a 12-bit 75MS/s ADC. Fig. 3-14 illustrates the architecture of the 4-bit pipeline stage used in [60].

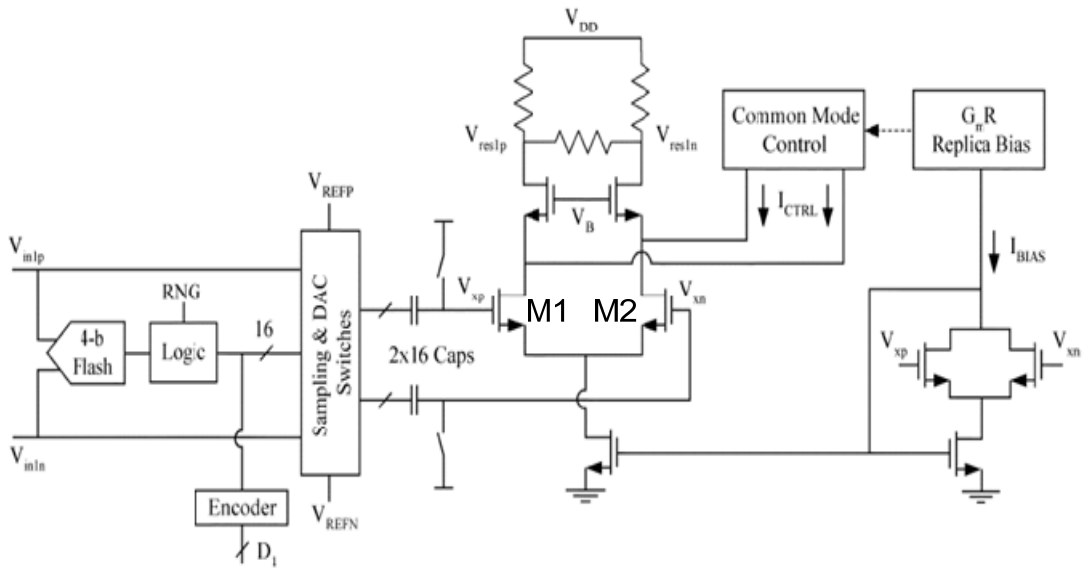

Fig. 3-14: 4-bit pipeline stage using open-loop amplifier (output of stage taken at nodes V_{res1p}, V_{res1n}) – figure taken from [60]

Rather than using an opamp in closed-loop, an open-loop differential amplifier is used to achieve the desired gain of 8x in the pipeline stage, where the gain is set approximately by

$g_m R_{load}=8$, where g_m is the transconductance of the differential pair M1-M2, and R_{load} the load impedance seen by the differential pair.

The unity gain frequency of the amplifier in Fig. 3-14 is given by:

$$\omega_t = \frac{g_m}{C_L} \qquad \textit{(eqn. 3-1)}$$

where ω_t is the unity gain frequency of the differential pair M1-M2, and C_L the load capacitance seen at the output of the amplifier (i.e. nodes V_{res1p}, V_{res1n}). In contrast a closed-loop topology with a gain of 8x would have had at best a unity gain frequency $1/8^{th}$ of that in eqn. 3-1. As a result it was shown in [60] that an open-loop approach yielded a power savings of 60% compared to a closed-loop opamp based approach.

A significant advantage of using an open-loop topology is a simple differential pair topology can be used rather than a more time consuming to design opamp. With an open-loop amplifier the designer is free from worrying about stability criteria such as phase margin. Furthermore it is conceivable that with an open-loop amplifier it would be easier to operate with low supply voltages as strictly speaking no stacked devices (which reduce signal swing – e.g. output transistors in a folded cascode) are required to ensure large a DC gain as is the case in opamps.

The tradeoff in using an open-loop topology is that the linearization and process desensitization found in feedback topologies are eliminated. In closed-loop topologies a large open-loop gain from the opamp creates a virtual ground at the input of the opamp and thus significantly reduces the input signal swing to the input differential pair of an opamp, resulting in a very linear closed-loop transfer curve. In an open-loop topology however the input to the differential pair of e.g. Fig. 3-14 has a signal swing $1/8^{th}$ that of the output. As a result the transfer curve of an open-loop approach suffers from a significantly increased non-linearity as shown in Fig. 3-15

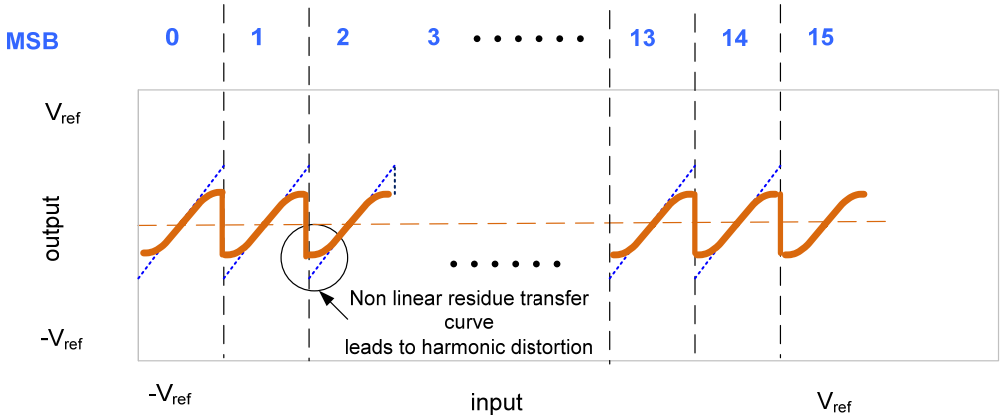

Fig. 3-15: Non-linear transfer curve for residue transfer function – ideal transfer curve shown in dashed lines

Furthermore in a closed-loop system, the gain is typically set by the ratio of capacitors which can be determined to a high precision, whereas in an open-loop approach the gain is set by the product of g_m and R_{load} – a product which varies significantly with process and temperature.

To compensate for the nonlinearity and gain variation in [60], a calibration scheme was developed which corrects the nonlinearity of the open-loop amplifier in the digital domain by using an inverse nonlinearity function, $f^{-1}(x)$, as shown in Fig. 3-16.

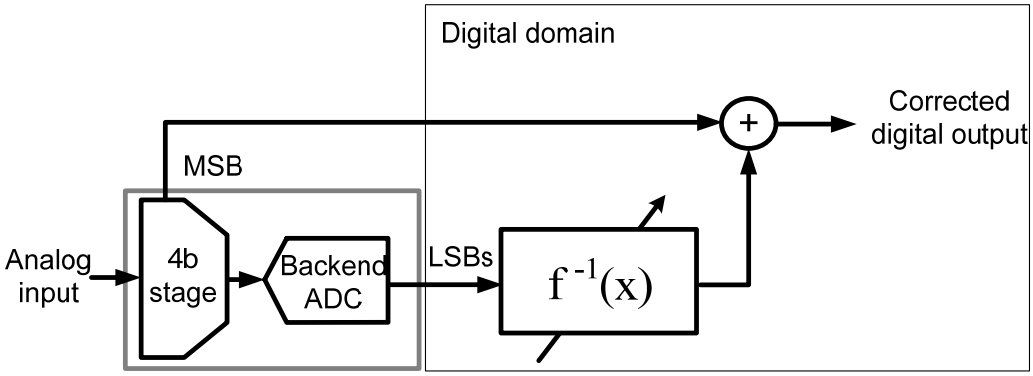

Fig. 3-16: Nonlinearity correction scheme

In [60] a statistics based digital background calibration scheme was used to estimate the value of $f^{-1}(x)$. The specific details of this calibration scheme are not discussed here, however it is noted that the calibration engine was mathematically intensive due to the large number of calculations involved in digitally estimating the nonlinear function $f^{-1}(x)$ to a precision higher than the 12-bits desired from the ADC. Furthermore, the digital calibration consumed $1/3^{rd}$ the overall power reported in [60], which was fabricated in a 0.35µm CMOS process. Thus while the unity gain frequency of the amplifier in an open-loop approach is dramatically increased for a given power, a large amount of power in the digital domain due to the complex DSP is subsequently required to enable the bandwidth enhancement. It is noted however that process scaling favors digital devices from a power and area perspective, hence it could be argued that trading analog power and complexity with digital power and complexity is a favorable one. It is noted that since the calibration scheme in [60] was a statistics based approach, a very large number of clock cycles were required in the calibration (on the order of 10^7 clock cycles to achieve 12-bits of resolution).

3.3.4.5 : COMPARATOR BASED SWITCHED CAPACITOR CIRCUITS

Like open-loop amplifier based pipeline stages, Comparator Based Switched Capacitor (CBSC) circuits [61] overcome limitations from technology scaling by replacing the opamp in closed-loop with a more power efficient structure. In [61], the CBSC technique was used in the MDACs to develop a 2.5mW, 10-bit 8MS/s pipelined ADC.

In switched capacitor circuits the fundamental job of the opamp is to provide a virtual ground to ensure accurate charge transfer from sampling to feedback capacitors. Instead of using an opamp to generate a virtual ground, CBSC circuits use a comparator in a feedback loop to effectively emulate the functionality of an opamp with a large DC gain. Fig. 3-17 illustrates a CBSC circuit which provides a gain of two.

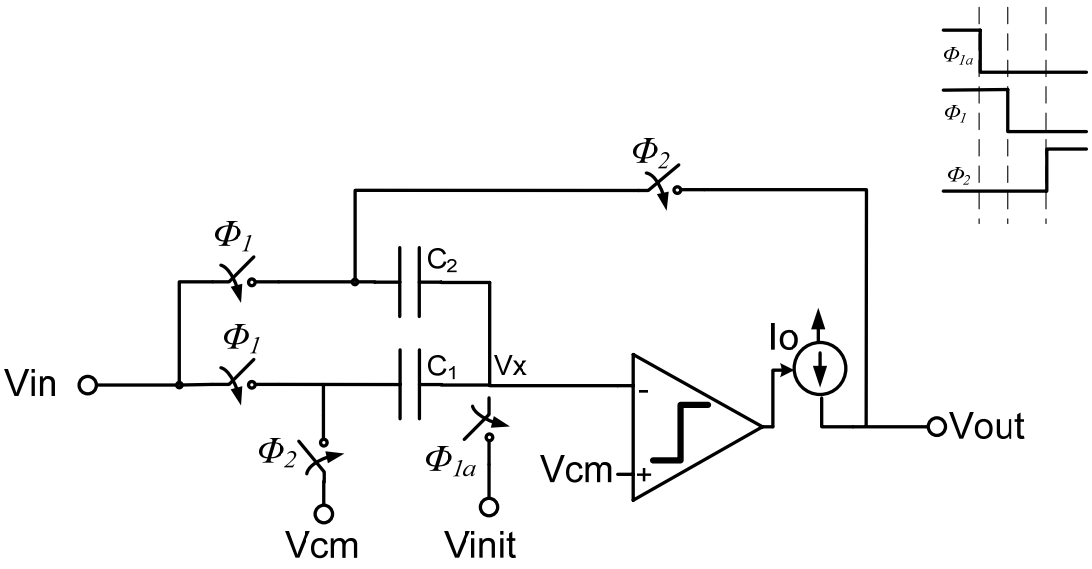

Fig. 3-17: CBSC gain of 2x circuit

In Fig. 3-17, during Φ_1, the input is bottom plate sampled on capacitors C_1 and C_2, and node Vx is initialized to a voltage below the value Vx would be if an opamp were used instead of a comparator, i.e. Vx is initialized to be below Vcm. During Φ_2, the current source at the output, Io, is turned on and Vout and node voltage Vx subsequently increase in voltage. When node Vx equals Vcm, the comparator toggles and turns off the current source Io, and the node voltages at Vx and Vout appear the same as if an opamp were used to create a virtual ground at Vx. Thus the CBSC arrangement implements the same functionality as an opamp based arrangement (i.e. gain of 2x). It is noted that in [61] an additional clock phase is

used to ensure a precise charge transfer, however the general principle of the circuit in [61] is the same as shown in Fig. 3-17.

The significant advantage of CBSC is that the topology does not depend on an opamp which would otherwise require a large DC gain and/or large supply voltage to implement a gain of 2x in a switched capacitor circuit. As comparators can be easily designed even with low supply voltages and with transistors that have low intrinsic gain, CBSC is well suited for implementation in deep submicron technologies. Furthermore in [61] it is shown that the CBSC approach has less inherent thermal noise than an opamp based approach and thus is able to use smaller sampling capacitors, hence have lower power consumption. One of the drawbacks of CBSC is thus far the technique has only been shown in silicon in single-ended form, thus is susceptible to common mode noise which could be very large in an integrated system.

3.4: SUMMARY

This chapter reviewed several state of the art pipelined ADC enhancement techniques. Areas of research most relevant to this work were covered in most detail, namely: rapid calibration, sample-and-hold removal techniques, and non-opamp based pipeline stages.

CHAPTER FOUR: RAPID CALIBRATION OF DAC AND GAIN ERRORS IN A MULTI-BIT PIPELINE STAGE

4.1: OVERVIEW

In section 3.2.7 techniques which attain calibration in short time intervals were shown to be of great interest in industry. Furthermore it was discussed that thus far rapid calibration techniques only correct gain error in 1.5b/stage based topologies. In this chapter a technique to digitally rapidly correct for both DAC and gain errors in the multi-bit first stage of an 11-bit pipelined ADC is presented. The proposed architecture [2], [3] uses a split-ADC based approach. A proof-of-concept prototype of the digital based background scheme was fabricated in a 1.8V 0.18μm CMOS process, where measured results show the calibration scheme is able to improve the peak INL of an 11-bit 45MS/s ADC from 6.4 LSB to 1.1 LSB after calibration. Measured results also show the SNDR/SFDR of the ADC improves from 46.9dB/48.9dB to 60.1dB/70dB after calibration. It will be seen that the calibration is achieved in approximately 10^4 clock cycles, which is more than 100x faster than previously published techniques which correct for both gain and DAC errors.

The organization of key sections in this chapter is as follows: Section 4.2 discuses the motivations of the work. Section 4.3 discusses the architecture used to implement rapid calibration of both DAC and gain errors in a multi-bit stage of a pipelined ADC. Section 4.4 discusses the circuits used to implement the ADC of this work. Section 4.6 discusses measured results of a prototype fabricated in a 1.8V 0.18um CMOS process. Section 4.7 concludes the chapter.

4.2: MOTIVATION

In sections 2.8.2 and 2.8.5 it was detailed how ADC accuracy was limited by imperfections such as low opamp DC gain and capacitor mismatch. In section 3.2 various techniques to measure and correct gain and capacitor mismatch were described. In particular the split-ADC topology was discussed as a useful approach to significantly reduce calibration time. However the published reports of the split-ADC approach thus far had only showed successful implementation in 1.5b/stage based topologies. In this work the split-ADC technique is expanded to also rapidly correct DAC errors in a multi-bit pipelined ADC stage. As discussed in section 3.2.4, since gain errors appear as a constant DAC error, the DAC calibration scheme proposed in this dissertation also corrects gain errors in a pipelined stage. Hence the proposed topology of this work is a more general or superset solution to those discussed previously in literature.

4.2.2: WHY ARE DAC ERRORS IMPORTANT TO CORRECT?

As discussed in section 2.8.4 ADCs with a low thermal noise floor require large capacitors to suppress thermal noise. Since capacitor mismatch decreases with increasing capacitor area, it could be argued that in high resolution ADCs the DAC errors are automatically addressed by virtue of using large sampling capacitors. In many applications however, the ADC does not need a very low thermal noise floor, however does need a very linear transfer characteristic – as discussed in section 2.2, OFDM applications are a prime example of this. From section 2.2 it was shown that for each individual signal in an OFDM spectrum, the SNDR was limited primarily by in-band harmonic distortion, rather than the thermal noise floor, hence small sampling capacitors in theory could be used to minimize power consumption. On the other hand, as shown in section 2.8.5 large sampling capacitors are still required in OFDM type applications to suppress the effects of capacitor mismatch. From section 3.2.4 it was shown that calibration techniques could be used to compensate for the effect of capacitor mismatch and thus enable a designer to use small sampling capacitors for an ADC in an OFDM type application. However it was shown in chapter three that present state of the art DAC

calibration techniques all use statistical based schemes and thus suffer from lengthy calibration times, which from a production standpoint was shown in section 3.2.7 to be unfavorable. Thus this the goal of this work was to develop a DAC calibration scheme, hence enabling low power ADCs in certain applications, where the calibration could be achieved in a very short time interval, and thus be very useful in an industrial environment.

4.3: PROPOSED RAPID DAC + GAIN CALIBRATION ARCHITECTURE

The calibration scheme of this work uses a dual-ADC approach (ADC A and ADC B), to simultaneously process in parallel the same analog input as shown in Fig. 4-1. The final ADC output is generated by the average of the two ADC outputs, thus each ADC is designed with half the total capacitance, hence half the power and area of the overall ADC to meet thermal noise requirements [28].

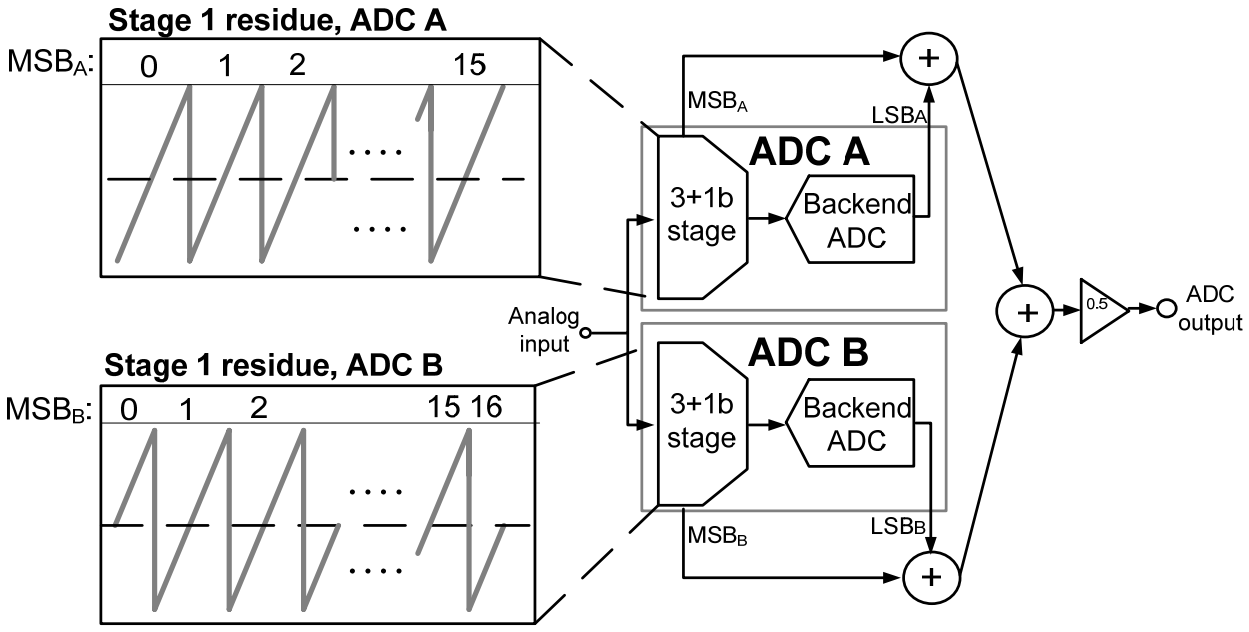

Fig. 4-1: Dual-ADC approach of this work

From Fig. 4-1 ADC A and B are identical except in each ADC the residue transfer function in the first stage is horizontally offset from the other by approximately ½ MSB.

4.3.2: PROPOSED MEASUREMENT OF MISSING CODES DUE TO DAC AND GAIN ERRORS

Fig. 4-2 illustrates the transfer function of key outputs from each split-ADC. From Fig. 4-2 if the analog input to the ADC is such that MSB_A=i, then MSB_B is either i or i+1. The offset between the digital outputs of ADCs A and B for the range of analog inputs where MSB_A=i and MSB_B=i is denoted Δ_{i1}, and Δ_{i2} where MSB_A=i and MSB_B=i+1 respectively, as shown in Fig. 4-2. In an ideal ADC without DAC or gain errors the difference between split-ADC outputs is constant regardless of the analog input, thus $\Delta_{i1}=\Delta_{i2}$. With ideal ADCs A and B $\Delta_{i1}=\Delta_{i2}=0$, however the offset is shown as a constant in Fig. 4-2 for clarity of illustration.

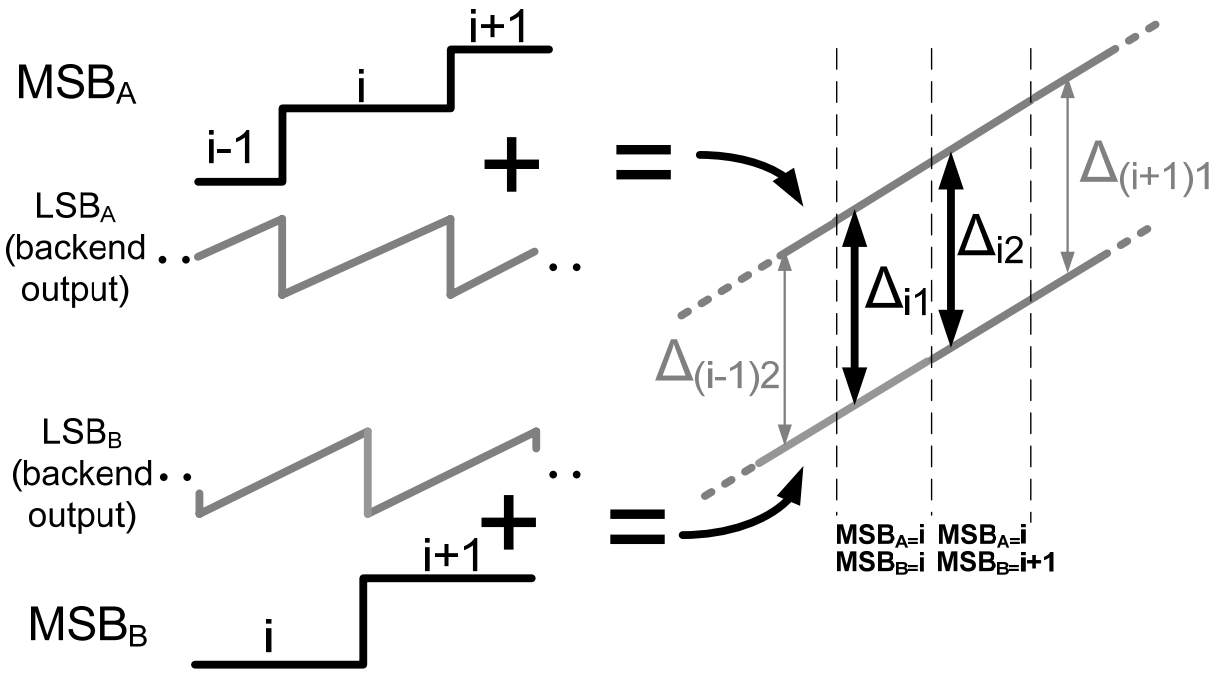

Fig. 4-2: Transfer curves of 1st stage (MSB), backend ADC (LSB) and total ADC outputs from each split-ADC with no errors

If however DAC and gain errors are included, as shown in Fig. 4-3, each split-ADC incurs unique missing codes wherever an MSB changes.

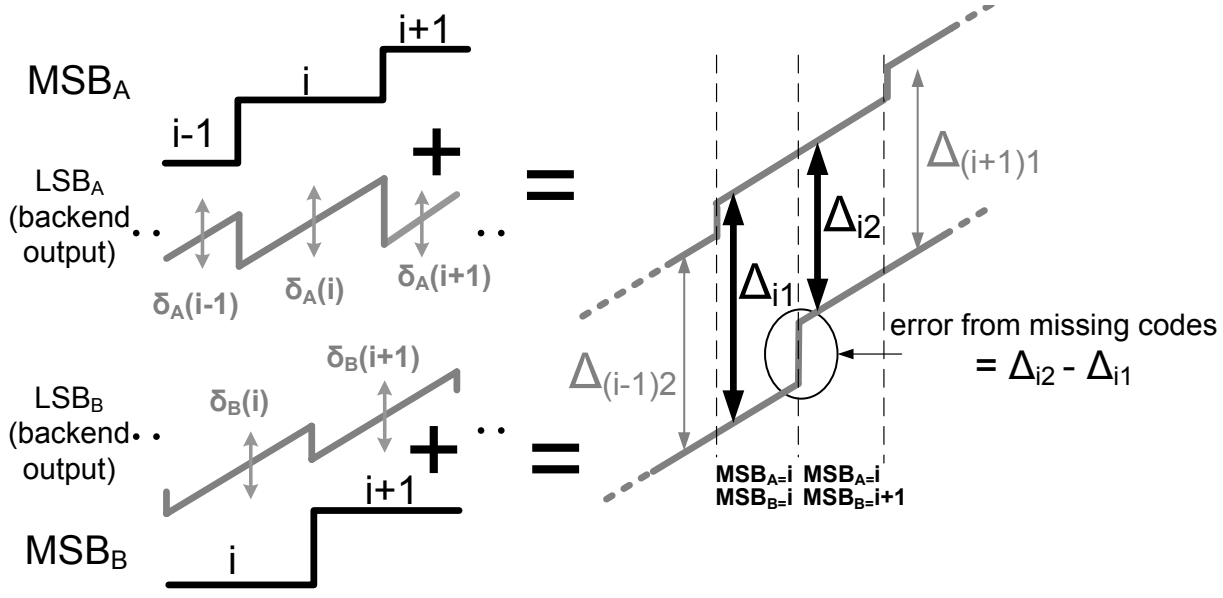

Fig. 4-3: Transfer curves of key ADC outputs with gain, DAC errors included

As the MSB transitions between each split-ADC are staggered however, ADC A does not incur an error in the first pipeline stage for the same range of analog inputs as ADC B. Thus the digital output of ADC A can be used as an *ideal reference* to measure the errors of ADC B. The difference between Δ_{i1} and Δ_{i2} precisely gives the error due to missing codes that occurs when MSB_B changes from i to i+1 as shown in Fig. 4-3. To minimize the effect of zero mean error sources such as thermal noise, Δ_{i1} and Δ_{i2} are averaged before subtraction using a simple first order IIR filter with transfer function $\mu/[1-(1-\mu)z^{-1}]$. In this work $\mu=1/64$ was used, where by implementing μ as a power of 2, multiplication by μ can be implemented using simple bit shifts. In a similar manner the unique error due to missing codes at all other MSB transitions can be measured for ADC B.

Errors due to missing codes for ADC A are measured by noting that Δ_{i2}-$\Delta_{(i+1)1}$ is the error due to missing codes in ADC A when MSB_A changes from i to i+1 as shown in Fig. 4-3. Hence the error due to missing codes in ADC A can be determined using already measured values $\overline{\Delta}_{i2}$ and $\overline{\Delta}_{(i+1)1}$. Errors due to missing codes at all other MSB transitions in ADC A are measured using an identical extension as done for ADC B.

4.3.3: CORRECTION OF MISSING CODES

With the errors from missing codes at each MSB transition measured, each ADC is corrected by shifting each ADC's digital output as a function of the MSB such that overall transfer function of each ADC is free from missing codes due to errors in the first stage as shown in Fig. 4-4 (same done for ADC A).

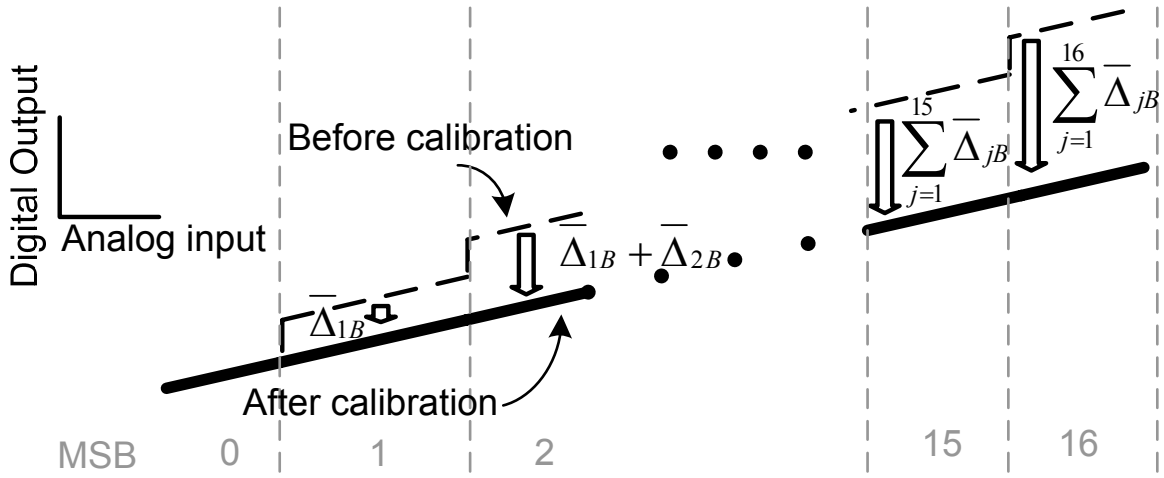

Fig. 4-4: Illustration of how correction terms for ADC B are derived from estimates of missing codes (correction topology of ADC A is similar)

Rapid calibration is achieved as $\overline{\Delta}_{iB}$ is highly correlated with the number of missing codes; only a small number of clock cycles are required to average out the effects of zero-mean noise. As long as the input is sufficiently busy to generate a sufficient number of estimates of Δ_{i1}, Δ_{i2}, for all i, there is no constraint on the shape of the input signal to the ADC. In contrast

statistical techniques use statistical correlations which require many output samples to extract similar information. The full topology of the ADC including digital calibration is shown in Fig. 4-5.

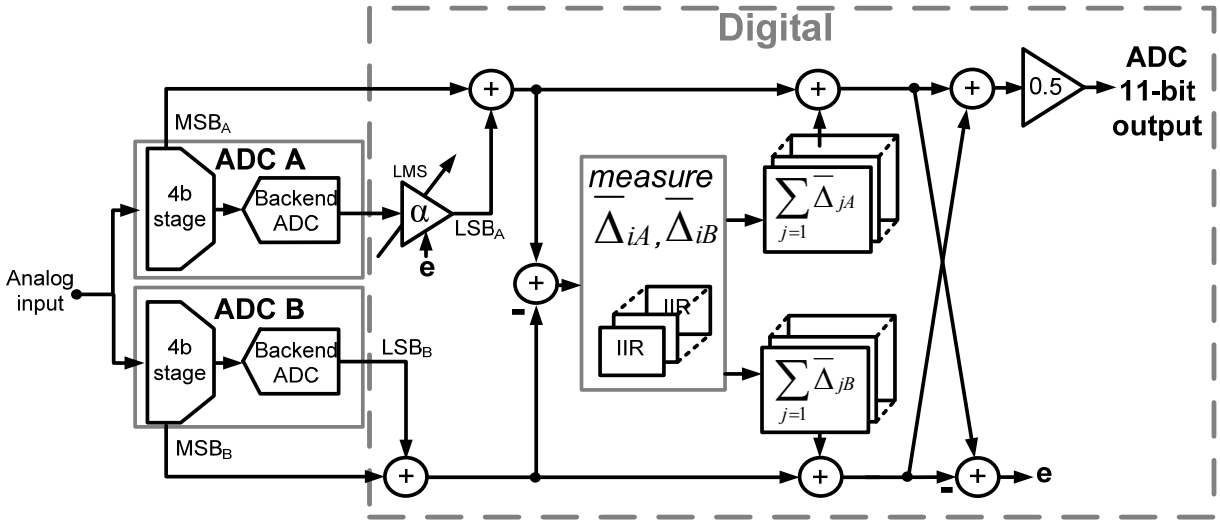

Fig. 4-5: Full ADC topology of this work

It is noted that the approach of this work is similar to background calibration techniques where a more accurate but slower ADC is used in parallel with the ADC under calibration (e.g.: [34]), where by adaptively equalizing the outputs of the fast ADC to the output of the slow ADC calibration can be achieved. In this work since the residue transfer function of one of the split-ADCs is offset from the other, ADC A does not suffer an error in the first stage for the same range of inputs as ADC B, thus one ADC can be used as an ideal reference for the other eliminating the need for one of the ADCs to be more accurate than the other. Hence there is no need to trade higher accuracy with lower sampling rates in the second ADC; both ADCs can operate at the same speed and both ADCs can be used to digitize the analog input. Thus the power of the additional ADC also goes towards lowering the noise floor in the digital output, unlike [34] where the additional ADC (since it operates slower) only aids the correction scheme. Furthermore using the technique outlined in this work, both ADCs are calibrated whereas in [34] only one ADC is.

4.3.4: MISMATCH BETWEEN ADCS

Due to random mismatches between the split-ADCs, the first stage of ADC A in Fig. 4-1 will have MSB transitions that are not exactly ½ MSB apart, as well as a different input referred offset and different stage gain than ADC B.

The calibration scheme of this work only requires the MSB transitions between ADC A and B to be non-overlapping, thus a precise definition of the offset between ADCs is not necessary. As long as the comparators which define the MSB transitions for each split-ADC are designed to have an offset within ½ MSB no errors are produced by the calibration scheme.

Constant offset between the outputs of ADC A and B appears as a common-mode shift in both $\overline{\Delta}_{i1}$, and $\overline{\Delta}_{i2}$. Since the number of missing codes at each MSB transition is measured by subtracting $\overline{\Delta}_{i2}$ from $\overline{\Delta}_{i1}$, the common mode is eliminated and thus input-referred offsets of each split-ADC have no impact in the calibration scheme (under the practical assumption that the offsets are not large enough to saturate the output of the pipeline stages).

To account for an overall gain mismatch between the two ADCs an LMS adaptive gain term (α in Fig. 4-5) is also included [28] which scales the backend code of ADC A so as to keep the outputs of ADC A and B parallel. As ADC B provides an ideal reference for ADC A, the error signal used for the LMS adaptation (which is formed by the difference of the two ADC outputs) is highly correlated with the gain error between ADCs, thus steady state convergence of α occurs within a relatively short time interval. To ensure stability, the adaptation of α is designed to converge more slowly than the loop which measures and corrects the gain and DAC errors of the first pipeline stage.

Ultimately, the accuracy of estimation in errors due to missing codes in the first stage is limited by missing codes and distortion in the backend, which although not addressed in this work can be minimized by also calibrating the backend stages.

4.3.5: SIMULATION RESULTS

To evaluate the functionality of the proposed calibration technique, a model of the system was implemented in Matlab Simulink. Gain and DAC errors were modeled in the first 3+1-bit pipeline stage, along with thermal noise, as shown in Fig. 4-6. The thermal noise floor was designed such that the input referred noise floor was approximately at the 11-bit level.

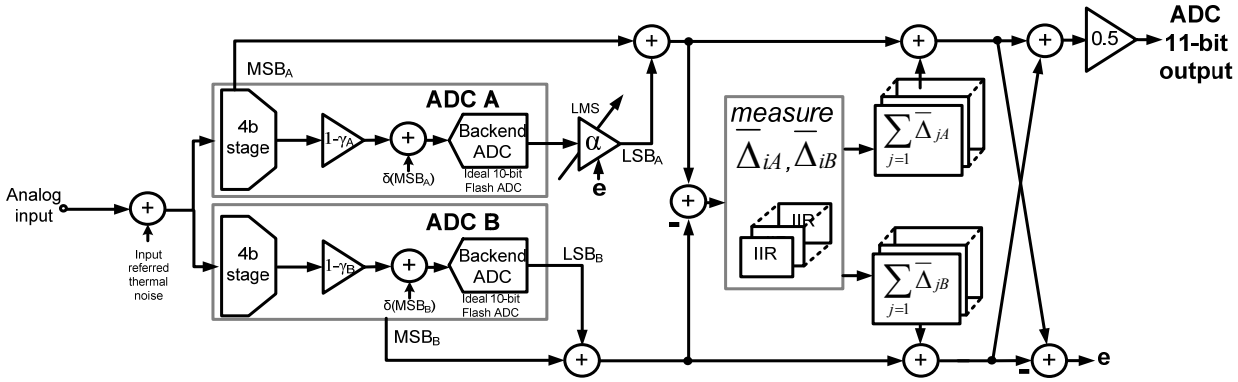

Fig. 4-6: Matlab Simulink test setup for simulation verification

Each dual-ADC in Fig. 4-6 was initialized with a different static gain error γ, and unique DAC errors δ at startup, such that each ADC path was not identical. To concentrate on the effectiveness of the calibration scheme, the backend ADC was modeled as an ideal 10-bit Flash ADC. The digital calibration engine implemented in Simulink emulated the topology shown in Fig. 4-5, where care was taken to ensure all digital correction terms had no more than 14-bits.

Fig. 4-7 shows the FFT plot of the ADC (before calibration) when the first stage in each pipeline stage was designed to have a relative gain error of ~1% and a relative capacitor mismatch ratio also of ~1%. The input to the ADC was a sinusoid at 1/200th the sampling rate so that all key harmonics could be shown without being aliased.

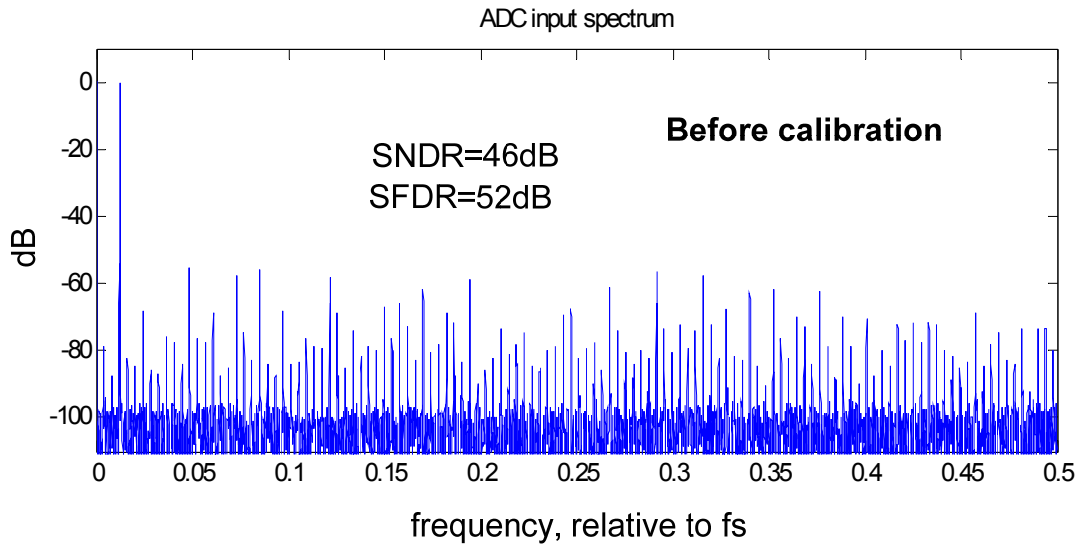

Fig. 4-7: Simulation results with 1% gain and DAC error – before calibration

Fig. 4-8 illustrates the FFT after calibration where the SNDR/SFDR is seen to dramatically improve to 66/85dB.

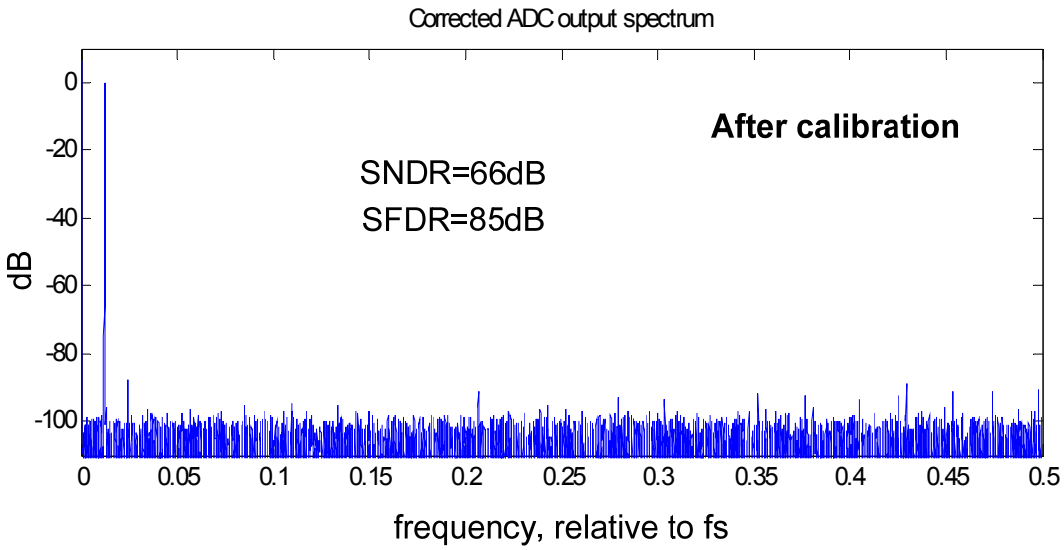

Fig. 4-8: Simulation results with 1% gain and DAC error – after calibration

Fig. 4-9 shows the evolution of ADC SNDR and SFDR with the number of clock cycles, where it is seen that a 23dB in SFDR and 20dB in SNDR improvement can be achieved in less than $1.5x10^4$ clock cycles.

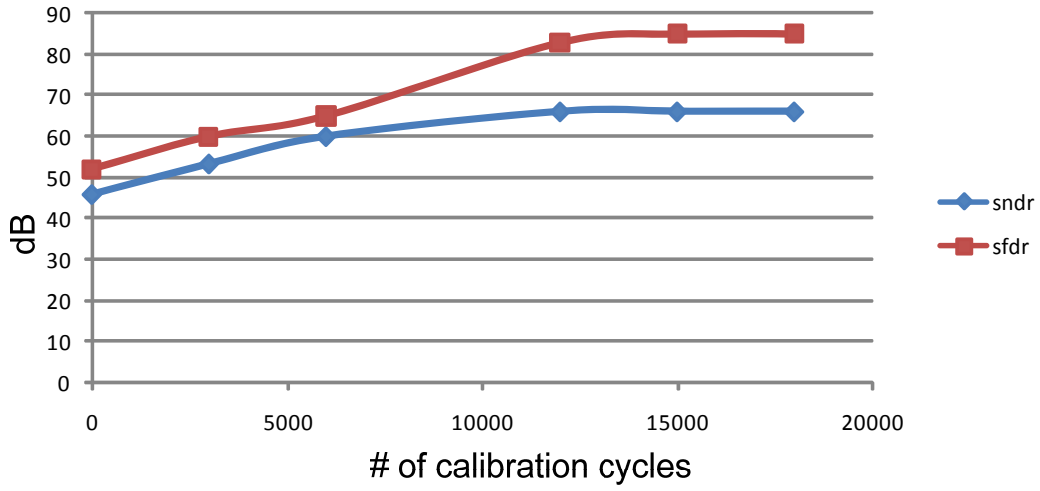

Fig. 4-9: SNDR/SFDR improvement with calibration cycles in Simulink model

To verify that the system worked with random inputs, an additional simulation was performed where a uniform random input was applied to the ADC for $1.5x10^4$ clock cycles and the error terms measured. After $1.5x10^4$ clock cycles the error measurement was frozen and a sinusoid applied to the input of the ADC where an FFT was subsequently taken of the sinusoidal digital output. Fig. 4-10 demonstrates the FFT of the ADC after calibration where it is verified that calibration is achieved even with a random input.

