

AN ABSTRACT OF THE THESIS OF

David Patrick Gubbins for the degree of Doctor of Philosophy in
Electrical and Computer Engineering presented on December 9, 2008.

Title: Continuous Time Input Pipeline ADCs

Abstract approved:

Un-Ku Moon

Analog-to-digital converters (ADCs) convert analog continuous time signals into discrete time, digital format. One precondition that must be met for conventional nyquist rate ADCs is that the input signal must be suitably band-limited to an input bandwidth less than the nyquist frequency. This mandates expensive anti-alias filters which contribute to system noise and distortion degradation. By choosing an OSR of 2 and adopting simple linear phase filtering techniques, significant inherent anti-alias filtering is achieved, avoiding the need for an explicit anti-alias filter in many applications. Additionally the proposed continuous time input pipeline ADC eases a number of other challenges present in conventional switched capacitor ADCs:- sampled opamp noise folding, sampling distortion, reduced ADC area, switched-capacitor pipeline ADC input loading. Chapter 1 introduces the first continuous time input pipeline ADC in the literature. This ADC,

while providing significant benefits, does not provide the inherent filtering. Chapter 2 presents the first continuous time input pipeline ADC with inherent anti-alias filtering.

©Copyright by David Patrick Gubbins
December 9, 2008
All Rights Reserved

Continuous Time Input Pipeline ADCs

by

David Patrick Gubbins

A THESIS

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Doctor of Philosophy

Presented December 9, 2008

Commencement June 2009

Doctor of Philosophy thesis of David Patrick Gubbins presented on
December 9, 2008.

APPROVED:

Major Professor, representing Electrical and Computer Engineering

Director of the School of Electrical Engineering and Computer Science

Dean of the Graduate School

I understand that my thesis will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my thesis to any reader upon request.

David Patrick Gubbins, Author

ACKNOWLEDGEMENTS

Firstly I would like to thank Professor Moon for the opportunity to pursue my doctorate in his group at OSU. I am so lucky to have had a major advisor who was so generous with his knowledge and time. Much thanks is also due to my committee members, Professors Mayaram, Temes and Hanumolu. These Professors have also been so generous with their time and knowledge. Thanks also goes to Professor Brady Gibbons who also served as GCR on my committee.

Kind thanks goes to all the members of Professor Moons group while I was at OSU. I immensely enjoyed all the discussions/banter -both technical and non-technical. I learned so much off of you all. Thanks to Sunwoo, Nima, Pavan, Igor, Vova, Youn-Jae, Josh, Nima, Tawfik, Manideep, Rob, Jon, Ben, Skyler, Peter, Ho-young, Sasidhar, Omid, Hari, Ting, Won-Seok, Merrick, Min-Gyu, Gil-Cho. It was a privilege for me to work and collaborate with this group.

Thanks also to all of the dedicated Professors at OSU that I took classes with. Thanks to Ferne and the EECS staff for steering me in the right path with regard to forms/classes etc!

Thanks to the wonderful people of Corvallis who made me feel so welcome in your great town. What great and generous people there are here. In particular I want to thank the wonderful circle of friends that I met here: Hank, Peter, Mikele, Marissa, Nic, Joey, Barb ,Carb, Judy, Marshall, Conor, Colin, Danielle , Scott, Christy ,Mike, Monica, Jacob, Hope and Troy, Lou, Chris, Amy. Thanks also to all my international friends : Carole, Chadi , Loraine, Toby, Bea, Charles, Hannes,

Siavash.

Thanks to all my room-mates in the "old house" - Marissa, Matheus, Fadsai, Tari, Peter, Tony, Travis and James. It's always been a welcoming place to drink tea and shoot the breeze-long may it continue!

Kind thanks to National Semiconductor for sponsoring my research. Having access to Nationals regular 0.18um fabrication was a major blessing. Also getting feedback on my research from National engineers was also very welcome. Kind thanks especially goes to Bumha Lee and Dave Boisvert.

Thanks to all the people who came to visit me from Ireland- Eleanor and Paudie, Kieran, Celina and Pat, Evelyn and Joe, Pat Culhane, Paul Mallon.

Muchas gracias a Elena por su paciencia, apoyo, aliento y reflexin durante mi doctorado.

I especially want to acknowledge my family. Thanks to my mother and father, Eileen and Tom, sisters, Ita, Lucy, Marie, Elizabeth and Helen for your unfaltering love and encouragement. I am also grateful to John, Donagh and Lorcan for being such gentlemen and good craic.

TABLE OF CONTENTS

	<u>Page</u>
1 Introduction	1
1.1 Pipeline ADC Operation	2
1.2 Switched Capacitor Implementation of the pipelined ADC	8
1.3 The Problems with the State of the Art	14
2 A Continuous-time Input Pipeline ADC	17
2.1 Introduction	17
2.2 Architecture of the continuous-time Input Pipeline ADC	23
2.3 Prediction Filter	25
2.4 Circuit Details	31
2.5 Experimental Results	32
2.6 Conclusions	36
3 A Windowed Integration, Residue Sampling Pipeline ADC with Inherent Anti-alias Filtering	44
3.1 Introduction	44
3.2 The Anti-alias Filter Problem	46
3.3 Architecture of the continuous-time Input Pipeline ADC	50
3.4 Prediction Filter	58
3.5 Circuit Details	58
3.6 Experimental Results	61
3.7 Conclusions	70
4 Conclusions	71
4.1 What in a nutshell has been done so far?	71
4.2 Specific Silicon Improvements	72
4.3 Future Directions	74
Bibliography	76

LIST OF FIGURES

Figure	Page
1.1 Basic Functionality of an ADC	2
1.2 Basic Functionality of a pipeline ADC	3
1.3 Alignment and Combination of the stage outputs	4
1.4 Simplified view of a K-bit pipelined ADC stage followed by an ideal M-bit ADC	5
1.5 Graphical view of how ADC operates	6
1.6 Typical Switched Capacitor Implementation of a pipeline stage . . .	7
1.7 A survey of recent high performance pipelined ADCs	12
1.8 One page anti-alias picture for a conventional nyquist rate ADC . .	14
2.1 The continuous-time input pipeline ADC architecture	19
2.2 The first two stages of the continuous-time input pipeline ADC . .	19
2.3 Elimination of the drive circuitry for CT input pipeline ADC	20
2.4 Simplified diagram of the conventional pipeline ADC first stage with timing diagram	21
2.5 The continuous-time input pipeline ADC first stage with timing diagram	22
2.6 Mathematical model showing propagation of error V_e	25
2.7 Pole-zero plot	27
2.8 Measured and Simulated Frequency Response	28
2.9 Finite gain opamp used through the prediction filter	30
2.10 Prediction Filter Biquad	31
2.11 Transfer Functions of the pipeline stages	33
2.12 Simplified first stage MDAC circuitry ($R_{in}=3k\Omega$, $R_f=4k\Omega$)	34
2.13 Opamp used through the ADC	35

LIST OF FIGURES (Continued)

Figure	Page
2.14 DNL and INL Performance of the ADC versus output code	37
2.15 Single tone FFT plot at 1MHz input and 16MSPS	38
2.16 SNR, SNDR, SFDR over input frequency at 16MSPS	39
2.17 Photo of the die	40
2.18 Power Breakdown of the prototype ADC	41
3.1 A typical pipeline ADC application circuit [16]	47
3.2 Magnitude and Group Delay Response requirements of an anti-alias filter	48
3.3 The continuous-time input pipeline ADC architecture	50
3.4 The first two stages of the continuous-time input pipeline ADC	51
3.5 Simplified diagram of the conventional pipeline ADC first stage with timing diagram	52
3.6 The continuous-time input pipeline ADC first stage with timing diagram	53
3.7 Simplified first stage MDAC circuitry. This stage implements windowed integration residue sampling.	62
3.8 Circuitry used to match the IDAC fullscale current to the unit poly resistance, R	63
3.9 First Stage CT ADC Opamp with switched-capacitor common mode feedback	64
3.10 SNDR and SFDR over input frequency at 26MSPS	66
3.11 Measured Frequency Response of the pipeline ADC	67
3.12 Photo of the die	68
3.13 Power Breakdown of the prototype ADC	69
4.1 Optimized Frequency Response of the CT pipeline ADC	73

LIST OF TABLES

<u>Table</u>		<u>Page</u>
2.1	Prediction Filter Specifications Based on 2.5b 1st Stage	26
2.2	Performance Summary	42
3.1	Prediction Filter Specifications Based on 2.5b 1st Stage	57
3.2	Performance Summary	65

DEDICATION

This thesis is dedicated to the memory of my mother Eileen and my sister Liz.

Chapter 1 – Introduction

Switched-capacitor techniques are very prevalent in the integrated circuit industry. The main reason for this is that switched-capacitor methods can be used to implement analog-to-digital converters (ADCs), digital-to-analog converters (DACs) and filters in a robust fashion. This leads to high yield I.C. products that are insensitive to PVT variation and easy to test. One could say that switched-capacitor methods have been a key enabler in the mobile electronics space over the last 30 years. The focus of this phd thesis is pipeline ADCs. Switched-capacitor techniques have traditionally been used to implement pipeline ADCs.

An ADC is an electronic circuit that changes or converts a continuous analog signal into a digital signal without altering its critical content.

In the simplest terms, an analog-to-digital converter samples an analog waveform at uniform time intervals and assigns a digital value to each sample. An ADC carries out two processes. These processes are quantization and sampling.

The quantization process occurs when the ADC represents an analog signal as a digital string of 1's and 0's with finite resolution. The ADC outputs a finite number of digital values, equal to 2^N , where N is the resolution in bits of the ADC. Digital signals are not an exact representation of the analog signal.

Sampling is the process of measuring the continuous analog signal at discrete and standard intervals. ADC operation in the real world, though, is affected by




Figure 1.1: Basic Functionality of an ADC

imperfect conditions. Figure 1.1 illustrates the sampling action of the ADC. The ADC receives the input signal V_{in} in addition to a clock source Clk . The ADC samples the input signal upon the positive edge of Clk and subsequently converts that signal into a digital format with N bits. This N -bit word is output to the N -bit output bus at the sampling rate. The sampling frequency is denoted as F_s .

1.1 Pipeline ADC Operation

A generalized pipeline ADC architecture is shown in Figure 1.2. The ADC is made up of a pipeline of stages each of which resolves some digital bits. Stage 1 resolves k bits, stage 2 resolves L bits and so on. There can be as many stages as required used to achieve the required resolution. For example 15 1-bit stages can be used to resolve 15 bits and hence realize a 15 bit ADC. Similarly 5 2-bit stages can be used to resolve 10 bits and hence realize a 10 bit ADC. It is this pipelining




Figure 1.2: Basic Functionality of a pipeline ADC

of stages that gives the pipeline ADC its name. The digital outputs from the pipeline stages are combined in a digital alignment block as shown in Figure 1.3. Re-alignment is required because as an input sample is passed down the pipeline the digital outputs corresponding to one sample become out of sync with respect to each other.

In order to understand the functionality of a pipeline ADC it is helpful to simplify our view of the entire ADC. This we do by just looking at a single pipeline stage followed by an ideal 'back-end ADC'. Here the back-end ADC represents the back-end stages of the pipeline ADC. Such a setup is shown in Figure 1.4.

In this situation the first stage K-bit pipeline stage resolves K bits and creates the residue V_{res1} . The backend M-bit ADC takes the residue, V_{res1} and converts it to digital format. The K bits from the first stage and the M bits from the back-end ADC are then combined to yield the final value. As an example suppose that $K=1$ and $M=3$. This functionality is shown graphically in Figure 1.5. Figure 1.5




Figure 1.3: Allignment and Combination of the stage outputs

shows how

1. Stage 1 creates its 1-bit digital output code
2. Stage 1 creates its analog output voltage V_{res1}
3. Stage 2 takes V_{res1} and creates a 2-bit digital output code
4. The first and second stage digital output codes are combined to create the final ADC output code.