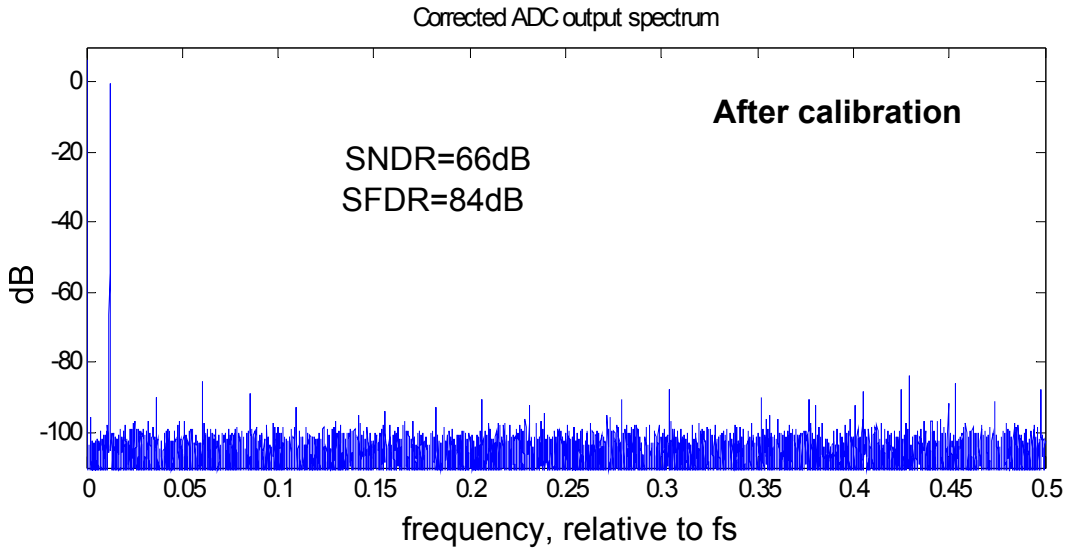

Fig. 4-10: FFT of ADC after calibration driven by uniform random input for $1.5x10^4$ clock cycles

Fig. 4-11 shows the evolution of ADC SNDR/SFDR when calibration correction terms are derived with a random input.

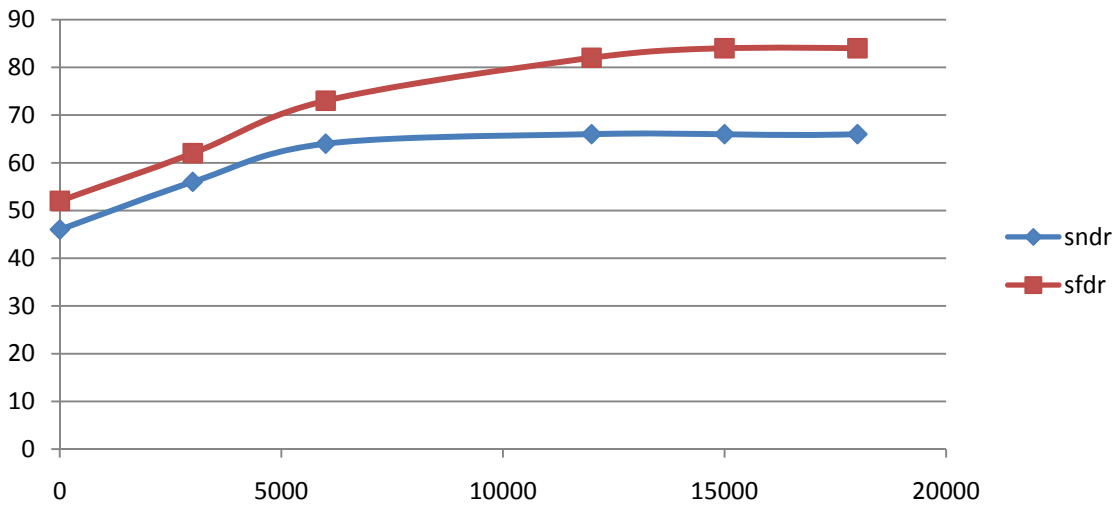

Fig. 4-11: SNDR/SFDR improvement with calibration cycles when ADC driven by a uniform random input

Similar experiments were conducted with different input sources (e.g. triangle wave, saturating input, etc.) and in all cases calibration was achieved so long as the input was sufficiently busy to excite each correction term.

From Fig. 4-7 to Fig. 4-11 it is clear that even with a very large capacitor and gain mismatch on the order of 1%, the proposed calibration scheme is able to provide both effective and rapid correction of circuit non-idealities, thus justifying a fabricated prototype.

It should be noted that the adaptation coefficients were determined heuristically. The value of μ ultimately selected in this work of 1/64 was found to adequately achieve a sufficient amount of error correction for the design specifications of this work for a variety of different inputs and reasonable mismatch levels. A detailed analysis of the impact of correction terms in the proposed architecture is deferred as future work.

4.4: CIRCUIT IMPLEMENTATION

To evaluate the architecture proposed in this work an 11-bit linear pipelined ADC was constructed in a 1.8V 0.18μm CMOS process. From section 4.3.5, although the proposed technique could enable much higher resolution ADCs, 11-bit resolution was ultimately chosen as to achieve a very high linearity for a mid-high speed ADC (i.e. >40MS/s), a costly and lengthy design of a PCB, and IC package would have to be undertaken to ensure sufficient suppression of noise sources (e.g. digital I/O noise, LC bond wire oscillations). Since only the first stage errors are corrected using the calibration technique presented in this chapter, a 3+1 bit first stage was used so that the non-idealities of the backend ADC would be small and not require calibration. The ADC was aggressively targeted to achieve a maximum sampling rate of 100MS/s. However as will be described in section 4.6.2, noise from the digital I/Os limited the maximum sampling rate to 45MS/s. The thermal noise floor of the ADC was designed to be near the 11-bit level for a 1.3V p-p input signal. Although a specific application was not targeted, a quick survey of industry data sheets of ADCs with 11-bit linearity with sampling rates around 45MS/s show for example: CCD imaging, video, portable instrumentation, IF/baseband communications, cable modems, wireless LAN, and medical ultrasound as potential applications.

Fig. 4-12 shows the architecture of the analog portion of the 11-bit pipeline ADC of this work. Scaling was used in the first three pipeline stages to reduce power consumption [13]. An additional three bits were added to the backend in each split-ADC to improve the accuracy in error estimation. As the last stages in a pipelined ADC consume only a small fraction of the total power, adding extra stages to reduce quantization noise has a minimal impact on ADC power.

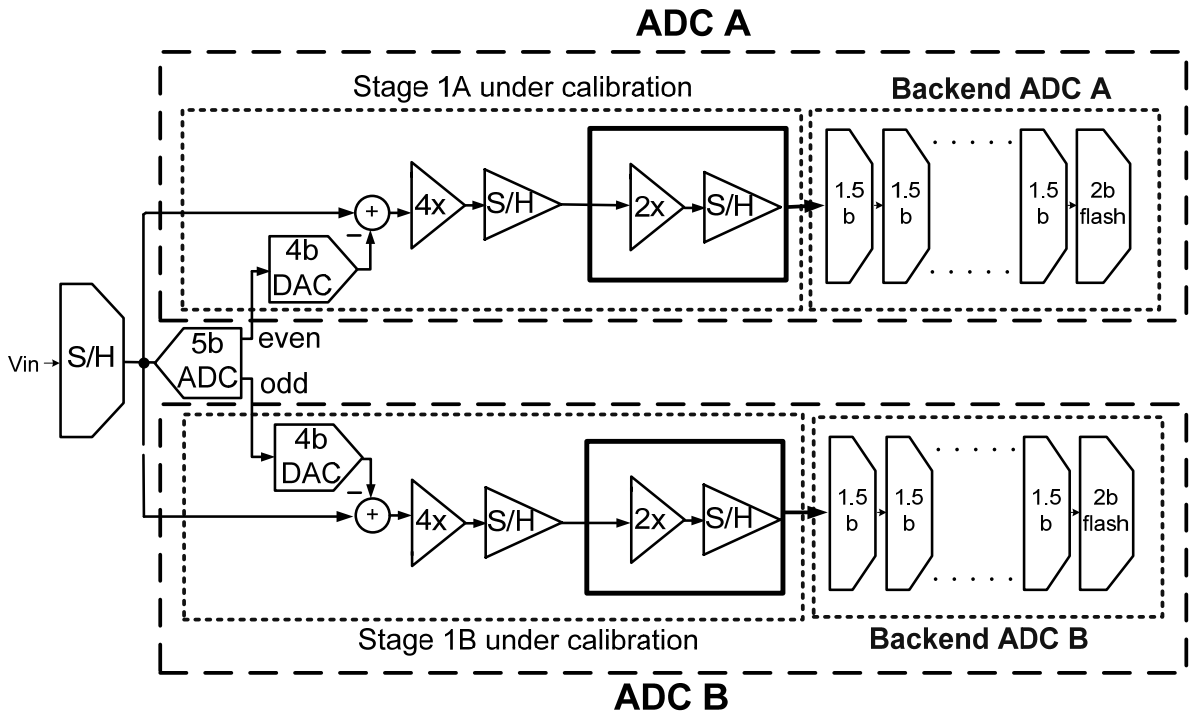

Fig. 4-12: Analog portion of ADC topology in detail

4.4.2: FRONT-END SAMPLE-AND-HOLD

To ensure both ADCs operate on the same analog input, a conventional flip-around front-end sample-and-hold [22] was used and is shown in Fig. 4-13. The flip-around topology has the advantage of a feedback factor near unity hence is power efficient. The sample-and-hold was designed such that the entire ADC had approximately an 11-bit input referred thermal noise floor relative to a 1.3V peak-to-peak input sinusoid.

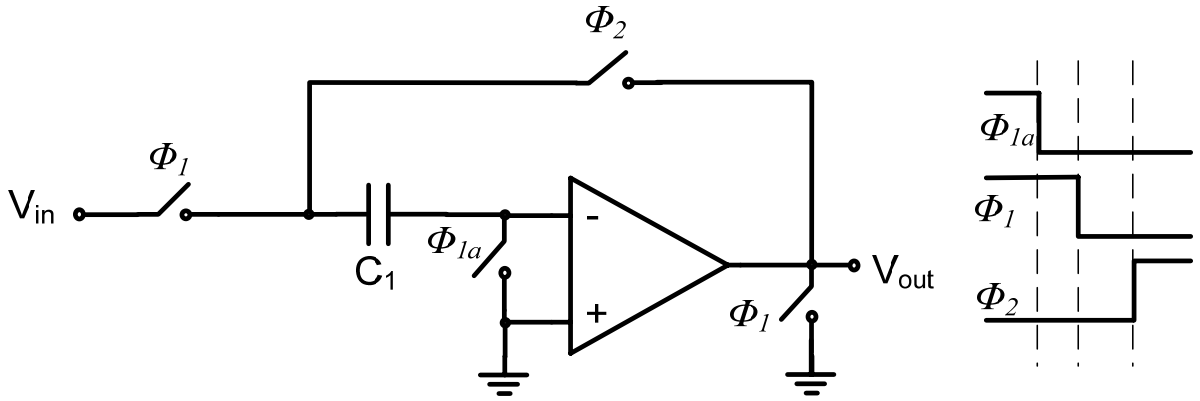

Fig. 4-13: Sample-and-hold topology (implemented fully-differentially in this work)

4.4.3: 5-BIT FLASH ADC

The MSB transitions for each split-ADC were generated by a single 5-bit flash ADC where the even numbered outputs were used for ADC A and odd numbered outputs for ADC B. The 5-bit flash ADC was implemented using an array of comparators where a single comparator is shown in Fig. 4-14.

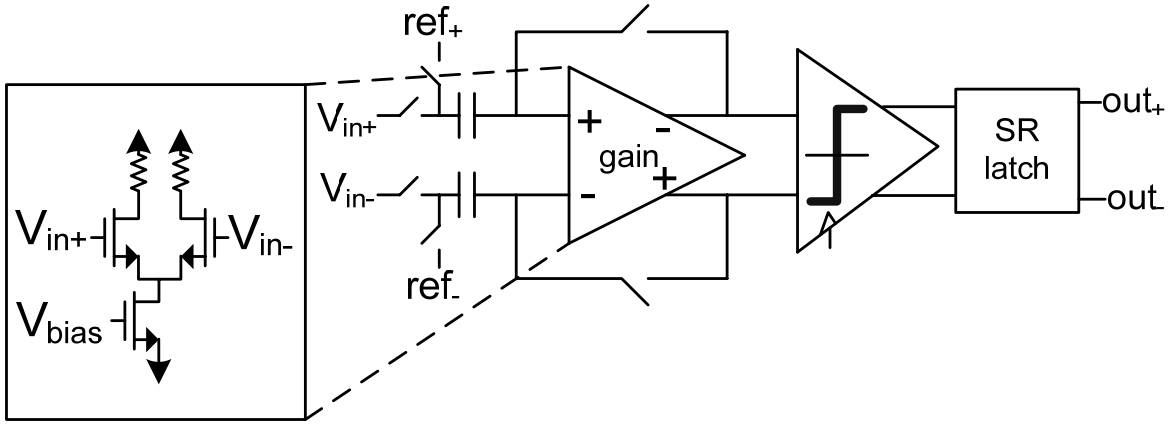

Fig. 4-14: Comparator topology used in 5-bit flash sub-ADC

Each comparator used a preamp before the slicer to reduce kickback noise from the slicer on to the reference voltages which were implemented with a resistor string. A switched capacitor based topology [7] was used to implement the threshold of the comparator as well as reduce offset from the preamplifier.

From Fig. 4-12 it is noted that one of the overheads of a split-ADC topology is an increase in resolution of the sub-ADC of the stage under calibration by 1-bit (to ensure the two ADCs do not have overlapping MSB transitions). However since the power of the sub-ADC is only a small fraction of overall ADC power, increasing the resolution of the sub-ADC by 1-bit results in only a small overall increase in analog power. Furthermore for low resolutions in the sub-ADC, as is the case in this work, the power of the sub-ADC is not limited by thermal noise - thus increased resolution can be obtained with only a marginal increase in power so as to sufficiently minimize comparator offsets.

4.4.4: 4-BIT MDAC

The topology of the multi-bit MDAC in the first stage of ADC A is shown in Fig. 4-15. The MDAC for ADC B is identical except slightly modified to account for a ½ MSB horizontal offset.

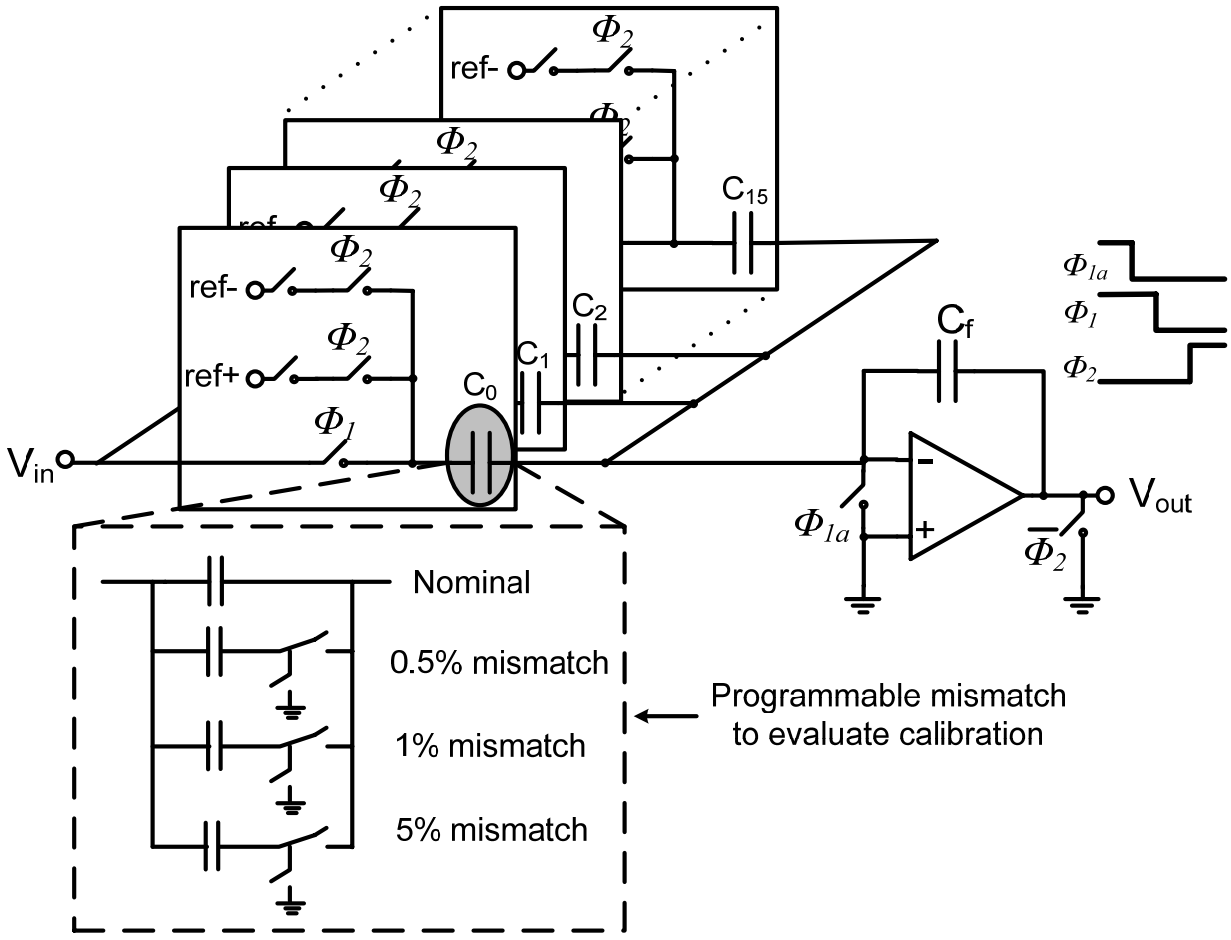

Fig. 4-15: Topology of 1st stage MDAC (implemented fully-differentially in this work)

As capacitor mismatch is a process variation, its impact can only be properly observed by testing thousands of chips for random variation. To minimize the number of ICs tested each sampling capacitor in the first stage of each split-ADC was made individually digitally programmable with 0.5%, 1% and 5% deviations from nominal values as shown in Fig. 4-15. Thus the effectiveness of the calibration scheme of this work could be very quickly evaluated by programming different mismatch profiles in the first stage of each split-ADC. The control bits to set the mismatch profile of the first stage in each ADC were set via an on-chip shift-register.

Since the calibration scheme of this work also corrected gain errors, the DC gain of the opamp in Fig. 4-15 was designed with the minimum DC gain required to adequately suppress distortion from the opamp. The opamp, which had a DC gain of ~50dB, was implemented using a one-stage p-input folded cascode topology, and is illustrated in Fig. 4-16. Common mode feedback was implemented using a switched capacitor based approach [7].

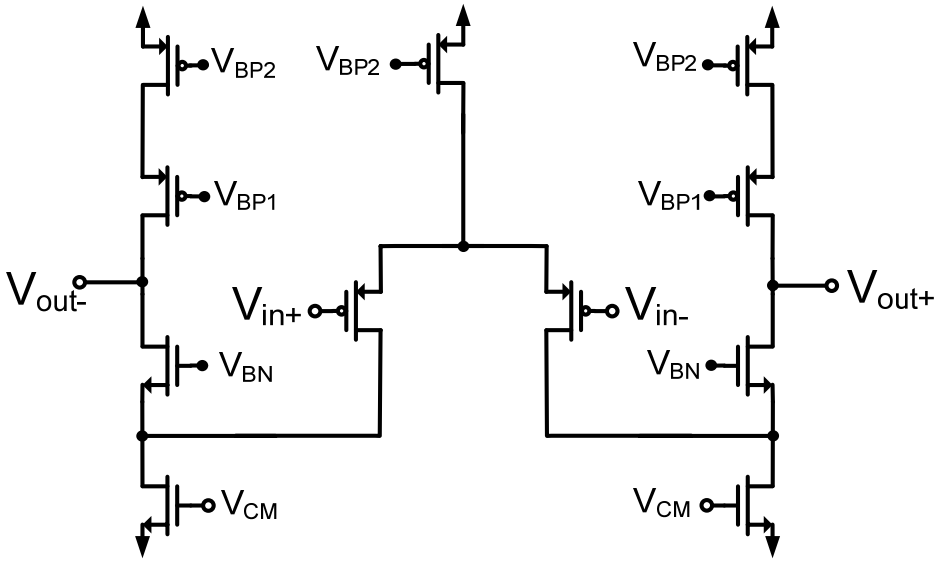

Fig. 4-16: Opamp used in 1st stage MDAC

In a 4-bit pipeline stage including 1-bit of redundancy, the ideal closed-loop gain of the MDAC is 8x. It is noted however that the feedback factor β of the MDAC of Fig. 4-15 is given by

$$\beta = \frac{C_f}{\sum_{i=0}^{15} C_i + C_p + C_f},$$

where C_p is the parasitic capacitance at the input of the opamp. The time constant of the settling time τ of the MDAC given by:

$$\tau = \frac{1}{\beta \omega_{ta}}$$

where ω_{ta} is the open-loop unity gain frequency of the opamp. Thus as the closed-loop gain becomes larger C_f becomes smaller and the feedback factor β becomes more sensitive to variations in parasitic capacitance at the input of the opamp. Furthermore, small feedback factors require the opamp to have a large open-loop unity gain frequency to achieve fast closed-loop settling as noted in section 2.8.3. Opamps requiring high unity gain frequencies can approach the limits of their technology, making the design less power efficient. In this design although measurements are presented at 45MS/s, the design was targeted for a ~2x higher sampling rate. Limitations in the test setup restricted measurements to half the target sampling rate. Thus in this design to reduce the open-loop opamp unity gain frequency (with the higher sampling rate in mind) as well as sensitivity to parasitic capacitors, the closed-loop gain of the MDAC was designed to be 4x rather than 8x. To ensure the backend ADC had a full scale analog input, a pipelined 2x gain stage was inserted between the 4-bit MDAC and backend ADC as shown in Fig. 4-12. The gain error of the 2x stage is effectively lumped into the gain error of the 4-bit MDAC, hence its non-idealities are also calibrated.

It is noted that for ease of implementation in this prototype the calibration scheme requires the input to be sufficiently busy to excite each sampling capacitor to achieve full calibration. However it is conceivable to relax the input signal swing requirement by rotating the sampling capacitors in a known sequence not correlated with the analog input, so that (e.g.) different sampling capacitors can be used to sample a DC input, hence exercise the range of different MSB transitions.

4.4.5: BACKEND PIPELINED ADC

The backend ADC was implemented using a cascade of nine standard 1.5b/stage pipeline stages followed by a 2-bit flash ADC as shown in Fig. 4-12. Gain-boosted [18] folded cascode opamps with a DC gain of 95dB were used in stages 2 and 3 to minimize missing

codes in the backend. The remaining pipeline stages used folded cascode based opamps with DC gains of 50dB.

4.4.6: DIGITAL CALIBRATION

To enhance flexibility in the test setup, the digital calibration engine was implemented off-chip, where the digital outputs of each split-ADC stage were taken off-chip and imported into Matlab via a logic analyzer. The digital outputs of the fabricated chip were input to a Simulink model of the digital calibration scheme outlined in this work. Since there is no feedback or direct interaction with the analog portions of the ADCs in the calibration scheme of this work, operating on the off-chip digital outputs verifies the calibration scheme without any loss of generality.

Since the missing codes at each MSB transition for both split-ADCs are required to be measured, 16+15=31 unique missing codes must be measured and estimated. Thus in a potential on-chip digital implementation each Δ_i is required to be passed through an IIR filter, and the digital output offset by each missing code as shown in Fig. 4-17 and Fig. 4-18.

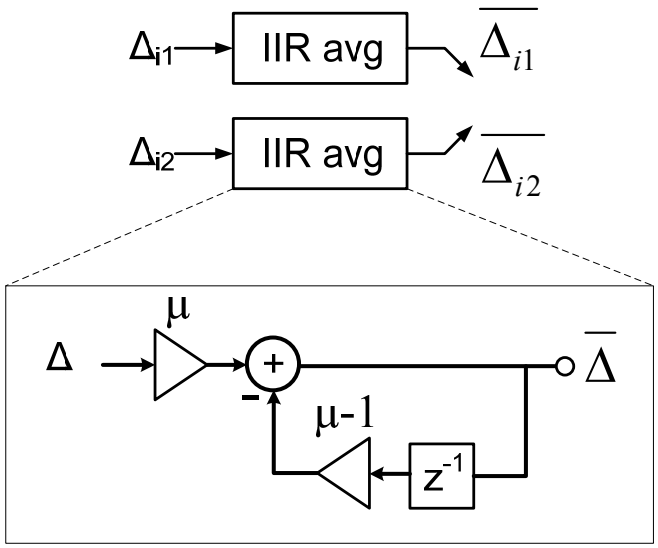

Fig. 4-17: Detail of IIR filter blocks used in 'estimate error' block of Fig. 4-18

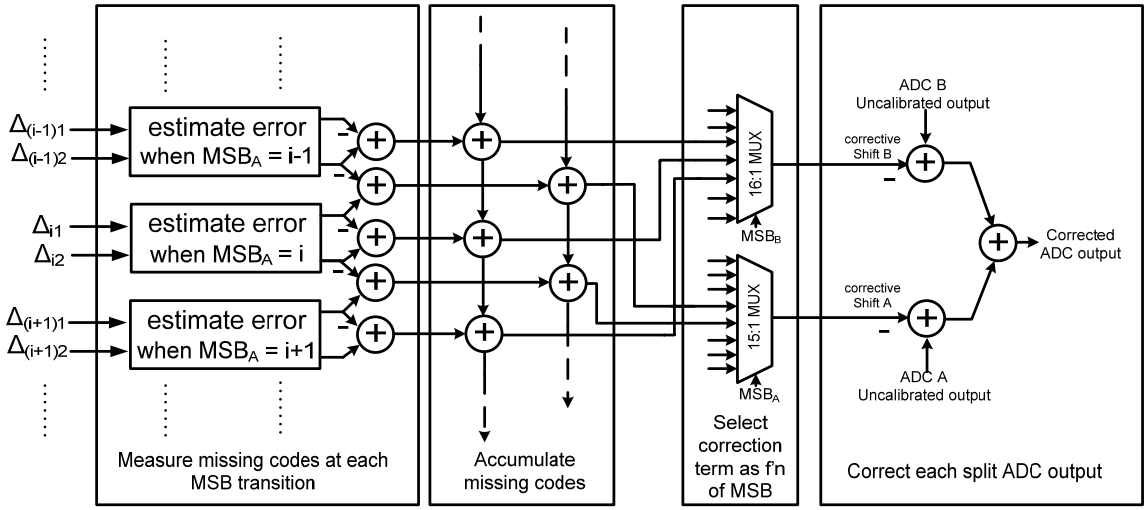

Fig. 4-18: Digital implementation of calibration

It is noted that each averaging block in Fig. 4-18 is only updated once every time its particular MSB bit is excited, i.e. has low activity. Thus the power of each path in Fig. 4-18 is fairly low since with a reasonably random input each MSB bit is excited at a small fraction of the overall sampling rate.

4.4.7: SPICE SIMULATIONS

The proposed ADC architecture was verified and found to be functional in Spice over all process corners and at a nominal temperature of 50°C. Fig. 4-19 illustrates an FFT of the ADC output after calibration at f_s=100MS/s. Calibration was performed by saving the output of the Spice simulation to a text file and parsing the Spice output in Simulink with a model of the calibration scheme.

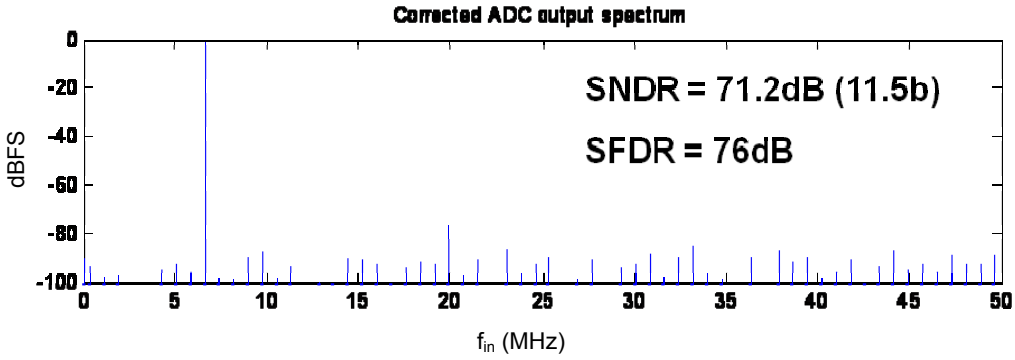

Fig. 4-19: Spice simulation of ADC output after calibration when f_s=100MS/s

It is noted that the Spice simulation did not include thermal noise sources. As the thermal noise floor was designed to be at ~11-bit level, the SNDR in the fabricated chip was expected to be lower. As ADC power is dominated by thermal noise considerations, in an attempt to minimize power consumption the ADC's resolution was designed to be limited by thermal noise, as opposed to quantization noise.

4.5: TESTING

Sections 4.5.2 and 4.5.3 discuss the test setup and measured results of the prototype fabricated in a 1.8V 0.18µm CMOS process.

4.5.2: PCB

A 4 layer FR4 dielectric PCB board with a minimum 6mil trace was designed and constructed for the device under test as shown in Fig. 4-20. Separate Power planes were used to isolate the analog, digital, I/O, and board power supplies. A differential input was generated using a 1:1 turns ratio Minicircuits transformer matched to 50Ω. Reference voltages were generated by passing the output of a resistive voltage divider through an opamp (LM7301) in a unity gain buffer configuration. To maintain constant supply voltages, all voltage supplies for each power plane were generated through regulators (LM337, LM1117), and heavily decoupled with capacitors. As the ADC utilized a constant current biasing scheme, an off-chip adjustable resistor was used as the master current source.

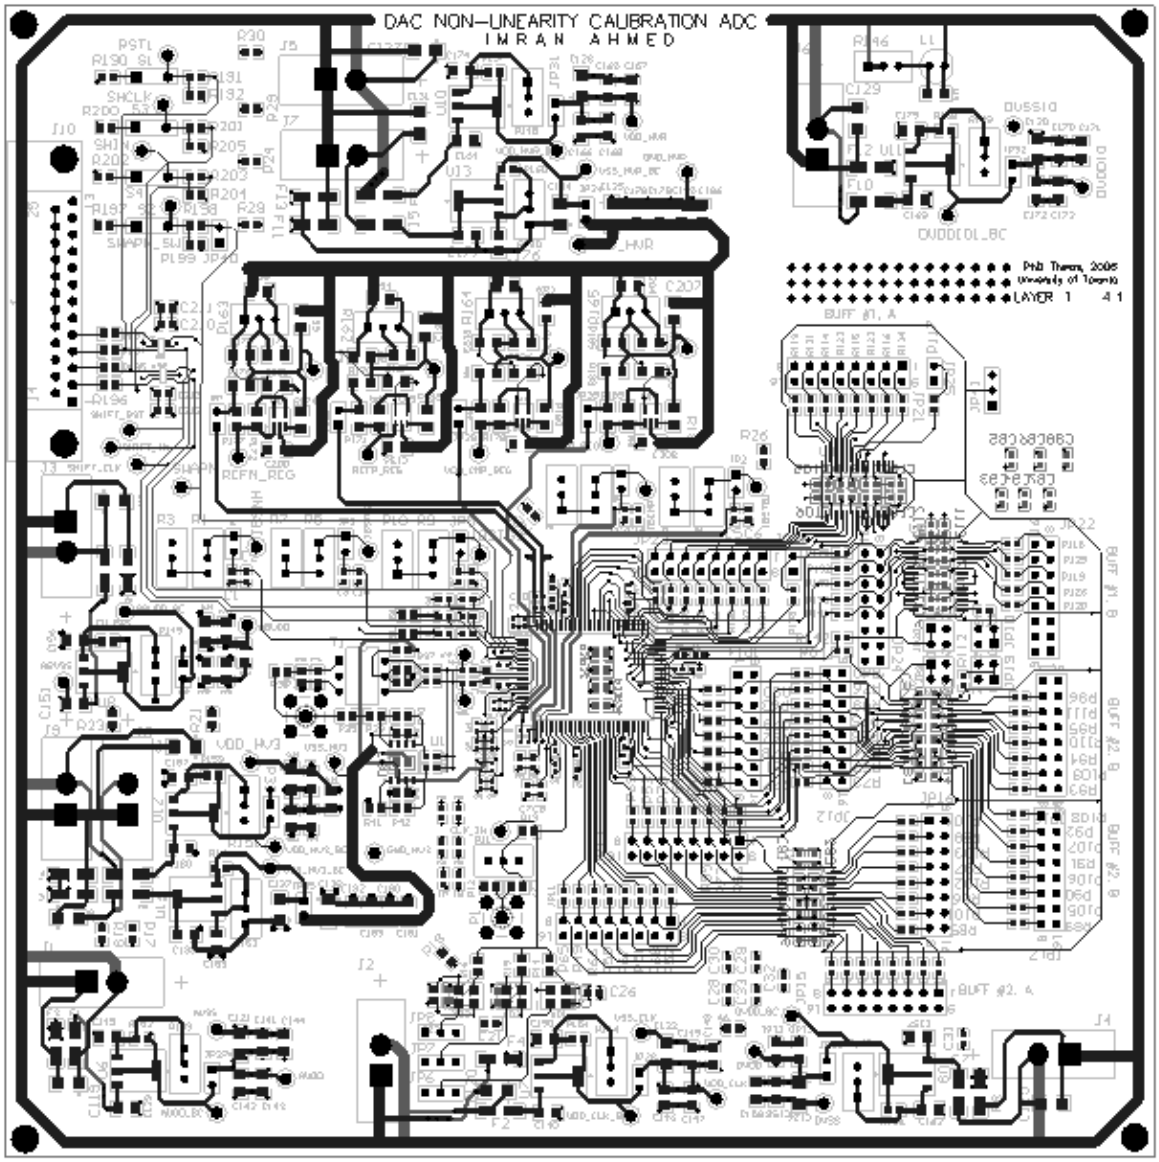

Fig. 4-20: Custom PCB layout

4.5.3: TEST SETUP

A test setup as shown in Fig. 4-21 was used. Sinusoidal inputs were generated using an HP 8664A function generator. A Minicircuits low pass filter was used to minimize harmonic distortion from the function generator such that the sinusoidal input to the ADC had an SNDR of well over 11-bits. A 45MHz clock was derived from a HP8130A pulse/pattern generator. The serial shift register was loaded via a parallel port connection to a PC, where a

Matlab script was executed to load the appropriate bits. The output bits of each pipeline stage were captured using a Tektronix TLA714 logic analyzer, capable of capturing 65,536 points at a time. An Agilent E3620A Dual output DC power supply was used to provide positive and negative voltages to the voltage regulators on the PCB. As discussed in section 4.4.6 the digital calibration was implemented off-chip in a behavioral Matlab model.

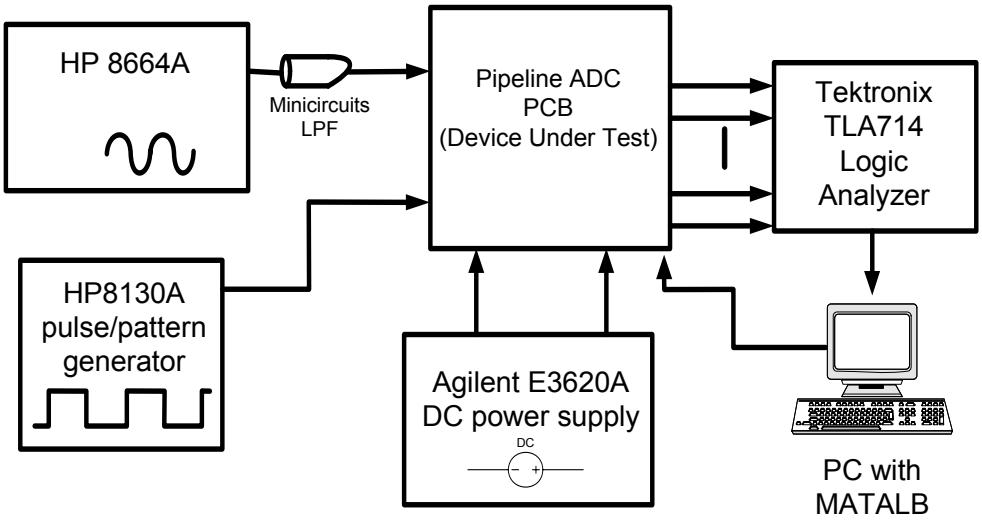

Fig. 4-21: Test setup for rapid DAC calibration ADC

4.6: MEASURED RESULTS

A prototype of the ADC architecture as shown in Fig. 4-22 was fabricated in a 1.8V 0.18µm CMOS process. The area of the analog core was 2.1mm x 1.7mm = 3.57mm^2. It is noted that the area of the first pipeline stage in each ADC had a large overhead due to the programmability of each sampling capacitor. In a practical design which does not have the additional test circuitry, the area of the first pipeline stage could be significantly reduced.

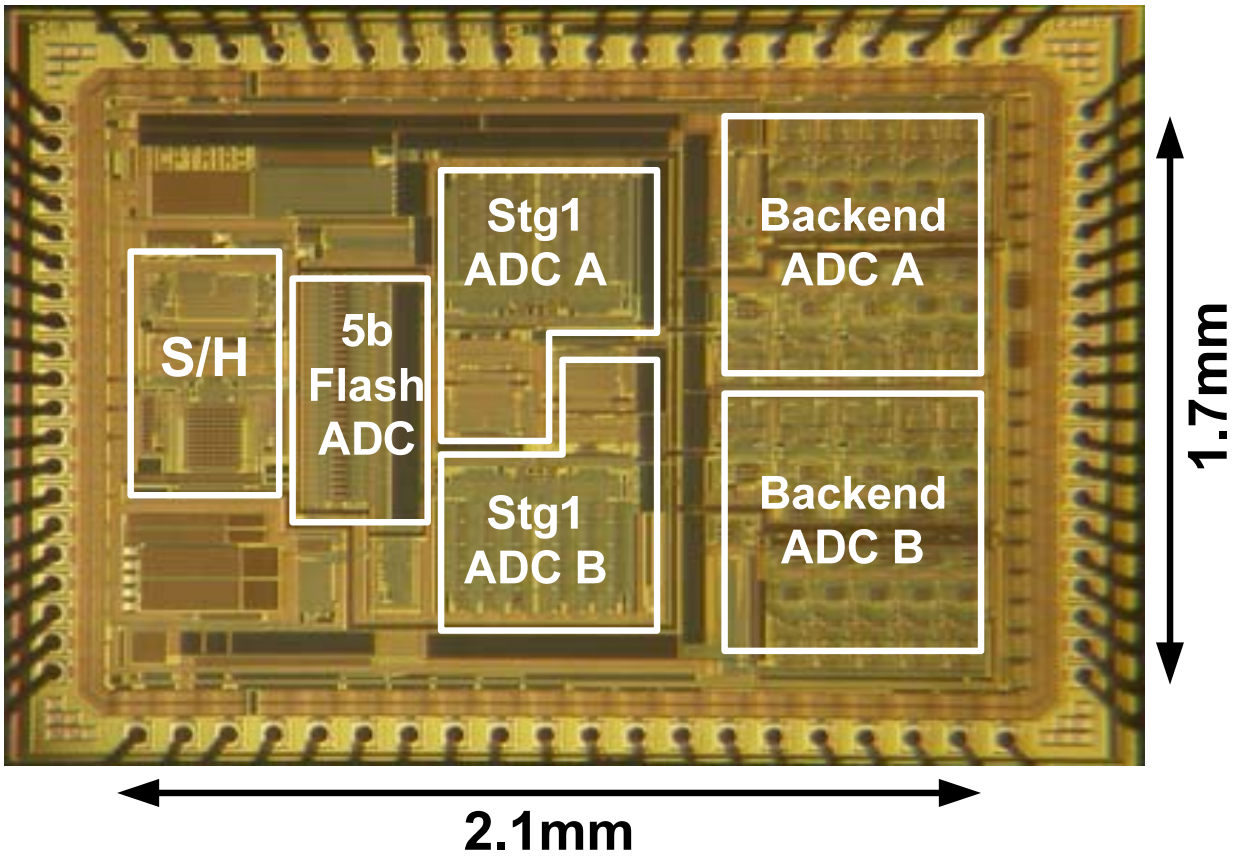

Fig. 4-22: Micrograph of fabricated IC in 1.8V 0.18µm CMOS

4.6.2: SAMPLING RATE REDUCTION

The ADC was operated at f_s=45MS/s, whereas the design was optimized for 100MS/s. Measured results showed the SNDR of ADC decreased for sampling rates larger than 45MS/s as shown in Fig. 4-23.