The first stage transfer function is on the left with V_{in} on the x-axis and V_{res1} on the y-axis. V_{res1} then feeds directly into the back-end ADC whose characteristics are shown on the right of the plot. Note how the ADC output code is generated by aligning the stage 1 and stage 2 bits as shown in Figure 1.5

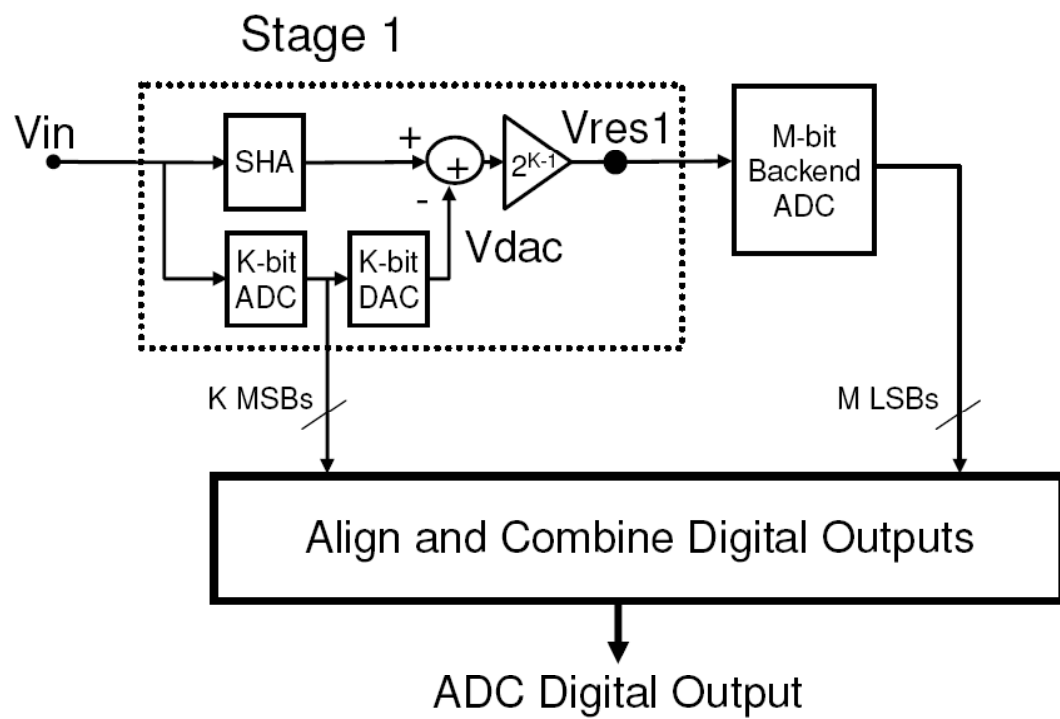


Figure 1.4: Simplified view of a K-bit pipelined ADC stage followed by an ideal M-bit ADC

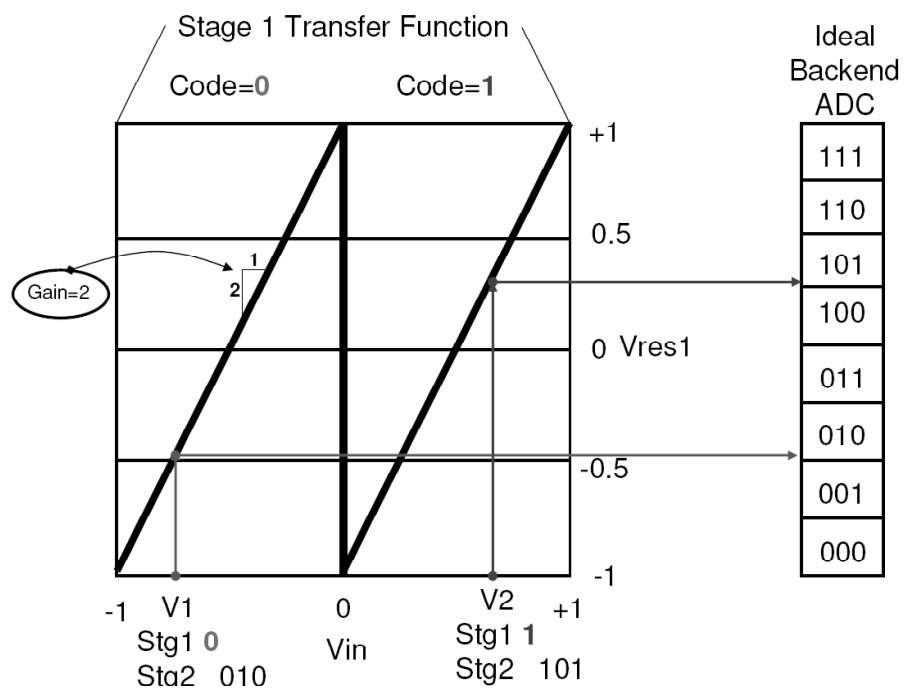


Figure 1.5: Graphical view of how ADC operates

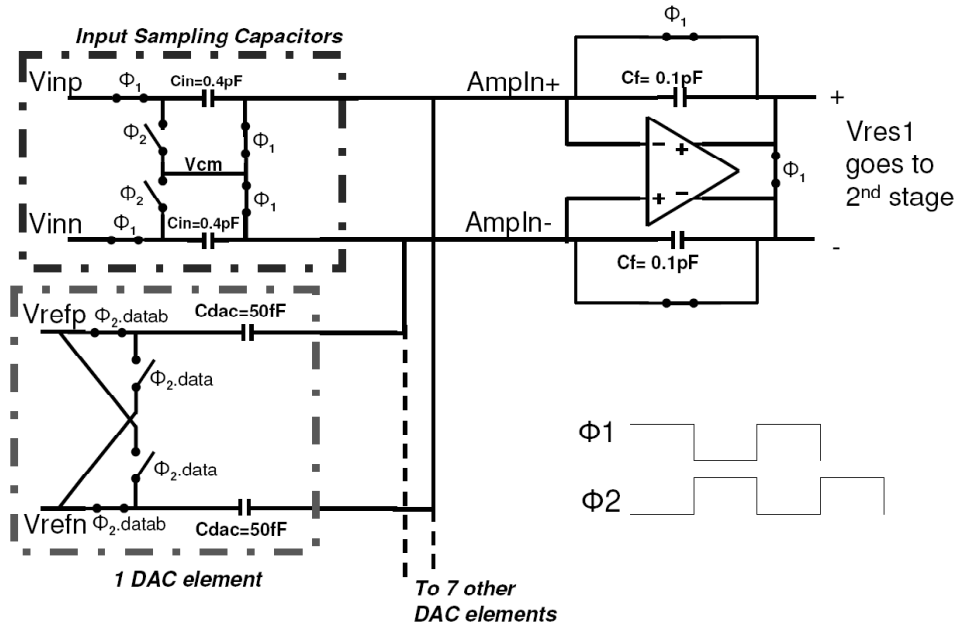


Figure 1.6: Typical Switched Capacitor Implementation of a pipeline stage

and simply adding together. As an example of how the pipeline stage works two example voltage conversions are shown, V_1 and V_2 . When V_1 is applied to the stage 1 it generates a digital code of 0 and creates a V_{res1} voltage which is applied to the backend ADC. The backend ADC then yields the digital code 010. It is shown in the slide how these are aligned and added to create the final value of 0010. The same procedure is used to explain the conversion of V_2 from the analog to digital domain. Another detail of Figure 1.5 is that stage 1 applies a gain of 2 to the input signal. This is clearly seen from the diagram. Gain is applied in order to maintain the same peak-to-peak signal level down the pipeline. This makes sure that the signal is always large enough so that it is easy to resolve with cheap comparators. Looking at Figure 1.5 it can be seen that the input voltage range to

the first stage is $\pm 1V$. The output voltage range is also $\pm 1V$. This is continued through the rest of the pipeline ADC.

Mathematically the following equation describes what is happening in the first pipeline ADC stage shown in Figure 1.5.

$$V_{res1} = 2(V_{in} \pm \frac{1}{2}) \quad (1.1)$$

There are two aspects to this equation. Firstly a dac voltage ($\pm \frac{1}{2}$) is subtracted from the input voltage. This dac voltage is an approximation to the input voltage and this causes the component in brackets in Equation 1.1 to be smaller than V_{in} . By applying gain this reduced signal level can be restored to the $\pm 1V$ voltage range.

1.2 Switched Capacitor Implementation of the pipelined ADC

Present day pipelined ADCs are implemented using switched capacitor methods. One example of switched capacitor pipeline circuitry is shown in Figure 1.6. This circuit implements the sample-and-hold (S/H), DAC and amplifier of Figure 1.4. This circuit is typically referred to as the MDAC in the pipeline stage- that is the pipeline stage excluding the K-bit ADC of Figure 1.4. The circuit of Figure 1.6 is composed of an input sampling network, a switched capacitor DAC and an amplifier. The input sampling capacitor samples the input signal on the negative edge of $\phi 1$. During $\phi 2$ V_{res1} is created as follows. The sampled V_{in} charge and DAC charge are injected into the virtual earth of the amplifier. The amplifier

converts the charge into the residue voltage V_{res1} as per Equation 1.1. A tradeoff arises between noise, distortion and power in such switched capacitor circuits. In order to achieve a certain noise performance a certain minimum capacitance is required in the MDAC stage of Figure 1.6. This stems from the fact that noise power in the switched capacitor network is largely proportional to kT/C . Thus in order to achieve low noise large capacitors are required. Linear capacitors occupy a lot of area on integrated circuits due to low capacitance per area numbers. Large sampling capacitance mandates the use of large switches with low R_{on} in order that the sampling network settles in the allotted time period. Sometimes t-gates are used as sampling switches. However this leads to limited distortion performance (roughly 75dB THD) due to signal dependent R_{on} in the sampling switches. Bootstrapped switches are often used to achieve a switch R_{on} that stays constant over the input voltage range. However this adds complexity and it is difficult to keep voltages across the MOS devices less than the process limit- this poses a reliability challenge. Driving the large switches at the clock rate can mean that the clock driver power is substantial. Looking at Figure 1.6 it is clear that when the residue V_{res1} is being created the amplifier output voltage is required to slew and settle in the allotted time (roughly half the clock period). There are two aspects to this. Firstly the amplifier is driving a substantial capacitive load. Such a load causes a pole at the output node of the opamp. Assuming a single-pole operational amplifier the closed-loop bandwidth is given by :-

$$BW = G_m/C_L \times f \quad (1.2)$$

A minimum closed loop bandwidth is required such that the opamp output V_{res1} settles adequately within the allotted time period (typically $1/(2 * F_s)$). If the opamp does not settle correctly then this creates distortion in the V_{res1} value sampled by the backend-ADC and the digitised signal. Lower feedback factor degrades the closed loop bandwidth. Large amplifier load capacitance degrades the closed loop bandwidth. Large closed loop bandwidth mandates a certain G_m in the opamp. This means that a certain amount of power needs to be dissipated in the opamp in order to achieve the required bandwidth. Clearly low noise and good distortion performance demands large power dissipation in the opamp. Clearly there are challenges in the input network with regard to sampling distortion.

Another problem with such ADCs is that the input sampling capacitor is difficult to drive- either from the outside world or from on-chip. The sampling network needs to be refreshed with charge every clock-cycle. This causes a disturbance in the driving amplifier which needs to settle by the end of the sampling time. This time-domain phenomenon is commonly referred to as kickback. The bonding wire inductance often plays a role in the capacitor settling transient. A common solution to this problem is to place an on-chip buffer in front of the ADC. However the on-chip buffer can then often limit the noise and distortion performance of the ADC.

Another problem with the input sampling network is that the input impedance of the ADC is frequency dependent and it changes from ϕ_1 to ϕ_2 . This creates a challenge when trying to properly terminate the ADC when driving from the outside world-this is a frequency domain phenomenon.

The first stage MDAC of the pipeline ADC typically dominates the power for the entire ADC. Due to the gain through the pipeline ADC the noise requirements of the backend-ADC get progressively relaxed. Thus the capacitor sizes can be scaled down in the backend ADC. This reduces the power in the latter stages of the ADC. This leaves the first stage as the most challenging stage to design. The challenges associated with the switched capacitor first stage of the pipeline ADC can be listed as follows:-

- Sampling distortion
- Bootstrapped switches required for good distortion performance
- Large Capacitance Area
- High performance pipeline ADCs occupy a lot of silicon area
- Limited input voltage range
- Driving pipeline ADCs from the outside world is difficult
- The pipeline ADC first stage burns most of the power due to op-amp requirements

High performance pipeline ADCs are used in applications that demand good noise performance in conjunction with good distortion performance. High performance high speed pipelined ADCs often require signal-to-noise-plus-distortion (SNDR) ratios of $\geq 70\text{dB}$. and spurious free dynamic ranges (SFDR) of $\geq 90\text{dB}$.

ADC	Stated # of bits	SNR [dB]	ERBW [MHz]	Power [mW]	SFDR [dB]	Technology	Sample Rate [MSPS]
Zanchi et al, JSSC June 05	16	78	40	1150	88	BiCmos	65
Ali et al JSSC Sept 06	14	75	500	1850	100	BiCmos	125
Siragusa et al JSSC Dec 04	15	72	20	400	90	Cmos	40
Gulati et al VLSI 04	12	60.6	180	1200	71	Cmos	180

Figure 1.7: A survey of recent high performance pipelined ADCs

However such ADCs also dissipate significant power. A survey of recent high performance pipelined ADCs shown in Figure 1.7 bears out this fact.

There are further issues with state of the art nyquist rate ADCs. Typically the input voltage range $V_{in_{rms}}$ of the switched capacitor network is limited to input supply range. This is mainly because various reverse biased diodes exist between the input voltage pins and the supply voltage. These diodes are in the input ESD protection circuitry and in the input sampling PMOS device. There also exist reverse biased diodes to ground. This input voltage limitation sets the maximum input peak-to-peak signal swing. When combined with the SNR specification one can calculate the input referred RMS noise requirement, $V_{n_{rms}}$ ($SNR = \frac{V_{in_{rms}}}{V_{n_{rms}}}$). This input referred RMS noise requirement strongly influences the power dissipation and area of the ADC. Thus for high performance ADCs any chance to increase the input signal swing is welcomed. This, of course, can increase the burden on circuitry before the ADC. For example in a radio receiver channel it would increase the burden on the amplifier before the ADC. Thus sometimes the ADC needs to achieve a given SNDR with a given input signal range that may be some fraction of the possible rail-to-rail input signal.

Another issue with high dynamic range ADCs is that there is the requirement for an anti-alias filter in front of the ADC. This filter prevents unwanted higher frequency noise or interference from aliasing down into the passband of the ADC as shown in Figure 3.1. In this figure F_b is the bandwidth of interest. It is assumed that once the input spectrum is sampled that an ideal "brick wall" digital filter with cutoff F_b is applied to the discrete time spectrum. Oversampling ($\frac{F_s}{2} > F_b$)




Figure 1.8: One page anti-alias picture for a conventional nyquist rate ADC

eases the design of the anti-alias filter. The anti-aliasing filter is specified to ensure that signal level at $F_s - F_b$ meets the dynamic range requirements of the ADC. This job is eased yet again if indeed there is not a lot of signal energy beyond $(F_s - F_b)$. Then assuming that the voltage range of the ADC is not exceeded it turns out that a normal nyquist rate ADC can tolerate a full scale sine wave between the frequencies F_b and $(F_s - F_b)$.

1.3 The Problems with the State of the Art

This section summarises the problems/challenges with state of the art pipeline ADCs.

- Sampling distortion
- Bootstrap switches required for good distortion
- Large Capacitance Area
- High performance pipeline ADCs occupy a lot of silicon area
- Limited input voltage range
- Driving pipeline ADCs from the outside world is difficult
- High performance pipeline ADCs burn a lot of power due to op-amp requirements
- An anti-alias filter is typically required in front of the ADC

This Phd thesis work addresses these challenges. The key principle proposed in this thesis is that the first stage residue should be processed in a continuous time (CT) fashion. By doing so the key limitations of switched-capacitor circuitry are very much reduced. By resolving ADC bits in the first stage in a continuous-time fashion and moving the sampling operation to the second stage input the back-end switched-capacitor circuitry can be scaled down appropriately because of the reduced back-end noise requirements. This eases the required switch resistances due to the smaller sampling capacitors and obviates the need for a front end sample-and-hold circuit. More importantly it allows the possibility that the proposed continuous-time first stage may be used to filter the input signal-thus allowing the possibility of inherent anti-aliasing.

This thesis is organized as follows. Chapter 2 is written as a stand-alone journal paper draft. It describes the first attempt at a CT input pipeline ADC. Chapter 3 is also written as a stand-alone journal paper draft. It describes the second CT input pipeline ADC that was designed as part of this Phd. This second ADC most importantly incorporates inherent anti-alias filtering. Chapter 4 draws some conclusions from this thesis work and makes some suggestions for further work.

Chapter 2 – A Continuous-time Input Pipeline ADC

2.1 Introduction

Pipeline ADCs have traditionally been implemented with switched-capacitor methods. However switched-capacitor methods present many challenges to ADC designers when more than 10 bit ENOB, high SFDR and low power are required [1]. These challenges are most acute in the pipeline ADC first stage, mainly due to the fact that the first stage is the main contributor to ADC noise, distortion and consequently power [3]. Sampling the input signal onto large sampling capacitors with low distortion poses a challenge to circuit designers. Such capacitors require large sampling switches with low on resistance, R_{on} , in order that the input sampling network settles in the allotted time period. Large sampling switches demand large switch logic drivers and this increases the digital power consumption. Large sampling switches also bring with them nonlinear junction capacitance which can cause the sampled signal to be distorted. Sampling is further complicated by input signal bondwire inductance which may cause ringing on the input signal sampling network. The large ADC capacitive load also presents a challenge for the reference buffers which need to settle to a required accuracy.

This work addresses these challenges by implementing a continuous-time(CT) first stage followed by a scaled down switched capacitor pipeline back-end ADC. For

every extra bit resolved in the CT first stage the back-end switched-capacitors can be scaled down by a factor of four [4]. This new architecture is shown in Fig. 3.3. Such an approach leads to lower power, reduced sampling distortion, lower capacitance area and allows rail-to-rail input swing. Switched-capacitor pipeline ADCs typically require a driver amplifier and an isolation circuit. This is especially important when the switched-capacitor input load is significant ($> 1pF$). A typical applications circuit is shown in Fig. 2.3(a). Such external circuitry increases the system cost and degrades the system noise and distortion and increases the system power. The resistive input load of the CT pipeline ADC allows the drive circuitry to be eliminated as shown in Fig. 2.3(b). In Fig. 2.3(b) the ADC resistive load acts as a termination for the anti-alias ladder filter. The key point here is that the ADC input current is a continuous-time entity whereas the switched capacitor input load requires discrete signal-dependent packets of charge every cycle. By resolving two bits in the first stage in a continuous-time fashion and moving the sampling operation to the second stage input the back-end switched-capacitor circuitry can be scaled down appropriately because of the reduced back-end noise requirements. This eases the required switch resistances due to the smaller sampling capacitors and obviates the need for a front end sample-and-hold circuit.

This paper is organized as follows. Section 3.3 introduces the proposed continuous time front end architecture. Section 3.4 describes the prediction filter used in the first stage. Section 3.5 provides the circuits used in the first stage. Section 3.6 presents the experimental results followed by conclusions.

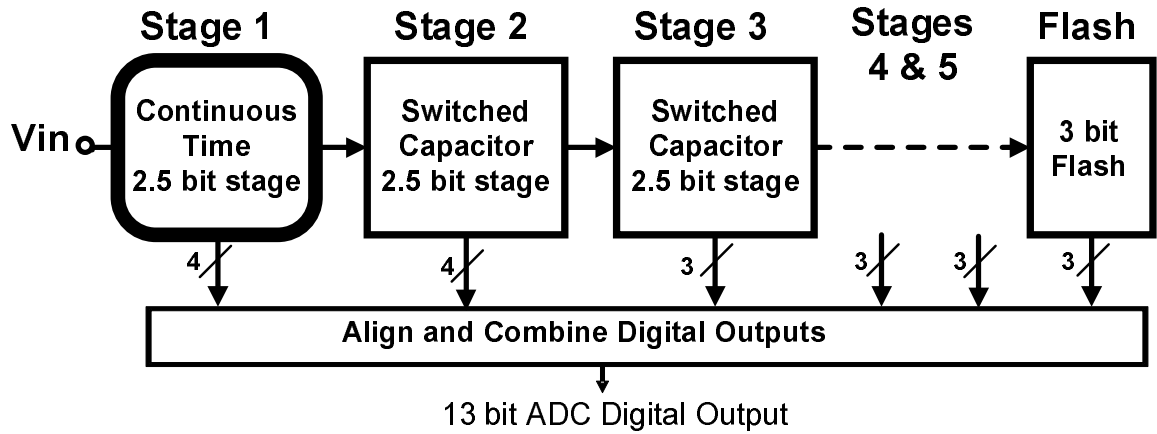


Figure 2.1: The continuous-time input pipeline ADC architecture

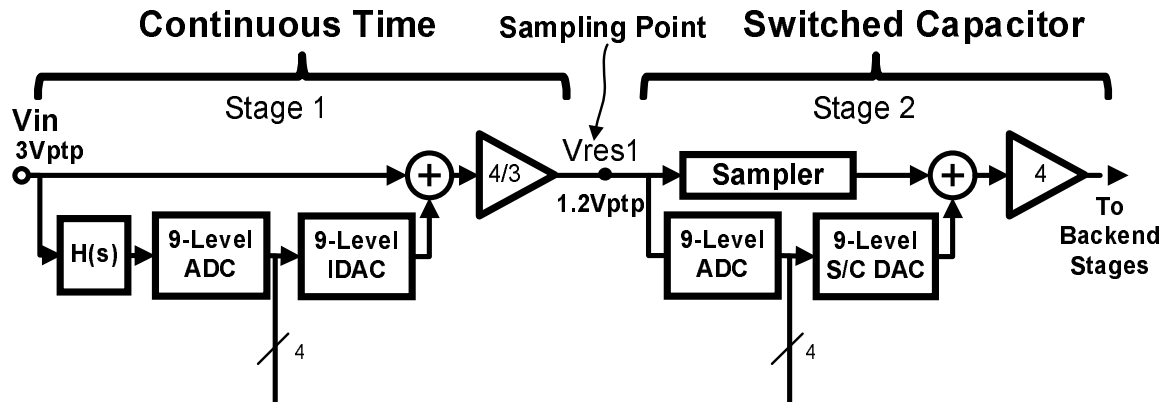
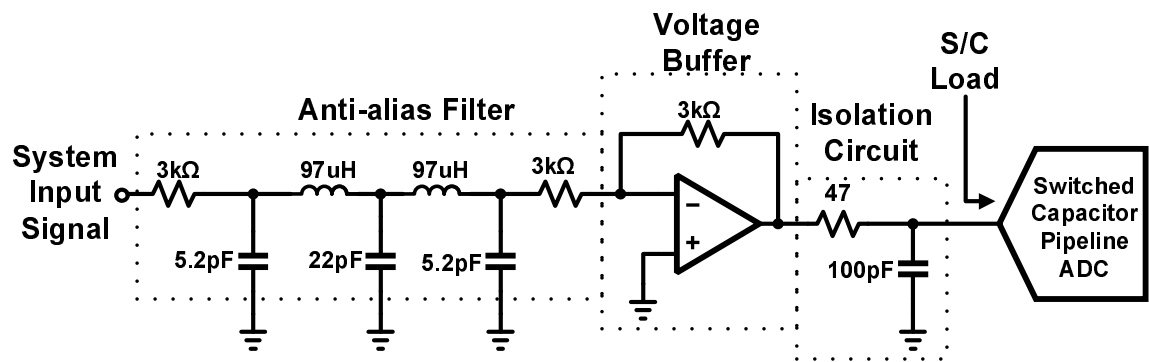



Figure 2.2: The first two stages of the continuous-time input pipeline ADC



(a) A switched-capacitor pipeline ADC applications circuit



(b) A continuous-time input pipeline ADC applications circuit

Figure 2.3: Elimination of the drive circuitry for CT input pipeline ADC

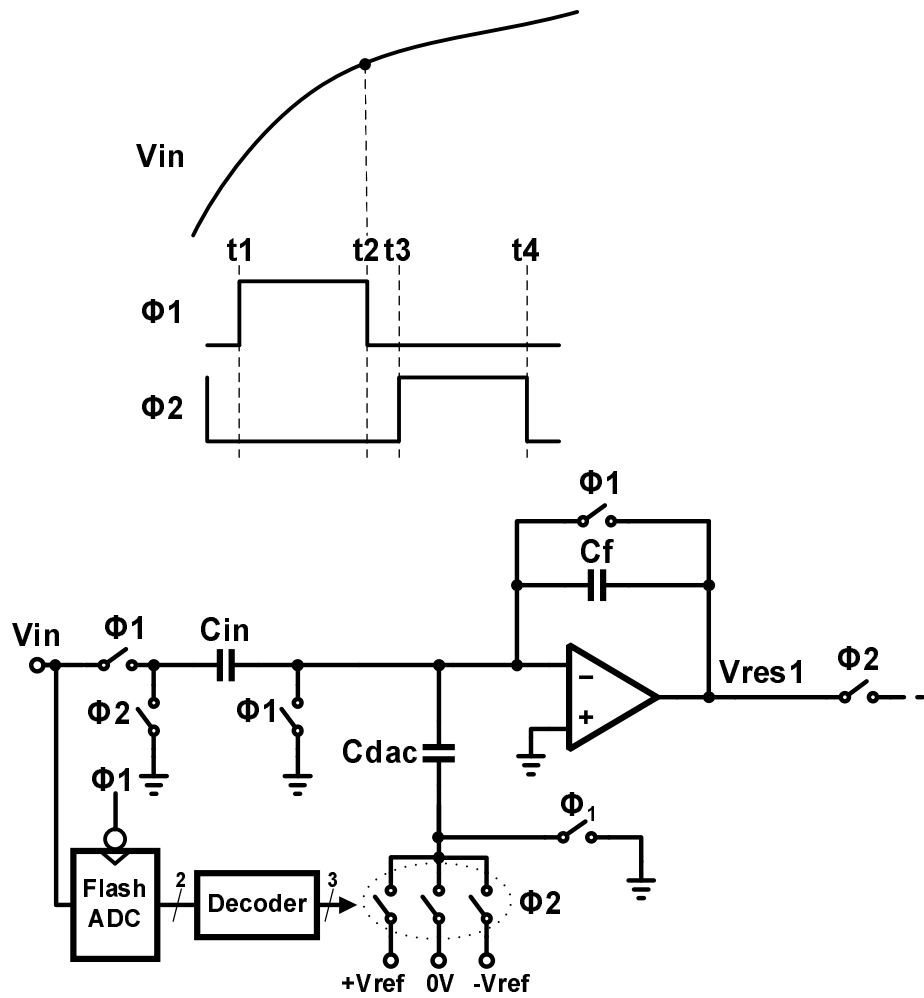


Figure 2.4: Simplified diagram of the conventional pipeline ADC first stage with timing diagram




Figure 2.5: The continuous-time input pipeline ADC first stage with timing diagram

2.2 Architecture of the continuous-time Input Pipeline ADC

Before discussing the new ADC architecture it is worth reviewing the basic operation and timing of a simplified 1.5 bit switched capacitor pipeline ADC first stage shown in Fig. 3.5. The continuous-time input signal is denoted by V_{in} . The input signal is sampled on to the capacitor C_{in} on the negative edge of ϕ_1 and the flash ADC quantizes V_{in} during the ensuing non-overlap period. Thus from Fig. 3.5 the sampling instant is given by time t_2 and the non-overlap time is between t_2 and t_3 . The flash ADC output propagates through to the multiplying DAC (MDAC) switches. At time t_3 , the positive edge of ϕ_2 , the charges on C_{in} and C_{dac} are injected into the virtual ground of the amplifier to create the amplified residual voltage, V_{res1} at the stage output. V_{res1} is then sampled by the following stage on the negative edge of ϕ_2 so that further bits can be resolved. The key point from this brief discussion is that there is a half sample delay through the first stage. This half sample delay allows the flash ADC to operate and allows the V_{res1} voltage to settle. We will now discuss the proposed continuous-time input stage that circumvents the explicit sampling of the input.