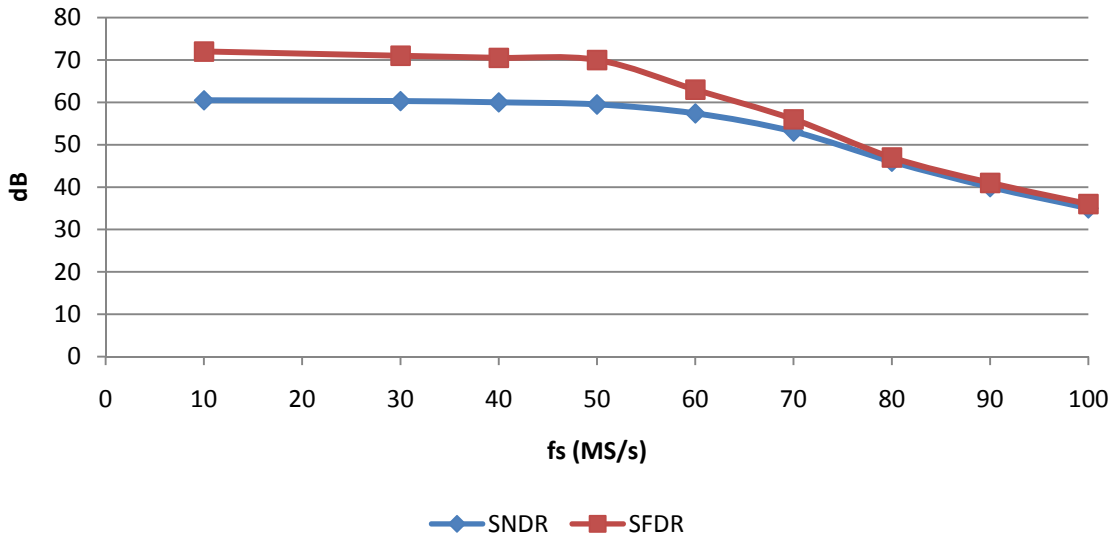

Fig. 4-23: SNDR/SFDR variation (after calibration) with sampling rate, with input tone at 2.39MHz

As the circuits in the ADC were designed to operate as high as 100MHz, it was inferred that the degradation of performance was due to noise sources not modeled in SPICE simulations of the ADC. Due to a lack of package models, the design was simulated without package parasitics. In [8] and [62] it is shown that parasitic inductance and capacitance on bond wires can cause oscillations on on-chip voltage nodes (typically digital nodes) due to large changes in current over a short time interval. As a result noise from digital signals can couple into the sensitive analog portion of the IC through coupling off-chip via ground loops and/or through coupling between bond wires via mutual inductance and parasitic capacitance, or on-chip via the substrate.

Upon investigating the source of performance degradation for high sampling rates in this work, it was found that several tones appeared in the ADC output spectrum which changed in magnitude whenever the supply voltage of the digital I/O drivers was varied, while keeping the rest of the setup unaltered. In Fig. 4-24, the I/O supply voltage is varied from 1.25V to 1.8V where it is seen that some tones increase in magnitude approximately by the increase in supply voltage.

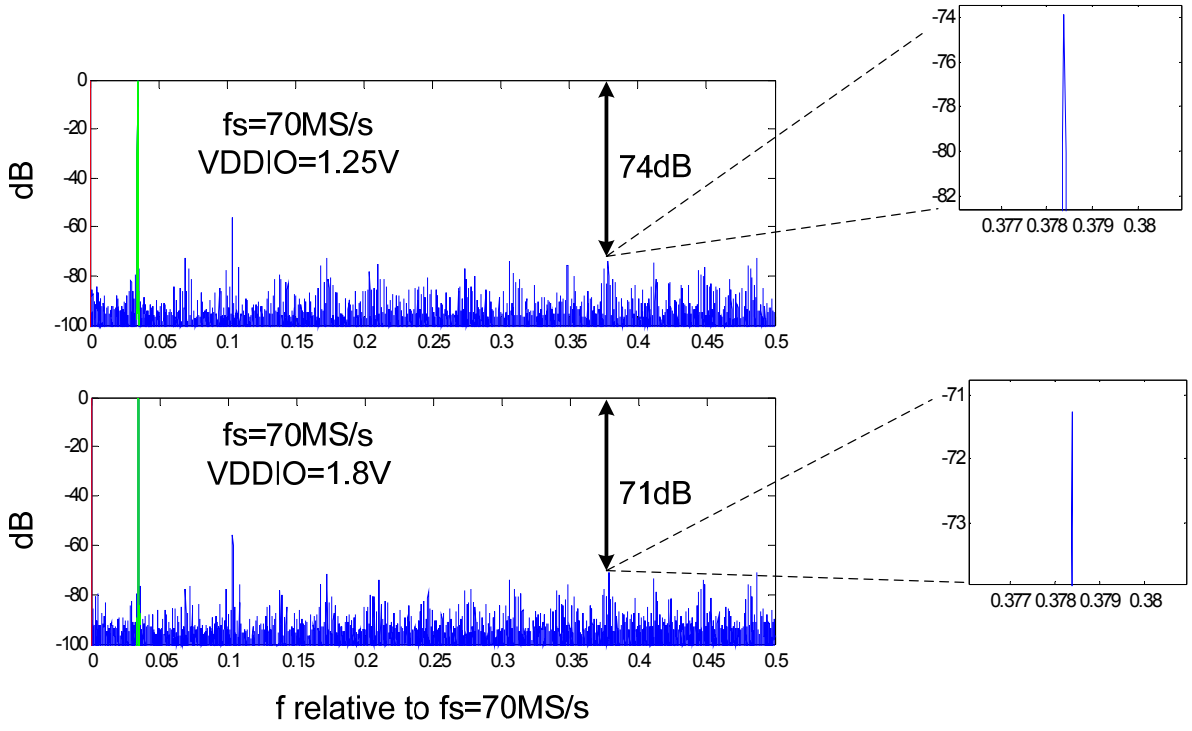

Fig. 4-24: Comparison of ADC FFTs with different digital I/O supply voltages

As changing the supply voltage to the digital I/O should have no impact on the ADC output spectrum, it was determined that unwanted digital noise was coupling into the system and was thus a major reason why performance degraded for higher sampling rates (since the effect of digital noise becomes more pronounced as the sampling rate increases since there is less time for the effects of the noise to settle.). As will be seen in sections 4.6.3 to 4.6.5 however, the performance of the ADC at half the designed sampling rate (45MS/s) was more than sufficient to successfully demonstrate the proposed calibration technique. Given the limited observability of internal nodes in the chip, and the promising results at half the sampling rate, a further analysis of the noise sources limiting high speed performance was halted, and the ADC was characterized at 45MS/s.

At 45MS/s the power of the fabricated ADC was 81mW plus 9.5mW for the reference voltage resistor string used in the 5-bit flash ADC. It is noted that since the design was optimized for a ~2x higher sampling rate the power could easily be reduced if optimized for 45MS/s.

4.6.3: INL/DNL PLOTS

Fig. 4-25 and Fig. 4-26 show the INL and DNL of the ADC before and after calibration, where each DAC element in each split-ADC was programmed with a mismatch of either: 0%, 0.5%, 1% or 5%.

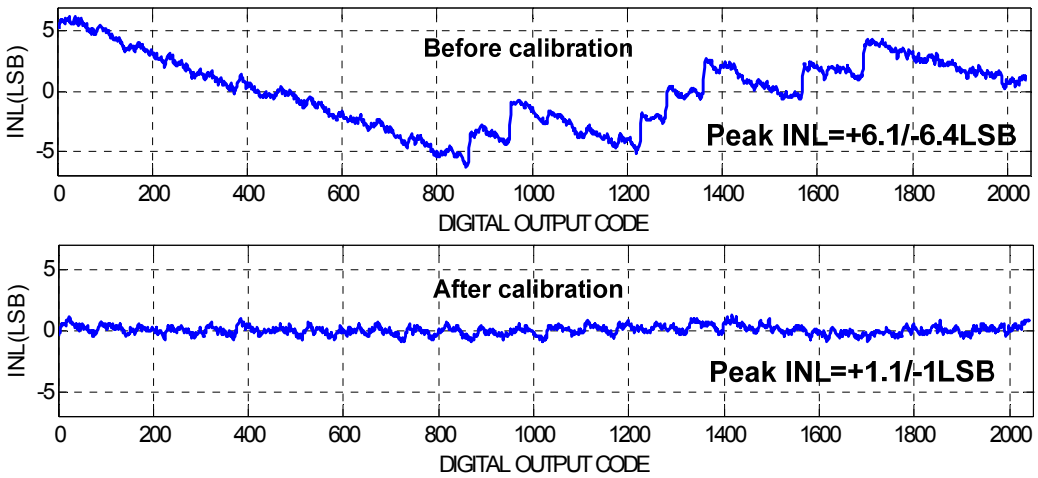

Fig. 4-25: INL before and after calibration, f_s=45MS/s (LSB @ 11-bit level)

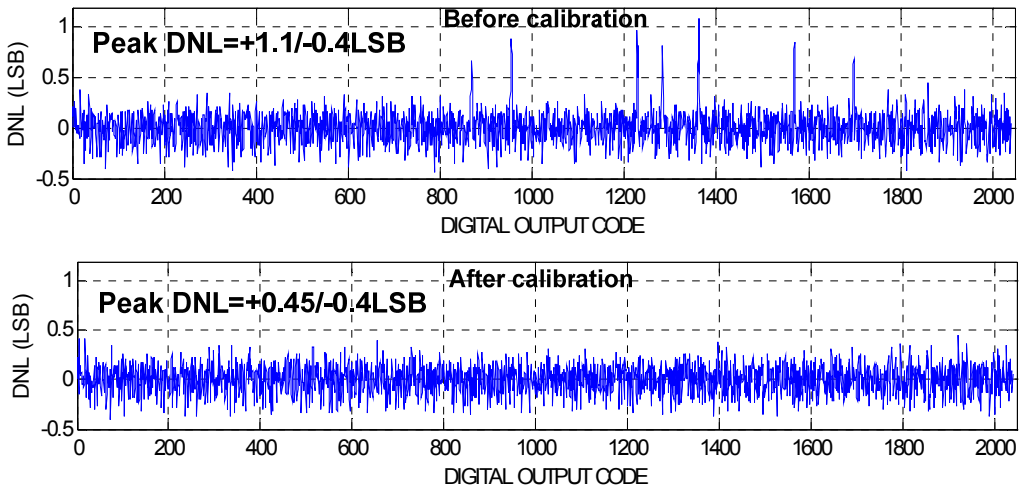

Fig. 4-26: DNL before and after calibration (LSB @ 11-bit level)

Using the calibration scheme of this work, the INL was improved from +6.1/-6.4LSB to +1.1/-1LSB after calibration, and DNL from +1.1/-0.4LSB to +0.45/-0.4LSB. It is noted that the residual INL errors after calibration are due primarily to distortion from the backend ADCs, as well as distortion from the front-end sample-and-hold which sets the best achievable linearity for both ADCs.

4.6.4: SNDR/SFDR PLOTS

Fig. 4-27 shows an FFT of the ADC output before and after calibration with a 1.3V p-p input sinusoid at 2.39MHz. The SNDR/SFDR of the ADC is improved from 46.9dB/48.9dB to 60.1dB/70dB.

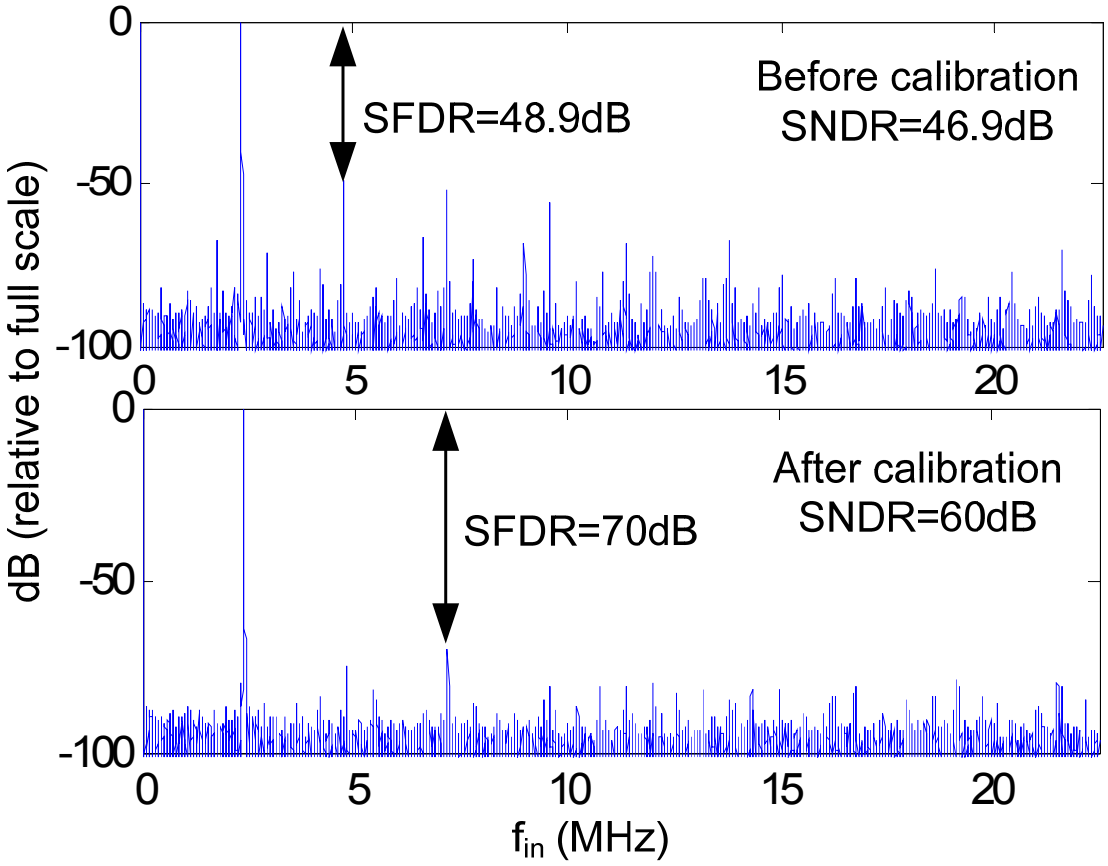

Fig. 4-27: FFT of ADC output before and after calibration

The calibration technique was verified with different programmed mismatches in the first stage as well as different full scale inputs (e.g. sinusoidal, random). In all cases successful background calibration was attained in a short time interval. Fig. 4-28 illustrates the variation of SNDR and SFDR with input frequency before and after calibration, where the fall-off of accuracy for higher input frequencies was attributed to the input switch in the front-end sample-and-hold.

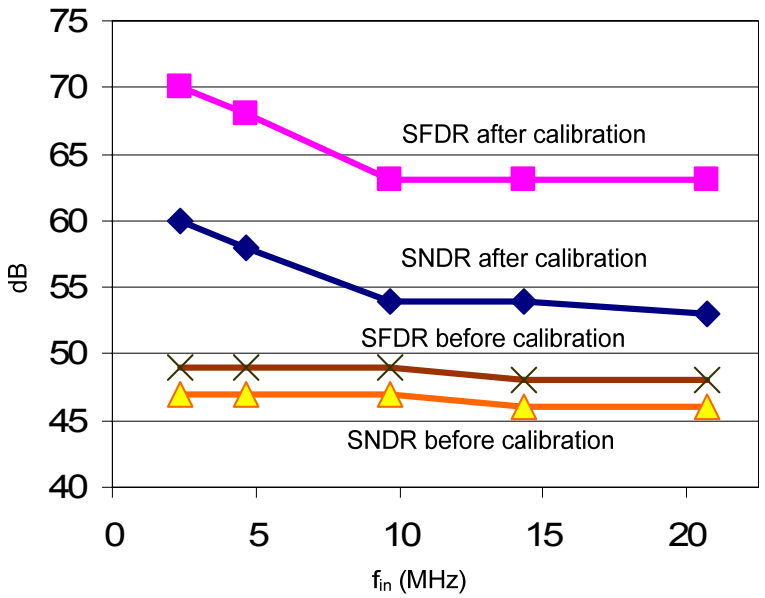

Fig. 4-28: Variation of ADC SNDR, SFDR with input frequency, before and after calibration

4.6.5: CALIBRATION TIME

The improvement of ADC SNDR and SFDR is plotted versus total number of calibration cycles in Fig. 4-29, where it is shown that steady state is attained within ~$1x10^4$ clock cycles or effectively 0.22ms (with μ=1/64). From [28] it is noted that if a statistics-based background calibration technique were used ~$4x10^6$ clock cycles would be required to correct only the gain error. In this work both gain *and* DAC errors are corrected in less than $1/100^{th}$ the total number calibration cycles of prior statistics based approaches. The plot of Fig. 4-29 was obtained by running the calibration for a given time and then freezing the calibration coefficients. Once the coefficients were frozen an FFT of the ADC output was taken with a sinusoid applied to the ADC input, and the SNDR/SFDR measured. This process was repeated for a range of total calibration cycles, yielding the plot shown in Fig. 4-29.

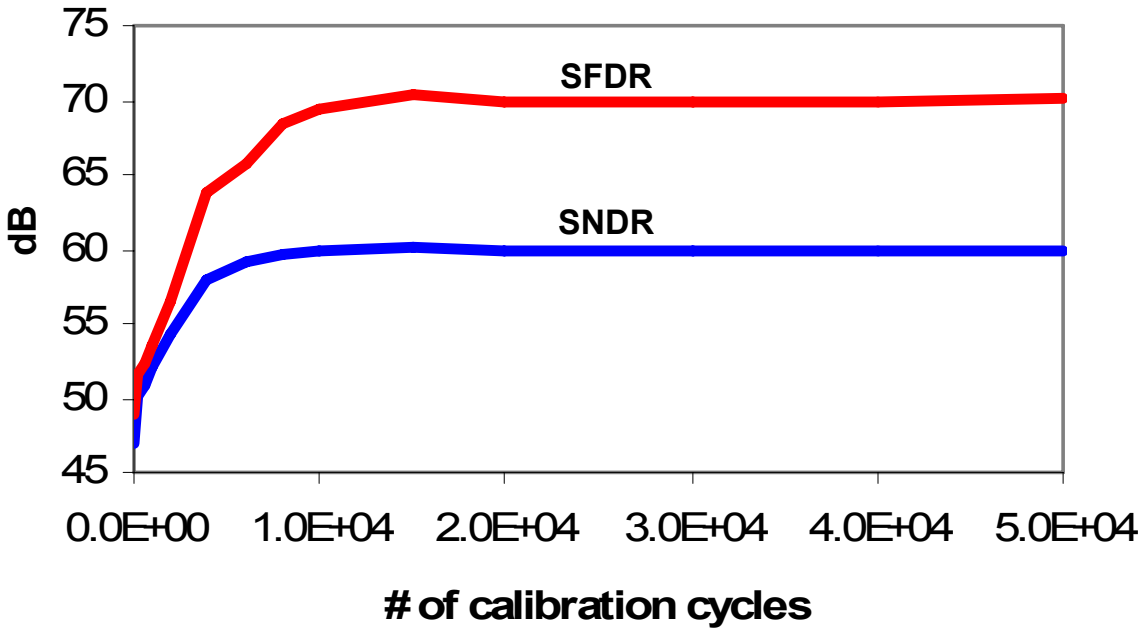

Fig. 4-29: ADC SNDR, SFDR improvement with # of calibration cycles

4.7: SUMMARY

A technique to rapidly measure and correct for missing codes introduced by a multi-bit first stage in a pipelined ADC was presented. Measured results from a prototype in 1.8V 0.18µm CMOS show an improvement in INL from +6.1/-6.4LSB to +1.1/-1LSB, and SNDR/SFDR from 46.9dB/48.9dB to 60.1dB/70dB using the calibration approach of this work. Calibration was achieved in ~10^4 clock cycles. Table 4-1 summarizes key measurements of the work.

Table 4-1: Summary of ADC performance

Technology		1.8V, 0.18µm CMOS
Sampling rate (f_s)		45MS/s
Input signal swing		1.3V p-p
Area		3.57mm^2
Power		81mW
# of calibration cycles		10^4 cycles (0.22ms)
	BEFORE CALIBRATION	**AFTER CALIBRATION**
SNDR	46.9dB	60dB
SFDR	48.9dB	70dB
INL	+6.4/-6.1 LSB	+1.1/-1 LSB
DNL	+1.1/-0.4 LSB	+0.45/-0.4 LSB

CHAPTER FIVE: A HIGH BANDWIDTH POWER SCALABLE SUB-SAMPLING ADC WITH EMBEDDED SAMPLE-AND-HOLD

5.1: OVERVIEW

In this chapter a pipelined ADC architecture for use in sub-sampled systems which is power scalable in relation to its down sampled bandwidth is presented. The ADC uses a novel technique which does not rely on carefully matched input paths, to eliminate the front-end sample hold, and thereby reduce power consumption. The technique allows for a power savings of >20% compared to a previous design. A method to improve the settling behavior of Rapid Power-on Opamps is also presented. Measured results in a 1.8V 0.18µm CMOS process verify the embedded sample-and-hold technique does not cause gross MSB errors for input frequencies higher than 267MHz. With f_s=50MS/s, for f_{in}=79MHz the SNDR is 51.5dB, and with f_s=4.55MS/s for f_{in}=267MHz the SNDR is 52.2dB.

The organization of the discussion in this chapter is as follows: section 5.3 describes the technique used to eliminate the front-end sample-and-hold, section 5.4 describes how power scalability is achieved in this work, section 5.5 describes the circuit implementation of the design, section 5.7 presents measurement results, and section 5.8 concludes the chapter.

5.2: MOTIVATION

All mobile communication systems consist of a receive path, in which it is typically required to down-convert a high frequency input down to an IF or baseband frequency. From section 2.10, it was shown that by using a sub-sampled ADC a front-end mixer could be eliminated from a receive system, and thereby potentially save power. Using a sub-sampled ADC however was shown to require an ADC which had a large input bandwidth.

As discussed in section 2.9 large input frequencies for a pipelined ADC require a power hungry front end S/H to ensure functionality. In the interest of reducing power consumption, techniques to eliminate the front end S/H have been proposed. From section 3.3.2 it is noted however that the present state of the art techniques to eliminate the front-end S/H rely on carefully matched signal paths and thus a lengthy and costly design procedure. Thus the goal of this work is to develop a sub-sampled ADC topology which enables the elimination of the front-end S/H to save power, yet does not require a carefully matched (thus costly) layout to ensure functionality.

5.3: PROPOSED EMBEDDED S/H TECHNIQUE

Fig. 5-1 illustrates a conventional 1.5b stage in a pipelined ADC. During Φ_1 the input is sampled on capacitors C_1 and C_2. During Φ_2 a gain of two is implemented by discharging the charge stored in C_1 to C_2, and DAC operation by connecting V_{DAC} to a voltage set by the sub-ADC.

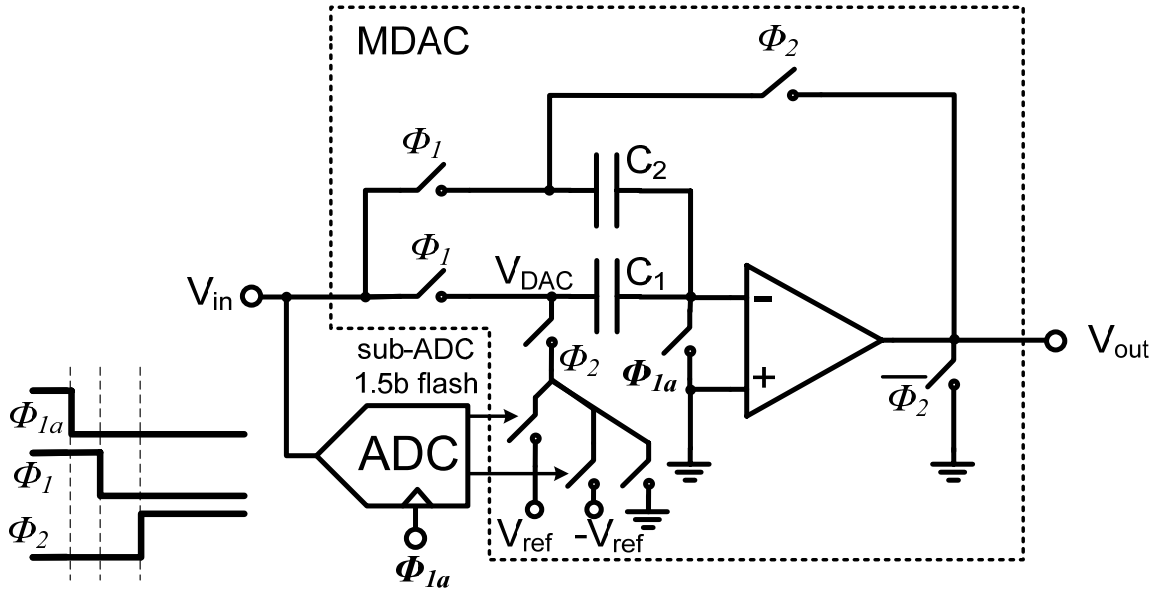

Fig. 5-1: Conventional 1.5 bit MDAC

In this work the first 1.5b pipeline stage is modified and an additional clock phase introduced as shown in Fig. 5-2. Like the conventional pipeline stage of Fig. 5-1, during Φ_1 the input is sampled on C_1 and C_2. However in this work, when Φ_2 switches from low to high and Φ_{2D} is low, V_{DAC} is set to a high impedance and C_2 is connected between the output and input of the opamp as shown in Fig. 5-3.

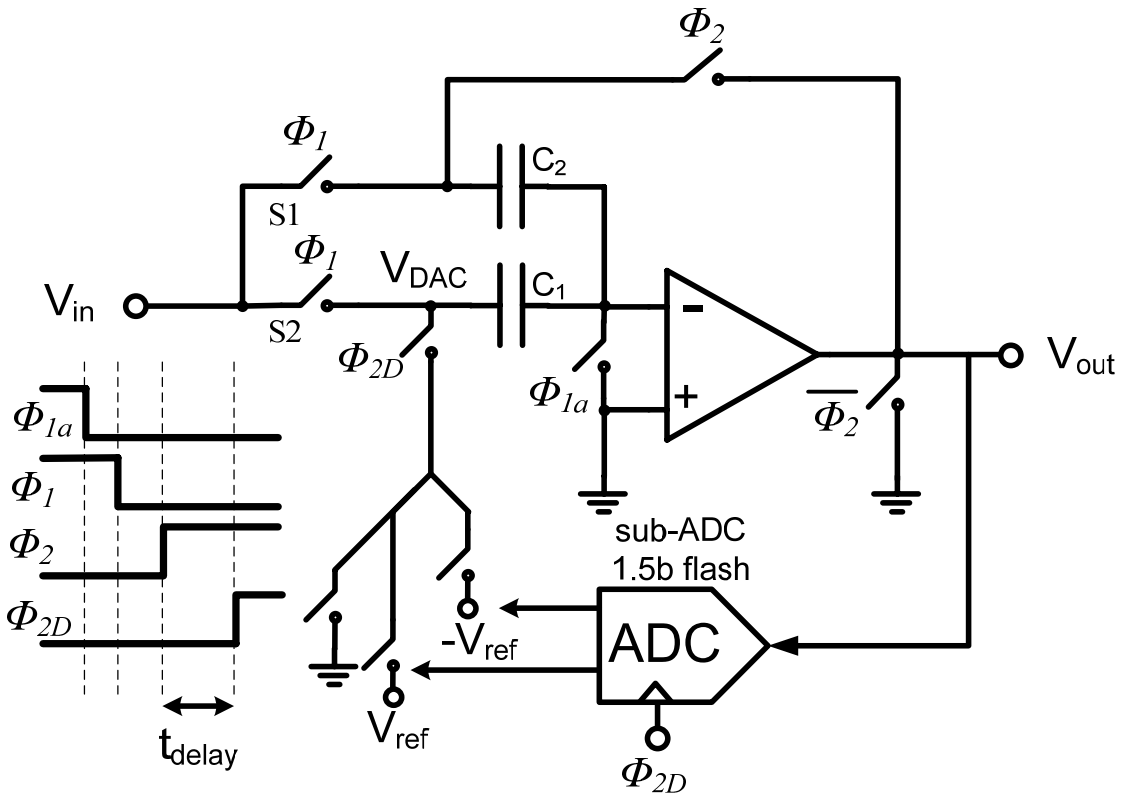

Fig. 5-2: MDAC of this work which enables elimination of front-end S/H (shown single-ended, but implemented fully-differentially)

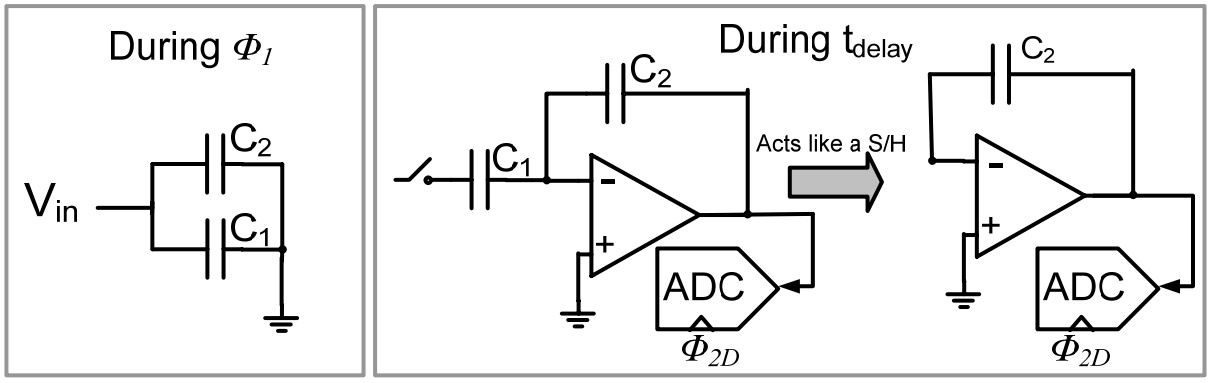

Fig. 5-3: Detailed illustration of MDAC functionality during t_{delay}

Connecting C_2 around the opamp produces a held value of the sampled analog input, and due to charge conservation at the negative input of the opamp the voltage across C_1 is preserved as shown in Fig. 5-3. Thus with the approach of this work, during the time labeled t_{delay} in Fig. 5-2, the output of the first stage can be connected to a 1.5b flash ADC. When Φ_{2D} subsequently goes high the 1.5b flash ADC resolves its input, sets V_{DAC} to the appropriate DAC voltage, and implements a gain of 2x by discharging the charge from C_1 into C_2. Therefore the first pipeline stage of this work implements the same functionality as a conventional pipeline stage, however is not sensitive to skew at the input as by using the embedded sample-and-hold of the MDAC, the sub-ADC operates on the same input that is sampled by the MDAC regardless of input frequency. Thus by using the proposed modified first pipeline stage a front-end sample-and-hold is not required to ensure functionality for high input frequencies, hence allowing for substantial power savings.

Since the flash ADC only requires an input that is >1.5-bit accurate, the output of the opamp during t_{delay} is only required to settle to >1.5-bit accuracy to generate the correct outputs from the flash ADC, hence t_{delay} is only a small fraction of the total available settling time as shown in Fig. 5-4.

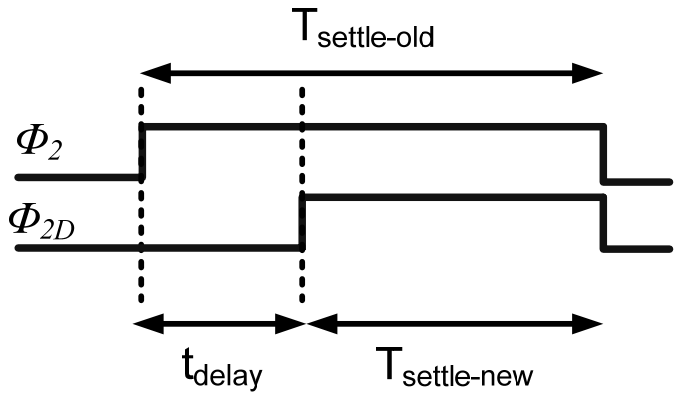

Fig. 5-4: Comparison of MDAC settling time of this work versus conventional MDAC

Although the power of the opamp in the first pipeline stage is slightly increased by the fraction of settling time taken by t_{delay}, the overall power of the ADC is significantly reduced as the power hungry front-end S/H is eliminated. The technique to eliminate the front-end S/H could also be applied to multi-bit pipeline stages, noting that prior approaches which relied on redundancy to eliminate the front-end S/H have an even smaller allowable skew thus demand an even more meticulously matched layout than discussed in section 3.3.2. It should be noted that the technique to use the embedded S/H, although independently derived for this work, is similar to that recently published in [63]. However the results of this work show: how the embedded S/H technique can be applied to a power scalable architecture, a much higher input bandwidth, and a 66% increase in sampling rate.

From Fig. 5-3 it is noted that during t_{delay} one end of C_1 is floating. Alternatively it is also possible to configure the MDAC such that during t_{delay} node V_{DAC} is also connected to the output of the opamp as shown in Fig. 5-5.

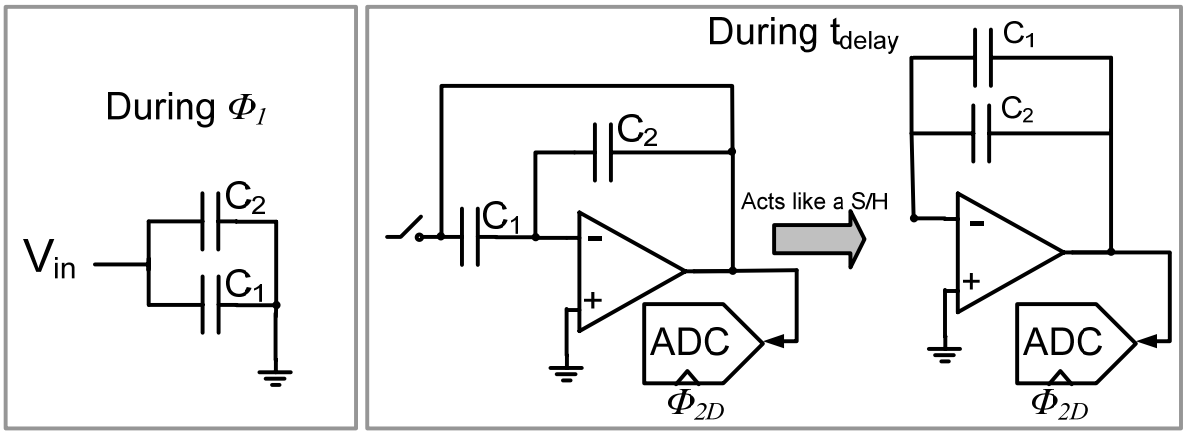

Fig. 5-5: Alternative configuration of MDAC during t_{delay} without floating capacitors

Using the alternative approach the functionality of the MDAC is the same as Fig. 5-2 however a floating capacitor is avoided. By using the approach of Fig. 5-5, the feedback factor of the MDAC during t_{delay} is increased; however the load capacitance seen by the

opamp is also increased. Depending on the relative size of C_1 compared to the opamp input capacitance, the approach of Fig. 5-5 could yield a faster or slower solution compared to that in Fig. 5-3. To evaluate both approaches, a programmable switch was included in this work to allow the ADC to operate as shown in Fig. 5-3 or as shown in Fig. 5-5, so that the viability of each approach could be validated.

5.4: POWER SCALABILITY WITH SAMPLING RATE

The ADC of this work achieves power scaling using the techniques of [45], i.e. CMPS, which was described in section 3.3.3.3.

5.5: CIRCUIT IMPLEMENTATION

The ADC topology of this work is based on a previous chip designed by the author [45]. All the circuits from [45] are reused in this work with the exception that the first stage in this work is implemented using the techniques described in section 5.3, and no front-end S/H is used in this work. Furthermore in this work a clock delay generator circuit is added, as well as an additional (but not strictly necessary) bias circuit for the first stage to improve noise isolation between the first stage and subsequent pipeline stages. An advantage of basing this work on a previous circuit design is the improvement in power reduction can be made by directly comparing this work against the power consumption of the previous work [45]. As the goal of the prototype was proof-of-concept, a specific application was not targeted, however as the ADC is power scalable a variety of applications from low power sensors to wireless LAN could take advantage of the proposed 10-bit architecture.

5.5.2: ADC ARCHITECTURE

The architecture of the pipelined ADC of this work is shown in Fig. 5-6. The first stage is as shown in Fig. 5-2, and all remaining stages are standard 1.5b stages. Stages 3-9 are identical

to those used in [45]. Stage 2 is a standard 1.5b stage but has the same sized sampling capacitors and opamp as stage 1.

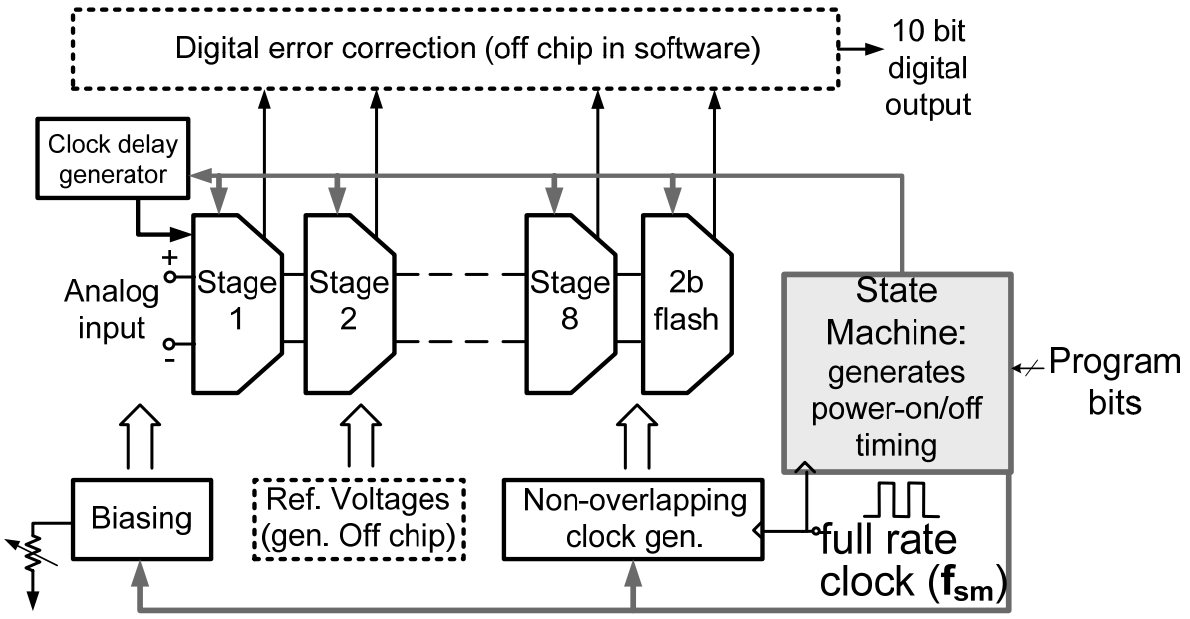

Fig. 5-6: Architecture of pipelined ADC

As there was no explicit front-end S/H block in this work, the total sampling capacitances of the first two pipelined stages in this work were reduced by 40% compared to [45] (from 940fF to 580fF in stages 1, and 2) while maintaining the same input referred noise as [45]. In this work the opamp of stage 1 was conservatively designed such that t_{delay} was 20% of the total settling time. However t_{delay} was made tunable so that the lowest t_{delay} possible without gross MSB errors could be measured. To meet the same settling accuracy as [45], the opamps of stage 1 and 2 were made 33% smaller than those of [45]. Thus although less settling time is available in the first pipeline stage, overall power is still reduced without a front-end S/H since the thermal noise floor of the pipeline stages can be made higher (with smaller sampling capacitors) to maintain a fixed input referred noise floor. To enable a large input bandwidth, bootstrapped switches [64] were used as the input switches S1, S2 in Fig. 5-2.

5.5.3: RAPID POWER-ON OPAMP

A Rapid Power-on Opamp [45] was used in each pipeline stage to realize a power scalable architecture and is shown in Fig. 5-7.

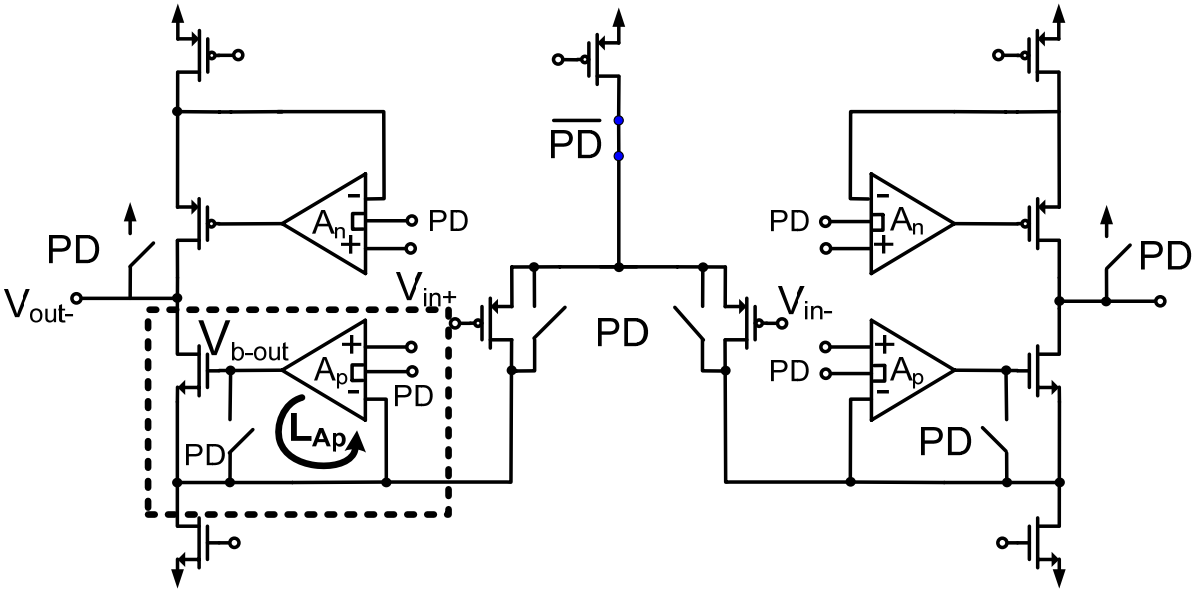

Fig. 5-7: Rapid Power-on Opamp

Like conventional gain-boosted opamps, the Rapid Power-on Opamp requires the unity gain frequency of the gain booster opamps (opamps A_p, A_n in Fig. 5-7) to be higher than the 3dB frequency of the closed-loop but lower than the second pole of the main opamp [18]. Furthermore the loop formed by the gain boosters (labeled L_{Ap} in Fig. 5-7) is also required to be stable. Rapid Power-on Opamps power on and completely off each clock cycle, hence the inputs to the gain booster opamps A_p, A_n effectively see a step function every clock cycle. Thus to ensure good settling, loop L_{Ap} requires a high phase margin. In [45] standard folded cascode PMOS-input opamps as shown in Fig. 5-8 were used for the gain booster opamps A_p.