The continuous-time input pipeline ADC first stage is shown in Fig. 3.6. The continuous-time signal path from V_{in} to V_{res1} consists of an inverting amplifier whose gain is given by the ratio of resistors R_{in} and R_f . An NMOS current digital-to-analog converter, labelled as I_{dac} , completes the MDAC. Looking at Fig. 3.6 it is clear that the sampling instant occurs at the output of the first stage at time t_4 . However in order to have an accurate residue voltage that is within the second

stage input range the Idac component of the continuous-time residue, V_{res1} must be settled. The Idac thus needs to be updated at time t_3 . Therefore at time t_2 an estimate of the input voltage at time t_4 must be presented to the flash ADC. A prediction filter is used to provide this estimate within a certain error budget that is consistent with the first stage resolution. This error budget determines the filter specification which is shown in Table 3.1. A fundamental trade-off exists between bits-per-stage and the prediction filter requirements for a given signal bandwidth. As the number of bits resolved in the continuous time first stage increases the back-end switched-capacitor ADC benefits from scaling [2]. However this also causes the filter accuracy requirements to become more challenging. It was found by trial and error that resolving two and a half bits in the first stage yields a balance between prediction filter accuracy and back-end ADC scaling. A trade-off also exists between the tolerable input signal bandwidth and the required filter order. Higher input bandwidth requires a higher order filter to maintain the equiripple magnitude and phase response over the bandwidth. It was found that a fourth order filter was required to accurately predict the input signal over 80% of the nyquist band-20MHz for this ADC. The details of the prediction filter will be dealt with in Section 3.4. The obvious question at this point is why not instead introduce a phase lag in the continuous-time signal path from V_{in} to V_{res1} . By placing the prediction filter before the flash ADC the filter accuracy requirements are much reduced due to the presence of redundancy in the ADSC path [3]. Also artifacts of the prediction filter frequency response do not appear in the ADC frequency response. Mathematically this is explained concisely in Fig. 2.6 where

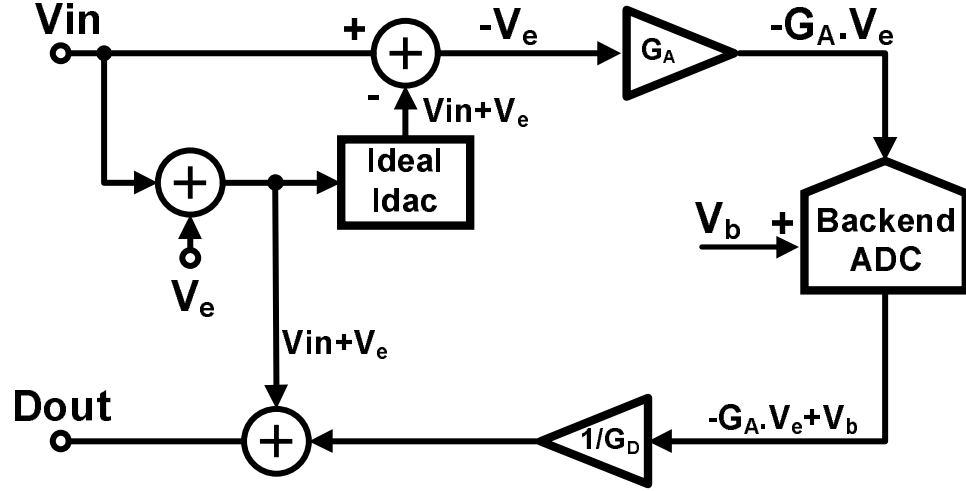


Figure 2.6: Mathematical model showing propagation of error V_e

G_A is an analog gain, G_D is a digital gain, V_b is the back-end ADC quantization error and V_e is a term containing both coarse quantization and prediction filter error. It can be seen that if $G_D = G_A$ then

$$D_{out} = V_{in} + \frac{V_b}{G_D} \quad (2.1)$$

Eq. 2.1 shows that any frequency shaping introduced by the prediction filter will not be present in the final digital representation of the input signal.

2.3 Prediction Filter

The motivation of this work is to develop a CT input architecture that efficiently digitizes lowpass wideband signals. It is assumed that the input signal is bandlimited by an anti-aliasing filter, so as to ensure that the signal level at $F_s - F_b$

Table 2.1: Prediction Filter Specifications Based on 2.5b 1st Stage

Tuning Range	$\pm 40\%$
Frequency Accuracy	$\pm 5\%$
THD	$40dB$
Noise (Output Referred)	$2mV_{rms}$
Magnitude Ripple	$\pm 0.3dB$
Phase Ripple	$\pm 2^\circ$

meets the dynamic range requirements of the ADC [6], where F_b is the signal bandwidth. For this ADC F_b can be as high as $0.8 \times \frac{F_s}{2}$. Given that the input is a lowpass, bandlimited signal an analog prediction filter at the ADC input can use continuous-time information that is usually ignored by conventional discrete time ADCs to estimate the next sample ahead of time. In this work the prediction filter zeros and poles are optimized to provide an approximate half-sample time advancement and constant gain over the bandwidth of interest. This phase advancement is valid for both sinusoidal tones and general transient signals provided they are lowpass and bandlimited, in this case to $0.8 \times \frac{F_s}{2}$.

The pole-zero plot of the minimum phase prediction filter is shown in Fig. 2.7. The pole and zero locations were optimised in Matlab to fit the desired magnitude and phase response of the prediction filter in a mean-squared sense. By weighting the mean squared error of the magnitude response and phase response portions appropriately an acceptable frequency response was achieved. The frequency response of such a filter is shown in Fig. 2.8.




Figure 2.7: Pole-zero plot

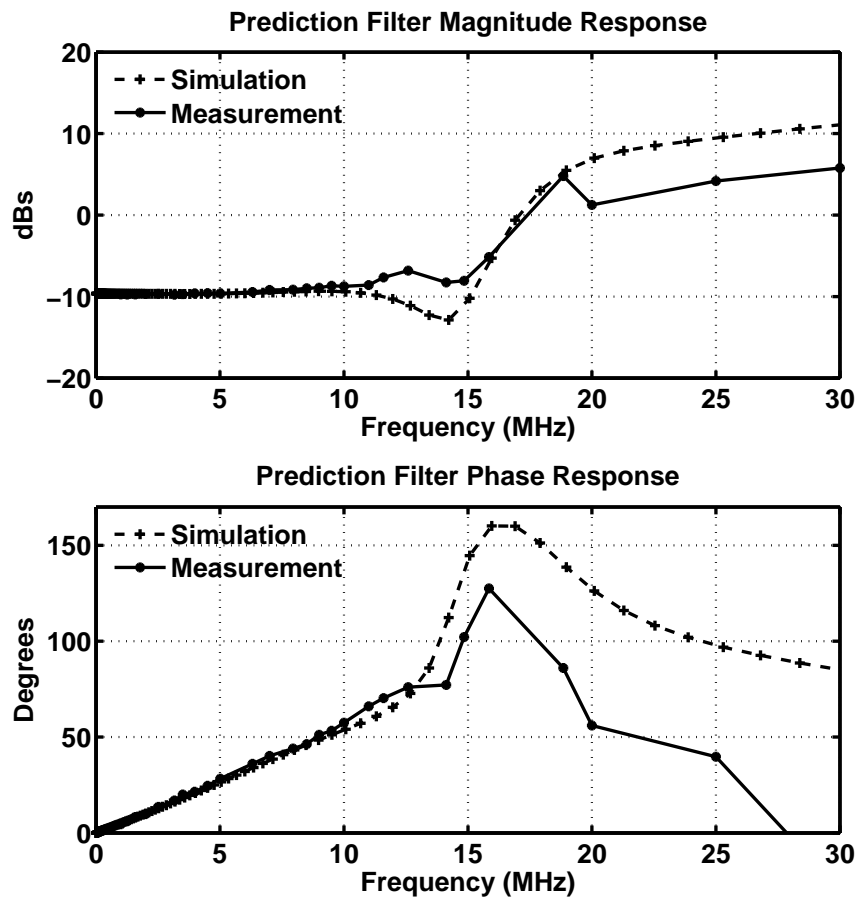


Figure 2.8: Measured and Simulated Frequency Response

The zeros are more dominant at lower frequency. This provides the required phase lead. Poles are necessary not only to make the filter physically realizable but also to limit the out of band gain that can be seen in Fig. 2.8. This out of band gain is an unavoidable consequence of the in band phase lead [7]. Fortunately, due to the practical finite gain bandwidth of the amplifier, the filter gain rolls off at higher frequencies. The prediction filter was implemented with two biquads using active-RC filter techniques. Such a realisation allows a rail-to-rail input signal. The capacitors of the continuous time prediction filter were trimmed digitally to account for resistor and capacitor component variation. The temperature variation of the RC product post calibration was within the filter frequency accuracy specification. All the flash ADCs in the pipeline ADC operate with the same reference voltages. For this pipeline ADC these voltages are 0.6V and 1.2V. Therefore in order to scale the input signal down to be consistent with the first stage flash ADC the filter has a gain scaling factor of 0.33. From Table 3.1 the noise and distortion requirements of the filter are easily achieved. One of the key challenges in the design of wideband active-RC filters is to ensure accurate pole and zero placement. Traditionally filter OTAs have been designed to have a sufficiently high DC gain and bandwidth to mitigate their effect on the filter pole and zero positions. The approach taken here was to design simple filter OTAs with finite, precise($\pm 1dB$) DC gain [10] [11] [12]. The prediction filter is implemented with two active-RC low-Q biquads shown in Fig. 2.10. Such an approach leads to low filter power and ensures that the filter is not a bottleneck in terms of sampling speed. The maximum biquad capacitor ratio is 5 and the maximum resistor ratio is 19. The prediction filter

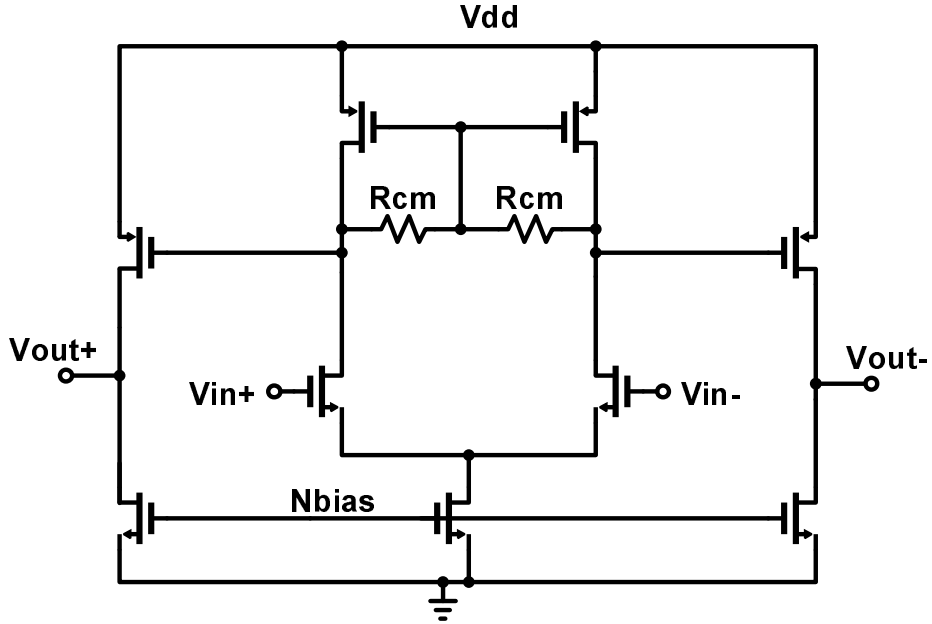


Figure 2.9: Finite gain opamp used through the prediction filter

OTA is a 2 stage OTA, shown in Fig. 2.9. The common mode of the first stage is regulated by the R_{cm} resistors and the diode connection they form. No common mode regulation was put in place for the second stage due to the low DC gain and low impedance connection to the common mode voltage. No internal OTA compensation was used. The filter OTAs were output compensated by the output node capacitance [12]. Such circuit simplifications in conjunction with relaxed requirements simplify the prediction filter circuit design.