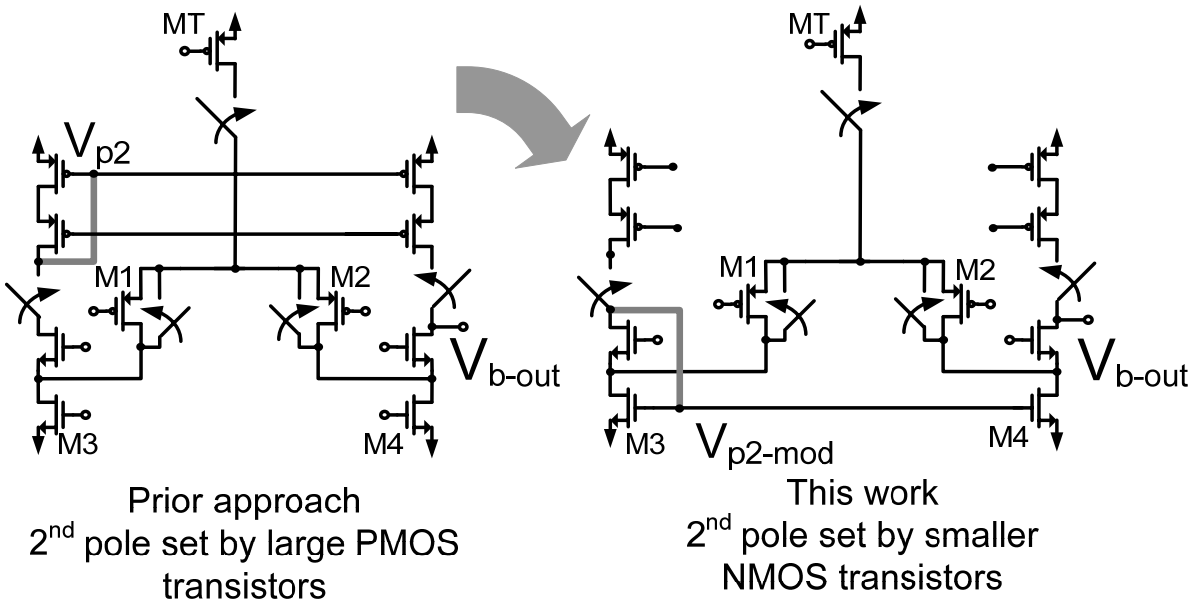

Fig. 5-8: Gain booster opamp used previously (left), and in this work (right)

The phase margin of the loop L_{Ap} is limited by the second pole of the loop which occurs at the PMOS current mirror node labeled V_{p2} in Fig. 5-8. The same bias voltages used for the main opamp were also used for the gain booster opamps, thus PMOS transistors were sized ~5x larger than the NMOS transistors to maintain a similar overdrive voltage as the NMOS transistors for a fixed current density. As such a large parasitic gate-source and gate-drain capacitance exists on node V_{p2} and sets the second pole of the loop to a low frequency which degrades the phase margin of the loop L_{Ap}. To improve the phase margin of loop L_{Ap} the p-input gain booster opamp was modified by implementing the current mirror in the opamp with NMOS transistors instead of PMOS transistors [65] as illustrated in Fig. 5-8. By performing the mirror operation with NMOS transistors, the capacitance on the mirror node is reduced by ~2.5x and the second pole pushed to a much higher frequency - significantly improving the phase margin. Simulation results show that the phase margin of loop L_{Ap} is improved from 56º to 72º by only changing the location of the mirror node. By switching the location of the current mirror node from PMOS to NMOS transistors the slew rate of the gain

boosting opamp is reduced by 2x (i.e. when M2 is cutoff, in the prior approach the current through MT is mirrored by the PMOS current mirror to the output node, however in the approach of this work when M2 is cutoff the output current is set by M4 which is biased with half the current of MT). Although a lower slew rate increases the opamp's power-on time, the significant benefit of a more stable power-on transient makes it a favorable trade-off. Also it is noted that since the power-on time of the opamps of this work only require a small percentage of the settling time to power-on, the total available settling time of the opamp is not significantly reduced with a lower slew rate in the gain booster opamps.

5.5.4: GENERATION OF DELAYED CLOCK Φ_{2D}

The delayed clock edge Φ_{2D} of Fig. 5-4 was generated using a chain of four current starved inverters as shown in Fig. 5-9. Current starved inverters were chosen to allow t_{delay} to be widely tuned for different sampling rates by changing the bias currents. Also by varying the off-chip reference current I_{ref}, the dependence of ADC SNDR vs. t_{delay}, could be measured.

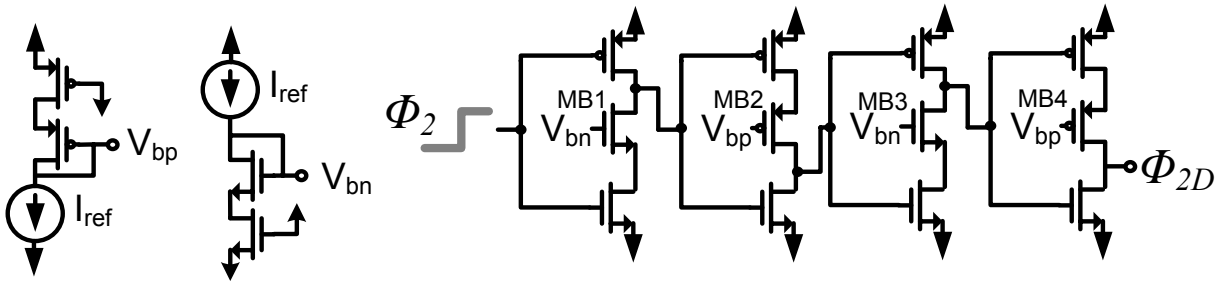

Fig. 5-9: Clock delay block to generate Φ_{2D} from Φ_2

The current starved inverters were designed such that only one clock edge was delayed. Commonly in a current starved inverter, the current sources are connected to the source nodes of the inverters (e.g.: [62], [66], [67]) as shown in Fig. 5-10.

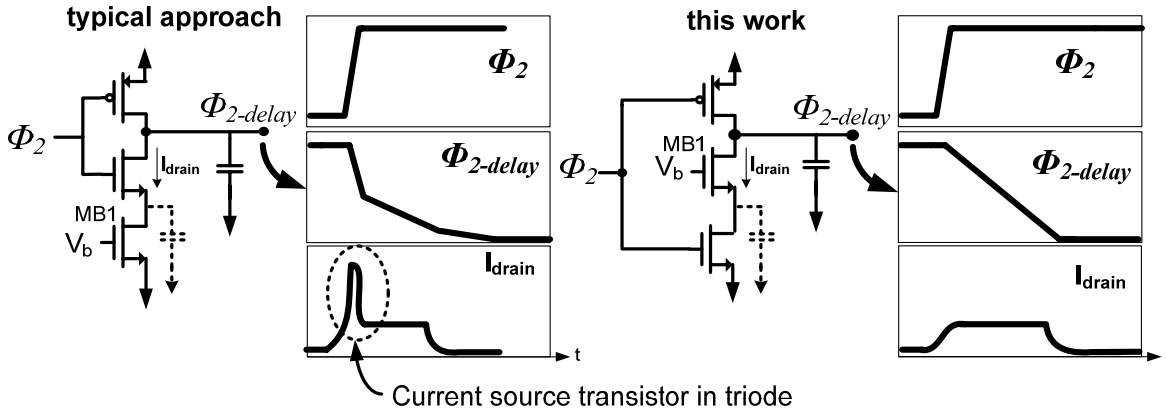

Fig. 5-10: Comparison of different current starved inverter topologies

When the input clock switches from low to high in e.g. the first inverter, the current source initially starts off in triode causing a large drain-source current, which is different than the desired bias current, to flow until the pre-charged output discharges to the point where the current source is biased in the active region as shown in Fig. 5-10. As a result the discharge rate of the output also becomes a function of the rise time of the input, reducing the control I_{ref} has on the delay of the current starved inverter, hence reducing range of delay values possible by varying I_{ref}.

In this work, the current source transistors were connected between the drains of the PMOS and NMOS inverter transistors as shown in Fig. 5-10. This was done so that when the input to e.g. the first inverter switches from low to high, the current source transistor MB1 switches from operating in cut-off to the active region, forcing the inverter output to discharge with a rate set by the bias current of MB1. Hence the discharge rate at the output of the inverter is a strong function of the biasing current of MB1 allowing wide variation in the delay of the inverter.

The value of t_{delay} was set by an off-chip reference current I_{ref}. In a practical system for use in industry where a precise definition of the duration of t_{delay} would be required to ensure sufficient settling times for the opamp over process and temperature variations, a DLL can be used to ensure a fixed t_{delay}. Given a sufficient settling time, the clock edge of Φ_{2D} can be allowed to have a large amount of jitter without any significant impact on ADC performance. Hence a potential DLL solution could be implemented with very low power. Furthermore since the most power consuming block of a DLL is typically the delay cell, which is already implemented on-chip in this work, the additional power of a loop filter, charge pump, and phase detector to complete the DLL loop would be relatively small.

5.5.5: SPICE SIMULATIONS

The proposed architecture was designed and verified using Spice. Simulations were carried out over all process corners, and temperatures of 0°C, 80°C and 120°C, where at all corners the system was found to be fully functional. Monte Carlo analysis was used to verify functionality with device and capacitor mismatch in the ADC. Fig. 5-11 shows an FFT of the ADC's output when the ADC operates at 50MS/s with an input tone at 215.625MHz.

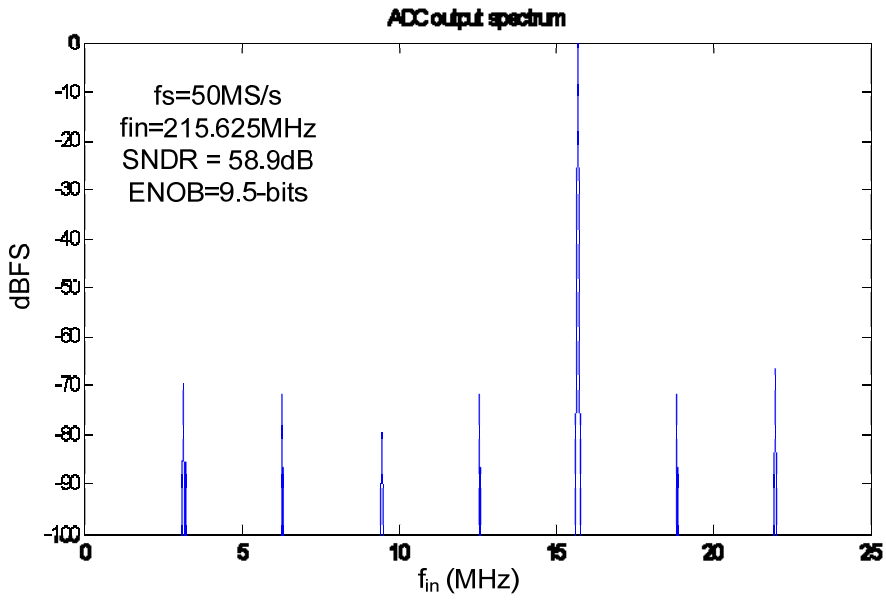

Fig. 5-11: FFT of Spice simulation output

From the Spice simulations with f_{in}=215MHz, the ENOB was 9.5-bits. It is noted that the Spice simulations did not include thermal noises sources and parasitics from the packaging, in an attempt to simplify the simulations. As many of the circuits used in this work are identical to that used in [45], it was expected that the ENOB realized in a fabricated chip will be similar to the 8.5 to 9 bits realized in [45].

5.6: TESTING

Sections 5.6.2 through 5.6.3 discuss the test setup and measured results of the prototype fabricated in a 1.8V 0.18µm CMOS process.

5.6.2: PCB

A 4 layer FR4 dielectric PCB board with a minimum 6mil trace was designed and constructed for the device under test as shown in Fig. 5-12. Separate Power planes were used to isolate the analog, digital, I/O, and board power supplies. A differential input was

generated using a 1:1 turns ratio Minicircuits transformer matched to 50Ω. Reference voltages were generated by passing the output of a resistive voltage divider through an opamp (LM7301) in a unity gain buffer configuration. To maintain constant supply voltages, all voltage supplies for each power plane were generated through regulators (LM337, LM1117), and heavily decoupled with capacitors. As the ADC utilized a constant current biasing scheme, an off-chip adjustable resistor was used as the master current source. The resistance was a series combination of 1kΩ, 10kΩ, 200kΩ, 1MΩ, and 3MΩ potentiometers such that the biasing current could be accurately controlled over a wide range to facilitate the evaluation of wide range current scaling.

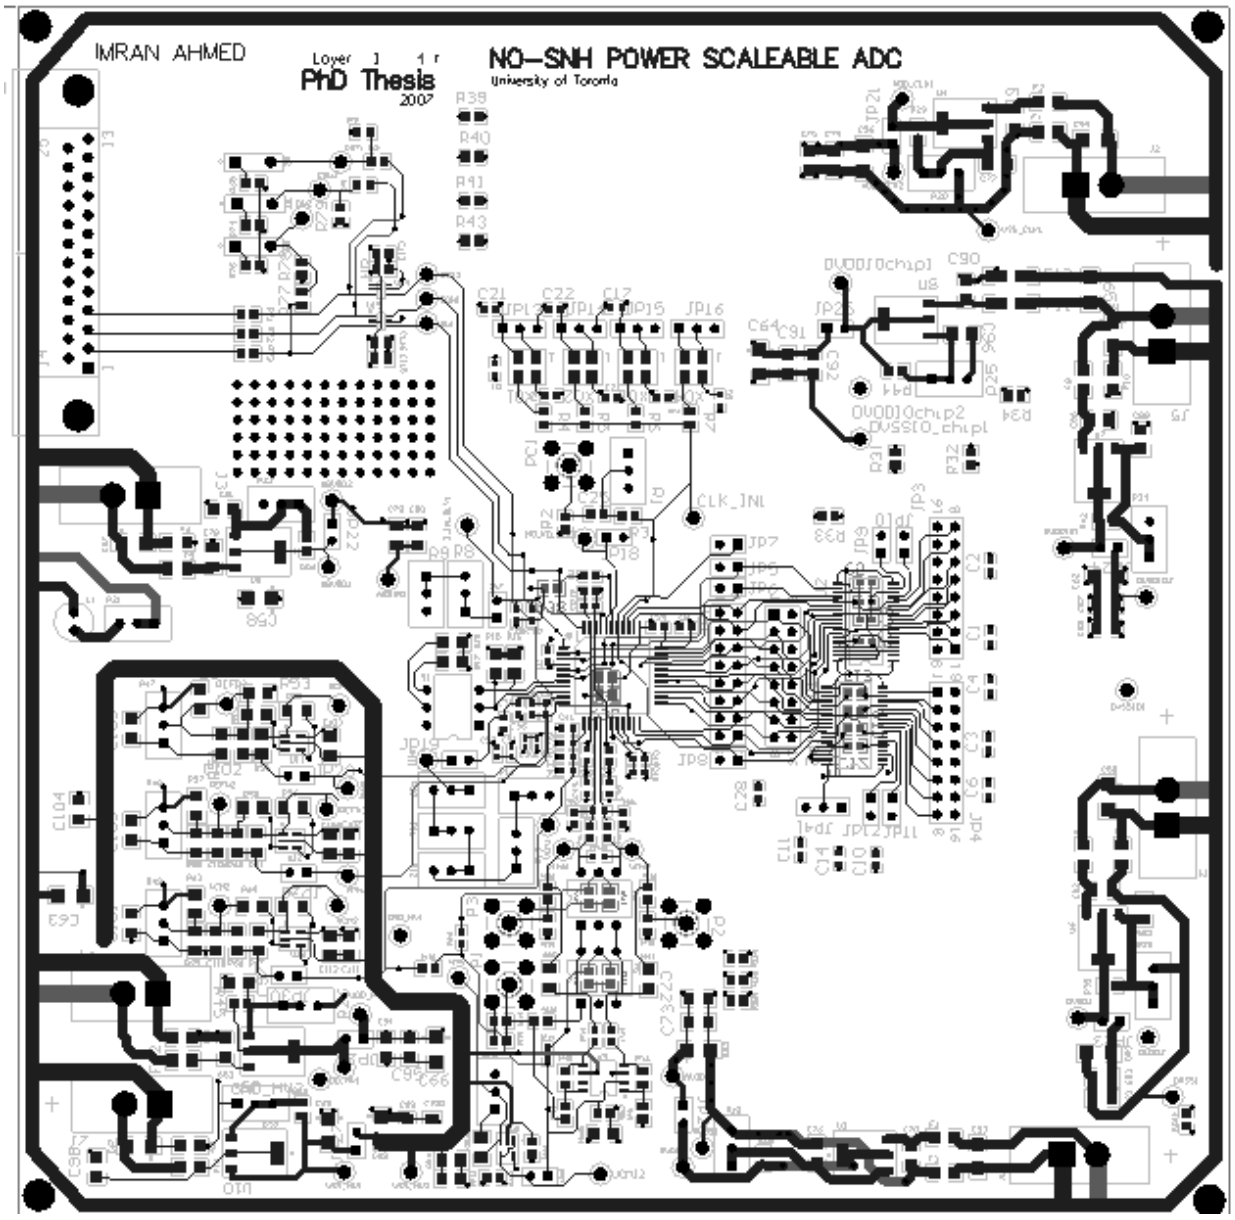

Fig. 5-12: Custom PCB layout

5.6.3: TEST SETUP

A test setup as shown in Fig. 5-13 was used. Sinusoidal inputs were generated using an HP 8664A function generator. Several Minicircuits filters were used to minimize harmonic distortion from the function generator such that the sinusoidal input to the ADC had an SNDR of well over 62dB for input frequencies larger than 270MHz. A 50MHz crystal oscillator (ECS3518-XO) on the PCB was used to generate the low jitter clock for the ADC.

The serial shift register was loaded via a parallel port connection to a PC, where a Matlab script was executed to load the appropriate bits. The output bits of each pipeline stage were captured using a Tektronix TLA714 logic analyzer, capable of capturing 65,536 points at a time. An Agilent E3620A Dual output DC power supply was used to provide positive and negative voltages to the voltage regulators on the PCB. The 10-bit output word from the 10-bit ADC was determined via a Matlab script written to emulate the operation of a digital error correction circuit.

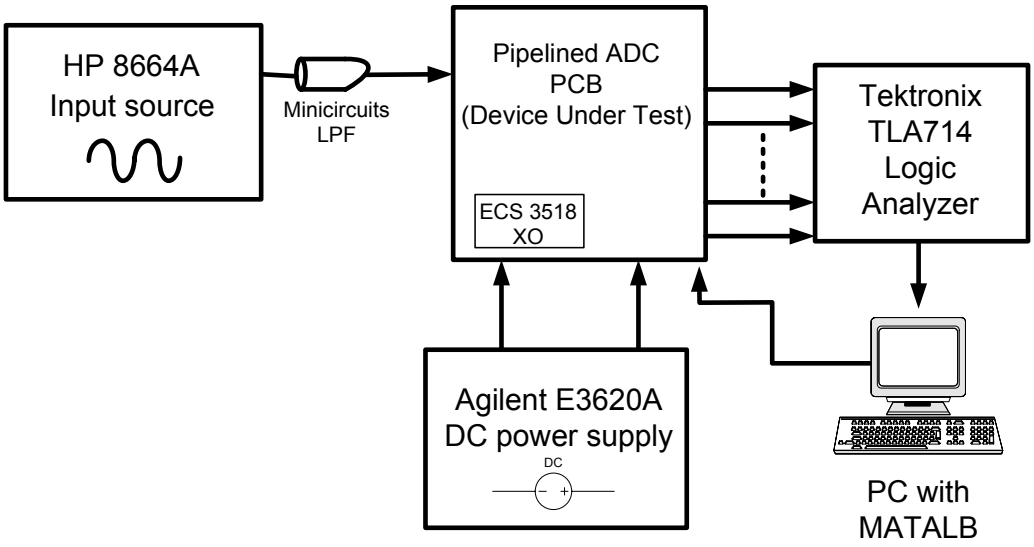

Fig. 5-13: Test setup for pipelined ADC without front-end S/H

5.7: MEASURED RESULTS

A prototype of the ADC of this work as shown in Fig. 5-14 was fabricated in a 1.8V, 0.18µm CMOS process. The core area was 1.1 mm x 1.1 mm, and the maximum input signal swing was 1.6V p-p.

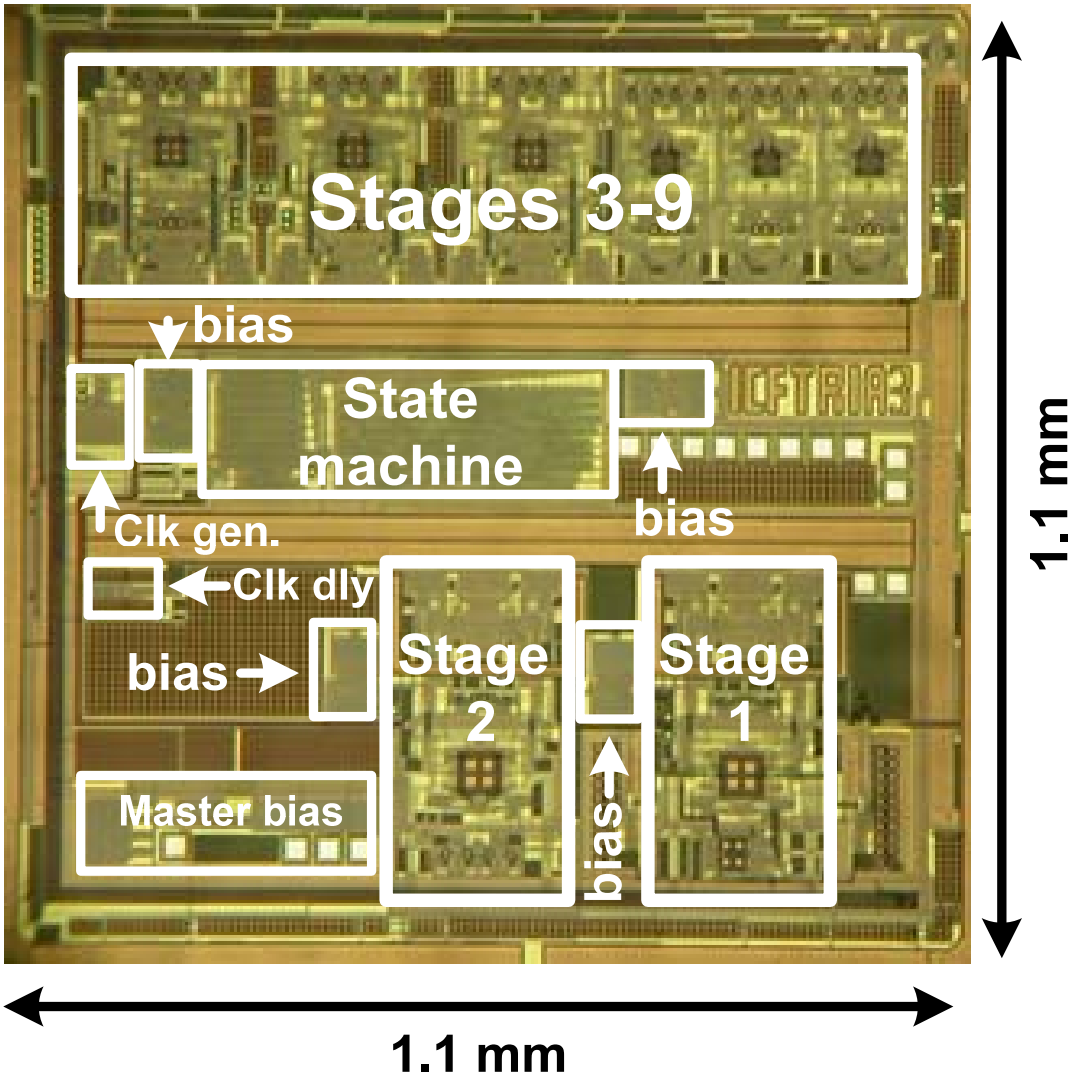

Fig. 5-14: Micrograph of fabricated chip in 1.8V, 0.18µm CMOS

5.7.2: SNDR VERSUS INPUT FREQUENCY

Fig. 5-15 and Fig. 5-16 show the SNDR of the ADC versus input frequency for various power scaled sampling rates. The ADC remains fully functional for input frequencies larger than 267MHz – frequencies which prior techniques using the redundancy of the first pipeline stage would require a very well matched layout (<<140ps skew).

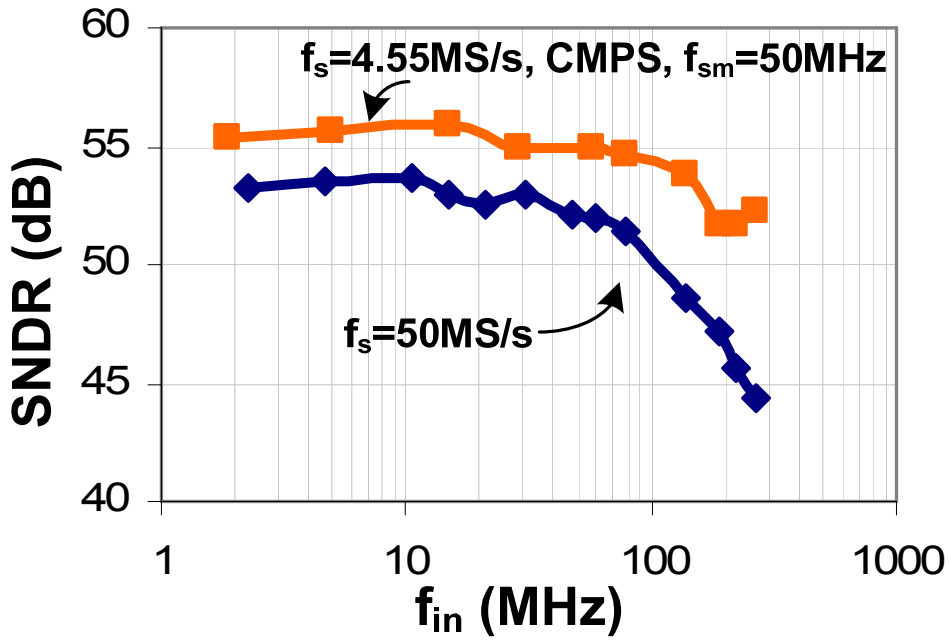

Fig. 5-15: SNDR vs. input frequency for f_s=50, 4.55MS/s

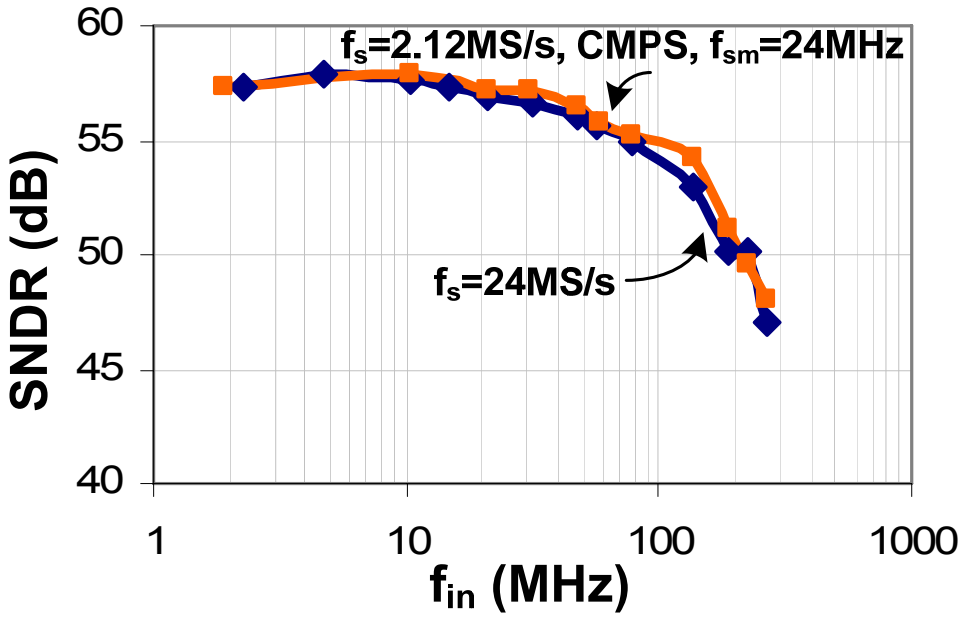

Fig. 5-16: SNDR vs. input frequency for f_s=24MS/s, 2.12MS/s

For f_s=50MS/s and f_{in}=79MHz the SNDR is 51.5dB. For lower sampling rates the input bandwidth is increased, e.g. for f_s=4.55MS/s, the SNDR is 52.2dB for f_{in}=267.1MHz.

Fig. 5-17 shows the FFT of the ADC output for f_s=50MS/s and 4.55MS/s. We note in this work the embedded S/H technique is demonstrated for input frequencies ~4x larger than in [63]. It is noted that the measurements for f_s=4.55MS/s are for the case when CMPS is enabled such that the settling time in each opamp as well as t_{delay} is the same as the case when f_s=50MS/s [45]. Also note that although a single sinusoidal input was used, distortion products are captured in the results due to aliasing from sub-sampling.

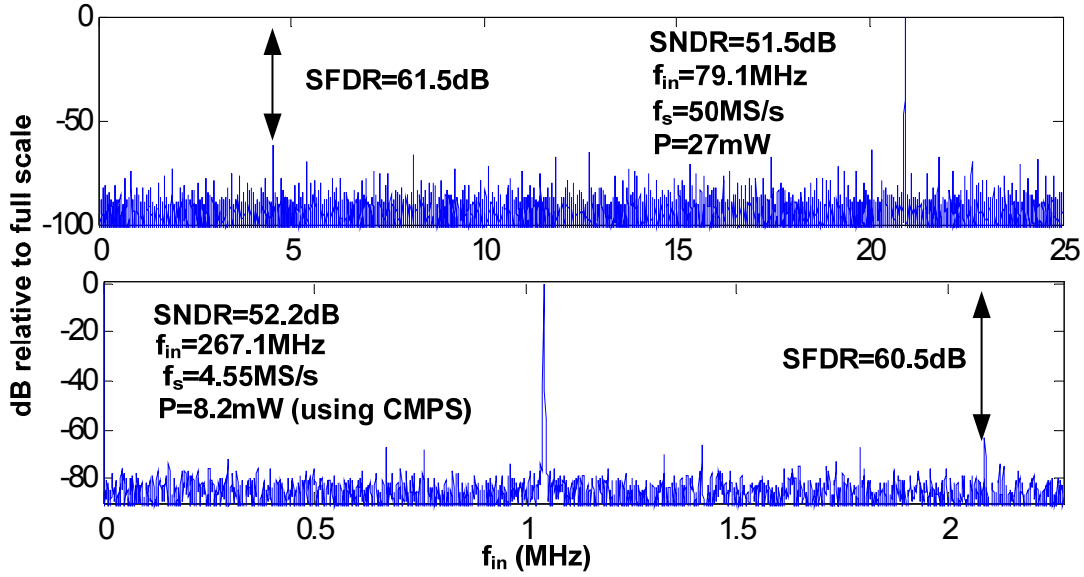

Fig. 5-17: FFT of ADC output at f_s=50, 4.55MS/s

5.7.3: POWER VERSUS SAMPLING RATE

Fig. 5-18 shows the power of the ADC versus sampling rate. The power of the ADC at f_s=50MS/s was 27mW, >20% lower than the 35mW of [45] while maintaining a similar

SNDR. It is noted in this work the 27mW includes an additional bias circuit (1mW) to improve the robustness of the system, as well as clock delay generator (0.5mW) to generate Φ_{2D}.

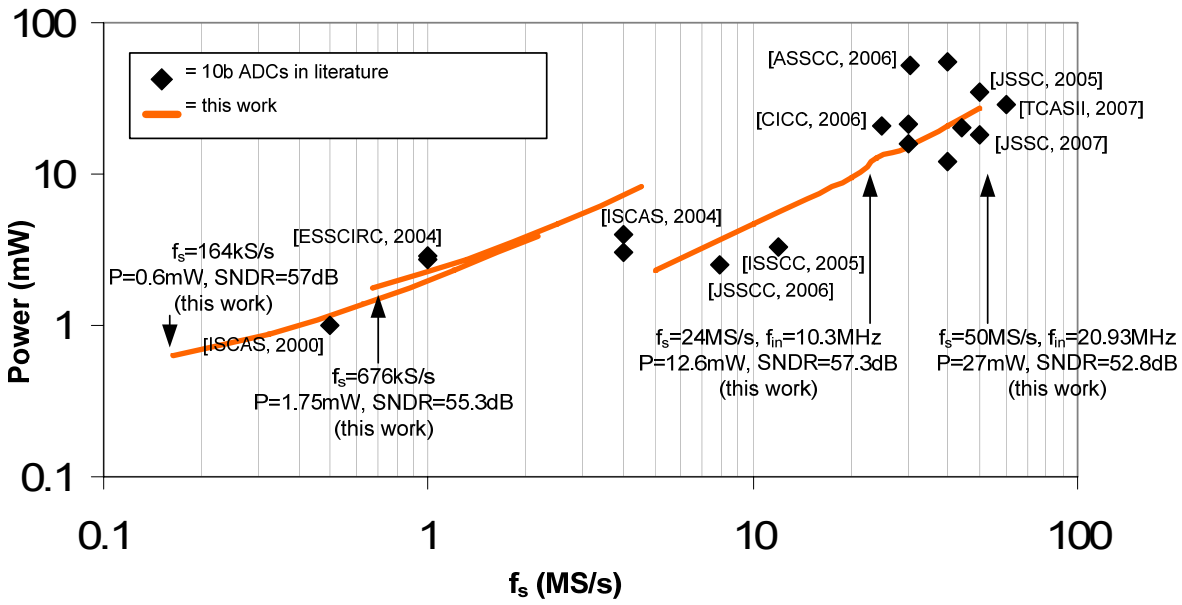

Fig. 5-18: Power vs. sampling rate

Lower sampling rates with correspondingly lower power can be realized by increasing t_{ON} while current scaling the ADC [45]. From Fig. 5-18 the power of the ADC of this work is shown to compare favorably to other recently published 10-bit ADCs [45], [40], [63], [68] - [79] over a wide range of sampling rates.

5.7.4: T_{DELAY} VS. SETTLING TIME – ROBUSTNESS OF PROPOSED TECHNIQUE

In Fig. 5-19 and Fig. 5-20 the SNDR of the ADC is plotted versus the percentage t_{delay} is of the total available settling time when the first pipeline stage is configured using the approach of Fig. 5-3 and Fig. 5-5 respectively.

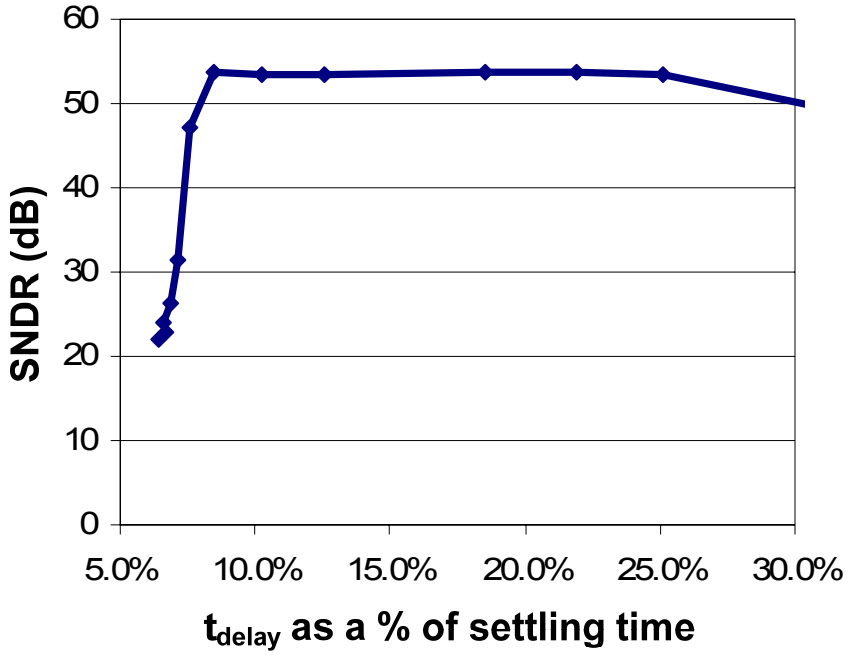

Fig. 5-19: Variation of ENOB with t_{delay} using approach of Fig. 5-3

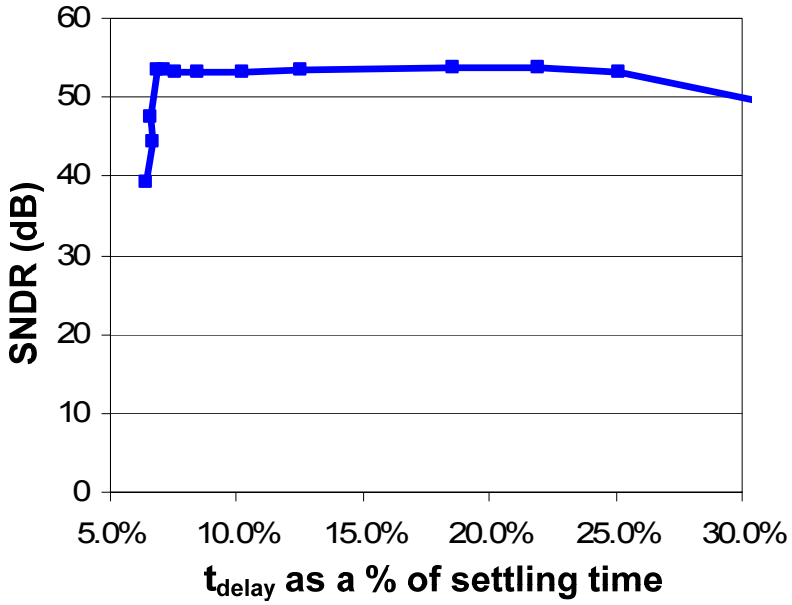

Fig. 5-20: Variation of ENOB with t_{delay} using approach of Fig. 5-5

From both plots it can be concluded that MSB errors only occur when t_{delay} is <10% of the available settling time. Thus the technique to use the embedded sample-and-hold described in this work does not require the first stage opamp to be significantly increased in power to maintain settling accuracy. Furthermore the fact that the SNDR only degrades for t_{delay} larger than 25% of the settling time indicates the power of the first stage opamp could easily be further reduced. Both approaches exhibit similar performance for different t_{delay}, where the approach of Fig. 5-5 shows better performance for slightly smaller values of t_{delay}. This can be attributed to the fact that although the load capacitance of the opamp in the first stage is increased using the approach of Fig. 5-5, for the opamp used in the first stage of this work, the larger feedback factor ultimately allows for faster settling.

5.8: SUMMARY

A power scalable ADC for use in sub-sampled systems with a large input bandwidth was described. Using a technique to remove an explicit front-end sample-and-hold block, a power

savings of >20% was realized. A method to improve the settling behavior of Rapid Power-on Opamps was also presented. Measured results from a 1.8V 0.18µm CMOS prototype show the ADC to achieve more than 51dB SNDR for input frequencies larger than 79MHz for f_s=50MS/s and 267MHz for f_s=4.55MS/s. Table 5-1 summarizes key performance metrics of the results presented in this chapter.

Table 5-1: Summary of key results

Technology	1.8V, 0.18µm CMOS
Area	1.21mm^2
Sampling rate (f_s)	<164 kS/s – 50MS/s
Input signal swing	1.6V p-p
Power	<0.6mW – 27mW
SNDR	>51.5dB for all f_s
SFDR	>60.5dB for all f_s
Input frequency range	0-267MHz
Power of [45] @ 50MS/s	35mW
Power of this work at 50MS/s	27mW

CHAPTER SIX: A LOW POWER PIPELINED ADC USING PASSIVE GAIN AND SOURCE FOLLOWERS

6.1: OVERVIEW

In this chapter a pipelined topology which significantly reduces ADC power consumption and does not require opamps is presented. The gain of each pipeline stage in this work is set by a technique which allows for a stage gain close to 2x to be achieved passively and fully-differentially, where digital foreground calibration is used to compensate for linear gain errors in each pipeline stage. The only active circuit in each pipeline stage is a unity gain buffer which drives the subsequent pipeline stage. In this work, the unity gain buffer was implemented with an NMOS source follower, which was designed to have sufficient linearity to meet the desired 10-bit requirements of this work for a signal swing of 1V p-p. The open-loop technique of this work allows for power savings on the order of those presented in [60], however has a much simpler calibration architecture as non-linearities from the gain stage are not required to be calibrated. Furthermore unlike [61], the approach of this work allows for great power savings but with a topology which is fully-differential. Measured results from a prototype fabricated in a 1.8V, 0.18um CMOS process show the 50MS/s ADC to achieve a peak SNDR/SFDR of 58.2/66 dB after calibration, where the power consumption of the ADC was only 9.9mW for a figure of merit of 0.3 pJ/step.

Section 6.2 discusses the motivation of the low power ADC. Sections 6.3 to 6.5 discuss the implementation of the novel low power ADC. Section 6.6 analyzes the thermal noise of the proposed approach. Section 6.7 reviews the foreground calibration technique used in this work. Section 6.8 presents a theoretical analysis of the improvement afforded with the proposed approach. Section 6.9 outlines the design specifications of this work, and section

6.10 discusses the circuit implementation of the ADC. Section 6.11 discusses simulation results, and sections 6.12 to 6.13 detail the test setup and measurement results of a prototype fabricated in a 1.8V 0.18µm CMOS process. Section 6.14 concludes the chapter.

6.2: MOTIVATION

The primary motivation of the work presented in this chapter is to significantly reduce power consumption in pipelined ADCs. ADC power reduction enables longer battery life in mobile applications, and lower cost packaging in wired applications. In section 2.8, it was shown that the opamp (which is used to achieve a linear stage gain) was the dominant power consumer in pipelined ADCs. As a result, the majority of prior published reports on power optimization techniques in pipelined ADCs ultimately involve power optimizations in the opamp. Since the opamp is the limiting factor in MDAC power consumption, the goal of this work is to fundamentally rethink the design of the MDAC and investigate alternative topologies which do not require power hungry opamps to achieve the desired stage gain.

6.3: ARCHITECTURE – PASSIVE GAIN

In this work, rather than obtaining stage gain from capacitors in a negative feedback network around an opamp, a passive charge sharing technique is used to obtain a 2x gain for a 1.5-bit pipeline stage. The approach of this work is inspired by capacitive charge pumps where successively larger voltages can be obtained by sampling voltages on different capacitors, and subsequently connecting each capacitor in series [80] to yield a total voltage which is the sum of the voltages sampled on each capacitor (e.g. Dickson Charge pump [81]). Fig. 6-1 illustrates an example of a passive voltage doubler using switches.

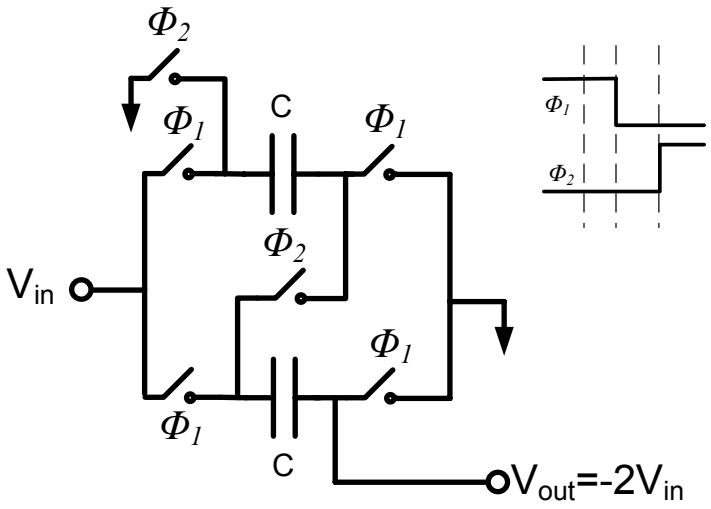

Fig. 6-1: Example of a voltage doubler

From section 2.7 it is recalled that the gain of a 1.5-bit pipeline stage is 2x, thus it is conceivable that a charge sharing technique similar to that shown in Fig. 6-1 can be employed to yield the stage gain in a pipelined ADC stage. Fig. 6-2 illustrates a potential pipeline stage topology which uses passive gain to achieve a stage gain of 2x, where C_1=C_2.

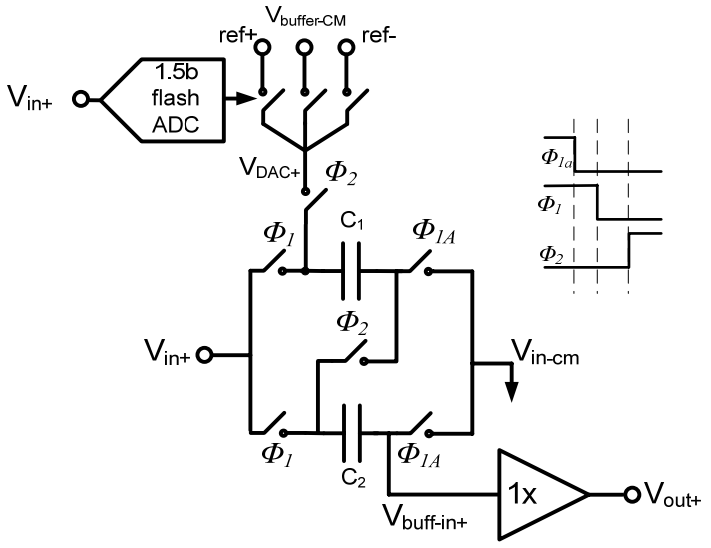

Fig. 6-2: Pipeline stage using passive gain (C_1=C_2)

Assuming the subsequent pipeline stage is constructed similarly to Fig. 6-2, the unity gain buffer is added so that charge sharing does not take place between capacitors from one pipeline stage to another as shown in Fig. 6-3.

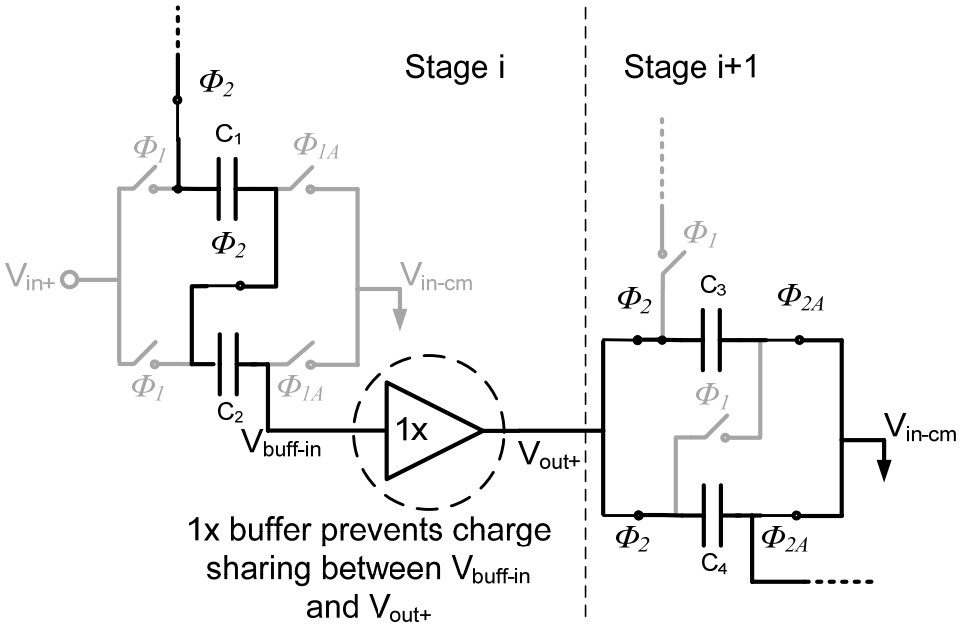

Fig. 6-3: Illustration of how 1x buffer prevents charge sharing

One of the drawbacks of the approach of the pipeline stage shown in Fig. 6-2 is that there is no common-mode rejection for a differential input signal. For example consider the case where each input V_{in+}, V_{in-} is offset from the desired common mode voltage by Δ. As shown in Fig. 6-4 this results in the common mode of the output also being doubled in addition to the desired analog input, i.e. the topology of Fig. 6-4 is pseudo-differential. Thus the pipeline stage shown in Fig. 6-2 has very poor input Common Mode Rejection (CMR)

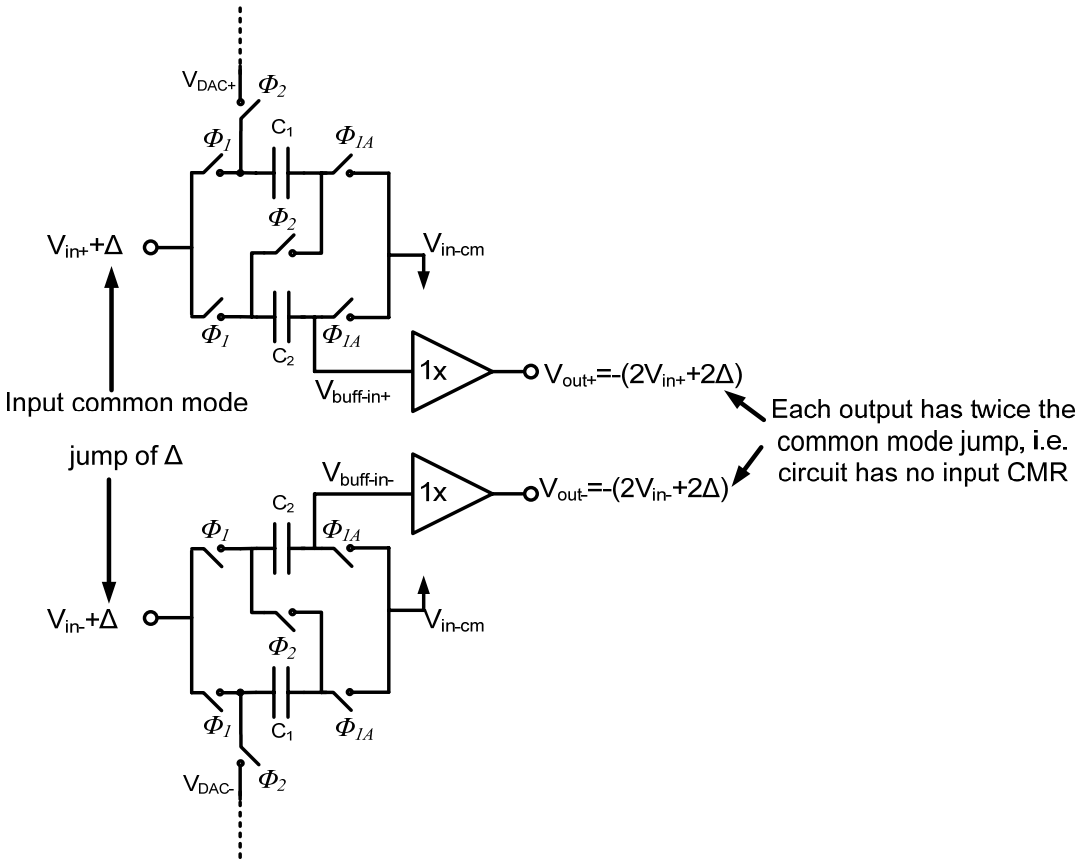

Fig. 6-4: Illustration of poor input CMR for pipeline stage shown in Fig. 6-2

In a pipelined ADC consisting of many stages it is quite likely that with such a structure a small common mode offset at the input of one of the pipeline stages could quickly multiply along the pipeline, making the absolute voltage of the input to a latter pipeline stage approach the supply voltage, thus eliminating the signal swing and hence rendering subsequent pipeline stages unusable. To avoid the poor common mode rejection problem in the architecture of Fig. 6-2, an alternative sampling topology is proposed as shown in Fig. 6-5, and is the topology used to implement each pipeline stage in the ADC discussed in this chapter.

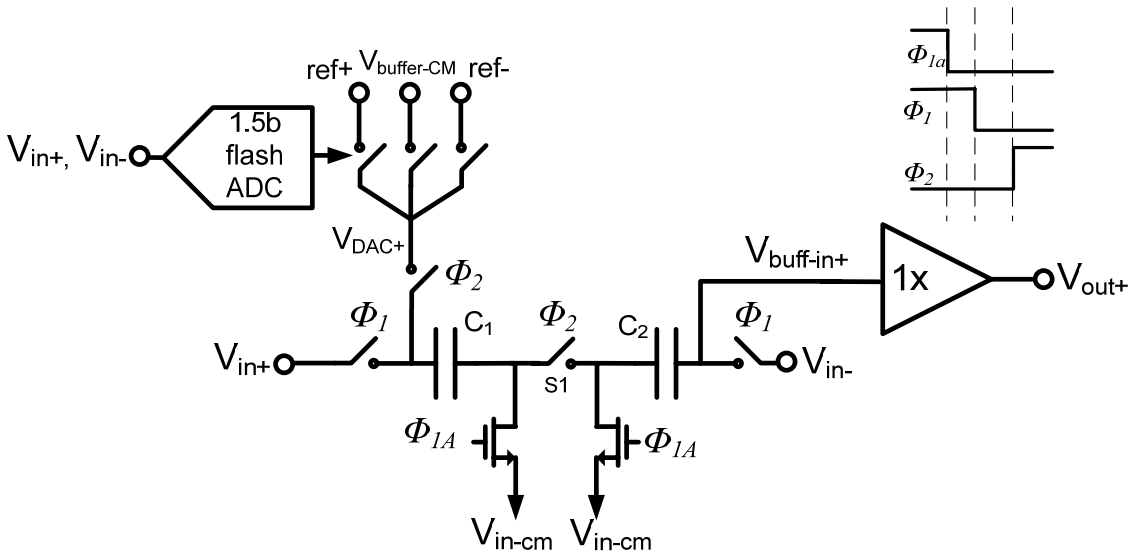

Fig. 6-5: Proposed pipeline stage used in the work of this chapter, where C_1=C_2

In the proposed topology of Fig. 6-5, since the ***differential*** input is effectively sampled across C_1 and C_2, variations in the common mode voltage are not captured by the sampling capacitors. Switch 'S1' is added to enable bottom plate sampling when Φ_{1A} switches from high to low. To generate the negative output V_{outp-} an identical circuit is used where the role of V_{in+} and V_{in-} are reversed as shown in Fig. 6-6.

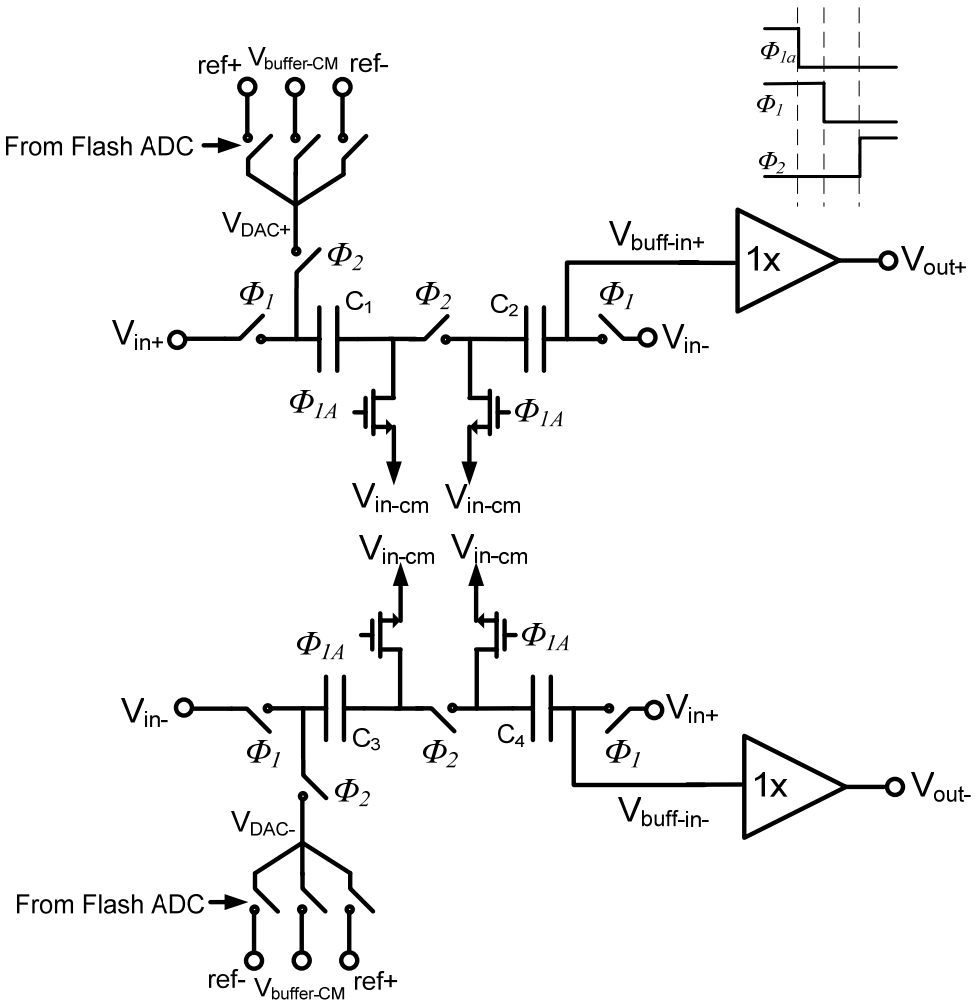

Fig. 6-6: Fully-differential implementation of proposed MDAC

From Fig. 6-6 it can be seen that common mode variations in the input V_{in+}, V_{in-} are rejected by the sampling configuration, and common mode variations in the output V_{out+}, V_{out-} are rejected by the sampling network in the subsequent pipeline stage.

In the proposed MDAC, since the gain of the pipeline stage is derived using a charge sharing technique (i.e. with no virtual ground), the gain of the pipeline stage will not be precisely set by only the ratio of the sampling capacitors C_1 and C_2. As will be seen in the following section, the gain is a function of parasitic capacitors in addition to the sampling capacitors. As noted in section 3.2 however, calibration can be used to measure and compensate for the

effect of a non-ideal stage gain, and is thus used in this work. Furthermore, as will be seen in the subsequent sections, using the topology of Fig. 6-5, the stage gain can be designed to be linear enough that only simple linear calibration techniques (i.e. gain error correction) are required to improve the accuracy of the ADC to the target resolution of this work (10-bits). This is contrasted to the open-loop approach of [60] where by virtue of the gain being derived from a differential pair (which has a non-linear input/output characteristic), a complicated non-linear calibration scheme was required.

The significant advantage of the proposed passive gain technique is the only active circuitry required is a unity gain block. In general all *linear* unity gain buffers use some form of internal feedback, where since the closed-loop gain of the buffer is 1x the value of 1/β=1. Recall from section 2.8.3 that the unity gain frequency of a closed-loop system is given by: $\omega_t = \beta\omega_{ta}$, and that the value of β is at best 0.5 for the traditional MDAC shown in Fig. 5-1. Since ω_{ta} is linearly related to the power of the amplifier [7], it can be deduced that the proposed MDAC topology can achieve the same speed as the traditional MDAC with half the power consumption, due to the larger value of β in the proposed MDAC. In practice it is noted that in high speed ADCs the traditional MDAC suffers from a β smaller than 0.5 due to large parasitic capacitors (typically β is between $1/4$ to $1/3$). Thus the proposed MDAC in practice can achieve the same speed as the traditional MDAC, but with *less than half* the power consumption.

An additional advantage of the proposed technique is: since the buffer comes after the passive gain block, the power of the buffer's noise when referred to the input of the pipeline stage is divided by the amount of passive gain squared – unlike the traditional MDAC where the noise of the opamp is not divided by the stage gain when referred to the input. Thus noise from active circuitry in the proposed topology contributes less to the overall noise floor than in the traditional MDAC, hence enabling further reduction in ADC power.

6.4: EFFECT OF PARASITIC CAPACITORS

In this section the effect of all parasitic capacitors on the gain of the proposed pipeline stage of Fig. 6-5 is examined, where Fig. 6-7 illustrates the key parasitic capacitors.

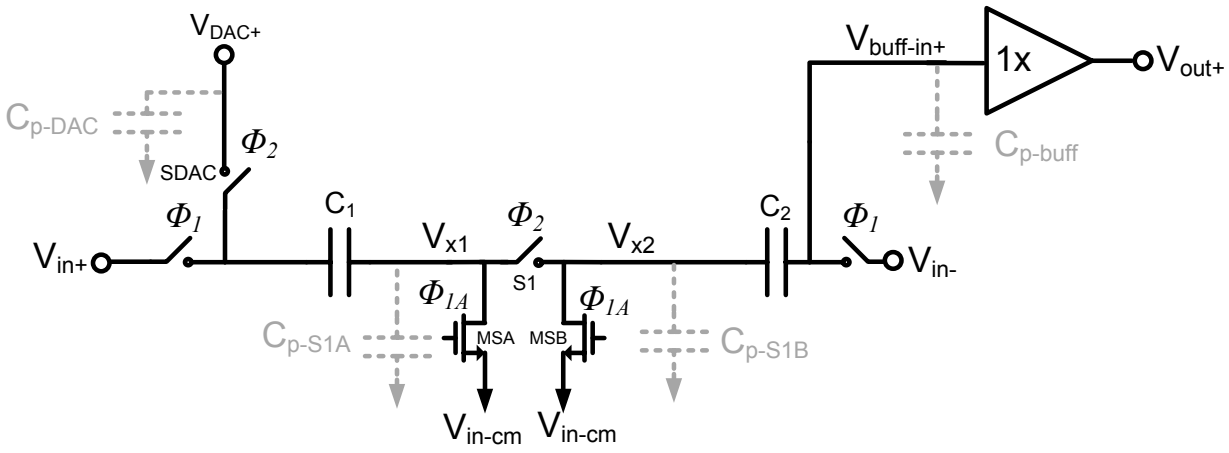

Fig. 6-7: Illustration of parasitic capacitors

From Fig. 6-7, during Φ_1 the charge sampled on C_1 and C_2 respectively are:

$$Q_1 = C_1(V_{in+} - V_{in-cm}), \qquad \textit{(eqn. 6-1)}$$

$$Q_2 = C_2(V_{in-cm} - V_{in-}). \qquad \textit{(eqn. 6-2)}$$

Thus the total charge sampled on C_1 and C_2 is given by:

$$Q_1 + Q_2 = Q_{\Phi 1} = C_1 V_{in+} - C_2 V_{in-} + V_{in-cm}(C_2 - C_1). \qquad \textit{(eqn. 6-3)}$$

Since nodes V_{x1}, V_{x2} are connected to constant voltage sources during Φ_1, parasitic capacitors $C_{p\text{-}S1A}$, $C_{p\text{-}S1B}$ have no impact during Φ_1. By using bottom plate sampling [7], charge injection from the sampling switches MSA, MSB is significantly minimized, and the input is hence linearly sampled on C_1, C_2.

During Φ_2, the total charge on the capacitive network is:

$$Q_{\Phi 2} = C_1(V_{DAC+} - V_x) - C_{p-S1}V_x + C_2\left(V_x - V_{buff-in+}\right) - C_{p-buff}V_{buff-in+} \quad \textit{(eqn. 6-4)}$$

where V_x is the node voltage at V_{x1}, V_{x2} when S1 is closed during Φ_2, and $C_{p\text{-}S1}=C_{p\text{-}S1A}+C_{p\text{-}S1B}$ the total parasitic capacitance due to the sampling switches. Using Kirchoff's Current Law (KCL), an expression for V_x interms of V_{DAC+} and $V_{buff\text{-}int+}$ can be found as:

$$V_x = V_{DAC+}\frac{C_1}{C_1+C_2+C_{p-s1}} + V_{buff-in+}\frac{C_2}{C_1+C_2+C_{p-s1}}, \quad \textit{(eqn. 6-5)}$$

Thus:

$$Q_{\Phi 2} = V_{DAC+}\left[C_1 - \frac{C_1(C_1-C_2+C_{p-s1})}{C_1+C_2+C_{p-s1}}\right] - V_{buff-in+}\left[C_2 + C_{p-buff} + \frac{C_2(C_1-C_2+C_{p-s1})}{C_1+C_2+C_{p-s1}}\right]. \quad \textit{(eqn. 6-6)}$$

Noting that ideally $C_1=C_2$, and that due to charge conservation $Q_{\Phi 1}=Q_{\Phi 2}$, thus:

$$C_1V_{in+} - C_2V_{in-} = V_{DAC+}\left[C_1 - \frac{C_1C_{p-s1}}{C_1+C_2+C_{p-s1}}\right] - V_{buff-in+}\left[C_2 + C_{p-buff} + \frac{C_2C_{p-s1}}{C_1+C_2+C_{p-s1}}\right]. \quad \textit{(eqn. 6-7)}$$

Defining $V_{in+} = V_{in}/2$ and $V_{in}- = -V_{in}/2$, the expression for the voltage at the output of the unity gain buffer is:

$$V_{out+} = A_{buff}\left[-(C_1+C_2)\frac{1}{C_2+C_{p-buff}+\frac{C_2C_{p-s1}}{C_1+C_2+C_{p-s1}}}V_{in} + \frac{C_1-\frac{C_1C_{p-s1}}{C_1+C_2+C_{p-s1}}}{C_2+C_{p-buff}+\frac{C_2C_{p-s1}}{C_1+C_2+C_{p-s1}}}V_{DAC+}\right], \quad \textit{(eqn. 6-8)}$$

where A_{buff} is the gain of the unity gain buffer (which is ~ 1x). When the parasitic capacitors are zero, $C_1=C_2$, and $A_{buff}=1$:

$$V_{out+} = -[2V_{in} - V_{DAC+}] \quad , \quad \textit{(eqn. 6-9)}$$

which is precisely the residue transfer characteristic of a 1.5-bit pipeline stage (negative residue). It is noted that since node V_{DAC+} is driven by the DAC voltage during Φ_2, parasitic capacitors at node V_{DAC+} have no impact, and thus the size of switch SDAC can be made large without impacting the gain or DAC voltage used in the pipeline stage. However as is clear from eqn. 6-8 parasitic capacitors at nodes V_{x1}, V_{x2} and $V_{buff\text{-}in+}$ do affect the gain of the MDAC as well as the DAC voltage subtracted in the output. As noted earlier, calibration is used to compensate for the effect of an MDAC gain which is not exactly equal to 2 as well as the fact that the DAC voltage subtracted is not exactly V_{DAC+}. Furthermore it is noted that that switches MSA, MSB can be sized relatively small (hence have small parasitic capacitance) as they connect to constant DC voltage sources. What is of concern however is the effect of non-linearity in the parasitic capacitors $C_{p\text{-}S1}$ and $C_{p\text{-}buff}$.

It is well known that the total gate capacitance of a MOS transistor varies with gate-source voltage with a shape as shown in Fig. 6-8.

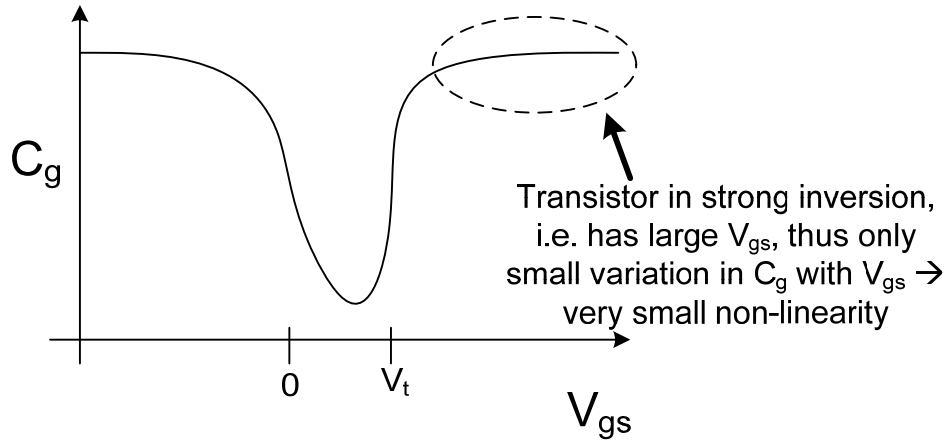

Fig. 6-8: Variation of gate capacitance with gate-source voltage

From Fig. 6-8, if it can be guaranteed that the gate-source voltage of the switches and input transistor in the buffer are in strong inversion during Φ_2, the amount of variation in gate capacitance (which is the dominant parasitic capacitance) with signal swing (i.e. gate and/or source voltages on transistors) can be significantly reduced. Since all switches in the switched capacitor circuit are designed to be strongly inverted to minimize 'on' resistance, and the input transistor of the buffer also designed to be strong inversion, the impact of non-linearities from capacitor variation with signal swing is minimal in the proposed topology. In section 6.11 it will be seen that simulation results showed that more than 10-bit linearity could be achieved at 50MS/s using the proposed topology. It is noted however that the effect of capacitor non-linearity could be a limiting factor in preventing the proposed topology from being useful in highly linear ADCs (i.e. >>10-bit linear).

6.5: UNITY GAIN BUFFER TOPOLOGY

One of the key advantages of the proposed pipelined topology is that the only active circuitry required is a circuit which has a gain of only 1x. From section 2.8.3 it was shown that in feedback systems the unity gain frequency of the closed-loop was maximized for a large feedback factor β. Thus the proposed MDAC topology allows one to achieve a gain larger than one, yet has the benefits of active circuitry which only requires a gain of 1x. Furthermore it is noted that since in many linear unity gain buffers the feedback network can be implemented using a simple wire, β truly equal to one can be achieved – unlike capacitive feedback based systems (e.g. traditional MDAC topology using opamps), where parasitic capacitors reduce the value of β.

To determine the best buffer topology for the proposed MDAC, several commonly used unity gain buffer stages were surveyed. The following figures and tables summarize the benefits and disadvantages of each buffer surveyed. Of note - each buffer topology is shown such that the total current is sufficient for each topology to have the same unity gain frequency of ω_t, where it is assumed each buffer has the same load capacitance C_L.

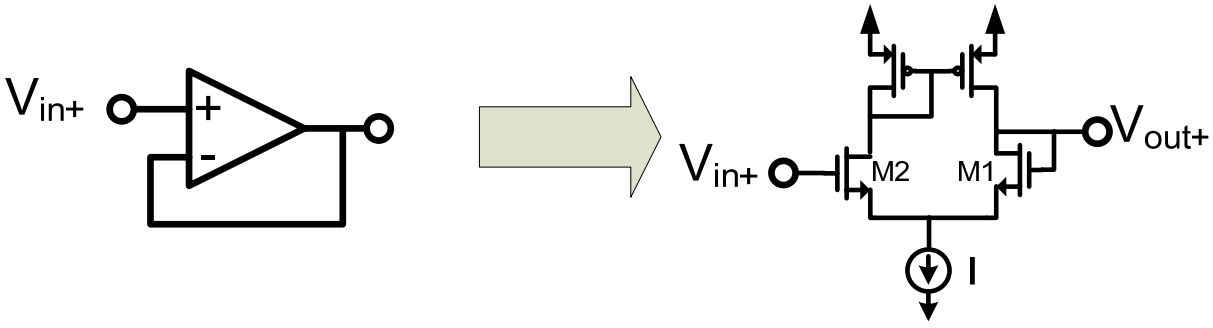

Fig. 6-9: Opamp in unity gain configuration

Opamp in unity gain configuration	
Advantages	**Disadvantages**
✓ 1/β=1 ✓ NMOS diff pair → large g_m ✓ Body effect does not have a huge impact ✓ Output CM tracks input CM ✓ Simulations show >10-bit linearity for 1V-pp input ✓ $Power = I \cdot V_{DD}$	x noise of M1 **and** M2 contribute to single-ended output noise x NMOS diff pair → output/input CM must be high → sampling switch of next stage is PMOS

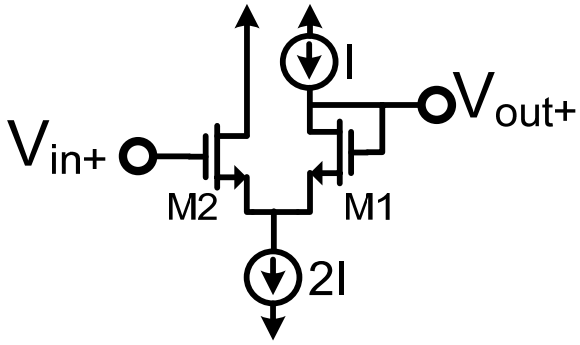

Fig. 6-10: Compound source follower

Compound source follower	
Advantages	**Disadvantages**
✓ NMOS diff pair → large g_m ✓ Body effect does not have a huge impact ✓ Output CM approx tracks input CM ✓ Simulations show >10-bit linearity for 1V-pp input	x noise of M1 **and** M2 contribute to single-ended output noise x $Power = 2I \cdot V_{DD}$ x NMOS diff. pair → output/input CM must be high → sampling switch of next stage is PMOS

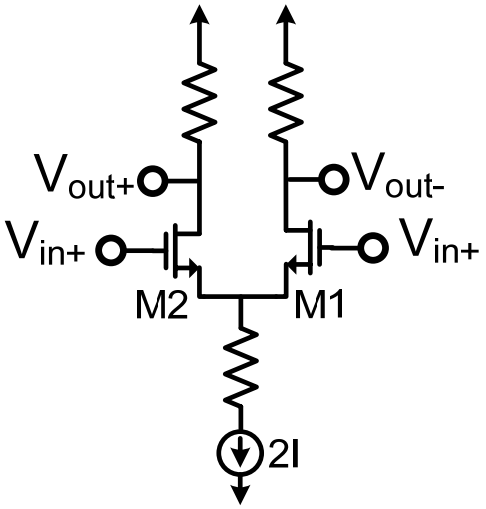

Fig. 6-11: Resistively degenerated differential pair

Resistively degenerated differential pair	
Advantages	**Disadvantages**
✓ NMOS diff pair → large g_m ✓ Can use CMFB to guarantee output CM ✓ Fully-differential – Input CM rejection is high ✓ Only noise of single transistor M1 or M2 contributes to single-ended output noise ✓ $Power = I \cdot V_{DD}$ (single-ended)	x Degeneration reduces bandwidth x Sensitive to process variation x NMOS diff pair → output/input CM must be high → sampling switch of next stage is PMOS x Simulations show <10-bit linearity for 1V-pp input

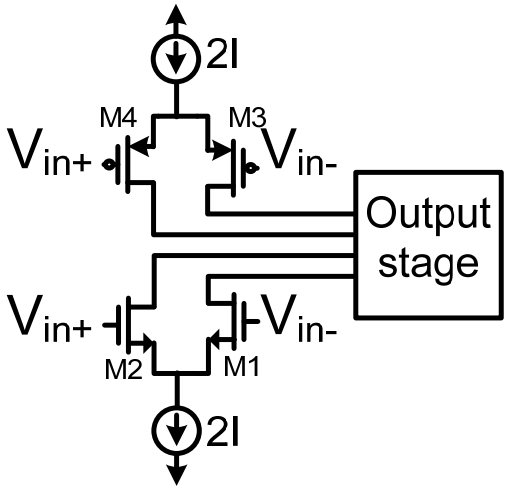

Fig. 6-12: Unity gain buffer which has an N-P complimentary input stage

Unity gain buffer which has an N-P complimentary input stage	
Advantages	**Disadvantages**
✓ Can have large input/output CM variation ✓ Larger swing, smaller sampling caps → lower power	x g_m varies with signal swing x Input offset varies as a function of input swing, leading to poor linearity x Simulations including offsets on reasonably sized input differential pairs show < 10-bit linearity x $Power > I \cdot V_{DD}$ x Large noise

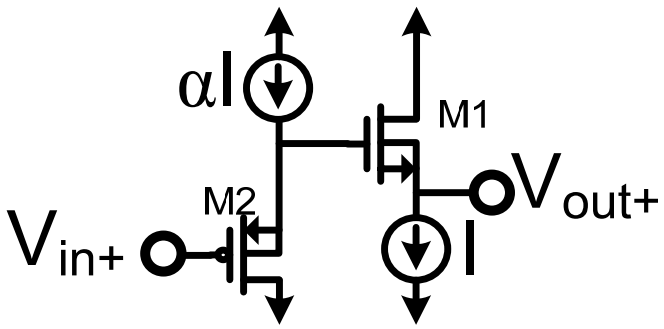

Fig. 6-13: Cascade source follower

Cascade source follower	
Advantages	**Disadvantages**
✓ Output common mode close to input common mode ✓ Simple, compact design	x Output CM varies over process corners x Less headroom x Gain <1 (cascade of two stages with gain <1) x Need DNW x $Power = (1+\propto)I \cdot V_{DD}$ x Large noise

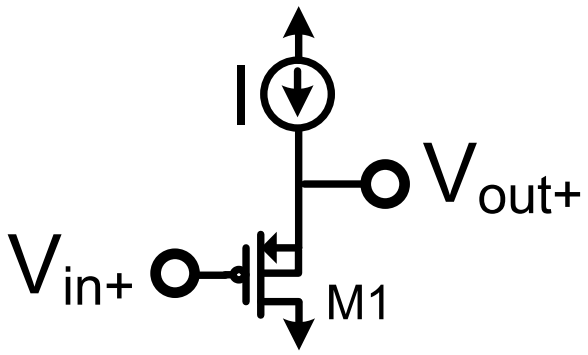

Fig. 6-14: PMOS source follower

PMOS source follower	
Advantages	**Disadvantages**
✓ Simple, compact design ✓ Body effect can be easily eliminated ✓ Can get 10b linearity with reasonable signal swing ✓ $Power = I \cdot V_{DD}$ ✓ Single transistor – low noise	x Small g_m with PMOS input x Output CM not equal to input CM x Output CM is high → next stage needs PMOS sampling switch x Smaller output signal swing

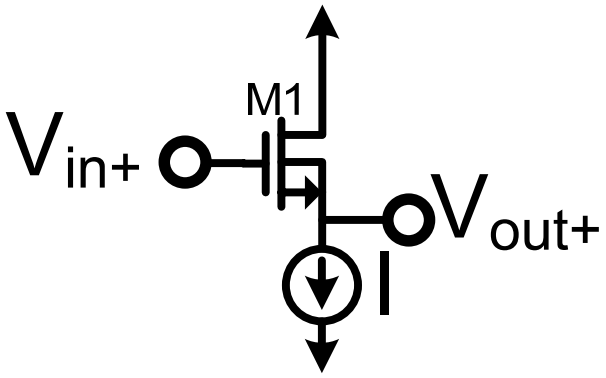

Fig. 6-15: NMOS source follower

NMOS source follower	
Advantages	**Disadvantages**
✓ Simple, compact design ✓ NMOS → high gm ✓ Output has low CM → next stage uses NMOS sampling switch ✓ Can get 10b linearity with reasonable signal swing ✓ Single transistor – low noise ✓ $Power = I \cdot V_{DD}$	x Output CM not equal to input CM x Need DNW to eliminate body effect x Smaller output signal swing

From the survey of unity gain buffers the most power efficient and simple buffer structure which meets the minimum 10-bit linearity requirements of this work is the NMOS source follower. The NMOS source follower had the lowest input referred noise spectral density amongst all surveyed amplifiers, and furthermore by using an NMOS device the g_m of the amplifier is maximized (noting that PMOS transistors usually have $1/4^{th}$ the g_m of an NMOS). An additional advantage is since the output common mode voltage of an NMOS source follower is closer to V_{SS}, it is possible to use NMOS sampling switches in the subsequent pipeline stage as shown in Fig. 6-16.

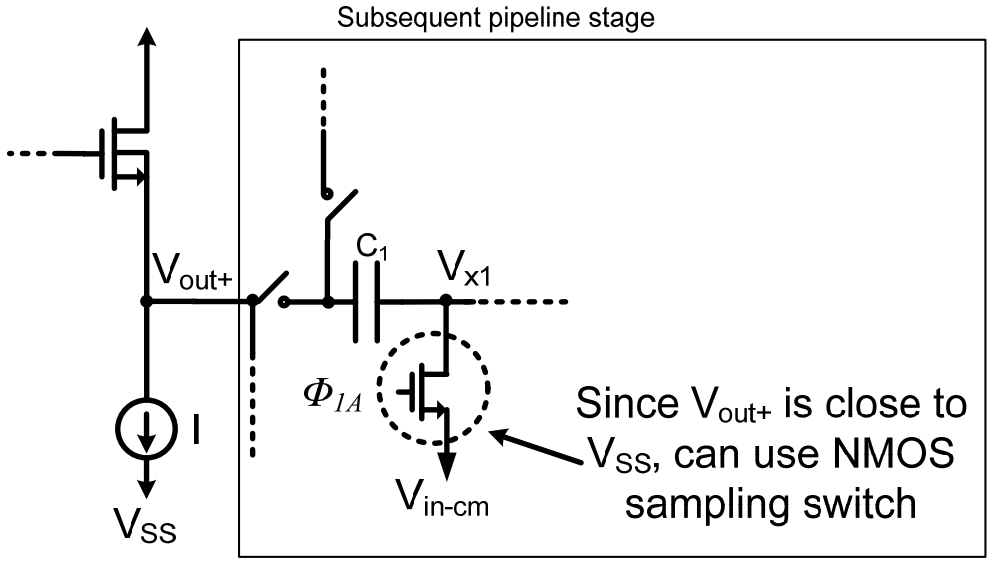

Fig. 6-16: Illustration of how an NMOS source follower enables NMOS sampling switches in subsequent pipeline stage

As NMOS transistors have lower 'on' resistances, they can be sized much smaller than PMOS switches (usually 1/4thsmaller), thus minimizing the parasitic capacitive load at node V_{x1}. NMOS switches also have less charge injection than PMOS switches, as PMOS switches are typically sized 4x larger than NMOS switches to achieve the same settling time. Hence using an NMOS source follower as opposed to a PMOS source follower also has the advantage of increased linearity.

One of the disadvantages of source followers is the different common-mode voltage between input and output nodes of the amplifier. However as noted in section 6.3, using the proposed topology the common mode at the input of the unity gain buffer is decoupled from the common mode at the output of the unity gain buffer. That is, since the input is differentially sampled on capacitors C_1, C_2, the common mode of node voltage $V_{buff\text{-}in+}$ in Fig. 6-5 is set by the common mode voltage of the DAC voltage V_{DAC+} which from Fig. 6-5 is $V_{buffer\text{-}CM}$, and hence the input of the unity gain buffer can have a different common mode voltage than the output. Thus in the proposed MDAC topology common mode voltage shifts between input

and output in the unity gain amplifier do not affect the common mode voltages in subsequent pipeline stages.

To minimize the body effect an NMOS source follower requires a Deep N-Well layer (DNW). It is noted however that in most digital processes used by industry, the DNW layer is readily available and thus there is no additional fabrication cost in using the DNW layer.

The parasitic input capacitance of the source follower is determined by the capacitance from gate to drain (C_{gd}) and gate to source (C_{gs}) of transistor M1 as shown in Fig. 6-17.

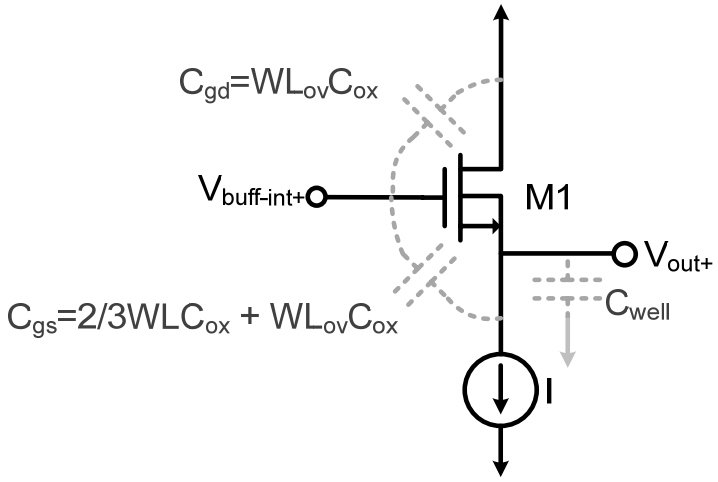

Fig. 6-17: Parasitic capacitance in NMOS source follower

From [7] it is noted that C_{gs} is the larger of the two input parasitic capacitors. A major advantage of the source follower topology is since the circuit has a unity gain between gate and source, the gate and source move approximately together, and thus the effect of the input capacitance C_{gs} is significantly reduced, leaving the input capacitance of the unity gain buffer to be dominated by the relatively small C_{gd} [20]. Thus a source follower topology has small

input parasitic capacitance, which hence enables a larger stage gain in the proposed MDAC topology.

Of note, the parasitic capacitance due to the DNW, C_{well}, slightly reduces the bandwidth, however as the output is discrete time, the non-linear nature of the well capacitance does not have a significant impact at the 10-bit level given a sufficient settling time.