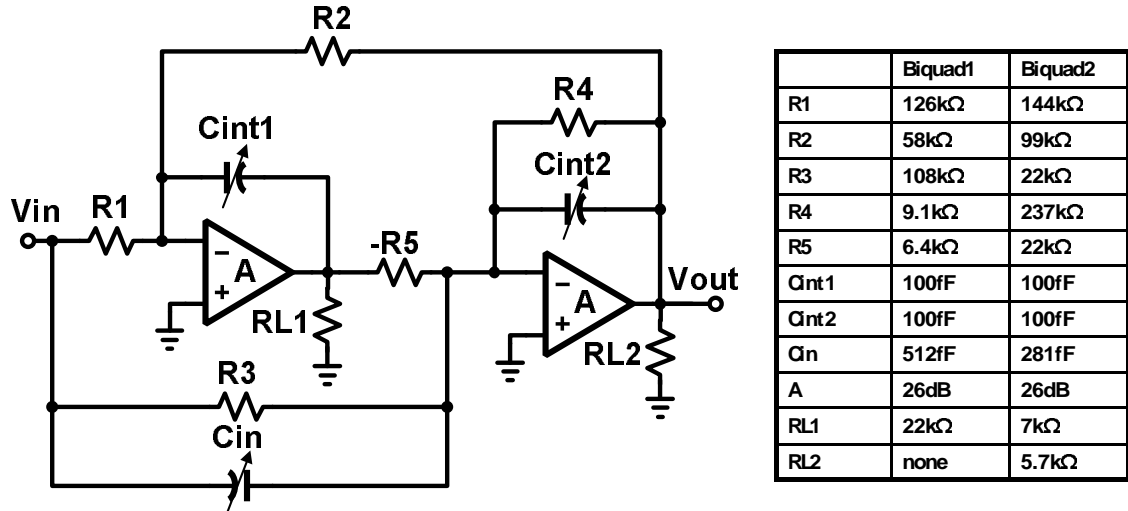


Figure 2.10: Prediction Filter Biquad

2.4 Circuit Details

The continuous-time input pipeline ADC architecture is shown in **Figs. 3.3 and 3.4**. The first two stages are implemented using two identical 9-level flash ADCs to limit the residue signal swing. This also eases the accuracy requirements of the first stage flash ADC and the prediction filter [8]. Stages 3, 4, and 5 are implemented using 7-level flash ADCs to reduce power in the remaining stages. The resulting stage transfer functions are shown in Fig. 2.11. The back end ADC is implemented using standard switched-capacitor techniques. Fig. 3.7 shows the first stage MDAC. The interstage gain is given by the I_{dac} fullscale current, the feedback resistor and the open-loop gain of the amplifier. This stage scales the input voltage range from $3V_{ptp}$ to $1.2V_{ptp}$. This improves the feedback factor of the stage. It is also worth noting that the I_{dac} does not significantly degrade the

feedback factor due to its higher impedance with respect to R_{in} and R_f . Simple digital foreground calibration similar to [9] was used to used to digitally measure the first stage analog jumps in Fig. 2.11. Once measured, these 9 digital values are used during normal operation to create the digitally calibrated output code. This scheme corrects for first stage interstage gain error and mismatch between the I_{dac} current elements. This allows the matching requirements of the I_{dac} devices to be relaxed and hence their size to be reduced. This allows a compact layout with lower parasitics at the virtual earth. A two stage miller compensated operational transconductance amplifier (OTA) , shown in Fig. 3.7 was designed for the CT first stage and reused throughout the switched-capacitor back-end ADC [4]. Such an opamp allows high DC gain and is suitable for driving the resistive load of the first stage MDAC. The OTA contains separate common mode loops for both the first and second stage. Comparators similar to [13] were used in the flash ADC.

2.5 Experimental Results

The converter was fabricated in a $0.18\mu\text{m}$ CMOS process. Fig. 3.12 shows the die photo with a total active area of about 1.6mm^2 , which contains a significant amount of white space. The most critical component from a precision point of view is the CT first MDAC stage, shown in the figure. A key point is that the CT approach to the first stage yields a compact MDAC layout. This contrasts with typical switched-capacitor ADC first stages which normally dominate the ADC layout due to the large capacitor requirements. Fig. 2.14 shows the DNL

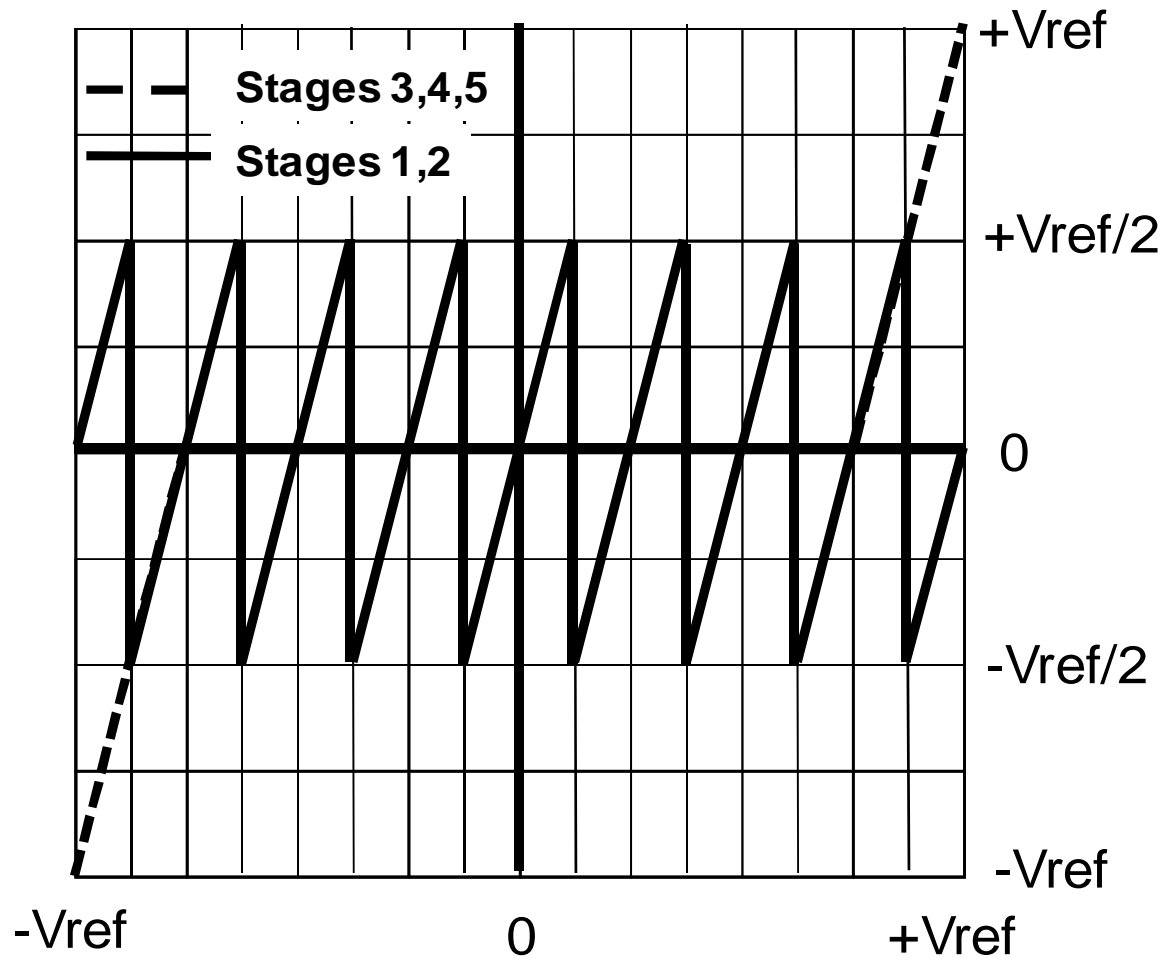


Figure 2.11: Transfer Functions of the pipeline stages

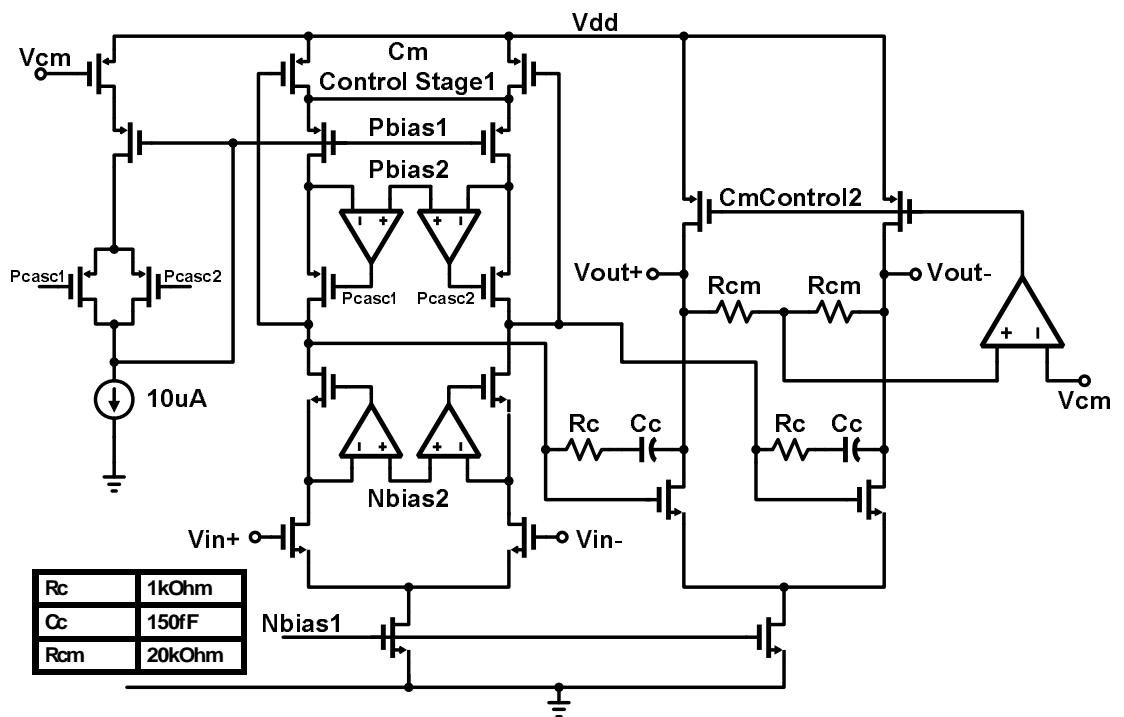


Figure 2.13: Opamp used through the ADC

and INL performance at the 11-bit level. Fig. 2.15 shows a typical single tone FFT performance plot at a sample rate of 16MSPS. Fig. 3.10 shows SNR, SNDR and SFDR over input frequency. The complete performance summary is given in Table 3.2. Unfortunately, results from this initial version of silicon were limited by a timing bug in the switched-capacitor back-end circuitry. This limited the sampling rate of the ADC and caused the power of the ADC to be excessive. This also increased the OSR from 1.25 in simulation to 4 on the bench. However the concept of a continuous time input pipeline ADC was verified. A pie chart in Fig. 3.13 shows the power breakdown of the ADC. The focus of this research was to verify the concept that the first stage residue can be processed in a continuous time fashion. The pipeline first stage presents a benign 3kOhm resistive load to the outside world. Simulated and measured prediction filter magnitude and phase frequency responses are shown in Fig. 2.8. There is reasonable agreement between the simulated and measured results in the bandwidth of interest.

2.6 Conclusions

A new ADC architecture employing a continuous-time input stage offers a more favorable trade-off between noise, distortion and amplifier power when compared to switched-capacitor methods. This architecture also results in a lower area due to back-end capacitor scaling. Consequently switch R_{on} requirements are very much eased obviating the need for bootstrap switches and sampling distortion is reduced. The benign resistive input load easily allows a rail to rail input signal to