6.5.2: LINEARITY OF SOURCE FOLLOWER IN A SAMPLED SYSTEM

The linearity of the source follower is limited primarily by the variation of the gain, and output impedance respectively with output signal swing. Since the source follower input is discrete time, non-linear settling times which plague source followers for sinusoidal inputs are not an issue for discrete time inputs assuming a sufficient settling time. If the body effect is eliminated by connecting the source and body together, the gain of a source follower as shown in Fig. 6-18 is given approximately by [20]:

$$A = \frac{g_{m1} r_{o1} || r_{ob}}{1 + g_{m1} r_{o1} || r_{ob}}, \qquad \textit{(eqn. 6-10)}$$

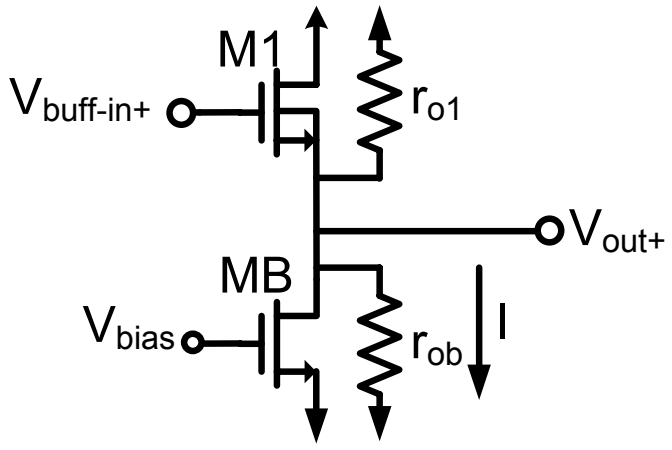

Fig. 6-18: NMOS source follower with output resistances labeled

where r_{o1} and r_{ob} are the output small signal resistances of M1 and MB respectively. Since $g_m \propto \sqrt{I}$ and $r_o \propto {}^{1}/_{I}$, the gain of the source follower is a weak function of $I^{-0.5}$, and thus changes in the bias current 'I' with variations in the output voltage V_{out+} result in small changes in the gain of the source follower with output signal swing, hence harmonic distortion. The bias current 'I' is set predominantly by the bias voltage V_{bias}, however due to short channel effects, variations in the drain source voltage of MB also changes the bias current 'I'. Thus the signal swing at the output V_{out+} modulates the bias current 'I', which in turn results in distortion. To minimize distortion in the source follower several techniques can be used (e.g. [53]), however a simple solution is to use a large length for transistor MB, which trades a slightly larger load capacitance with lower design complexity. In this work harmonic distortion in the source follower was minimized by using a large length (L=0.3µm) for MB. An added advantage of using a large length for the bias transistor is that the transconductance of MB is reduced, and as will be seen in eqn. 6-14 of section 6.6, results in lower noise in the source follower.

Thus in the proposed MDAC topology it is possible to achieve a linear stage gain using a source follower, and hence the linear gain calibration techniques as described in section 3.2.3 can be used to achieve a high resolution in the ADC. Unlike [60], complicated non-linear calibration is not required in the proposed approach, even though opamps are not used.

6.5.3: SIGNAL SWING OF SOURCE FOLLOWER

To maintain a linear output, the source follower is required to keep MB in the active region. Thus the minimum output voltage is given by the overdrive voltage of MB, i.e. $V_{out+} > V_{bias}-V_t$, (V_t is the threshold voltage). The maximum voltage is given by the fact that the source voltage of M1 in Fig. 6-19 is always V_t plus V_{ov1} below the gate (V_{ov1} is the overdrive voltage of M1). Thus if all voltages are limited to be no higher than V_{DD}, the maximum output voltage is $V_{DD}-V_t-V_{ov}$ as noted in Fig. 6-19.

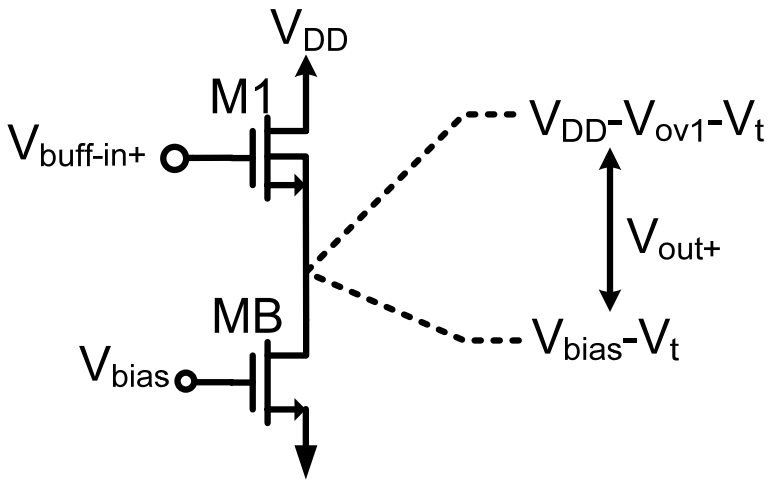

Fig. 6-19: NMOS source follower signal swing

It is noted however, that since the source of M1 tracks the gate of M1 (i.e. approximately constant V_{gs}), in theory a voltage higher than V_{DD} could be used at the gate of the source follower without having to worry about reliability issues from oxide break down which occur when $V_{gs} > V_{DD}$. Thus assuming care is taken at startup and at all other voltage nodes in the MDAC, it is conceivable that a source follower could be used and a large signal swing realized.

Alternatively, it is noted that in industry dual supply voltages are commonly used, where the analog portions of an ADC are designed with a large supply voltages and thick oxide devices, whereas the digital portions are designed with thin oxide low voltage devices. Thus in a deep submicron process (e.g. 45nm) where for example it is desired to keep all voltages below V_{DD}, a source follower could still be used as the unity gain amplifier so long as thick oxide devices and a second higher supply voltage are used for the source follower. Node voltages higher than V_{DD} and/or dual supply voltages were not used in this work; however an investigation of large signal swing (thus lower power) with the source follower is deferred as future research.

It is noted that if a low V_t NMOS (which is becoming more commonly available) was used in the source-follower, a significantly larger signal swing could be achieved, enabling smaller sampling capacitors and thus lower power consumption.

6.6: NOISE ANALYSIS OF PROPOSED MDAC

The input referred noise of the proposed MDAC shown in Fig. 6-5 can be evaluated by referring all noise sources to one of the single-ended inputs of the ADC (e.g. V_{in+}). During Φ_1, the noise sampled by the capacitors is as shown in Fig. 6-20, and given by:

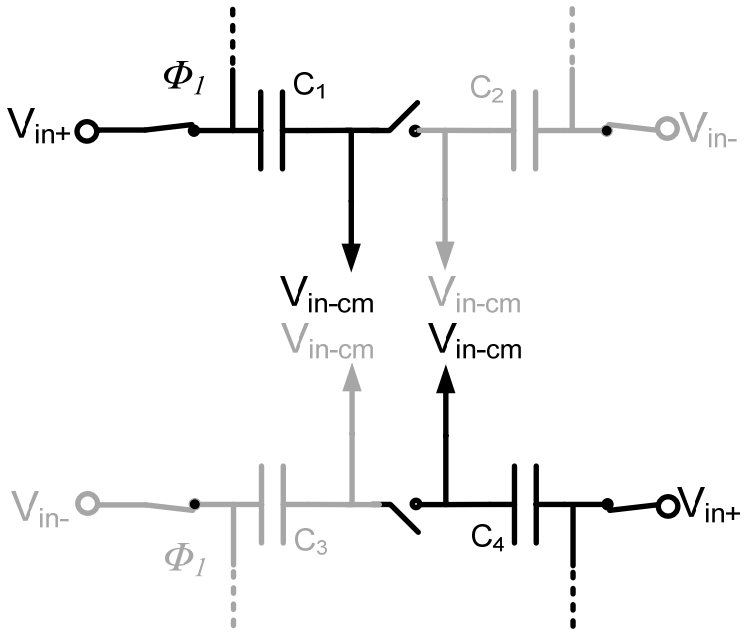

Fig. 6-20: Signal path of V_{in+} during Φ_1

$$\sigma^2_{noise-\Phi 1} = \frac{kT}{C_1+C_4}. \qquad \textit{(eqn. 6-11)}$$

As shown in Fig. 6-21, during Φ_2 the noise is determined by the noise from the switches with an 'on' resistance of R_{on}, where the total noise that appears at V_{out+} is shaped by the bandwidth of the unity gain buffer which has a unity gain frequency of $\omega_{ta}=g_{m1}/C_L$.

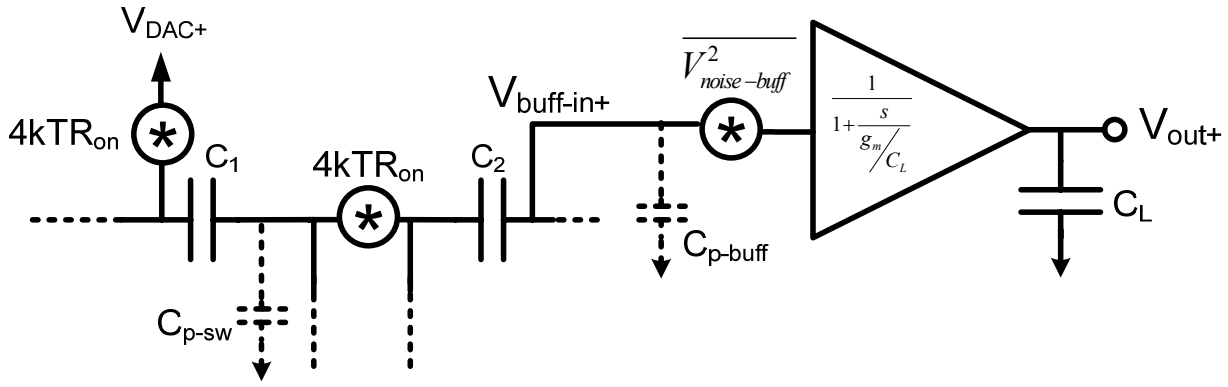

Fig. 6-21: MDAC configuration during Φ_2

Assuming the parasitic capacitances are only a small fraction of the sampling capacitors C_1, C_2, the noise during Φ_2 referred to the input V_{in+} can be found to be:

$$\sigma^2_{noise-\Phi 2} = \frac{1}{A^2}\left[2g_{m1}R_{on}\frac{kT}{C_L} + \sigma^2_{noise-buff}\right], \qquad \textit{(eqn. 6-12)}$$

where A is the gain of the pipeline stage, i.e.:

$$A = \frac{V_{out+}}{V_{in}} = -A_{buff}(C_1 + C_2)\frac{1}{C_2 + C_{p-buff} + \frac{C_2 C_{p-s1}}{C_1 + C_2 + C_{p-s1}}}. \qquad \textit{(eqn. 6-13)}$$

As $V_{buff-in+}$ is effectively a floating node during Φ_2, capacitors C_1 and C_2 have no effect on the noise during Φ_2 under the assumption that the unity gain buffer is the block which determines the noise bandwidth during Φ_2. This is a fair assumption as R_{on} is typically small and the parasitic capacitors C_{p-sw} and C_{p-buff} are typically only a fraction of the sampling capacitors C_1, C_2.

When a source follower as shown in Fig. 6-22 is used for the unity gain buffer, the total noise of the buffer [20] can be found to be:

$$\sigma^2_{noise-buff} = \frac{4}{3}\frac{kT}{C_L}\left(1+\frac{g_{m2}}{g_{m1}}\right) \qquad \textit{(eqn. 6-14)}$$

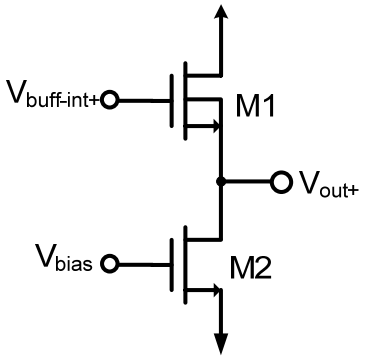

Fig. 6-22: NMOS source follower

Thus the noise in Φ_2 can be written as:

$$\sigma^2_{noise-\Phi 2} = \frac{1}{A^2}\left[2g_{m1}R_{on}\frac{kT}{C_L}+\frac{4}{3}\frac{kT}{C_L}\left(1+\frac{g_{m2}}{g_{m1}}\right)\right] \qquad \textit{(eqn. 6-15)}$$

Hence the total input referred noise of the proposed pipeline stage is given by the sum of eqn. 6-11 and eqn. 6-15:

$$\sigma^2_{noise-input} = \sigma^2_{noise-\Phi 1} + \sigma^2_{noise-\Phi 2} = \frac{kT}{C_1+C_4} + \frac{1}{A^2}\left[2g_{m1}R_{on}\frac{kT}{C_L}+\frac{4}{3}\frac{kT}{C_L}\left(1+\frac{g_{m2}}{g_{m1}}\right)\right] \quad \textit{(eqn. 6-16)}$$

Fig. 6-23 shows a Spice simulation of the power spectral density of the noise output of the circuit shown in Fig. 6-21 during Φ_2, when C_L=1.06pF, g_{m1}=2.2mA/V, g_{m2}=0.84mA/V, and T=300K.

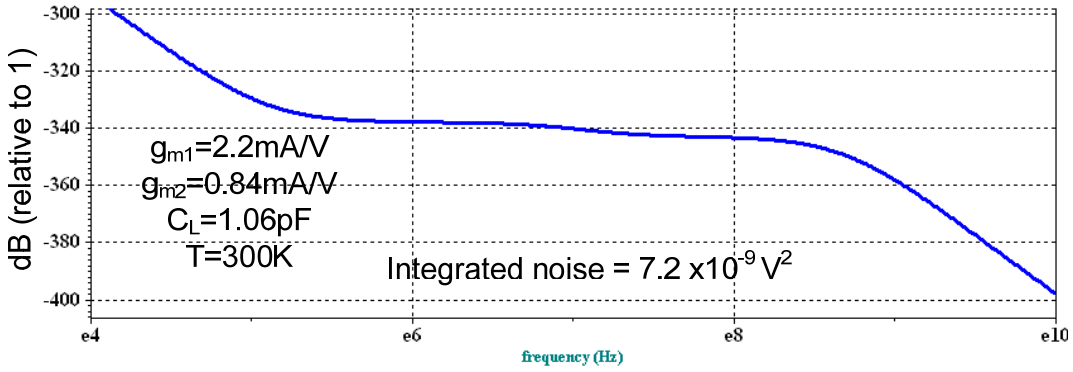

Fig. 6-23: Power spectral density of noise at V_{out+} during Φ_2

From eqn. 6-15, when the thermal noise from the switches is assumed small, the output referred noise (i.e. before dividing by the gain squared) is calculated to be:

$$\sigma^2_{\text{noise}-\Phi 2-\text{output}} = \frac{4}{3}\frac{\text{kT}}{C_L}\left(1+\frac{g_{m2}}{g_{m1}}\right) = \frac{4}{3}\frac{(1.38\text{x}10^{-23})(300\text{K})}{1.06pF}\left(1+\frac{0.84}{2.2}\right) = 7.2\ \times 10^{-9}\ V^2,$$

which agrees very well with the simulated results, thus verifying the derived noise analysis.

In general R_{on} can be made small, and $g_{m2}<g_{m1}$, thus the total input referred noise can be approximated by:

$$\sigma^2_{noise-input} \approx \frac{kT}{C_1+C_4} + \frac{1}{A^2}\left(\frac{4}{3}\frac{kT}{C_L}\right) \qquad \textit{(eqn. 6-17)}$$

In a traditional 1.5-bit based MDAC topology as shown in Fig. 5-1, during Φ_1, the noise is also given by $kT/(C_1+C_2)$. To find the noise during Φ_2, consider Fig. 6-24 which shows the basic configuration of the traditional MDAC topology during Φ_2.

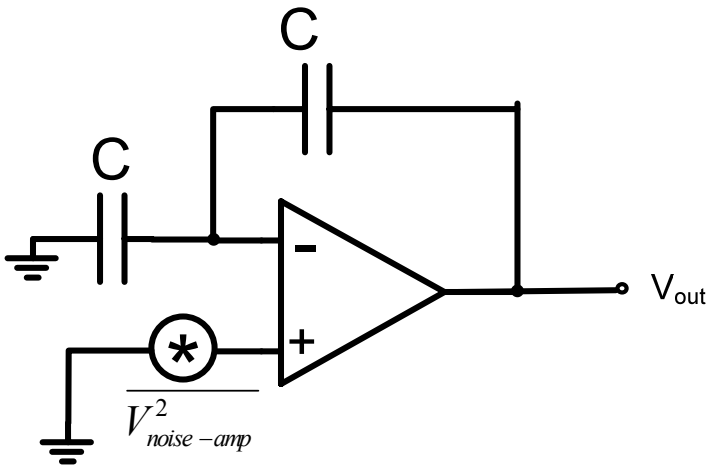

Fig. 6-24: Amplifier noise during Φ_2 in traditional MDAC

If it is reasonably assumed that the switch bandwidth is very large, the noise during Φ_2 will be dominated by the opamp. From Fig. 6-24, during Φ_2 the DC gain of the noise is given by [82]:

$$A_{noise-\Phi 2} = 1 + \frac{C}{C} = 2 \qquad \textit{(eqn. 6-18)}$$

However the signal gain is also 2, thus the thermal noise of the opamp directly contributes to the input referred noise of the pipeline stage. If it is assumed the noise of the opamp is dominated by the input differential pair, the opamp thermal noise referred to the input of the MDAC can be found to be [83]:

$$\sigma^2_{noise-buff} \approx \frac{4}{3}\frac{kT}{C_{L2}} \qquad \textit{(eqn. 6-19)}$$

where C_{L2} is the total load capacitance during Φ_2. Hence the total input referred noise for the traditional MDAC topology is given by:

$$\sigma^2_{noise-traditional} \approx \frac{kT}{C_1+C_2} + \frac{4}{3}\frac{kT}{C_{L2}} \qquad \textit{(eqn. 6-20)}$$

Comparing eqn. 6-17 and eqn. 6-20 it is clear that the proposed MDAC topology has a significant advantage in that since the unity gain buffer is connected to the output of the passive gain stage, the noise of the amplifier is reduced by A^2. Thus the proposed MDAC can be designed to achieve the same speed as an opamp based MDAC with smaller capacitors and thus lower power than the traditional MDAC.

6.7: CALIBRATION OF PIPELINE STAGES

In section 2.8.2 it was noted that gain errors in the MDAC resulted in missing codes which increase harmonic distortion in the ADC's output. From eqn. 6-13 it is clear that the gain of the 1.5-bit stage is a function of parasitic capacitance as well as the gain of the buffer – values which cannot be accurately estimated before fabrication. As discussed in section 3.2 however, calibration can be used to measure and correct the non-ideality. In this work foreground calibration was used to measure and correct the error of each stage. From eqn. 6-13 however, the gain of the proposed pipeline stage is a function of the unity gain amplifier's gain, which for a source follower is a function of temperature which varies with time. Thus in a practical implementation for use in industry it would be more desirable to use a background calibration scheme. Implementing a background calibration scheme however as discussed in chapter 4 is non-trivial, thus in the interest of reducing the complexities of an initial prototype, foreground calibration – which is much simpler to implement, was used. It is noted that the proposed low power MDAC topology is compatible with most background calibrations schemes, including the split-ADC approach discussed in chapter 4. Thus it is conceivable that in an industrial implementation background calibration could be used in the ADC.

6.7.2: FOREGROUND CALIBRATION IN DETAIL

Consider the ADC topology of Fig. 6-25, which shows a 1.5-bit first pipeline stage followed by an ideal backend Flash ADC.

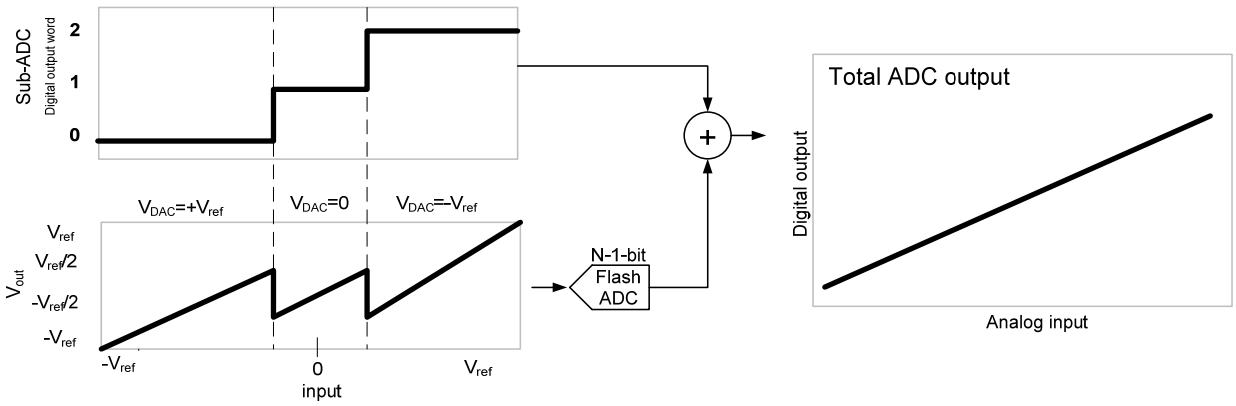

Fig. 6-25: ideal 1.5-bit first pipeline stage

If there is a gain error in the first pipeline stage, the output of the ADC is as shown in Fig. 6-26.

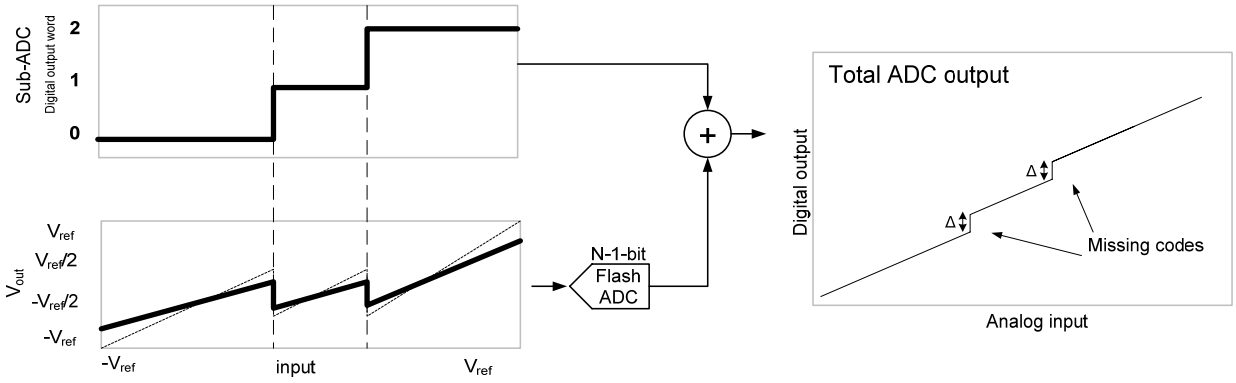

Fig. 6-26: 1.5-bit pipeline stage with gain error

Thus the objective of the calibration scheme is to estimate the number of missing codes, Δ.

Consider the residue transfer curve of a 1.5-bit stage as shown in Fig. 6-25. If the input to the pipeline stage is zero, the DAC voltage can be either 0, $+V_{ref}$, or $-V_{ref}$, and assuming no gain

errors in the pipeline stage, the output of the 1.5-bit pipeline stage will not saturate the output due to the 0.5-bit redundancy of the pipeline stage. Thus in an ideal 1.5-bit pipeline stage with zero input, the output of the ADC will be constant regardless of the DAC voltage (assuming the MSB bits reflect the value of the DAC voltage).

However if there is a gain error in the pipeline stage (such that the gain is less than the ideal value), the ADC will output different values when the DAC voltage is tied to $+V_{ref}$, 0, and $-V_{ref}$ respectively with the first pipeline stage set to have a zero input. Fig. 6-27 illustrates the ADC output when V_{in} is near zero without and with errors.

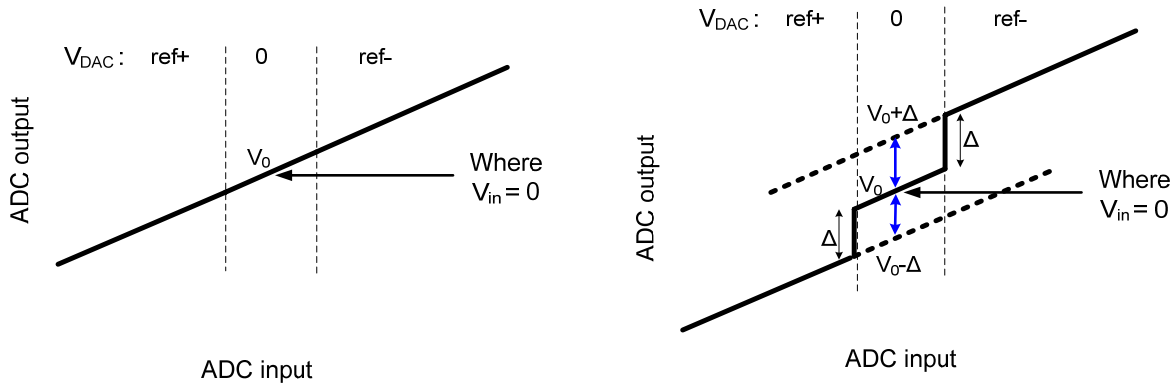

Fig. 6-27: Measure of missing codes when pipeline stage input (V_{in}) is zero – left is ideal, right is with errors

Thus the missing codes produced by a non-ideal stage gain can be corrected in the foreground by shorting the input of the pipeline stage under calibration to zero, and separately measuring the output of the ADC when the DAC voltage of the stage under calibration is connected to $+V_{ref}$, 0, $-V_{ref}$ respectively. By averaging out each value for a few clock cycles to suppress thermal noise, an accurate estimate of the error Δ can be found by subtracting the average ADC output when $V_{DAC}=V_{ref-}$ from the average output when $V_{DAC}=0$, and/or similarly by subtracting the average ADC output when $V_{DAC}=0$ from the average

output when $V_{DAC}=V_{ref+}$. The gain error is subsequently corrected by shifting the digital output by the amount of the missing codes as shown in Fig. 6-28.

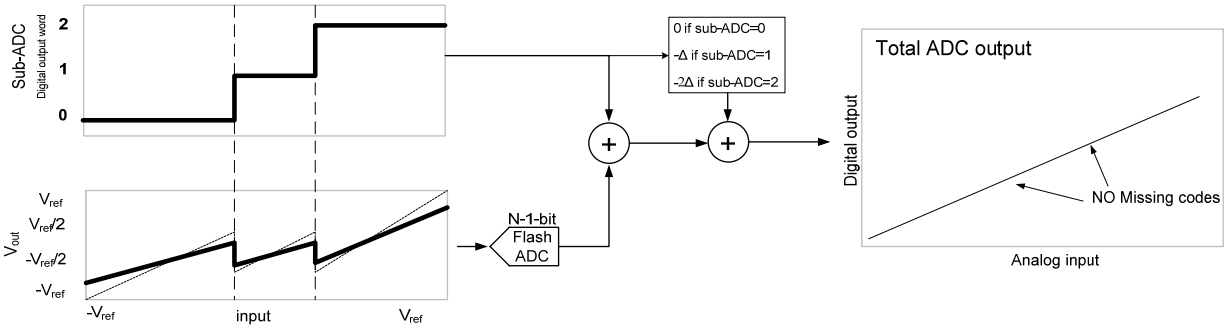

Fig. 6-28: Illustration of correction scheme

All pipeline stages are calibrated in this work using a foreground approach. Multiple stages are calibrated at startup by recursively using the method discussed this section initially on the last pipeline stage (while powering off all previous stages), then the second last, then the third last, etc., eventually calibrating the entire pipeline ADC as shown in Fig. 6-29 (e.g.: [84]).

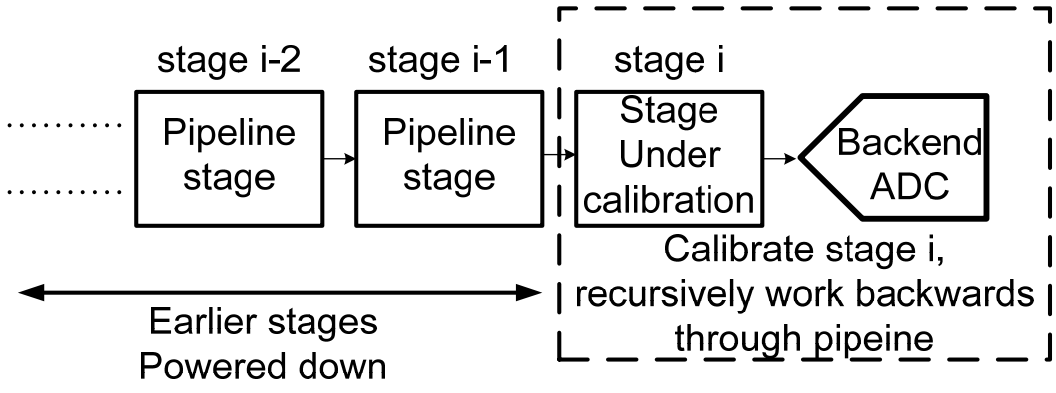

Fig. 6-29: Multistage foreground calibration

6.8: THEORETICAL POWER SAVINGS

In this section the theoretical power consumption of the proposed MDAC topology is compared against the power consumption of a traditional MDAC. The goal of the analysis is to examine under what circumstances the proposed topology offers a benefit of substantial power savings over the traditional MDAC architecture.

From section 2.8.3 it was noted that the unity gain frequency of an opamp in closed-loop is given by: $\omega_t = \frac{\omega_{ta}}{1/\beta}$, where β is the feedback factor. However in the proposed MDAC topology since the only active circuit (source follower in this work) requires only a unity gain, β=1. Hence on bandwidth considerations alone the proposed topology improves the settling time of the MDAC by feedback factor of the traditional MDAC topology. To make a fair comparison between the proposed and traditional MDAC topologies, the traditional MDAC topology should be designed with the same closed-loop gain as is achieved in the proposed MDAC, noting that since the proposed MDAC achieves a gain of less than 2x (due to parasitic capacitors as noted in section 6.4), the traditional MDAC can trade more bandwidth for less gain to match the proposed MDAC topology. Thus if the stage gain of the proposed MDAC is given by 'A', at best β=1/A for the traditional MDAC, hence considering only bandwidth the power savings of the proposed approach (η) is given by:

$$\eta = \frac{\omega_{t-proposed}}{\omega_{t-traditinalMDAC}} = \frac{\omega_{ta}}{\omega_{ta}/_{\beta}} = \beta = \frac{1}{A}. \qquad \textit{(eqn. 6-21)}$$

It was noted in section 6.6 that one of the benefits of the proposed approach was: since the active circuit comes after the gain stage, the noise of the active circuit is reduced by the gain when referred to the input. From eqn. 6-17 and eqn. 6-20, the ratio (ζ) of input referred noise between the proposed MDAC and the traditional MDAC of Fig. 5-1 is:

$$\zeta = \frac{\sigma^2_{T-proposedMDAC}}{\sigma^2_{noise-traditional}} = \frac{\sigma^2_{T\Phi 1} + \sigma^2_{T\Phi 2}/_{A^2}}{\sigma^2_{T\Phi 1} + \sigma^2_{T\Phi 2}} \qquad \textit{(eqn. 6-22)}$$

To simplify the analysis assume $\sigma^2_{T\Phi 1} = \sigma^2_{T\Phi 2}$. Thus:

$$\zeta \approx 0.5 + {}^{1}/_{2A^2} \qquad \textit{(eqn. 6-23)}$$

Hence to achieve the same input referred noise floor in both the proposed and traditional MDAC topologies the proposed MDAC can be designed with 1/ ζ smaller capacitors (assuming the same input signal swing). Since [7]:

$$\omega_{ta} = \frac{g_m}{C_{load}} = \frac{2I_D}{C_{load}V_{eff}}\ , \qquad \textit{(eqn. 6-24)}$$

it can thus be inferred that since the capacitors of proposed MDAC are reduced from the traditional MDAC by 1/ ζ, the power of the proposed MDAC can also be reduced by 1/ ζ while preserving the same speed. Thus the total factional reduction in power of the proposed MDAC versus a traditional MDAC (α) is given by:

$$\alpha = \eta\zeta, = \frac{1}{A}\left(0.5 + {}^{1}/_{2A^2}\right) \qquad \textit{(eqn. 6-25)}$$

Fig. 6-30 plots α versus the gain 'A' of the proposed pipeline stage.

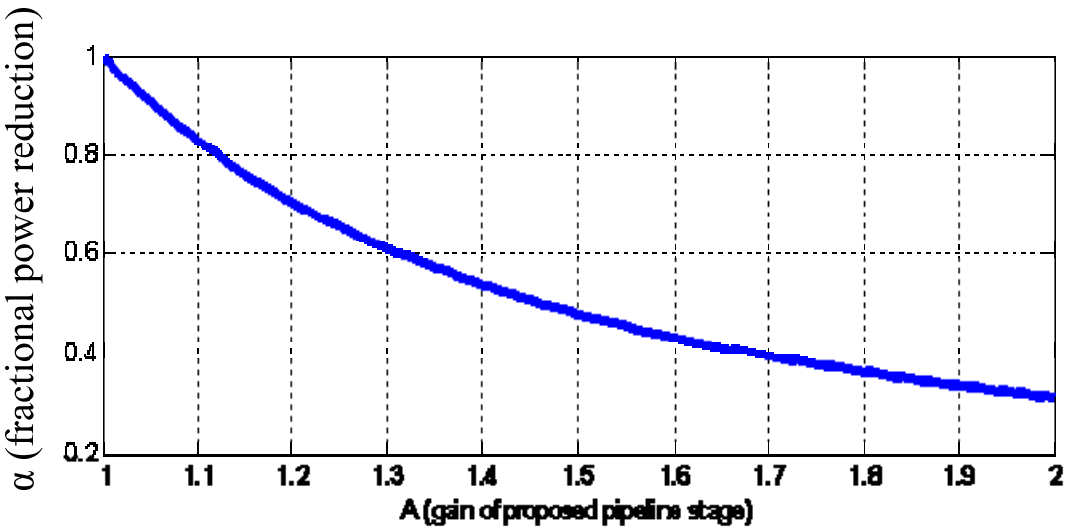

Fig. 6-30: Fractional reduction of power in proposed MDAC versus traditional MDAC

From Fig. 6-30 it is seen that the efficiency of power reduction in the proposed MDAC is maximized when the stage gain of the MDAC is maximized. It is noted that the efficiency

plot of Fig. 6-30 ignores the effect of parasitic capacitors in the opamp (which further reduce β), as well as the fact that in practice the noise of an opamp is usually larger than that of a source follower. These approximations were made as in practice the actual value of β and noise from the opamp varies from one implementation to another, thus the best case β and opamp noise were used in the analysis to show the minimum amount of power reduction possible using the proposed approach. Hence the power reduction of the proposed MDAC is even larger than that shown in Fig. 6-30.

The power consumption of the additional digital calibration circuitry required in the proposed ADC of this work has been ignored in the analysis, as it is noted that strictly speaking foreground calibration adds only a small amount of additional power since it only powers on periodically, and the power of digital adders can be made small. It is noted that in general digital background calibration of linear errors in deep sub micron processes typically only add 10-20% extra power, where the specific amount of additional power varies depending on which background calibration technique is used, and how frequently the background calibration is powered on.

6.9: DESIGN SPECIFICATIONS

From section 6.8 it was noted that the proposed MDAC topology showed the greatest improvement in power consumption over the typical MDAC topology when the stage gain was maximized. To maximize operating speed, all switches need to be sized large to minimize the settling time, and the unity gain buffer needs to be large to maximize bandwidth. However, from eqn. 6-13 it is seen that the larger the parasitic capacitors (thus larger the switches, and larger the unity gain amplifier) the smaller the stage gain. Hence there is a tradeoff in the proposed topology of higher speed with lower power efficiency. The design specifications of a 10-bit 50MS/s ADC were ultimately selected as Spice simulations showed at 50MS/s stage gains on the order of ~1.75x could be achieved using the proposed MDAC topology. Furthermore as this dissertation already deals with the design of a power

efficient opamp based 10-bit 50MS/s ADC in Chapter five, a fair comparison can be made between the proposed topology of this chapter and the ADC discussed in chapter five which uses a traditional MDAC approach for many of the pipeline stages (except the first stage). Although a specific application was not targeted with this work as it was a proof-of-concept prototype, a quick survey of 10-bit ADC with sampling rates on the order of 50MS/s show a variety of potential applications from medical imaging to digital communications.

6.10: CIRCUIT DESIGN

In this section the circuits used in the design of a prototype of a 10-bit 50MS/s ADC in 1.8V 0.18µm CMOS which used the power efficient MDAC topology proposed in this chapter are described.

6.10.2: ADC TOP LEVEL TOPOLOGY

Fig. 6-31 illustrates the top level topology of the ADC designed using the power reduction techniques discussed in this chapter.

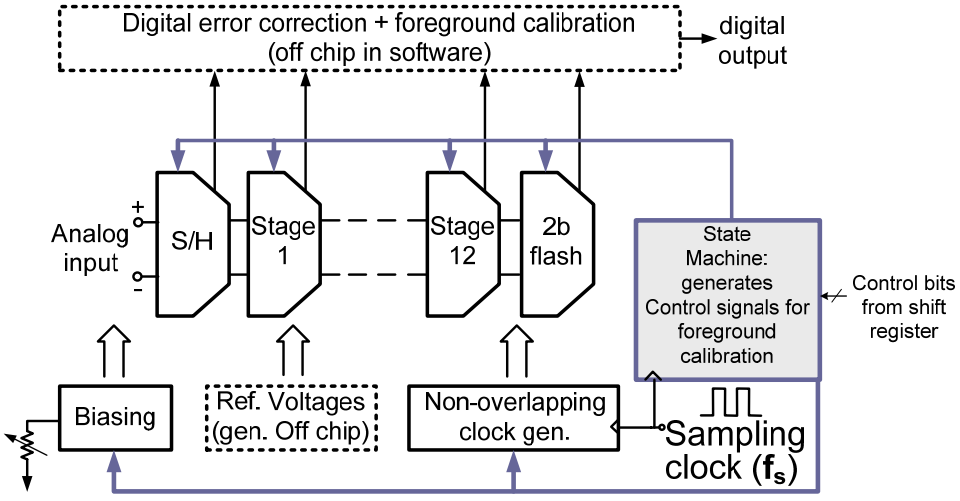

Fig. 6-31: Top-level topology of ADC used in the work of this chapter

Simulation results showed each pipeline stage to have a stage gain of ~1.75x i.e. $\log_2(1.75)=0.8$ true bits resolved per stage (where 'true bits' denotes the number of bits resolved which reduce the quantization noise floor, i.e. ignoring the redundant bits). Thus with 12 total stages followed by a 2-bit Flash ADC, approximately 11.6 true bits are resolved ignoring the thermal noise. As ADC power is dominated by thermal noise considerations, the thermal noise floor was designed to be at the ~10-bit level at the input of the ADC. To minimize power, the first three pipeline stages were scaled approximately by their respective stage gains [13].

6.10.3: FRONT-END SAMPLE-AND-HOLD

To simplify the design of the ADC and focus all design efforts on the novel power reduction techniques, no attempt was made to eliminate the front end S/H. Further power reduction by elimination of the front end S/H is deferred as future research.

The front-end S/H was also realized in this work without opamp-based capacitive feedback as shown in Fig. 6-32, where rather than using a passive voltage gain technique as used in the MDAC, the sampled input is simply buffered by a source follower [85].

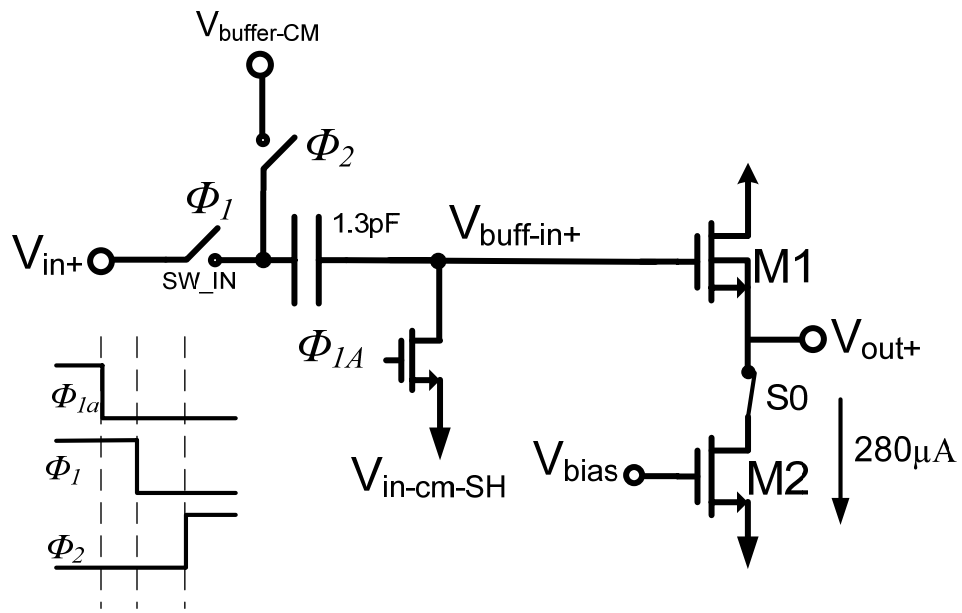

Fig. 6-32: Front-end sample-and-hold using unity gain buffer (shown single-ended, implemented pseudo-differentially in practice)

Switch S0 is included so that during Φ_1 the source follower can be powered off, and thus save additional power. $V_{in\text{-}cm\text{-}SH}$ was set as 0.4V and $V_{buffer\text{-}CM}$ was set as 1.4V, where each voltage was generated off-chip.

Although not done in this work it is possible that a passive gain of two could be used in the S/H as shown in Fig. 6-33 so as to relax the signal swing requirements and thus linearity requirements of the input sampling switches S_IN1, S_IN2.

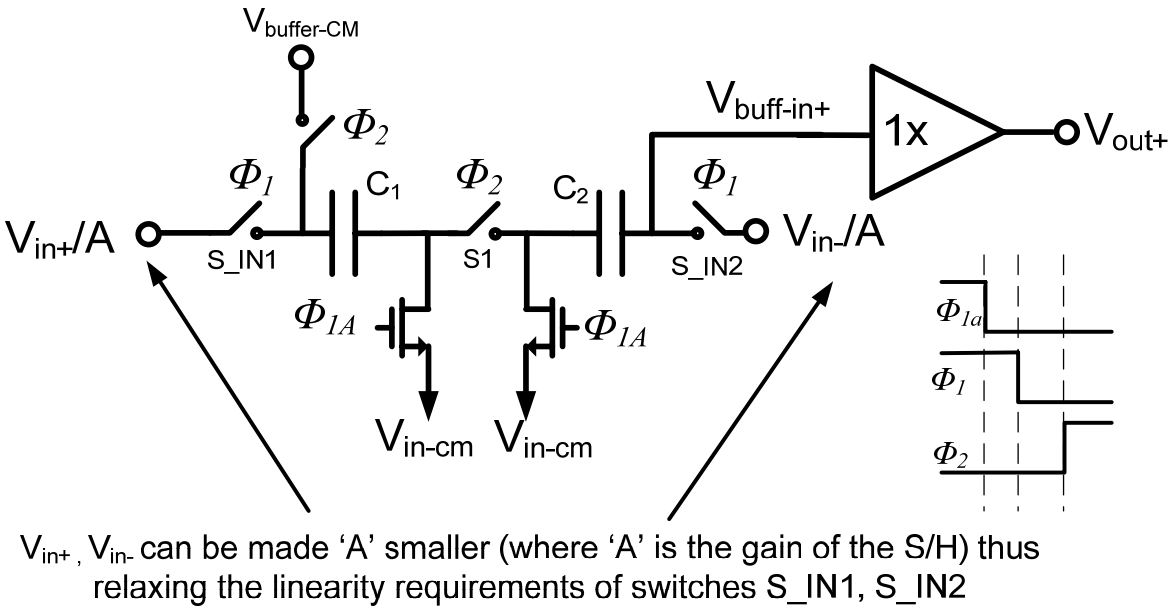

Fig. 6-33: Modified S/H which has a gain of 'A'

6.10.4: MDAC AND UNITY GAIN AMPLIFIER

Fig. 6-34 shows the full circuit topology of the first stage MDAC used in this work. Subsequent pipeline stages were identical where the first three pipeline stages were scaled by the gain of the pipeline stage. An NMOS source follower using a DNW layer was used to implement the unity gain amplifier. To further reduce power consumption a switch S0 was added to power off the unity gain amplifier during Φ_1 as it is not required to operate during Φ_1.