Figure 2.14: DNL and INL Performance of the ADC versus output code




Figure 2.15: Single tone FFT plot at 1MHz input and 16MSPS




Figure 2.16: SNR, SNDR, SFDR over input frequency at 16MSPS

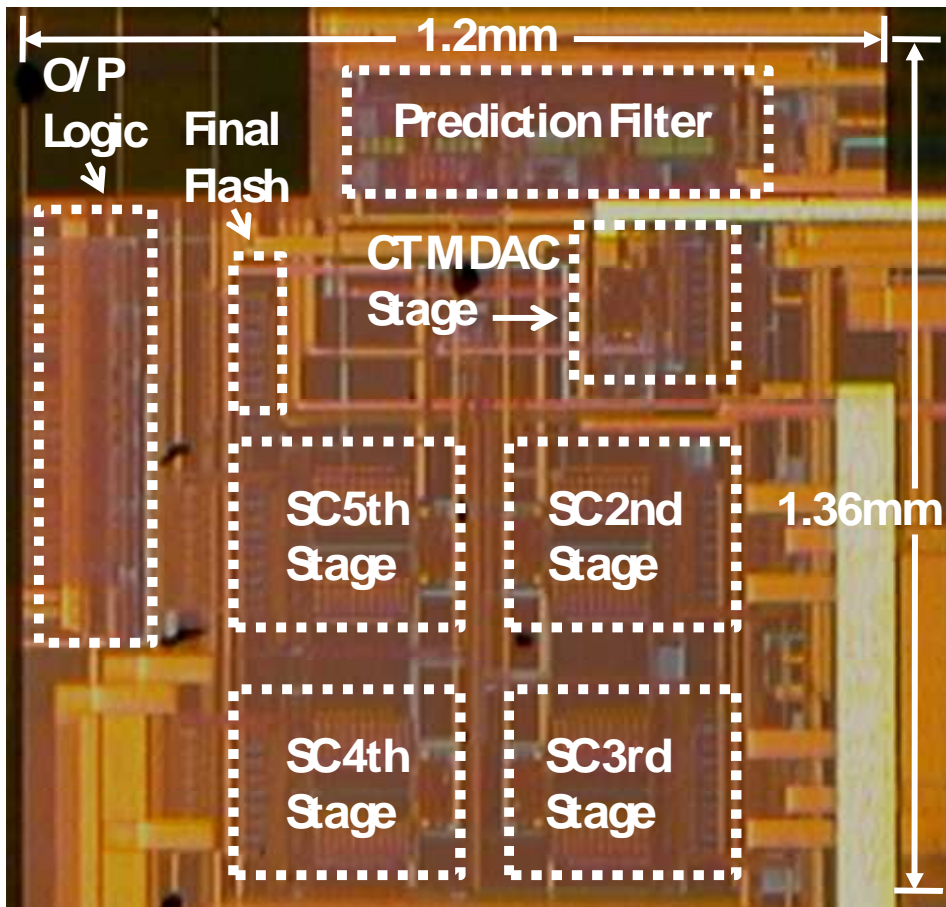


Figure 2.17: Photo of the die

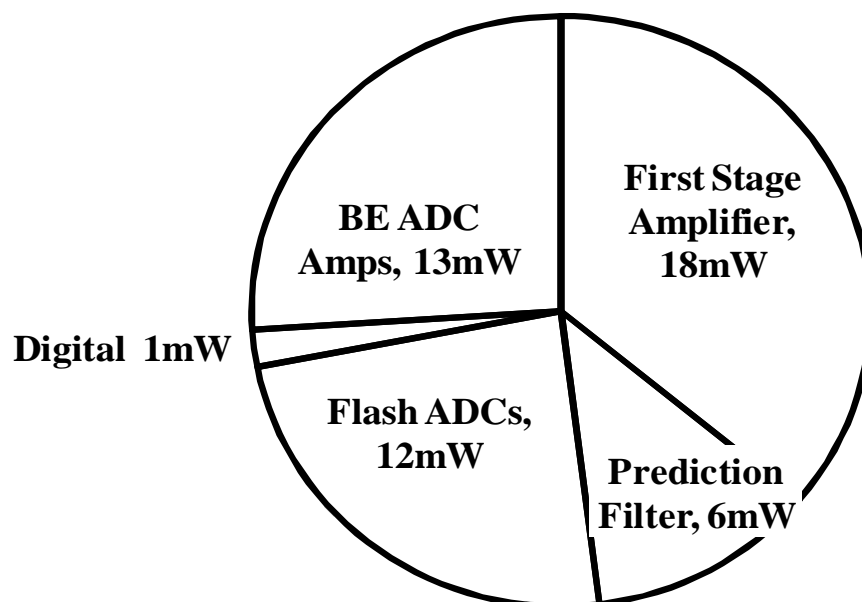
Power Consumption (50mW Total)

Figure 2.18: Power Breakdown of the prototype ADC

Table 2.2: Performance Summary

Resolution	11 bits
Conversion Rate	16MSPS
Input Range	3 V _{ptp}
SNDR	56.51dB @ Fin=1MHz
SNR	57dB @ Fin=1MHz
THD	65.95dB @ Fin=1MHz
SFDR	66dB @ Fin=1MHz
DNL	0.43/-0.68 LSBs
INL	0.78/-1.3 LSBs
Total Chip Power Consumption	50mW
Supply voltage	1.8V
Die area	1.2 × 1.36 mm ²
Technology	0.18μm CMOS

be coupled into the ADC and eliminates the requirement for ADC drive circuitry. This continuous-time input architecture has been enabled by the introduction of signal prediction within the ADC first stage. Future possibilities include moving the prediction filter into the digital domain and incorporating anti-alias filtering into the CT first stage.

ACKNOWLEDGMENT

The authors would like to thank National Semiconductor for supporting this work and providing chip fabrication. Much thanks is due to Professor Gabor Temes, Sunwoo Kwon, Naga Sasidhar Lingam and Nima Maghari for useful discussions

and layout help.

Chapter 3 – A Windowed Integration, Residue Sampling Pipeline ADC with Inherent Anti-alias Filtering

The continuous-time input pipeline nyquist rate ADC architecture is introduced. Such an approach overcomes many of the challenges associated with a pure switched-capacitor architecture. Inherent anti-alias filtering allows the possibility of eliminating costly anti-alias filters. The effect of switched-capacitor sampling distortion is reduced. Opamp noise folding is reduced for the first stage. This architecture eases the jitter requirements of the ADC clock when compared to switched capacitor pipeline ADCs. 9.85 ENOB is achieved with 21.4mW analog power from a 1.8V supply at 26MSPS. The proof of concept test chip fabricated in a 0.18 μ m CMOS process validate the effectiveness of proposed techniques.

3.1 Introduction

Pipeline ADCs have traditionally been implemented with switched-capacitor methods. Switched-capacitor methods present many challenges to ADC designers when more than 10 bit ENOB, high SFDR and low power are required [1]. These challenges are most acute in the pipeline ADC first stage, mainly due to the fact that the first stage is the main contributor to ADC noise, distortion and consequently power [3]. Sampling the input signal onto large sampling capacitors with

low distortion poses a challenge to circuit designers. Such capacitors require large sampling switches with low on resistance, R_{on} , in order that the input sampling network settles in the allotted time period. Large sampling switches demand large switch logic drivers and this increases the digital power consumption. Large sampling switches also bring with them nonlinear junction capacitance which can cause the sampled signal to be distorted. Sampling is further complicated by input signal bond-wire inductance which may cause ringing on the input signal sampling network. The large ADC capacitive load also presents a challenge for the reference buffers which need to settle to a required precision. A fundamental requirement of switched capacitor circuitry is that a continuous time anti-alias filter must precede the switched-capacitor sampler. This problem, while well-known is restated in Section 3.2 for completeness.

Another downside of switched capacitor circuitry is that of opamp noise folding. The key point is that switched-capacitor opamp bandwidth requirements are typically a multiple of the clock frequency, depending on the settling precision requirements. This causes the opamp noise to fold down into the nyquist band-degrading the ADC noise performance [14]. The continuous time approach also significantly eases the issue of opamp noise folding because the first stage opamp bandwidth requirements are very much eased.

This work addresses the above challenges by implementing a continuous-time(CT) 2.5b pipeline first stage with inherent anti-aliasing filtering. By choosing an OSR of 2 and adopting simple linear phase filtering techniques significant anti-alias filtering is achieved, avoiding the need for an explicit anti-alias filter in many applica-

tions. The CT first stage is followed by a scaled down switched capacitor pipeline back-end ADC. For every extra bit resolved in the CT first stage the back-end switched-capacitors can be scaled down by a factor of four [4]. This new architecture is shown in Fig. 3.3. Such an approach leads to lower power, reduced sampling distortion, lower capacitance area and easily allows rail-to-rail input swing. The resistive input load of the continuous-time first stage is also more favorable when compared to a dynamically changing switched-capacitor input load [5]. By resolving two bits in the first stage in a continuous-time fashion and moving the sampling operation to the second stage input the back-end switched-capacitor circuitry can be scaled down appropriately because of the reduced back-end noise requirements. This eases the required switch resistances due to the smaller sampling capacitors and obviates the need for a front end sample-and-hold circuit.

This paper is organized as follows. Section 3.2 summarizes the challenges associated with anti-alias filters. Section 3.3 introduces the proposed continuous time front end architecture. Section 3.4 describes the prediction filter used in the first stage. Section 3.5 provides the circuits used in the first stage. Section 3.6 presents the experimental results followed by conclusions.

3.2 The Anti-alias Filter Problem

The switched-capacitor sampler at the front end of traditional switched-capacitor pipeline ADCS is an impulse sampler. A key property of an ideal impulse sampler is that all frequencies are aliased with equal fidelity and gain down to the

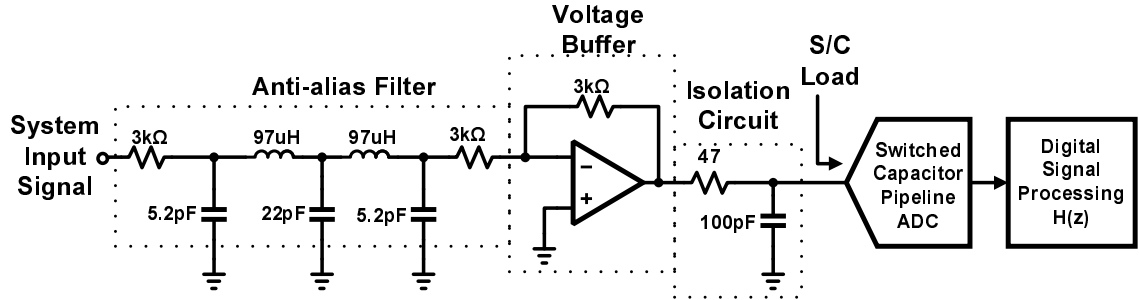


Figure 3.1: A typical pipeline ADC application circuit [16]

nyquist bandwidth [15]. This means that an anti-aliasing filter must precede the switched-capacitor impulse sampler. Additionally the anti-alias filter must be able to drive the pipeline switched-capacitor load without degrading the noise and distortion performance of the system. These requirements lead to a typical pipeline applications circuit as shown in Fig. 3.1.

The applications circuit is composed of an anti-alias filter, a voltage buffer, an isolation circuit and the switched-capacitor ADC itself. For this example it is assumed that the ADC is followed by a brick wall digital filter denoted by $H(z)$. The digital filter ideally removes signal content from Fb to $F_s/2$. Any out of band signal that is aliased into the band from $0Hz$ to Fb cannot be removed once it is sampled because it cannot be distinguished from the wanted signal. Working back from the ADC input the isolation circuit provides the discrete packets of charge required by the ADC input load and isolates the opamp output from the switched-capacitor load. The opamp buffers the signal from the anti-alias filter. The job of the anti-alias filter is as follows: the pass-band must pass the desired

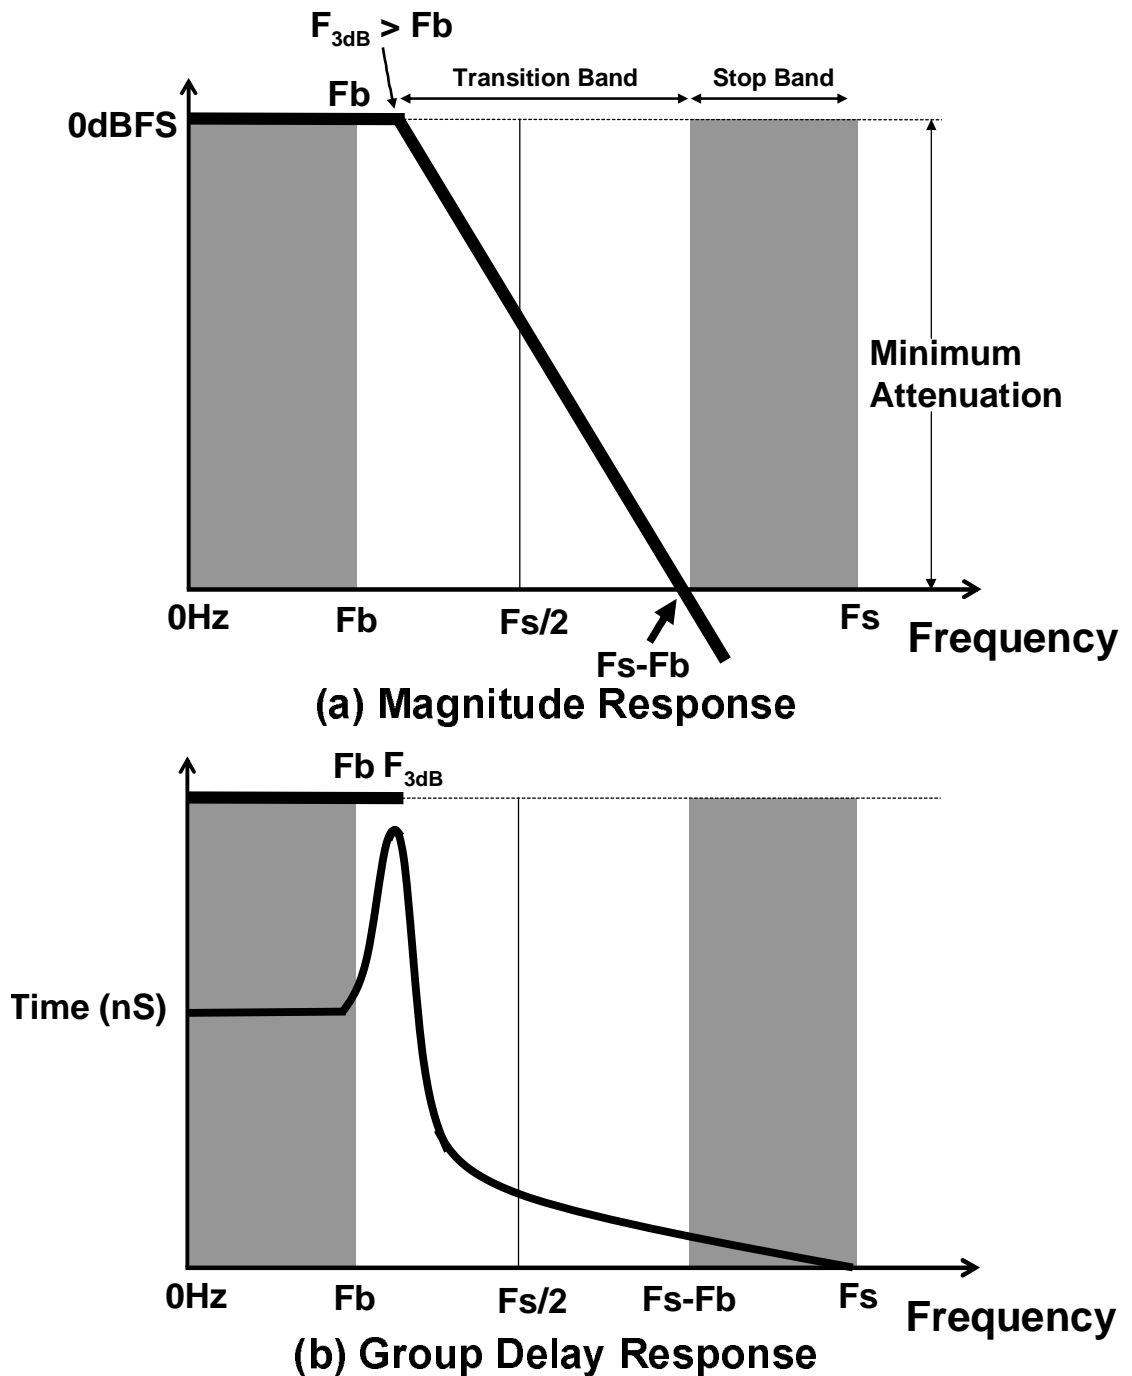


Figure 3.2: Magnitude and Group Delay Response requirements of an anti-alias filter