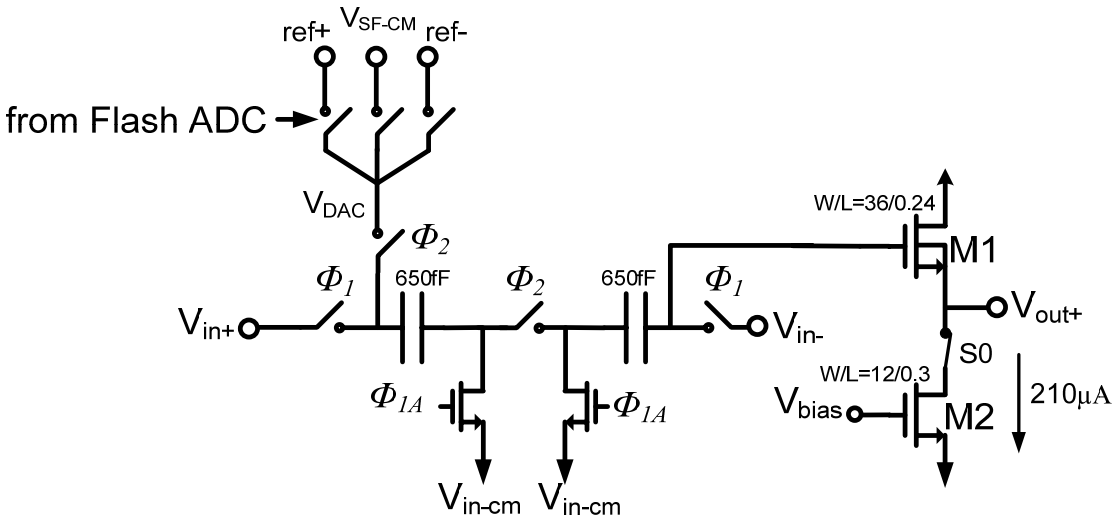

Fig. 6-34: First stage MDAC circuit

As noted in section 6.5.2, to maximize linearity and minimize noise, the length of M2 was 0.3μm. To maintain linearity beyond the 10-bit level, the output signal swing of the source follower was designed to be 0.5V p-p single-ended (i.e. 1V p-p when considering the differential output). It is noted that the ADC proposed in this chapter has a signal swing smaller than the differential 1.6V p-p of the ADC discussed in chapter five. This is because simulations show a signal swing larger than 1V p-p results in increased nonlinearity from the buffer such that more than 10-bits linearity cannot be achieved. Thus a clear tradeoff in using a simple source follower buffer is reduced signal swing, thus increased capacitor sizes in the MDAC, hence increased power consumption. However as will be seen in section 6.13, even with a ~40% smaller signal swing, the proposed topology is so power efficient that a significant power reduction still occurs when comparing the work of this chapter versus that of chapter five.

The sizes of all the switches were optimized based on simulation results, where all the switches were large enough to achieve the desired sampling rate of 50MS/s, but small enough to minimize distortion and parasitic capacitances on critical nodes.

The values of the reference voltages for the MDAC, which were generated off-chip, were: V_{ref+}=1.55V, V_{ref-}=1.05V, $V_{cm-buff}$=1.3V, V_{in-cm}=0.5V.

6.10.5: SUB-ADC

The sub-ADC was designed using the same dynamic comparators used in the ADC of chapter five. Dynamic comparators have the advantage of low power consumption, but at the cost of increased offset. This however is a favorable tradeoff as the 1.5-bit topology has a large amount of redundancy to trade with comparator offset [12]. Fig. 6-35 illustrates the comparators used in this work.

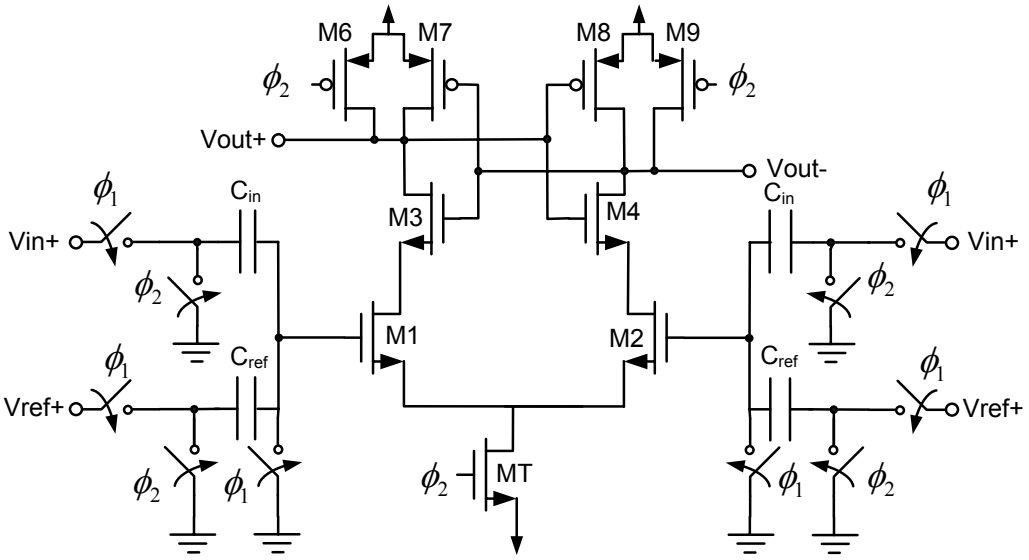

Fig. 6-35: Dynamic comparator used in Flash sub-ADC

As the sub-ADC connects to the output of a source follower which has a common-mode voltage near V_{SS}, the comparators required different reference voltages than those supplied to the DACs of each pipeline stage. Since the pipeline topology can tolerate offsets from the comparator, different values of the reference voltages can be provided to the sub-ADC provided they do not incur an effective offset in each comparator larger than $V_{ref}/4$.

Furthermore by separating the reference voltages in the comparators from the DAC, the amount of switching noise on the reference voltages is reduced. The comparator's reference voltages which were generated off-chip were: V_{ref+}=0.75V and V_{ref-}=0.25V

6.10.6: DIGITAL STATE MACHINE

A digital state machine was used to generate the control signals for each pipeline stage during foreground calibration. The state machine was only powered on during foreground calibration and powered completely off subsequently. The state machine was ultimately programmed so that complete calibration was completed within 10^4 clock cycles.

6.10.7: ANALOG TEST-MUX

To enhance the testability of the ADC each pipeline stage was equipped with three analog test muxes: two to enable viewing of the differential input to each pipeline stage, and a third test mux to observe the bias voltage of the source follower. The goal of the test mux was to be able to verify basic functionality without adding a significant additional capacitive load. As such, each analog mux was implemented as a transmission gate which was sized approximately the same as the transmission gates for the sampling switches. Thus basic functionality of the ADC can be verified by running the ADC at low sampling rates, avoiding otherwise non-trivial wafer probing.

The three analog test-muxes in each pipeline stage were all connected in parallel to three I/O pins in the chip as shown in Fig. 6-36. Control bits loaded via a shift register were used to select which pipeline stage connected to the I/O pins.

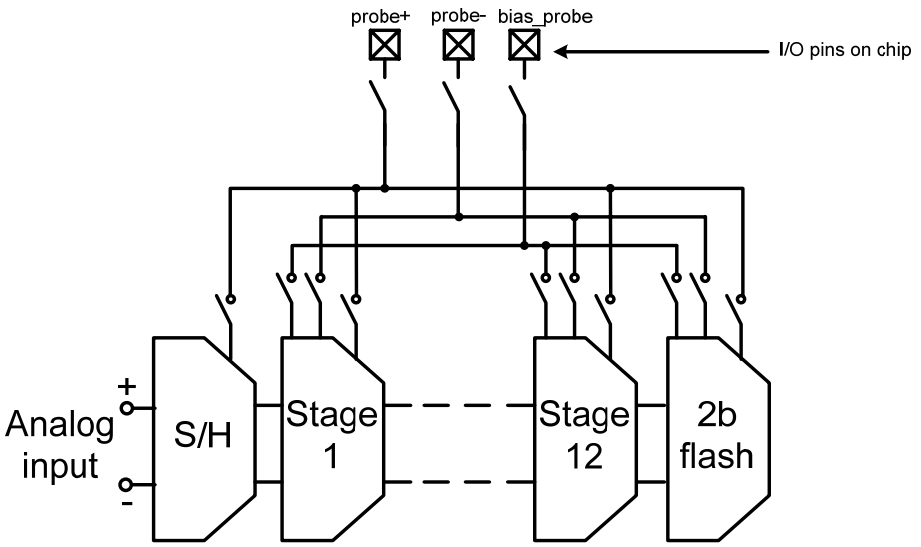

Fig. 6-36: Analog test mux configuration

6.11: SIMULATIONS

The proposed topology was simulated in Spice using TSMC models for a 1.8V 0.18µm CMOS process. The design was simulated at all process corners and at temperatures of 0, 80, and 120 °C. The design was also verified with Monte Carlo analysis and at supply voltages as low as 1.7V. Foreground calibration and digital error correction was implemented by taking the digital outputs generated in Spice and passing them via a text file to a model of the calibration and error correction in Matlab. Fig. 6-37 shows an FFT of the ADC's output after calibration; the total power consumption of the 50MS/s ADC was ~9mW.

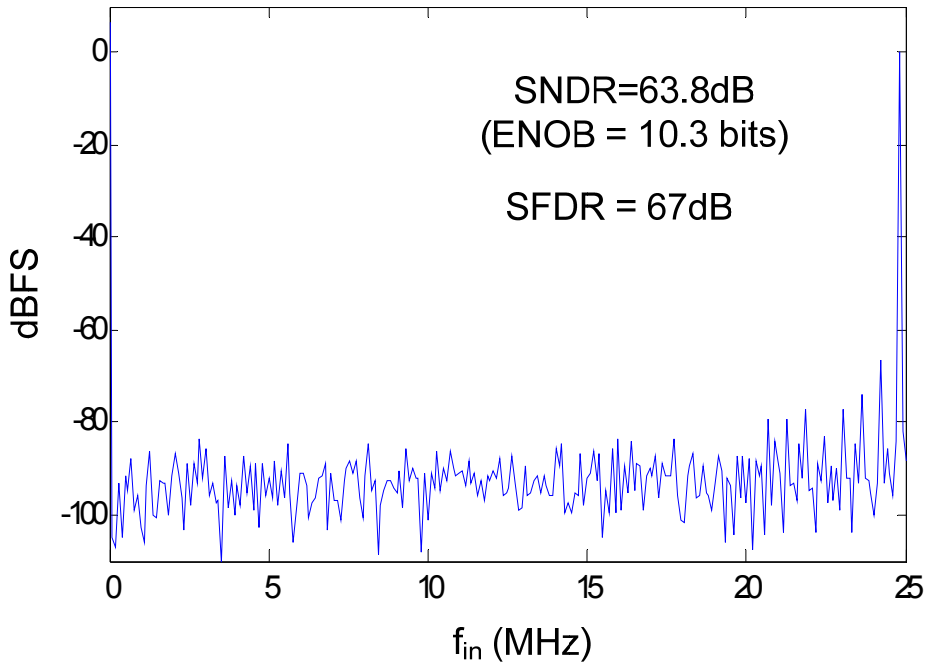

Fig. 6-37: FFT of digital output from Spice simulation – simulation excludes thermal noise

It is noted that the FFT in Fig. 6-37 excludes thermal noise - the SNDR is only indicative of quantization noise floor. Including the thermal noise, which is approximately at the 10-bit level, the expected peak resolution of the ADC is ~9.5 bits.

6.12: TESTING

Sections 6.12.2 and 6.12.3 discuss the test setup and measured results of the prototype fabricated in a 1.8V 0.18µm CMOS process.

6.12.2: PCB

A 4 layer FR4 dielectric PCB board with a minimum 6mil trace was designed and constructed for the device under test as shown in Fig. 6-38. Separate Power planes were used

to isolate the analog, digital, I/O, and board power supplies. A differential input was generated using a 1:1 turns ratio Minicircuits transformer matched to 50Ω. Reference voltages were generated by passing the output of a resistive voltage divider through an opamp (LM7301) in a unity gain buffer configuration. To maintain constant supply voltages, all voltage supplies for each power plane were generated through regulators (LM337, LM1117), and heavily decoupled with capacitors. As the ADC utilized a constant current biasing scheme, an off-chip adjustable resistor was used as the master current source.

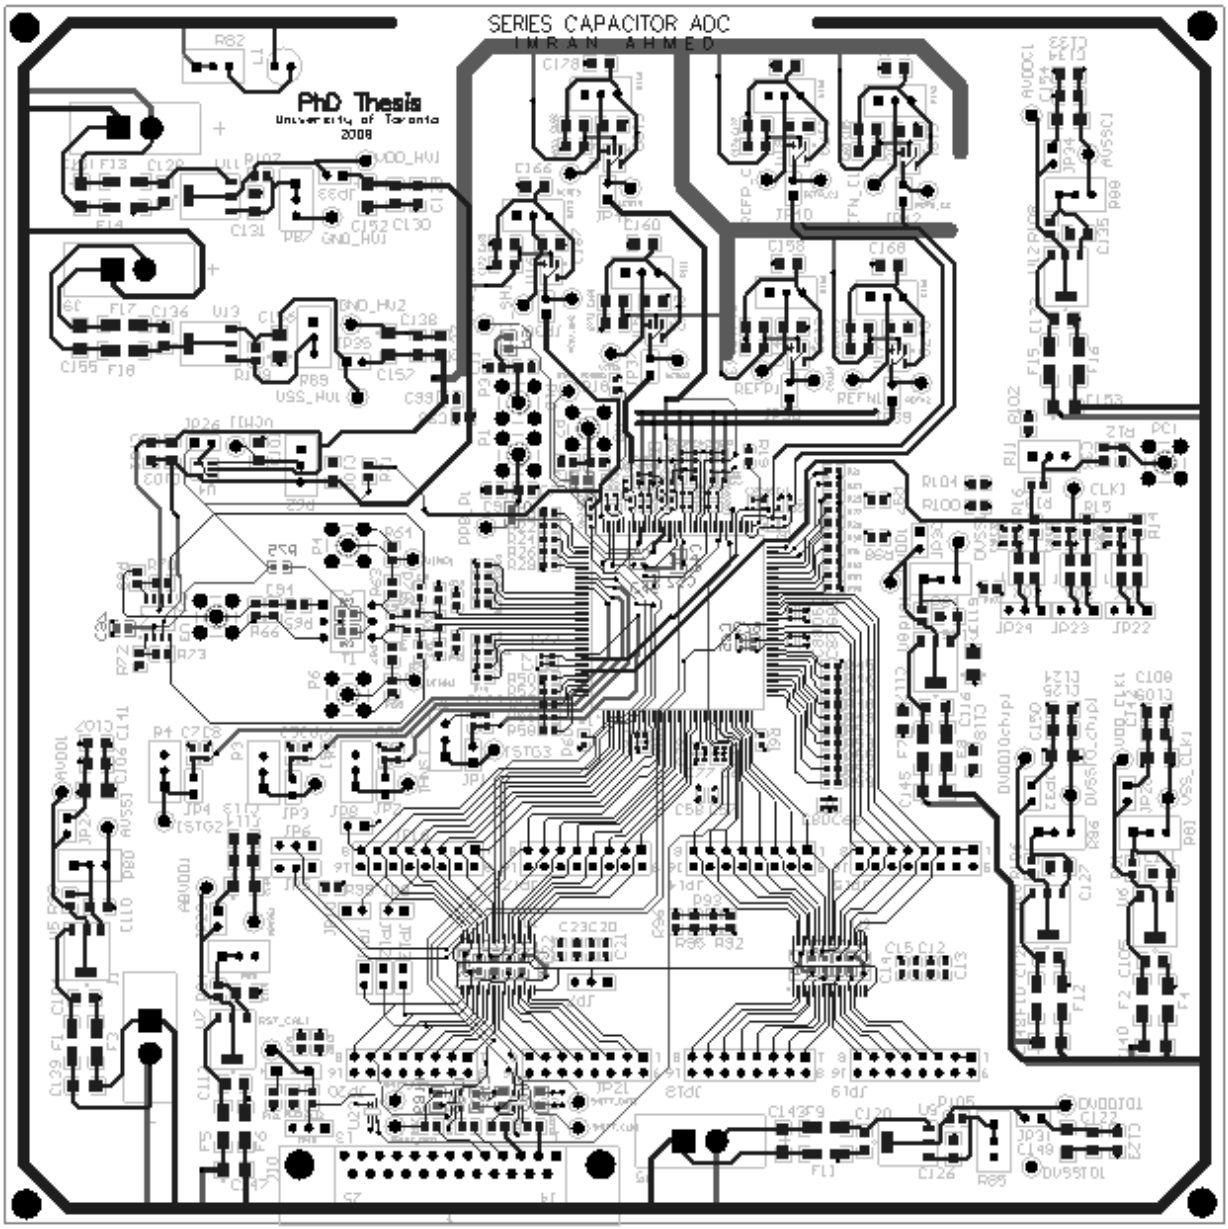

Fig. 6-38: PCB used to in test setup for ADC described in this chapter

6.12.3: TEST SETUP

A test setup as shown in Fig. 6-39 was used. Sinusoidal inputs were generated using a HP 8664A function generator. Several Minicircuits filters were used to minimize harmonic distortion from the function generator such that the sinusoidal input to the ADC had an SNDR of well over 62dB for input frequencies larger than 21MHz. A 50 MHz crystal

oscillator (ECS3518-XO) on the PCB was used to generate the clock for the ADC. The serial shift register was loaded via a parallel port connection to a PC, where a Matlab script was executed to load the appropriate bits. The output bits of each pipeline stage were captured using a Tektronix TLA714 logic analyzer, capable of capturing 65,536 points at a time. An Agilent E3620A Dual output DC power supply was used to provide positive and negative voltages to the voltage regulators on the PCB. The 10-bit output word from the 10-bit ADC was determined via a Matlab script written to emulate the operation of a digital error correction circuit and foreground calibration.

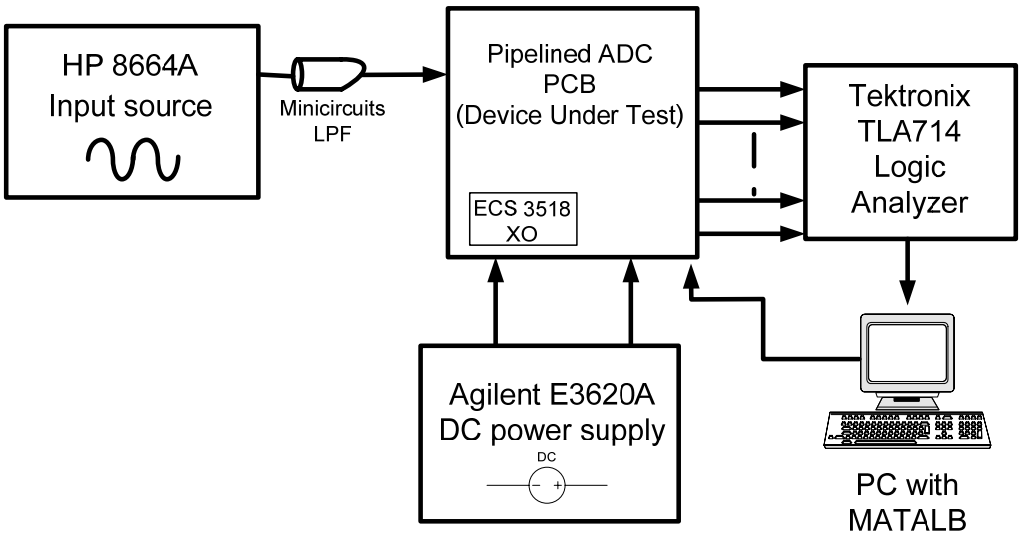

Fig. 6-39: Test setup of ADC

6.13: MEASURED RESULTS

A prototype of the proposed ADC of this chapter was fabricated in a 1.8V 0.18µm CMOS process as shown in Fig. 6-40; the core area was 2.0mm x 0.7mm (1.4 mm^2). The fabricated IC was packaged in a 120-pin CQFP package, where 78 of the total 120 pins were used. A large number of pins were used to allow many pins to be used for DC power supplies and reference voltages, so as to reduce the impact ringing due to the parasitic inductance of bond wires.

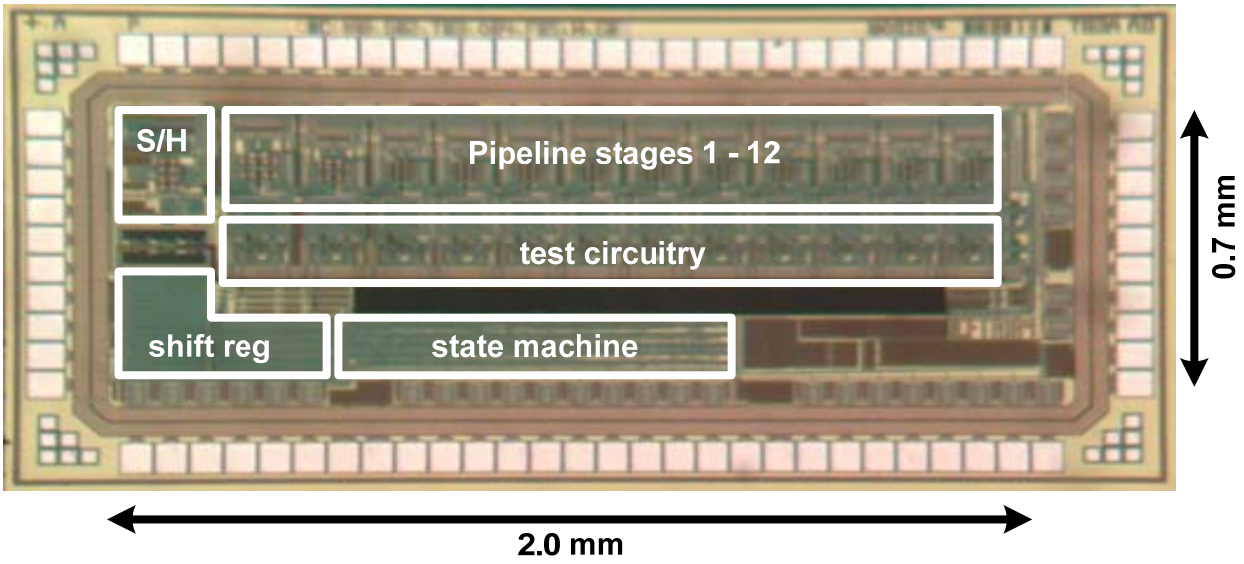

Fig. 6-40: Micrograph of low powered pipeline ADC

As noted in Fig. 6-40 a significant amount of area in the chip was dedicated to test circuitry such as digital muxes which enabled each digital signal in each pipeline stage to be configured in different modes for testability. In a practical implementation much of the test circuitry can be removed hence allowing for a reduction in area consumption.

6.13.2: MEASURED ADC SNDR VARIATION

Fig. 6-41 shows the variation of the SNDR and SFDR of the ADC at 50MS/s for input frequencies between 2.4MHz and 25MHz, where it is seen that the ADC is capable of achieving an SNDR/SFDR as high as 58.2/66 dB.

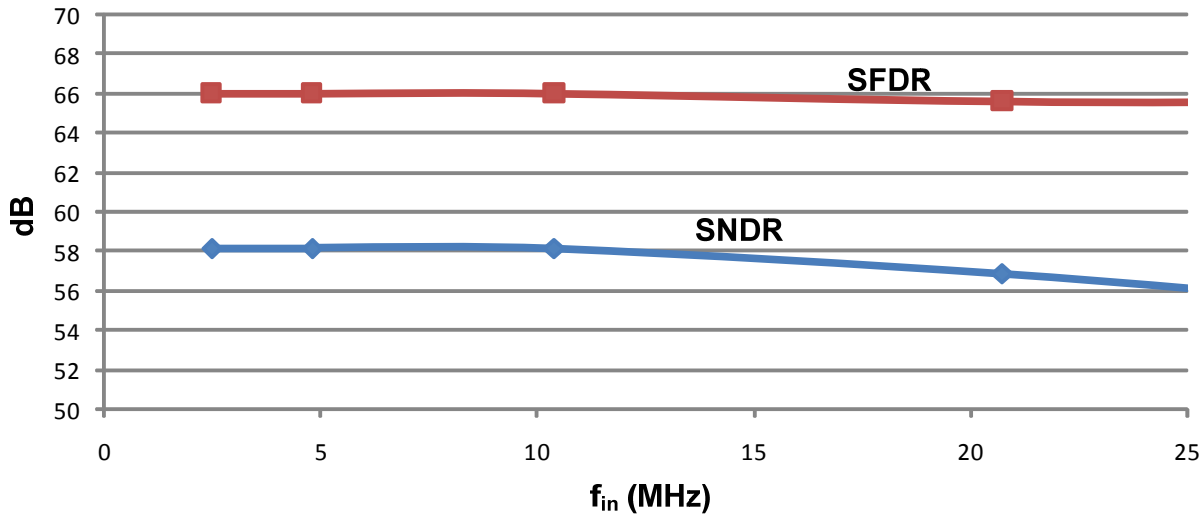

Fig. 6-41: SNDR/SFDR variation with input frequency, f_s=50MS/s

Fig. 6-42 shows the variation of ENOB with input frequency where it is seen the ADC can achieve as high as 9.4-bits.

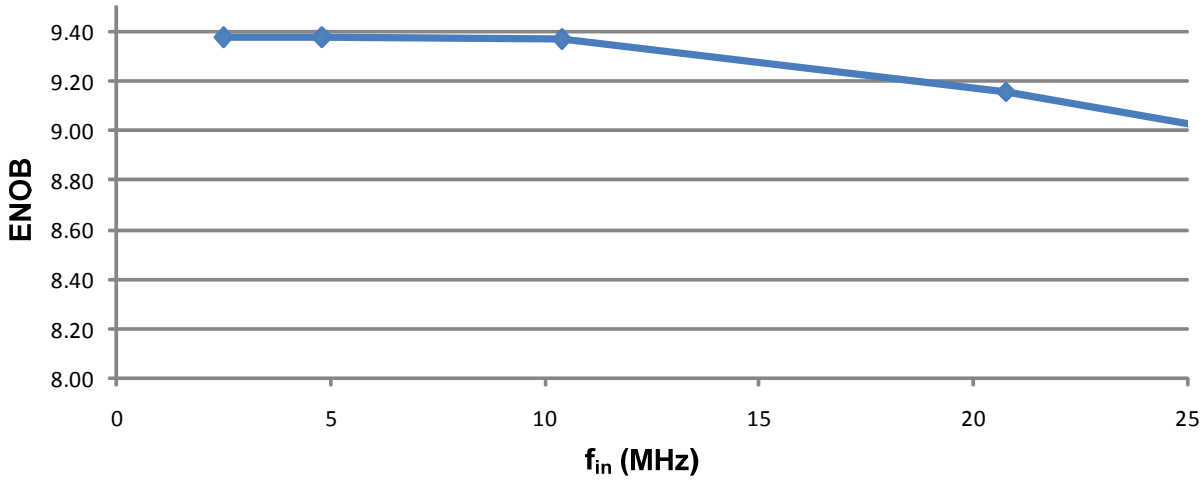

Fig. 6-42: ENOB variation with input frequency, f_s=50MS/s

The ADC consumed only 3.9mW for the active amplifiers and 6mW for all the clocking circuits, resulting in a total power consumption of only 9.9mW. Although the power of the reference voltages is not included, it is noted that the total average current sourced by the reference voltages was only 0.34mA. The ADC discussed in this chapter consumed less than half the power of the ADC discussed in chapter five (9.9mW versus 27mW), even though the ADC of this chapter had a smaller signal swing (1.0V p-p versus 1.6V p-p) and included a front-end sample-and-hold. It is conceivable if a larger signal swing was used and the front end sample-and-hold removed, the ADC discussed in this chapter could have a further reduction in power. It is noted that the clocking power was large due to generous clock buffer sizing, and long clock lines (due to the additional area taken by the test circuitry) which could be made much shorter in a practical implementation. It is also noted that with faster technologies (e.g. 65nm, 45nm etc.) the clocking power would be significantly lower due to a lower supply voltage and smaller parasitic capacitors.

Fig. 6-43 compares the power of the ADC of this work versus other recently published 10-bit ADCs where it is seen that amongst 10-bit pipeline ADCs the proposed architecture of this chapter has the lowest power consumption at 50MS/s.

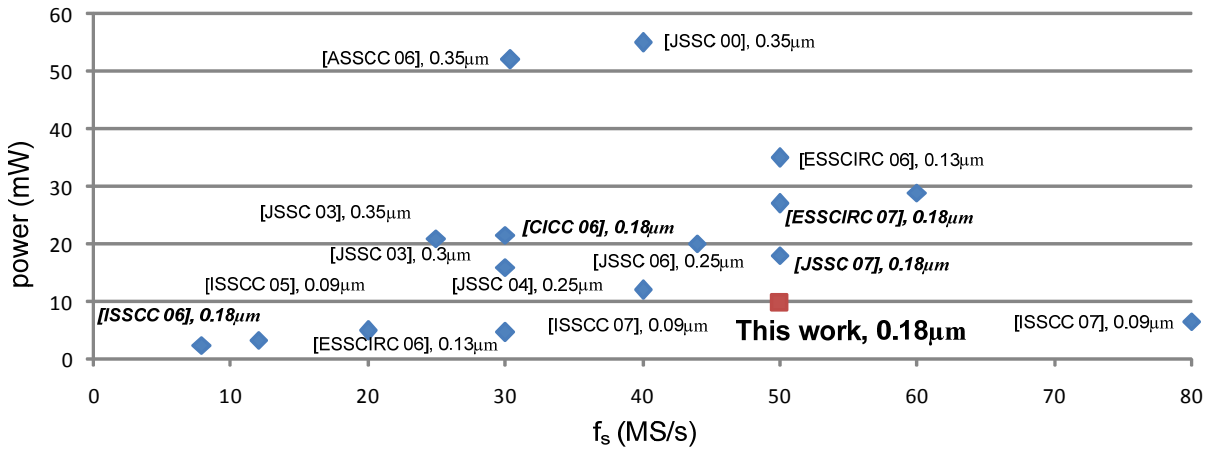

Fig. 6-43: Comparison of power of ADC of this work versus other 10-bit ADCs

Fig. 6-44 compares the figure of merit (eqn. 2-1) for 10-bit pipelined ADCs where it is seen that the ADC of this work has amongst the best published figure of merits for 10-bit ADCs in the 10-80MS/s range. The few publications which have a slightly better figure of merit have the benefit of a faster technology (0.18µm publications are shown in bold italics). Furthermore it is noted that amongst 0.18µm ADCs, this work achieves the best figure-of-merit.

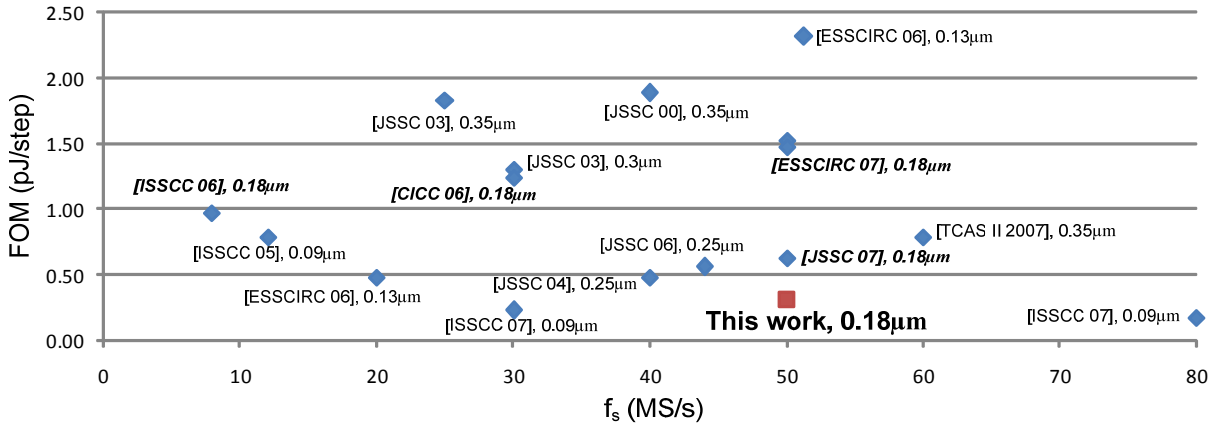

Fig. 6-44: Comparison of FOM of ADC of this work versus other 10-bit ADCs

Fig. 6-43 and Fig. 6-44 clearly display the significant power savings afforded by the proposed ADC topology.

6.13.3: ADC FFTs

Fig. 6-45 to Fig. 6-48 show FFTs of the ADC output for input frequencies of 2.4MHz and 20.7MHz before and after foreground calibration for f_s=50MS/s.

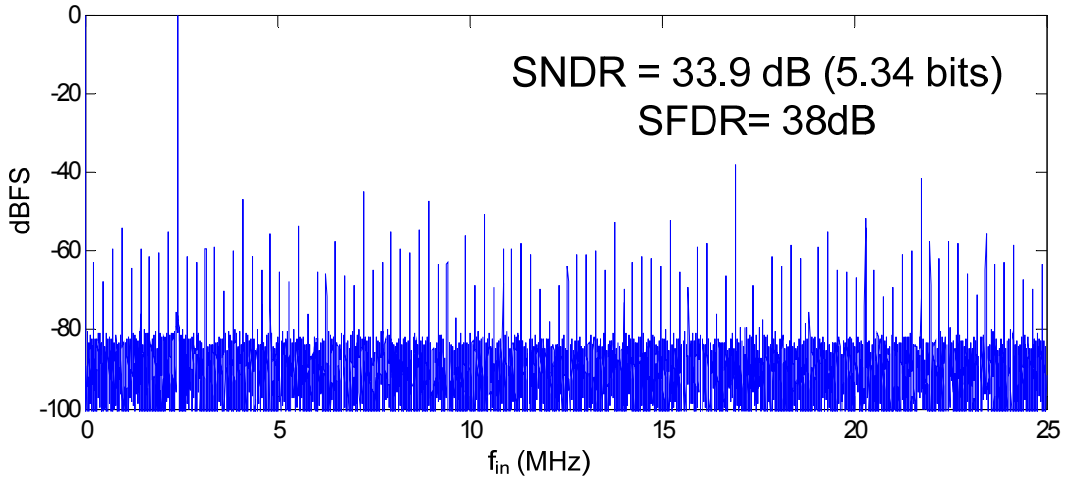

Fig. 6-45: FFT of ADC output before calibration, f_{in}=2.4MHz, f_s=50MS/s

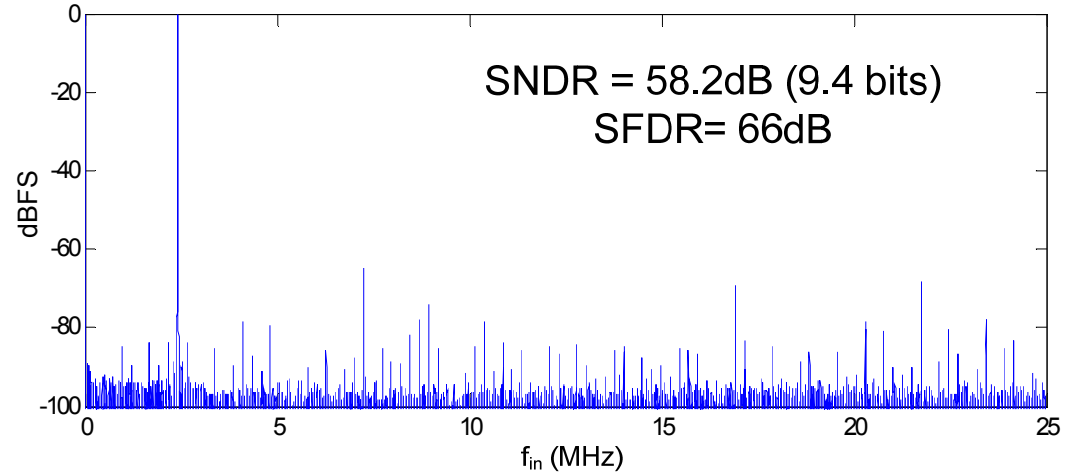

Fig. 6-46: FFT of ADC output after calibration, f_{in}=2.4MHz, f_s=50MS/s

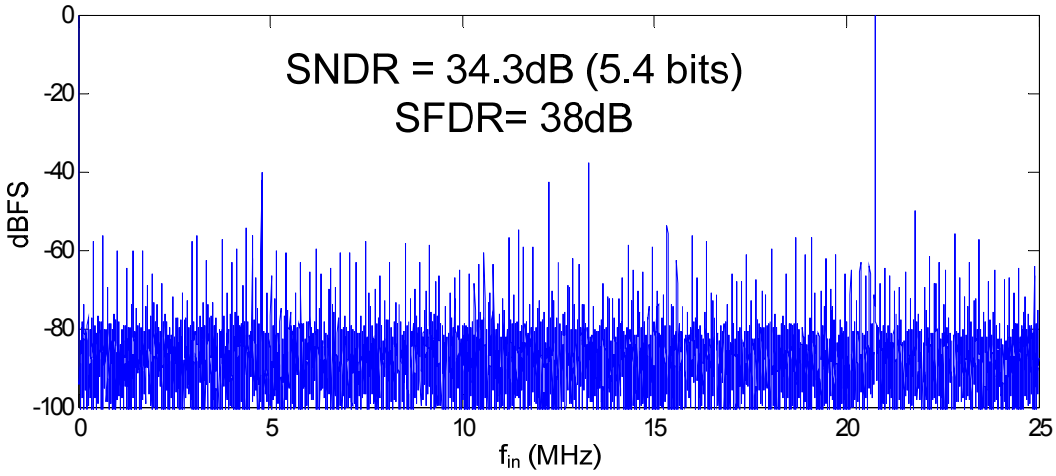

Fig. 6-47: FFT of ADC output before calibration, f_{in}=20.7MHz, f_s=50MS/s

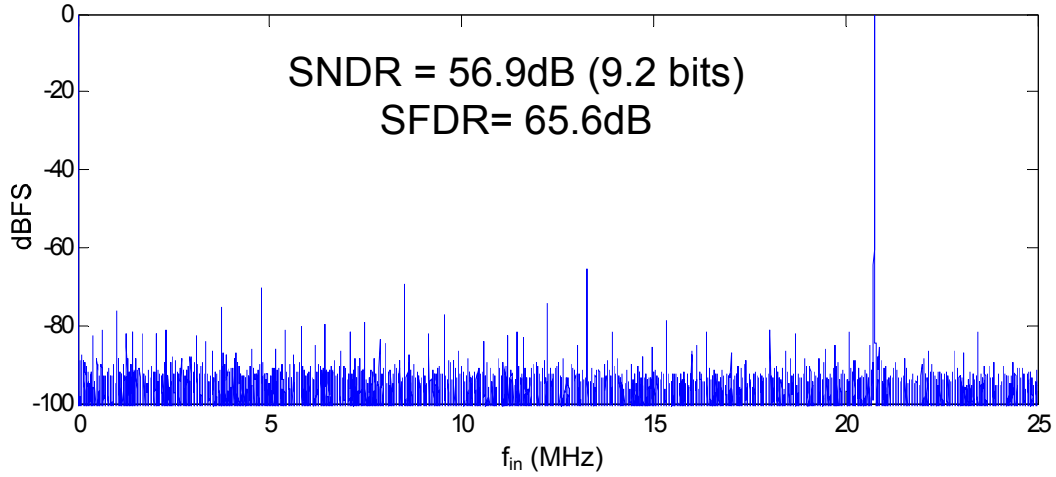

Fig. 6-48: FFT of ADC output after calibration, f_{in}=20.7MHz, f_s=50MS/s

From the FFTs it is clear that calibration is clearly heavily leveraged to achieve a significant improvement in ADC performance – almost 4-bits.

It is noted that when the bias currents of the ADC were varied by more than +/-20%, the ADC SNDR varied by less than 0.1 bits when using correction terms extracted when the ADC was biased with nominal bias currents. These results indicate that the gain of each pipeline stage was set primarily by capacitor ratios – which do not change over time. Hence in a practical implementation, the interval between foreground calibrations could be very large, allowing for a minimal impact on normal ADC operation. Drift of correction terms is also negligible, as measured results from the ADC separated by one week show insignificant (less than 0.01 bits) difference in ADC ENOB when the same correction terms are used for both measurements.

From the FFT plots after calibration, it can be seen that the even order distortion terms are heavily attenuated, verifying that the proposed sampling topology shown in Fig. 6-5 achieves fully-differential functionality. The degradation of ADC accuracy for higher sampling rates was attributed to distortion from the input sampling switch (switch SW_IN in Fig. 6-32), and increased digital noise from the I/O buffers. The distortion due to the sampling switch could be reduced with only a small increase in power by either using larger transistors in the input switch and/or using bootstrapping techniques [64].

6.13.4: INL/DNL PLOTS

Fig. 6-49 to Fig. 6-52 show INL/DNL of the ADC before and after calibration, where it seen that calibration significantly improves the INL of the ADC from +15.7/-17.9 LSB to +0.7/-0.8 LSB and DNL from +1.6/-1 LSB to +0.35/-0.35 LSB.