input bandwidth. The stop-band must attenuate any interferer beyond $F_s - F_b$ sufficiently such that it will not affect the system performance when aliased into the pass-band due to sampling by the A/D converter. F_b is usually a lot less than the $F_s/2$ rate. This over-sampling is required to provide the requisite transition band where the filter can roll off to some minimum attenuation at $F_s - F_b$. This is shown in Fig. 3.2(a). Lower over-sampling ratio increases the required filter order because the transition band becomes smaller. This is bad for two reasons. Firstly it increases the filter cost. Secondly it is a further challenge to design the anti-alias filter to have a linear phase response. Linear phase is quantified by how constant the group delay is over the band of interest. A linear phase response is the same as having constant group delay. Linear phase is extremely important to maintain the shape of pulses and pulse type signals being passed through the filter. Video applications for example require a linear phase response. It turns out that peaking in the group delay occurs at the pass-band edge for analog filters due to the fast roll-off. As a result the anti-alias filter 3dB frequency is often placed well beyond the signal band of interest to avoid phase distortion in the wanted signal band. This further reduces the frequency range for the filter to roll off. This can be seen in Fig. 3.2(b) where group-delay peaking coincides with the 3dB point which has been placed beyond F_b . Thus, there is a tightly coupled trade-off between OSR, anti-alias filter order, group delay response and system cost. Fig. 3.1 shows the anti-alias filter implemented external to the chip. Such external circuitry is expensive and takes up PCB area. An alternative would be to integrate this circuitry with the pipeline ADC. Such an approach would require

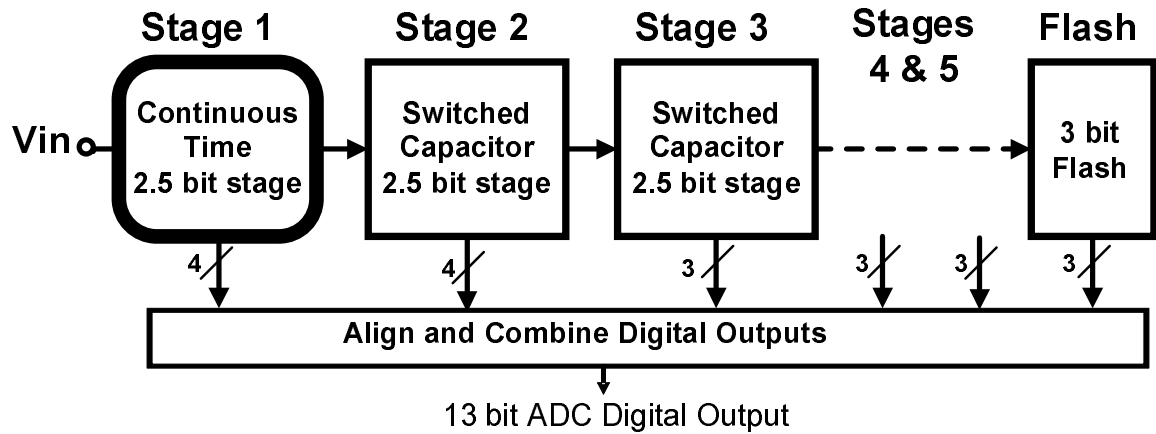


Figure 3.3: The continuous-time input pipeline ADC architecture

active circuitry and would contribute to a degradation in noise and distortion and an increase in system power. Another key problem would be that of driving the pipeline ADC with active circuitry without introducing significant distortion and noise. The anti-alias filtering problem further motivates the move to a CT pipeline ADC first stage with inherent anti-alias filtering.

3.3 Architecture of the continuous-time Input Pipeline ADC

Before discussing the new ADC architecture it is worth reviewing the basic operation and timing of a simplified 1.5 bit switched capacitor pipeline ADC first stage shown in Fig 3.5. The continuous-time input signal is denoted by V_{in} . The input signal is sampled on to the capacitor C_{in} on the negative edge of ϕ_1 and the flash ADC quantizes V_{in} during the ensuing non-overlap period. Thus from Fig. 3.5 the sampling instant is given by time t_2 and the non-overlap time is

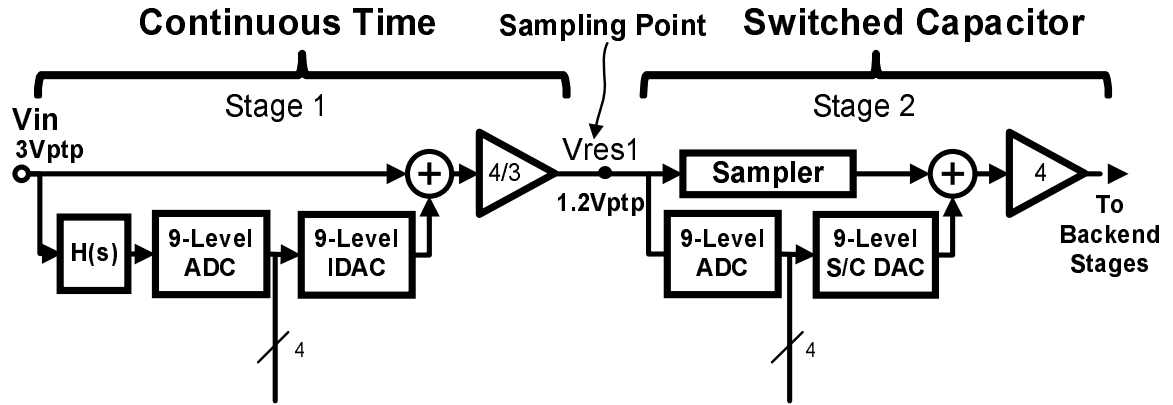


Figure 3.4: The first two stages of the continuous-time input pipeline ADC

between t_2 and t_3 . The flash ADC output propagates through to the multiplying DAC (MDAC) switches. At time t_3 , the positive edge of ϕ_2 , the charges on C_{in} and C_{dac} are injected into the virtual ground of the amplifier to create the amplified residual voltage, V_{res1} at the stage output. V_{res1} is then sampled by the following stage on the negative edge of ϕ_2 so that further bits can be resolved. One key point from this brief discussion is that there is a half sample delay through the first stage. This half sample delay allows the flash ADC to operate and allows the V_{res1} voltage to settle. Secondly it is important to note that any jitter on the negative sampling edge of ϕ_1 will directly translate to jitter induced noise on V_{in} . This jitter induced noise contribution is directly input referred to the ADC input V_{in} . We will now discuss the proposed continuous-time input stage that circumvents the explicit impulse sampling of the input.

The continuous-time input pipeline ADC first stage is shown in Fig. 3.6. Similar to the SC implementation shown in Fig. 3.5, the proposed CT stage also consists of two paths: (a) the input signal path denoted by CT path and (b) the




Figure 3.5: Simplified diagram of the conventional pipeline ADC first stage with timing diagram

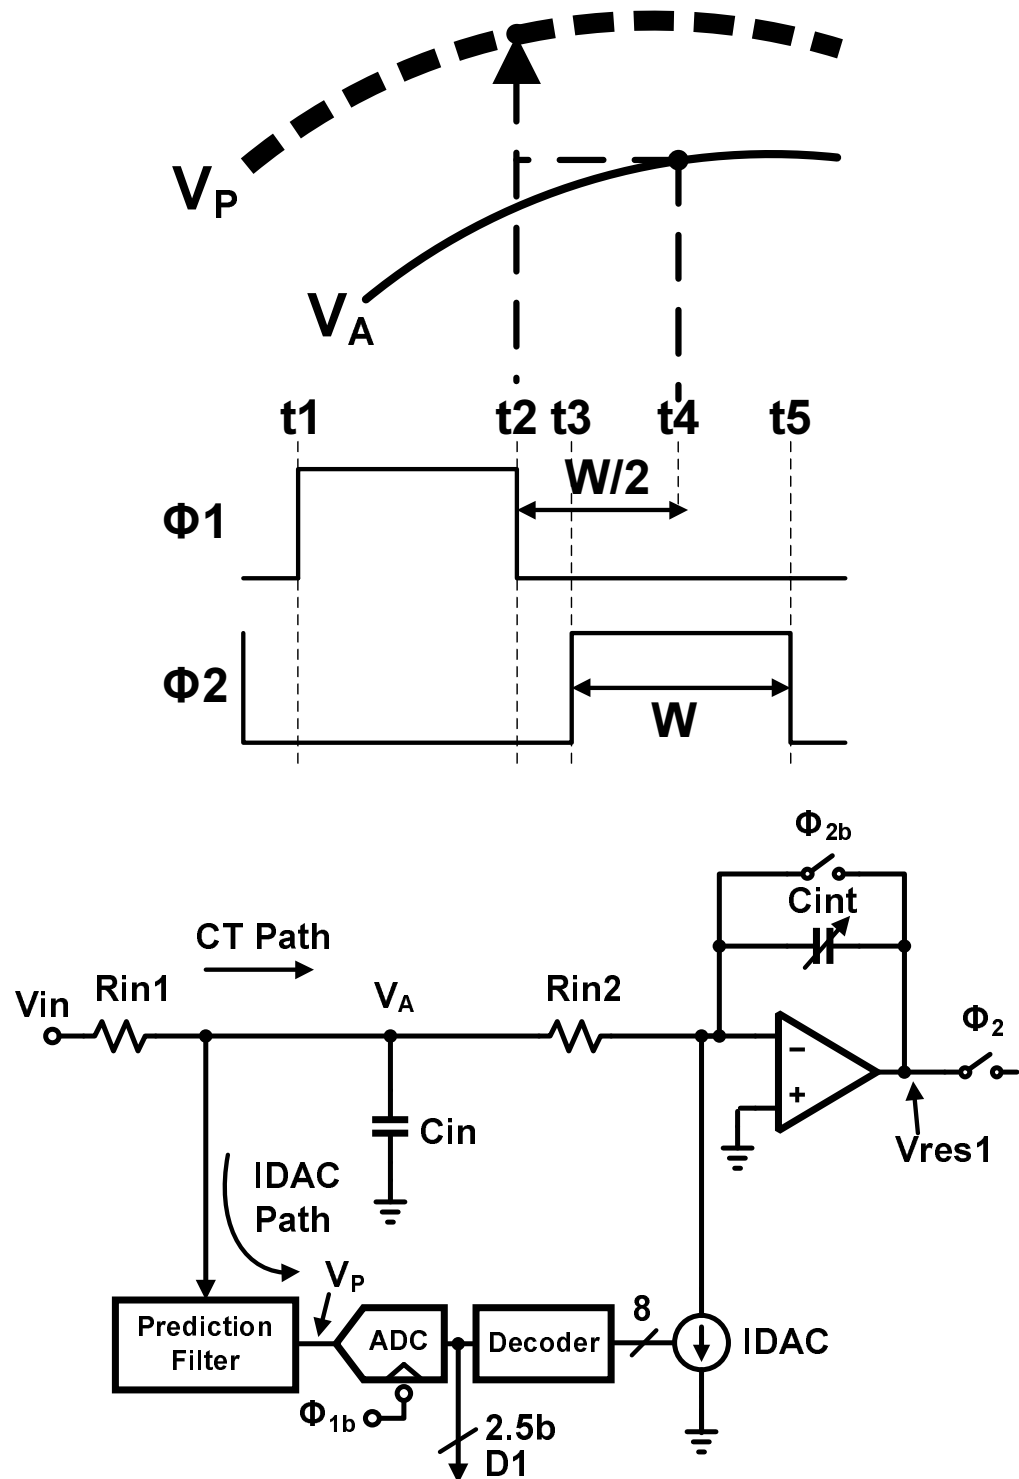


Figure 3.6: The continuous-time input pipeline ADC first stage with timing diagram