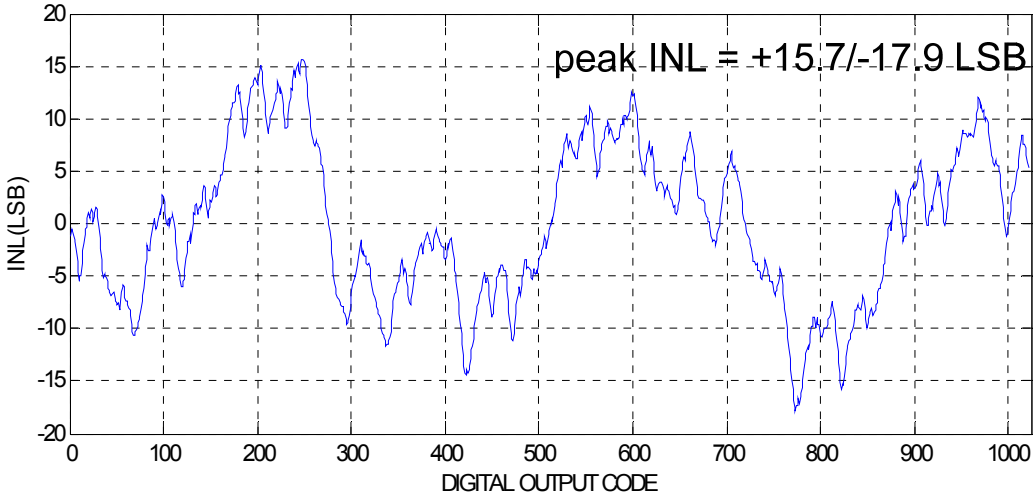

Fig. 6-49: INL before calibration (LSB @ 10-bit level)

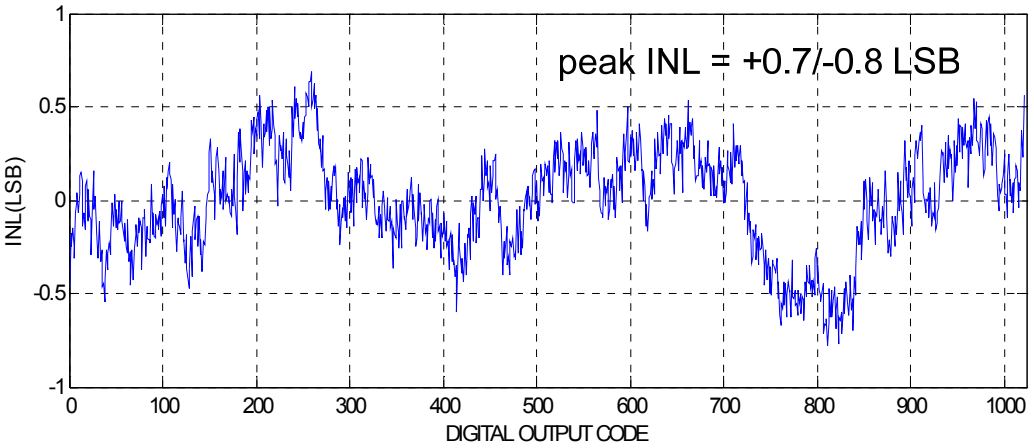

Fig. 6-50: INL after calibration (LSB @ 10-bit level)

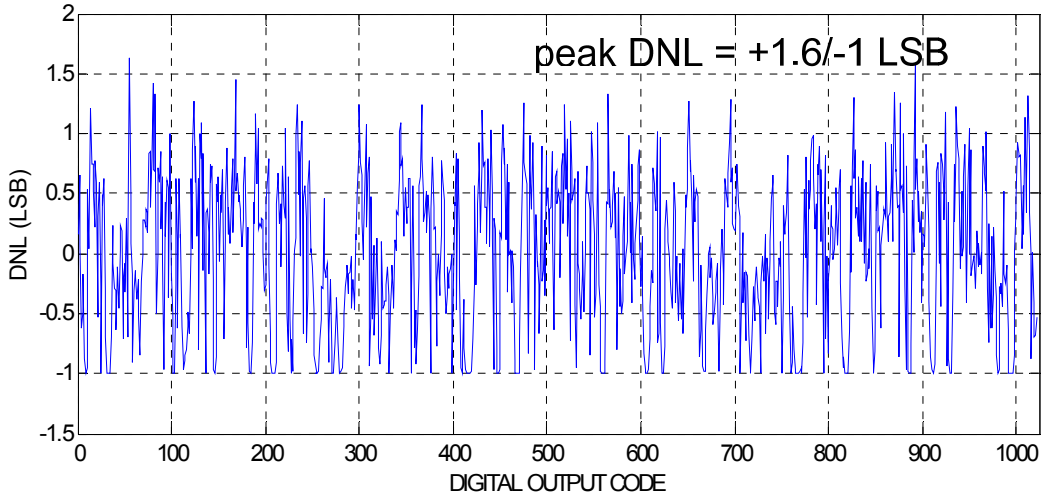

Fig. 6-51: DNL before calibration

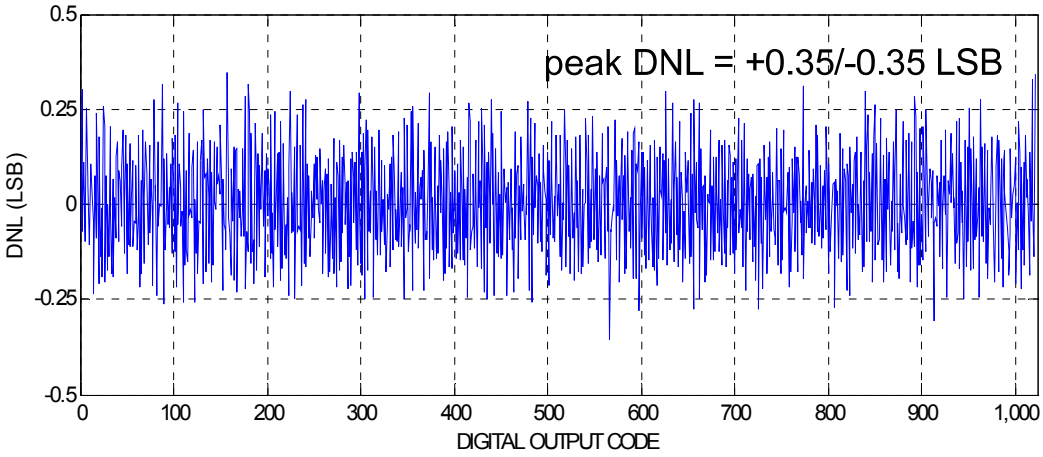

Fig. 6-52: DNL after calibration

6.14: SUMMARY

In this chapter a technique to significantly reduce MDAC power was discussed. A fully-differential charge sharing technique in combination with a simple source follower as a unity gain buffer and foreground calibration was shown to enable lower power consumption over

prior opamp based MDAC topologies. A summary of key measurement results is presented in Table 6-1.

Table 6-1: Summary of measured results

Technology	1.8V, 0.18µm CMOS
Input signal swing	1.0V p-p
Area	1.4mm^2
Sampling rate	50MS/s
SNDR	58.2 dB (peak)
SFDR	66dB (peak)
ENOB	9.4 (peak)
Power	9.9 mW
FOM	0.3 pJ/step

CHAPTER SEVEN: CONCLUSIONS

7.1: SUMMARY

In this dissertation several techniques to improve pipelined ADC performance were detailed. The first third of the work dealt with a novel technique to enable rapid background calibration of both gain and DAC errors in a multi-bit pipeline stage using a dual-ADC based approach. Measured results from a 1.8V, 0.18µm CMOS prototype showed a more than 100x faster calibration than previous published reports, where a peak SNDR/SFDR of 60/70dB was achieved in only 10^4 clock cycles, making the proposed technique very useful in an industrial setting. The last two thirds of the dissertation discussed techniques in which the power of a pipelined ADC could be reduced. The second third presented a novel technique where the front end sample-and-hold of a pipelined ADC for sub-sampling applications was eliminated, thereby enabling a large reduction in power consumption. Unlike prior published reports the technique did not require a carefully matched (and thus costly) layout to achieve functionality for very high input frequencies. Measured results from a prototype fabricated in a 1.8V, 0.18µm CMOS prototype showed better than 51dB SNDR for input frequencies larger than 267MHz. The last third of the dissertation presented a novel MDAC topology which heavily exploits charge sharing and calibration to achieve a very low power pipelined ADC design without the use of opamps. Measured results of a prototype in 1.8V, 0.18µm CMOS show the 50MS/s ADC can achieve an SNDR/SFDR has high as 58.2/66 dB while only consuming 9.9mW, yielding a figure of merit which is amongst the best 10-bit medium-high speed ADCs published.

Table 7-1 to Table 7-3 summarizes the performance of the results presented in chapters four, five, and six.

Table 7-1: Summary of measurement results from chapter four

Technology		1.8V, 0.18µm CMOS
Sampling rate (f_s)		45MS/s
Input signal swing		1.3V p-p
Area		3.57mm^2
Power		81mW
# of calibration cycles		10^4 cycles (0.22ms)
	BEFORE CALIBRATION	**AFTER CALIBRATION**
SNDR	46.9dB	60dB
SFDR	48.9dB	70dB
INL	+6.4/-6.1 LSB	+1.1/-1 LSB
DNL	+1.1/-0.4 LSB	+0.45/-0.4 LSB

Table 7-2: Summary of measurement results from chapter five

Technology	1.8V, 0.18µm CMOS
Area	1.21mm^2
Sampling rate (f_s)	<164 kS/s – 50MS/s
Input signal swing	1.6V p-p
Power	<0.6mW – 27mW
SNDR	>51.5dB for all f_s
SFDR	>60.5dB for all f_s
Input frequency range	0-267MHz
Power of [45] @ 50MS/s	35mW
Power of this work at 50MS/s	27mW

Table 7-3: Summary of measurement results from chapter six

Technology	1.8V, 0.18µm CMOS
Input signal swing	1.0V p-p
Area	1.4mm^2
Sampling rate (f_s)	50MS/s
SNDR	58.2 dB (peak)
SFDR	66dB (peak)
ENOB	9.4 (peak)
Power	9.9 mW
FOM	0.3 pJ/step

7.2: FUTURE RESEARCH

1. **Multistage calibration:** The calibration techniques discussed in chapter four only addressed the calibration of the first pipeline stage. Future research could investigate if it is feasible for the backend pipeline stages to also be rapidly calibrated using the techniques proposed in chapter four and/or a combination of other calibration techniques

2. **Digital calibration circuit:** The ADC presented in chapter four relied on a Matlab script to emulate the functionality of the digital calibration. To complete the study in the rapid calibration ADC an investigation of the detailed complexity and power consumption of the digital calibration circuit in different technologies could be done.

3. **Embedded S/H technique in multi-bit pipeline stages:** The proposed technique of eliminating the S/H in chapter five dealt with a pipelined ADC where the first stage was only a 1.5-bit stage. The proposed technique can also work with multi-bit pipeline stages – since multi-bit stages have an even tighter requirement of the maximum clock skew due to layout tolerable using prior S/H removal techniques, an implementation of the proposed technique with a multi-bit stage would be attractive.

4. **Embedded S/H technique combined with PGA**: In some applications pipelined ADCs are preceded by a PGA. In the case that the PGA is a switched capacitor circuit, there is no need for a front-end S/H, however the PGA becomes the dominant consumer of power. Future investigations could look into the possibility of using the techniques developed in chapter five towards eliminating a front-end PGA, yet somehow maintain gain programmability.

5. **Background calibration with ADC of chapter 6:** To simplify the ADC discussed in chapter 6, foreground calibration was used, however as noted in section 3.2.6, background calibration is generally more attractive. A future investigation could look at the various different calibration techniques and see which are best suited to the ADC. Alternatively a future investigation could look at the prospect of developing a new

background calibration scheme that exploits the fact that opamps are not used in the ADC topology.

6. **Implementing topologies in newer technologies:** The three prototypes discussed in this thesis all were implemented in 1.8V, 0.18μm CMOS. As the three contributions of this work are primarily at the architectural level, they could also be applied to deeper sub-micron technologies (e.g. 65nm, 45nm). It is expected that in newer technologies a potentially lower power consumption could be realized.

7. **Combining the ideas developed in chapters 4, 5, and 6:** The ultimate synthesis of this dissertation would be to develop a low powered pipelined ADC as discussed in chapter 6, where the front-end S/H is removed using the embedded S/H technique developed in chapter 5, and background calibration using the topology shown in chapter 4. Such a potential ADC would have the benefit of rapid background calibration, and even lower power than the ADC of chapter 6, as the power hungry front-end S/H would be eliminated.

7.3: KEY DEVELOPMENTS OF THIS WORK

1. A novel topology which enables rapid calibration of both DAC and gain errors (enabling high linearity with poorly matched capacitors) in a multi-bit pipelined stage [2], [3].

2. A novel approach to eliminating the front-end S/H (thus saving power) which does not rely on a carefully matched layout [4], [5].

3. A novel MDAC topology which does not require opamps with capacitive feedback to achieve gain, and hence enables substantial power reduction.

APPENDIX I: CALIBRATION OF NON-LINEAR ERRORS

Consider Fig. 8-1 which is the same as Fig. 4-2 with the exception that the residue amplifier has a non-linear transfer characteristic.

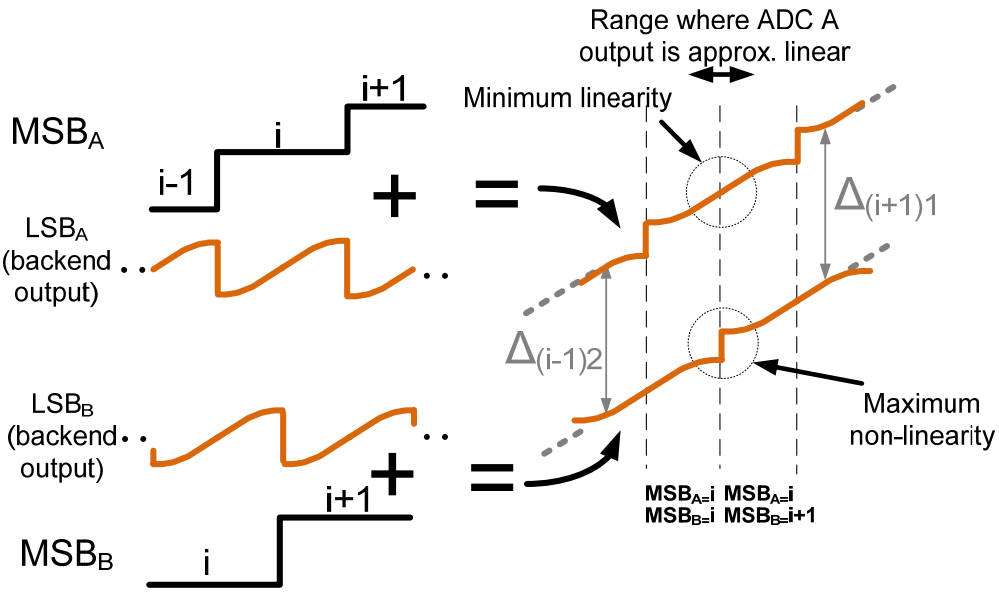

Fig. 8-1: Nonlinearities in residues of each split-ADC

From Fig. 8-1 we note that when the effect of nonlinearity on the output of the first stage in ADC B is largest, the effect of nonlinearity on the output of the first stage in ADC A is smallest. Thus if the nonlinearity affecting the residue amplifier is reasonably small [60], the output of ADC A around the MSB transition of ADC B can be used as an ideal reference to measure the effect of non-linearity in the first pipeline stage of ADC B.

As any input which produces an output around one of the MSB transitions of the ADC produces an output which can immediately generate an estimate of the nonlinearity, a rapid calibration of non-linear errors can thus also be performed assuming the analog input to the ADC is reasonably busy.

In the work presented in chapter four, the emphasis was placed on the correction of the linear error sources of DAC and gain errors, thus nonlinearity correction was not explored in great detail, however brief simulations were performed with the test setup in Simulink as shown in Fig. 8-2.

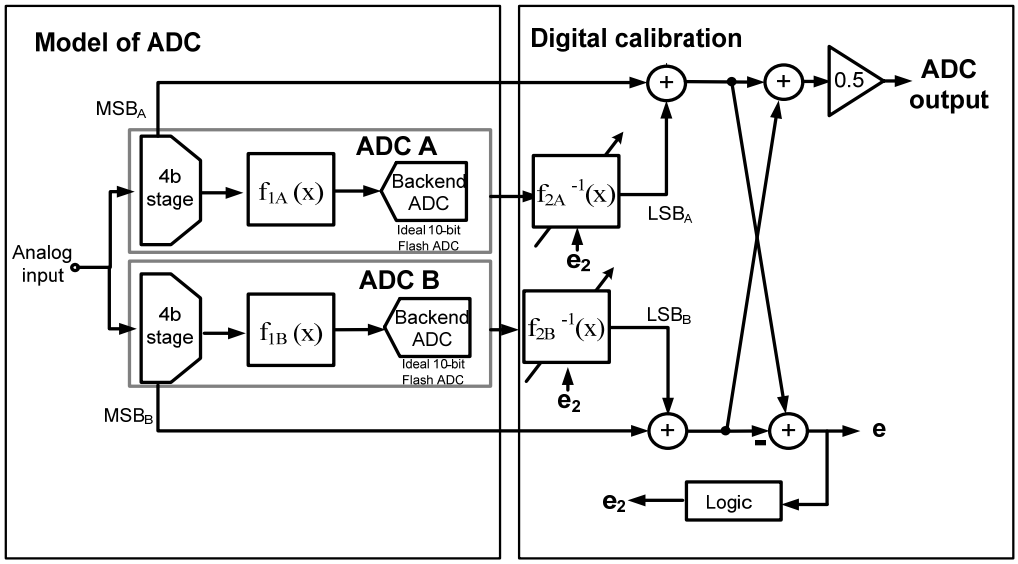

Fig. 8-2: Simulink setup to examine non-linearity correction

The nonlinearity of the residue amplifier $f_1^{-1}(x)$ was modeled by taking the transfer function of a differential pair which performs a gain of 16x using resistive loads in Spice and mapping the transfer function to Simulink by using a lookup table. The inverse function $f^{1}(x)$ was implemented with a simple adaptive neural network. Fig. 8-3 and Fig. 8-4 show an FFT of the output of the ADC before and after calibration.

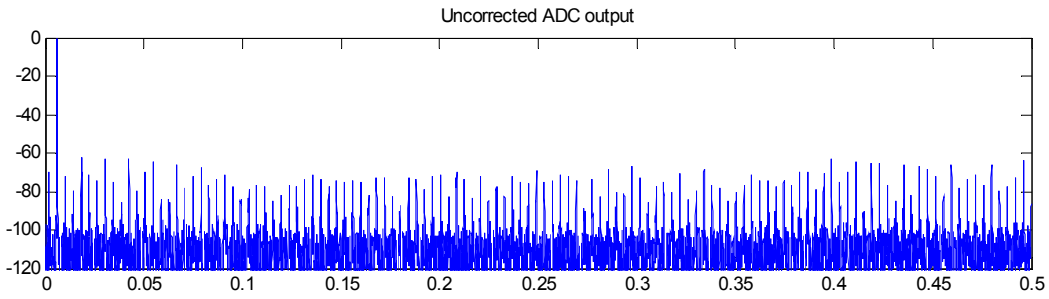

Fig. 8-3: FFT of ADC output without Calibration (SNDR = 48dB, SFDR = 60dB)

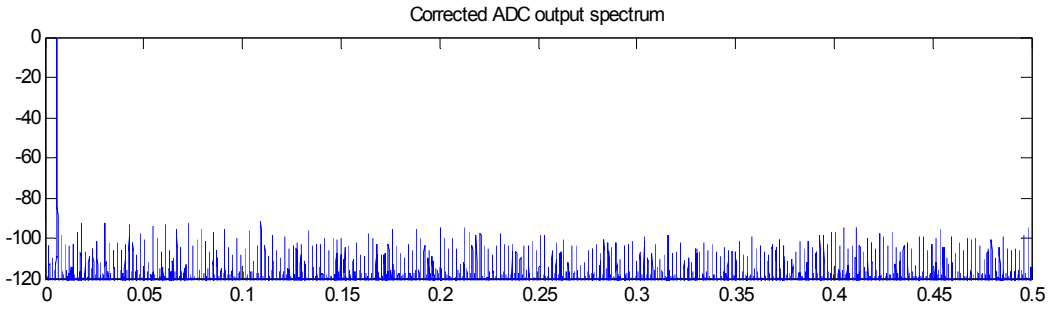

Fig. 8-4: FFT of ADC output after using calibration techniques of this work (SNDR = 80dB, SFDR = 92dB)

From the initial simulations it was found that calibration could be achieved within $1x10^6$ clock cycles, which is at least an order of magnitude faster than the calibration technique used in [60]. It should also be noted that while using a split-ADC technique to correct gain nonlinearity was independently derived for this work, publications have recently emerged which show similar architectures to achieve rapid nonlinearity correction [86].

It is also noted that some initial simulations were performed where gain, DAC, and nonlinearity errors were introduced and were all simultaneously corrected using both the techniques of section 4.3.2 and this section. The simulations where all three error sources were corrected showed promising results, opening the possibility of future work which for the first time could realize an ADC that calibrates all three major error sources.

REFERENCES

[1] http://www.maxim-ic.com/appnotes.cfm/appnote_number/1023/ , July 25th, 2008 9:39PM

[2] I. Ahmed, D.A. Johns, "An 11-bit 45MS/s pipelined ADC with rapid calibration of DAC errors in a multi-bit pipeline stage," *Solid-State Circuits Conference, 2007. ESSCIRC 2007. Proceedings of the 33rd European* , vol., no., pp. 147-150, 11-13 Sept. 2007

[3] Ahmed, I.; Johns, D. A., "An 11-Bit 45 MS/s Pipelined ADC With Rapid Calibration of DAC Errors in a Multibit Pipeline Stage," *Solid-State Circuits, IEEE Journal of* , vol.43, no.7, pp.1626-1637, July 2008

[4] I. Ahmed, D.A. Johns, "A high bandwidth power scalable sub-sampling 10-bit pipelined ADC with embedded sample and hold," *Solid-State Circuits Conference, 2007. ESSCIRC 2007. Proceedings of the 33rd European* , vol., no., pp. 159-162, 11-13 Sept. 2007

[5] Ahmed, I.; Johns, D. A., "A High Bandwidth Power Scalable Sub-Sampling 10-Bit Pipelined ADC With Embedded Sample and Hold," *Solid-State Circuits, IEEE Journal of*, vol.43, no.7, pp.1638-1647, July 2008

[6] Lathi, B.P. Modern Digital and Analog Commuincation Systems. Oxford University Press, New York, 1998

[7] Johns, David and Martin, Ken. Analog Integrated Circuit Design. John Wiley & Sons, Inc: New York, 1997.

[8] Analog Devices Inc, Kester, Walk. The Data Conversion Handbook. Newnes: New York, 2005.

[9] http://www.maxim-ic.com/appnotes.cfm/an_pk/1080 , July 25th, 2008 9:39PM

[10] Craninckx, J.; Van der Plas, G., "A 65fJ/Conversion-Step 0-to-50MS/s 0-to-0.7mW 9b Charge-Sharing SAR ADC in 90nm Digital CMOS," *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International* , vol., no., pp.246-600, 11-15 Feb. 2007

[11] Agnes, Andrea; Bonizzoni, Edoardo; Malcovati, Piero; Maloberti, Franco, "A 9.4-ENOB 1V 3.8µW 100kS/s SAR ADC with Time-Domain Comparator," *Solid-State*

Circuits Conference, 2008. ISSCC 2008. Digest of Technical Papers. IEEE International , vol., no., pp.246-610, 3-7 Feb. 2008

[12] S. H. Lewis, H. S. Fetterman, G. F. Gross Jr, R. Ramachandran, and T.R. Viswanathan, "A 10-bit 20-Msample/s analog-to-digital converter", *IEEE J. Solid-State Circuits*, vol. 27, no. 3, pp. 351–358, Mar. 1992.

[13] P.T.F. Kwok et al, "Power Optimization for Pipeline Analog-to-Digital Converters", *IEEE Transactions on Circuits and Systems--II: Analog and Digital Signal Processing*, vol 36, May 1999, pp. 549-553

[14] D. Cline, "Noise, Speed, and Power Trade-offs in Pipelined Analog to Digital Converters", Doctor of Philosophy in Engineering thesis, University of California Berkeley, 1995

[15] G. Chien, "High-Speed, Lower-Power, Low-Voltage Pipelined Analog-to-Digital Converter", Masters of Science thesis, University of California Berkeley, 1996

[16] Uyttenhove *et al*, "Speed-Power-Accuracy Tradeoff in High-Speed CMOS ADCs", *IEEE transactions on Circuits and Systems –II: Analog and Digital Signal Processing*, vol 49, April 2002, pp. 280-287

[17] I. Ahmed, "A power scaleable and low power pipeline ADC using power resettable opamps", Masters of Applied Science thesis, University of Toronto, 2004

[18] K. Bult, G.J.G.M. Geelen, "A Fast settling CMOS Op Amp for SC Circuits with 90-dB DB Gain", *IEEE Journal of Solid-State Circuits*, vol 25, pp.1379-1384, December 1990

[19] You, F.; Embabi, S.H.K.; Sanchez-Sinencio, E., "Multistage amplifier topologies with nested G_m-C compensation," *Solid-State Circuits, IEEE Journal of*, vol.32, no.12, pp.2000-2011, Dec 1997

[20] Razavi, Behzad. Design of Analog CMOS Integrated Circuits. McGraw-Hill, New York, 2000

[21] L. Sumanen et al, "CMOS dynamic comparators for pipeline A/D converters", *2002 IEEE International Symposium on Circuits and Systems (ISCAS)*, vol 5, pp. 157-160, May 2002

[22] Yang, W.; Kelly, D.; Mehr, L.; Sayuk, M.T.; Singer, L., "A 3-V 340-mW 14-b 75-Msample/s CMOS ADC with 85-dB SFDR at Nyquist input," *Solid-State Circuits, IEEE Journal of* , vol.36, no.12, pp.1931-1936, Dec 2001

[23] I. Galton, “Digital Cancellation of D/A Converter Noise in Pipelined A/D Converters,” IEEE TCAS-II. vol. 47, pp. 185–196, March 2000.

[24] Siragusa, E.; Galton, I., "A digitally enhanced 1.8-V 15-bit 40-MSample/s CMOS pipelined ADC," *Solid-State Circuits, IEEE Journal of* , vol.39, no.12, pp. 2126-2138, Dec. 2004

[25] S. Ray, B.-S Song, “A 13-b Linear, 40-MS/s Pipelined ADC with Self-Configured Capacitor Matching”, IEEE JSSC, vol 42, pp.463-474, March 2007

[26] D.Y. Chang, J. li, U.K moon, “Radix-based digitial calibration techniques for multi-stage recycling pipelined ADCs”, *IEEE Transactions on Circuits and Systems I*, vol. 51, pp.2133-2140, Nov. 2004

[27] C.R. Grace, P.J. Hurst, S.H. Lewis, “A 12b 80 MS/s pipelined ADC with bootstrapped digital calibration”, in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2004, pp. 460-539

[28] McNeill, J.; Coln, M.; Larivee, B., "A split-ADC architecture for deterministic digital background calibration of a 16b 1 MS/s ADC," *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International* , vol., no., pp. 276-598 Vol. 1, 6-10 Feb. 2005

[29] U.K Moon, B.S. Song, “Background Digital Calibration techniques for Pipelined ADC’s”, *IEEE Transactions on Circuits and Systems II*, vol. 44, pp.102-109, Feb. 1997

[30] Kwak, S.-U.; Song, B.-S.; Bacrania, K., "A 15-b, 5-Msample/s low-spurious CMOS ADC," *Solid-State Circuits, IEEE Journal of* , vol.32, no.12, pp.1866-1875, Dec 1997

[31] Erdogan, O.E.; Hurst, P.J.; Lewis, S.H., "A 12-b digital-background-calibrated algorithmic ADC with -90-dB THD," *Solid-State Circuits, IEEE Journal of* , vol.34, no.12, pp.1812-1820, Dec 1999

[32] I. Galton, “Digital Cancellation of D/A Converter Noise in Pipelined A/D Converters,” IEEE TCAS-II. vol. 47, pp. 185–196, March 2000.

[33] J. Li, UK Moon, “Background Calibration Techniques for Multistage Pipelined ADCs With Digital Redundancy,” IEEE *Transactions on Circuits and Systems II*, vol. 50, pp. 531-538, Sept. 2003

[34] Y. Chiu, C.W.Tsang, B. Nikolic, P.R.Gray, “Least Mean Square Adaptive Digital Background Calibration of Pipelined Analog-to-Digital Converters”, IEEE TCAS-I, vol 51, pp.38-46, Jan. 2004

[35] Xiaoyue Wang; Hurst, P.J.; Lewis, S.H., "A 12-bit 20-Msample/s pipelined analog-to-digital converter with nested digital background calibration," *Solid-State Circuits, IEEE Journal of* , vol.39, no.11, pp. 1799-1808, Nov. 2004

[36] Siragusa, E.; Galton, I., "A digitally enhanced 1.8-V 15-bit 40-MSample/s CMOS pipelined ADC," *Solid-State Circuits, IEEE Journal of* , vol.39, no.12, pp. 2126-2138, Dec. 2004

[37] Jipeng Li; Gil-Cho Ahn; Dong-Young Chang; Un-Ku Moon, "A 0.9-V 12-mW 5-MSPS algorithmic ADC with 77-dB SFDR," *Solid-State Circuits, IEEE Journal of* , vol.40, no.4, pp. 960-969, April 2005

[38] Hung-Chih Liu; Zwei-Mei Lee; Jieh-Tsorng Wu, "A 15-b 40-MS/s CMOS pipelined analog-to-digital converter with digital background calibration," *Solid-State Circuits, IEEE Journal of* , vol.40, no.5, pp. 1047-1056, May 2005

[39] S. Ray, B.-S Song, “A 13-b Linear, 40-MS/s Pipelined ADC with Self-Configured Capacitor Matching”, IEEE JSSC, vol 42, pp.463-474, March 2007

[40] I. Mehr and L. Singer, “A 55-mW, 10-bit, 40-Msample/s Nyquist-rate CMOS ADC,” *IEEE J. Solid-State Circuits*, vol. 35, no. 3, pp. 318-325, Mar. 2000.

[41] Ali, A.M.A.; Dillon, C.; Sneed, R.; Morgan, A.S.; Bardsley, S.; Kornblum, J.; Wu, L., "A 14-bit 125 MS/s IF/RF Sampling Pipelined ADC With 100 dB SFDR and 50 fs Jitter," *Solid-State Circuits, IEEE Journal of* , vol.41, no.8, pp. 1846-1855, Aug. 2006

[42] Gulati, K.; Peng, M.S.; Pulincherry, A.; Munoz, C.E.; Lugin, M.; Bugeja, A.R.; Li, J.; Chandrakasan, A.P., "A Highly Integrated CMOS Analog Baseband Transceiver With 180 MSPS 13-bit Pipelined CMOS ADC and Dual 12-bit DACs," *Solid-State Circuits, IEEE Journal of*, vol.41, no.8, pp. 1856-1866, Aug. 2006

[43] Rusu, A.; Rodriguez de Llera Gonzalez, D.; Ismail, M., "Reconfigurable ADCs enable smart radios for 4G wireless connectivity," *Circuits and Devices Magazine, IEEE* , vol.22, no.3, pp. 6-11, May-June 2006

[44] Gulati, K.; Hae-Seung Lee, "A low-power reconfigurable analog-to-digital converter," *Solid-State Circuits, IEEE Journal of* , vol.36, no.12, pp.1900-1911, Dec 2001

[45] Ahmed, I.; Johns, D.A., "A 50-MS/s (35 mW) to 1-kS/s (15 μW) power scalable 10-bit pipelined ADC using rapid power-on opamps and minimal bias current variation," *Solid-State Circuits, IEEE Journal of* , vol.40, no.12, pp. 2446-2455, Dec. 2005

[46] Anderson, M.; Norling, K.; Dreyfert, A.; Yuan, J., "A reconfigurable pipelined ADC in 0.18 μm CMOS," *VLSI Circuits, 2005. Digest of Technical Papers. 2005 Symposium on* , vol., no., pp. 326-329, 16-18 June 2005

[47] Audoglio, W.; Zuffetti, E.; Cesura, G.; Castello, R., "A 6-10 bits Reconfigurable 20MS/s Digitally Enhanced Pipelined ADC for Multi-Standard Wireless Terminals," *Solid-State Circuits Conference, 2006. ESSCIRC 2006. Proceedings of the 32nd European* , vol., no., pp.496-499, Sept. 2006

[48] Craninckx, J.; Van der Plas, G., "A 65fJ/Conversion-Step 0-to-50MS/s 0-to-0.7mW 9b Charge-Sharing SAR ADC in 90nm Digital CMOS," *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International* , vol., no., pp.246-600, 11-15 Feb. 2007

[49] Verma, N.; Chandrakasan, A.P., "An Ultra Low Energy 12-bit Rate-Resolution Scalable SAR ADC for Wireless Sensor Nodes," *Solid-State Circuits, IEEE Journal of* , vol.42, no.6, pp.1196-1205, June 2007

[50] Hsu, Cheng-Chung; Huang, Chen-Chih; Lin, Ying-Hsi; Lee, Chao-Cheng; Soe, Zaw; Aytur, Turgut; Yan, Ran-Hong, "A 7b 1.1GS/s Reconfigurable Time-Interleaved ADC in 90nm CMOS," *VLSI Circuits, 2007 IEEE Symposium on* , vol., no., pp.66-67, 14-16 June 2007

[51] Hernes, B.; Briskemyr, A.; Andersen, T.N.; Telste, F.; Bonnerud, T.E.; Moldsvor, O., "A 1.2V 220MS/s 10b pipeline ADC implemented in 0.13/spl mu/m digital CMOS," *Solid-State Circuits Conference, 2004. Digest of Technical Papers. ISSCC. 2004 IEEE International* , vol., no., pp. 256-526 Vol.1, 15-19 Feb. 2004

[52] Andersen, T.N.; Hernes, B.; Briskemyr, A.; Telsto, F.; Bjornsen, J.; Bonnerud, T.E.; Moldsvor, O., "A cost-efficient high-speed 12-bit pipeline ADC in 0.18-μm digital CMOS," *Solid-State Circuits, IEEE Journal of* , vol.40, no.7, pp. 1506-1513, July 2005

[53] Hernes, B.; Bjornsen, J.; Andersen, T.N.; Vinje, A.; Korsvoll, H.; Telsto, F.; Briskemyr, A.; Holdo, C.; Moldsvor, O., "A 92.5mW 205MS/s 10b Pipeline IF ADC Implemented in 1.2V/3.3V 0.13μm CMOS," *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International* , vol., no., pp.462-615, 11-15 Feb. 2007

[54] C. C. Enz and E. A. Vittoz, “CMOS low-power analog circuit design,” in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, ch. 1.2, Tutorials, 1996, pp. 79-133.

[55] J. Crols et al, “Switched-Opamp: An Approach to Realize Full CMOS Switched-Capacitor Circuits at Very Low Power Supply Voltages”, *IEEE Journal of Solid-State Circuits*, vol 29, pp. 936-942, August 1994

[56] M. Waltari et al “1-V 9-Bit Pipelined Switched-Opamp ADC”, *IEEE Journal of Solid-State Circuits*, vol 6, pp. 129-134, January 2001

[57] H.C. Kim et al, “A 30mW 8b 200MS/s Pipelined CMOS ADC Using a Switched-Opamp Technique”, *ISSCC Dig. Tech. Papers*, pp. 284-285, Feb. 2005

[58] R. Wang et al, “A 3.3mW 12MS/s 10b Pipelined ADC in 90nm Digital CMOS”, *ISSCC Dig. Tech. Papers*, pp. 278-279, Feb. 2005

[59] V. S. L. Cheung et al, “A 1-V 10.7-MHz Switched-Opamp Bandpass ΣΔ Modulator Using Double=Sampling Finite-Gain-Compensation Technique”, *IEEE Journal of Solid-State Circuits*, vol 37, pp. 1215-1225, October 2002

[60] B. Murmann et al, “A 12-bit 75-MS/s Pipelined ADC Using Open-Loop Residue Amplification”, *IEEE Journal of Solid-State Circuits*, vol 38, December 2003, pp. 2040-2050

[61] Fiorenza, J. K.; Sepke, T.; Holloway, P.; Sodini, C. G.; Lee, H.-S., "Comparator-Based Switched-Capacitor Circuits for Scaled CMOS Technologies," *Solid-State Circuits, IEEE Journal of* , vol.41, no.12, pp.2658-2668, Dec. 2006

[62] W. Dally and J. Poulton. Digital Systems Engineering. Cambridge University Press, Cambridge, UK, 2001

[63] Jian Li; Xiaoyang Zeng; Lei Xie; Jun Chen; Jianyun Zhang; Yawei Guo, "A 1.8-V 22-mW 10-bit 30-MS/s Subsampling Pipelined CMOS ADC," *Conference 2006, IEEE Custom Integrated Circuits* , vol., no., pp.513-516, 10-13 S

[64] A.M. Abo, P.R. Gray, “A 1.5 V, 10-bit, 14MS/s CMOS Pipeline analog-to-digital converter” *IEEE, VLSI Circuits*, June 1998, pp. 166-169

[65] R.J. Baker. CMOS Circuit Design, Layout, and Simulation. IEEE Press/ Wiley Interscience, New Jersey, 2005

[66] Maulik, P. C.; Mercer, D. A., "A DLL-Based Programmable Clock Multiplier in 0.18-µm CMOS With -70 dBc Reference Spur," *Solid-State Circuits, IEEE Journal of* , vol.42, no.8, pp.1642-1648, Aug. 2007

[67] Tai-Cheng Lee; Keng-Jan Hsiao, "The design and analysis of a DLL-based frequency synthesizer for UWB application," *Solid-State Circuits, IEEE Journal of* , vol.41, no.6, pp. 1245-1252, June 2006

[68] J. Park, H. J. Park, K. Jae-Whui, S. Sangnam, and P. Chung, "A 1 mW 10-bit 500 KSPS SAR A/D converter," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 5, May 2000, pp. 581-584.

[69] P. Confalonleri, M. Zarnprogno, F. Girardi, G. Nicollini, and A. Nagari, "A 2.7 mW 1 MSps 10 bit analog-to-digital converter with built-in reference buffer and ILSB accuracy programmable input ranges," in *Proc. Eur. Solid-State Circuits Conf. (ESSCIRC)*, Sep. 2004, pp. 255-258.

[70] L. Dorrer, F. Kuttner, A. Wiesbauer, A. Di Giandomenico, and T. Hartig, "10-bit, 3 mW continuous-time sigma-delta ADC for UMTS in a 0.12 μm CMOS process," in *Proc. Eur. Solid-State Circuits Conf. (ESSCIRC)*, Sep. 2003, pp. 245-248

[71] Fiorenza, J. K.; Sepke, T.; Holloway, P.; Sodini, C. G.; Lee, H.-S., "Comparator-Based Switched-Capacitor Circuits for Scaled CMOS Technologies," *Solid-State Circuits, IEEE Journal of* , vol.41, no.12, pp.2658-2668, Dec. 2006

[72] R. Wang, K. Martin, D. Johns, and G. Burra, "A 3.3 mW 12 MS/s 10 bit pipelined ADC in 90 nm digital CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2005, pp. 278-279.

[73] D. Y. Chang and U. K. Moon, "A 1.4-V 10-bit 25 MS/s pipelined ADC using opamp-reset switching technique," *IEEE J. Solid-State Circuits*, vol. 38, no. 8, pp. 1401-1404, Aug. 2003.

[74] D. Miyazaki, S. Kawahito, and M. Furuta, "A 10-bit 30-MS/s low-power pipelined CMOS A/D converter using a pseudodifferential architecture," *IEEE J. Solid-State Circuits*, vol. 38, no. 2, pp. 369-373, Feb. 2003.

[75] Yin, Wenjing; Jiang, Jie; Xu, Jun; Ye, Fan; Ren, Junyan, "An Undersampling 10-bit 30.4-MSample/s Pipelined ADC," *Solid-State Circuits Conference, 2006. ASSCC 2006. IEEE Asian* , vol., no., pp.343-346, Nov. 2006

[76] J. Arias, et al., "Low-power pipeline ADC for wireless LANs," *IEEE J. Solid-State Circuits*, vol. 39, no. 8, pp. 1338-1340, Aug. 2004.

[77] Bo Xia; Valdes-Garcia, A.; Sanchez-Sinencio, E., "A 10-bit 44-MS/s 20-mW configurable time-interleaved pipeline ADC for a dual-mode 802.11b/Bluetooth

receiver," *Solid-State Circuits, IEEE Journal of* , vol.41, no.3, pp. 530-539, March 2006

[78] Ryu, S.T.; Song, B.S.; Bacrania, K., "A 10-bit 50-MS/s Pipelined ADC With Opamp Current Reuse," *Solid-State Circuits, IEEE Journal of*, vol.42, no.3, pp.475-485, March 2007

[79] Lu, C.-C.; Lee, T.-S., "A 10-bit 60-MS/s Low-Power CMOS Pipelined Analog-to-Digital Converter," *Circuits and Systems II: Express Briefs, IEEE Transactions on [see also Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on]* , vol.54, no.8, pp.658-662, Aug. 2007

[80] Quinn, P.; Pribytko, M., "Capacitor matching insensitive 12-bit 3.3 MS/s algorithmic ADC in 0.25 /spl mu/m CMOS," *Custom Integrated Circuits Conference, 2003. Proceedings of the IEEE 2003* , vol., no., pp. 425-428, 21-24 Sept. 2003

[81] Dickson, J.F., "On-chip high-voltage generation in MNOS integrated circuits using an improved voltage multiplier technique," *Solid-State Circuits, IEEE Journal of* , vol.11, no.3, pp. 374-378, Jun 1976

[82] H.S. Lee, “Fundamental Limits and Practical issues of Analog-to-Digital Converters”, ISSCC 2006 Short Course

[83] Schreier, R.; Silva, J.; Steensgaard, J.; Temes, G.C., "Design-oriented estimation of thermal noise in switched-capacitor circuits," *Circuits and Systems I: Regular Papers, IEEE Transactions on [Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on]* , vol.52, no.11, pp. 2358-2368, Nov. 2005

[84] Lee, S.-H.; Song, B.-S., "Simplified digital calibration for multi-stage analog-to-digital converters," *Circuits and Systems, 1993., ISCAS '93, 1993 IEEE International Symposium on* , vol., no., pp.1216-1219 vol.2, 3-6 May 1993

[85] Hadidi, K.; Sasaki, M.; Watanabe, T.; Muramatsu, D.; Matsumoto, T., "An open-loop full CMOS 103 MHz -61 dB THD S/H circuit," *Custom Integrated Circuits Conference, 1998. Proceedings of the IEEE 1998* , vol., no., pp.381-383, 11-14 May 1998

[86] McNeill, J.A.; Goluguri, S.; Nair, A., ""Split-ADC" Digital Background Correction of Open-Loop Residue Amplifier Nonlinearity Errors in a 14b Pipeline ADC," *Circuits and Systems, 2007. ISCAS 2007. IEEE International Symposium on* , vol., no., pp.1237-1240, 27-30 May 2007