digital-to-analog converter (DAC) path shown as IDAC path. These two paths are combined in a way to generate the needed residue, $V_{res1} = G_A(V_{in} - D1.V_{ref})$, at the end of the $\phi2$ phase. G is the interstage gain. We will first describe the operation of the CT path with the aid of the timing diagram shown in Fig. 3.6. The CT path is made up of the windowed active-RC integrator that accumulates the input signal during $\phi2$ (of width W) and is reset during $\phi2b$. It can be shown that this windowed integrator has the following three properties: (1) The DC gain, G_A can be approximated to be $W/C(R_{in1} + R_{in2})$, (2) The input-output transfer function exhibits a sinc shape with spectral nulls at frequency multiples of $1/W$ and (3) the time delay is $W/2$. Such a sampler is referred to as a windowed integration sampler in the literature [17]. In our design the DC gain, G_A is nominally set to $4/3$ and is digitally calibrated, as explained later, to mitigate interstage gain errors. The sinc filtering and RC filtering combine to provide anti-aliasing and measurements indicate $14dB$ of attenuation at the sampling frequency. Normally for a 2.5b stage the gain, G_A is set to a value of 4. However in this architecture it is easy to couple the fullscale signal into the ADC through the input resistors. This yields an input range of $3.6V_{ptp}$ using a $1.8V$ supply. However the gain of the first stage was then scaled down to $4/3$ to yield an output range of $1.2V_{ptp}$ differential. This range is consistent with the amplifier linear output swing. Reducing the gain of the stage also improves the feedback factor of the stage. The details of the IDAC path will now be presented. The purpose of IDAC path is to coarsely quantize the input and generate the analog equivalent of the quantized digital code, $D1$. Because the windowed integrator introduces a delay of only $W/2$ in

the CT path, it is required to introduce a phase lead of $W/2$ to the input of the sub-ADC. Such a phase lead guarantees that residue, V_{res1} , corresponds to the amplified quantization error at time t_3 (see Fig. 3.6). To this end, the input signal is first processed by the prediction filter to introduce a phase lead of about $W/2$ and then quantized using the ADC at the end of phase ϕ_1 . The analog equivalent of the digital code is generated using the current-mode DAC, IDAC in Fig. 1. The IDAC value is updated during the non-overlapping time while the stage is still in reset mode. During ϕ_2 the IDAC code remains fixed and the residue is accumulated at V_{res1} . At the end of ϕ_2 V_{res1} is sampled by the back-end ADC. The specification of the prediction filter, detailed in Table 3.1, is relaxed due to the presence of redundancy in the ADC [3]. The input impedance of the prediction filter is relatively high which means that it can simply be tied to node VA without having any impact on the operation of the CT stage. An unavoidable consequence of in-band phase lead is the out of band filter gain. This out of band gain is bad because potentially an out of band input signal can excite the filter and cause the first stage output to saturate. This can happen because the out of band frequencies get gained by the IDAC path whereas they get attenuated by the CT path. This problem is mitigated by the ADC redundancy and the RC filter created by R_{in1} , R_{in2} and C_{in} , which band-limits the input signal so that out of band signals do not excite the prediction filter.

The response of this new ADC architecture to clock jitter will be described next. Fig. 3.6 shows that any jitter on either the positive or negative edge of ϕ_2 will cause the window width, W to vary from cycle to cycle. This causes a jitter

induced noise in the ADC output. The rms value of the jitter induced noise is proportional to the value of the window width rms jitter and the amplitude of the window integrated signal, which for this ADC is the residue, $(V_{in} - V_{dac})$. The exact reason for this is that W determines the interstage gain of the CT stage as follows $G_A = W/C(R_{in1} + R_{in2})$. Thus from Fig. 3.6 jitter on the ϕ_2 clock edges modulates the interstage gain of the first stage and thus the first stage residue is subject to this jitter. This is very different to a switched capacitor pipeline ADC. One downside is that this new residue sampling method is subject to the jitter on **both** clock edges. However a key advantage when compared to the switched-capacitor ADC is that it is the quantity $(V_{in} - V_{dac})$ that is effected- not the input signal V_{in} directly. $(V_{in} - V_{dac})$ is smaller in amplitude than V_{in} . Also the precision requirements of this means that the noise created by jitter will be smaller when referred to V_{in} .

A fundamental trade-off exists between bits-per-stage and the prediction filter requirements for a given signal bandwidth. As the number of bits resolved in the continuous time first stage increases the back-end switched-capacitor ADC benefits from scaling [2]. However this also causes the filter accuracy requirements to become more challenging. It was found by trial and error that resolving two and a half bits in the first stage yields a balance between prediction filter accuracy and back-end ADC scaling. A trade-off also exists between the tolerable input signal bandwidth and the required filter order. Higher input bandwidth requires a higher order filter to maintain the equiripple magnitude and phase response over the bandwidth. It was found that a fourth order filter was required to accurately

Table 3.1: Prediction Filter Specifications Based on 2.5b 1st Stage

Tuning Range	$\pm 40\%$
Frequency Accuracy	$\pm 5\%$
THD	$40dB$
Noise (Output Referred)	$2mV_{rms}$
Magnitude Ripple	$\pm 0.3dB$
Phase Ripple	$\pm 2^\circ$

predict the input signal over the full nyquist band. In fact experimental results show that this ADC is functional with an input signal up to 1.6 times the nyquist rate (see Fig. 3.10). In order to expedite the design and evaluation of the CT methods proposed herein the prediction filter has been reused from a previous published chip design [18]. At this point it is important to point out the key differences between this work and [18]. This work improves upon the work of [18] in two key ways. Firstly this work introduces inherent anti-alias filtering into the ADC. Secondly this work also offers lower jitter induced noise when compared to its impulse sampling switched-capacitor counterpart.

Reference [18] shows that any frequency shaping introduced by the prediction filter will not be present in the final digital representation of the input signal. The anti-alias filtering is provided solely by the CT path in Fig. 3.6.

3.4 Prediction Filter

Section 3.2 clearly showed that over-sampling is a clear strategy when trying to implement the anti-alias filter. The inherent anti-alias filtering strategy in this ADC is no different. A modest OSR of 2 is used. This is seen as a simple trade-off between usable bandwidth and sample-rate. The target sample rate for this CT ADC was 26MHz. This implies a 3dB bandwidth of 6.25MHz. The input RC filter of Fig. 3.6 (R_{in1} , R_{in2} and C_{in}) provides enough filtering to ensure that an out of band full-scale input signal at V_{in} does not saturate the residue voltage. Given that the signal at V_A is a lowpass, bandlimited signal an analog prediction filter at the ADC input can use continuous-time information that is usually ignored by conventional discrete time ADCs to estimate the next sample ahead of time. With this in mind the prediction filter that was described in Chapter 2 was reused in this ADC. The reason it can be reused is as follows. The filter of Chapter 2 was designed to give a half sample phase lead at 50MHz over a bandwidth of 20MHz. This half sample rough linear phase lead corresponds to 10nS. However the ADC described in this chapter requires a quarter sample linear phase lead at 25MSPS over a lower bandwidth. This in a nutshell allows the reuse of the same prediction filter of Chapter 2.

3.5 Circuit Details

The continuous-time input pipeline ADC architecture is shown in **Figs. 3.3 and 3.4**. The first two stages are implemented using two identical 9-level flash

ADCs to limit the residue signal swing. This also eases the precision requirements of the first stage flash ADC and the prediction filter [8]. Stages 3, 4, and 5 are implemented using 7-level flash ADCs to reduce power in the remaining stages. The back end ADC is implemented using standard switched-capacitor techniques. The primary contribution of this paper is the CT first stage. A telescopic operational transconductance amplifier (OTA) is used in the ADC CT first stage as shown in Fig. 3.9. The main motivations behind such an amplifier are power dissipation, signal swing, open loop gain and ease of compensation. The telescopic OTA is a power efficient amplifier mainly because it contains a single stage and thus the static power dissipation is limited to just one tail pair. The output signal swing is consistent with the $1.2V_{ptp}$ differential required for this ADC. This telescopic ADC was used without any gain boosting. This yielded an open loop DC gain of 52dB. This gain was sufficient to suppress nonlinearity in the amplifier. The closed loop interstage gain error was corrected for with the gain calibration technique to be described later. The telescopic amplifier is compensated by its own capacitive load. Choice of an NMOS differential pair maximizes the transconductance of the stage at the cost of increased flicker noise when compared to a PMOS differential pair. Fig. 3.7 shows the first stage MDAC. This stage scales the input voltage range from $3.6V_{ptp}$ to $1.2V_{ptp}$ at the output. This improves the feedback factor of the stage and thus also the noise gain of the amplifier. It is also worth noting that the IDAC does not significantly degrade the feedback factor due to its higher impedance with respect to R_{in} and R_f . An important consideration in pipeline ADC design is inter-stage gain accuracy. One option is for circuit designers to

ensure that the analog gain, G_A is as close to some desired nominal gain, typically 2^{M-1} , where M is the number of bits resolved in the stage. This is achieved by making the open-loop amplifier gain as high as possible [1]. Alternatively a foreground gain calibration technique can be used to, instead, measure the analog IDAC steps of Figure 2.11 and match the digital codes to these IDAC steps [19].

The MDAC precision requirements are most acute in the first stage where $V_{res1} = G_A(V_{in} - D1.V_{ref})$. G_A can be approximated to be $W/C(R_{in1} + R_{in2})$.

It is also further compounded by the PVT variation of the in this new architecture. The key issue is that the transfer function jumps due to the IDAC code changes need to be either calibrated in an analog or digital fashion.

Two foreground techniques are used to set this problem aside. Firstly the integration capacitor, C_{int} is calibrated digitally using switches to ensure the stage does not suffer from loss in dynamic range. This calibration centers the gain to $\pm 5\%$ of the nominal value. Secondly a simple digital calibration routine similar to [19] is used to digitally measure the IDAC jumps using the back-end ADC. A simple analog loop, shown in Fig. 3.8 is used to match the IDAC full-scale current to the poly input resistance.

Once measured these 9 digital values are used during normal operation to create the digitally calibrated output code. This scheme corrects for first stage interstage gain error and mismatch between the IDAC current elements. This allows the matching requirements of the IDAC devices to be relaxed and hence their size to be reduced. This allows a compact layout with lower parasitics at the virtual earth. No attempt was made in this research to make the gain temperature stable.

One additional technique that could be used to do this would be to make the pulse width, W to track temperature variations in the RC product- thus keeping $W/C(R_{in1} + R_{in2})$ constant over temperature. Comparators similar to [13] were used in the flash ADC.

3.6 Experimental Results

The converter was fabricated in a $0.18\mu\text{m}$ CMOS process. Fig. 3.12 shows the die photo with a total active area of about 1.9mm^2 . The most critical component from a precision point of view is the CT input stage, shown in the figure. A key point is that the CT approach to the first stage yields a compact first stage MDAC layout. This contrasts with typical switched-capacitor ADC first stages which normally dominate the ADC layout due to the large capacitor requirements. Fig. 3.10 shows a typical single tone FFT performance plot at a sample rate of 26MSPS along with SNDR, SFDR variation with respect to input signal frequency. It is clear that this ADC can process input signals well beyond the nyquist rate. The complete performance summary is given in Table 3.2. A pie chart in Fig. 3.13 shows the analog power breakdown of the ADC. The digital power was 5.3mW which included the digital output drivers. The focus of this research was to verify the concept that the first stage residue can be processed in a continuous time fashion. The pipeline first stage presents a benign 3kOhm resistive load to the outside world. Simulated and measured prediction filter magnitude and phase frequency responses are shown in Fig. 2.8. There is reasonable agreement between the simulated and measured

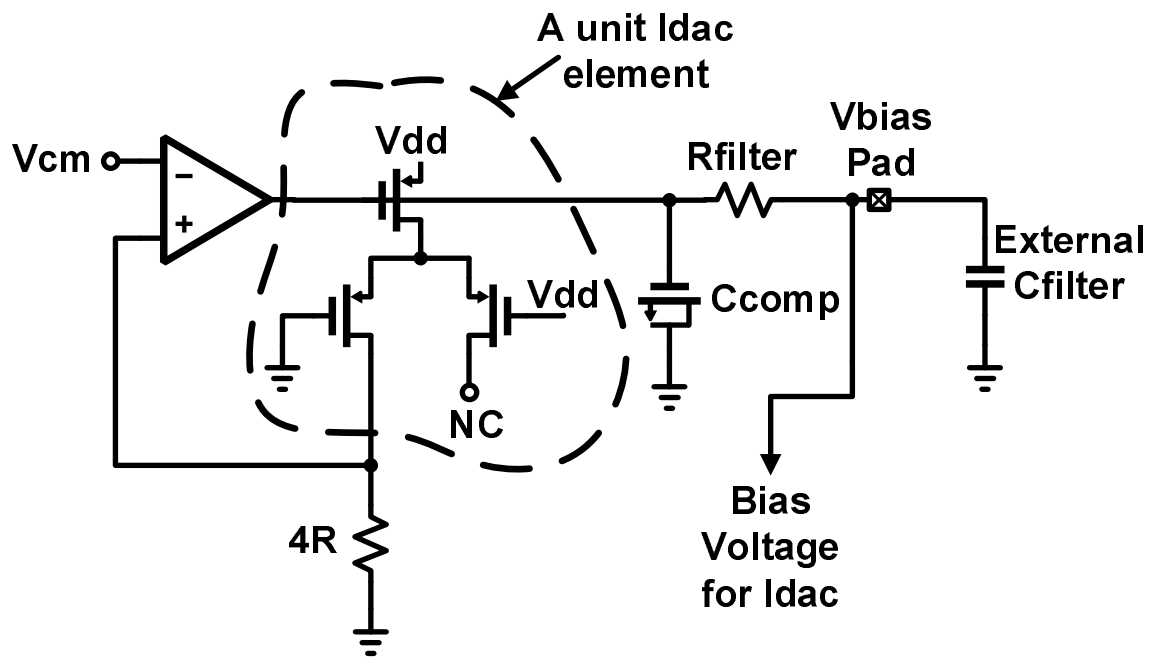


Figure 3.8: Circuitry used to match the IDAC fullscale current to the unit poly resistance, R

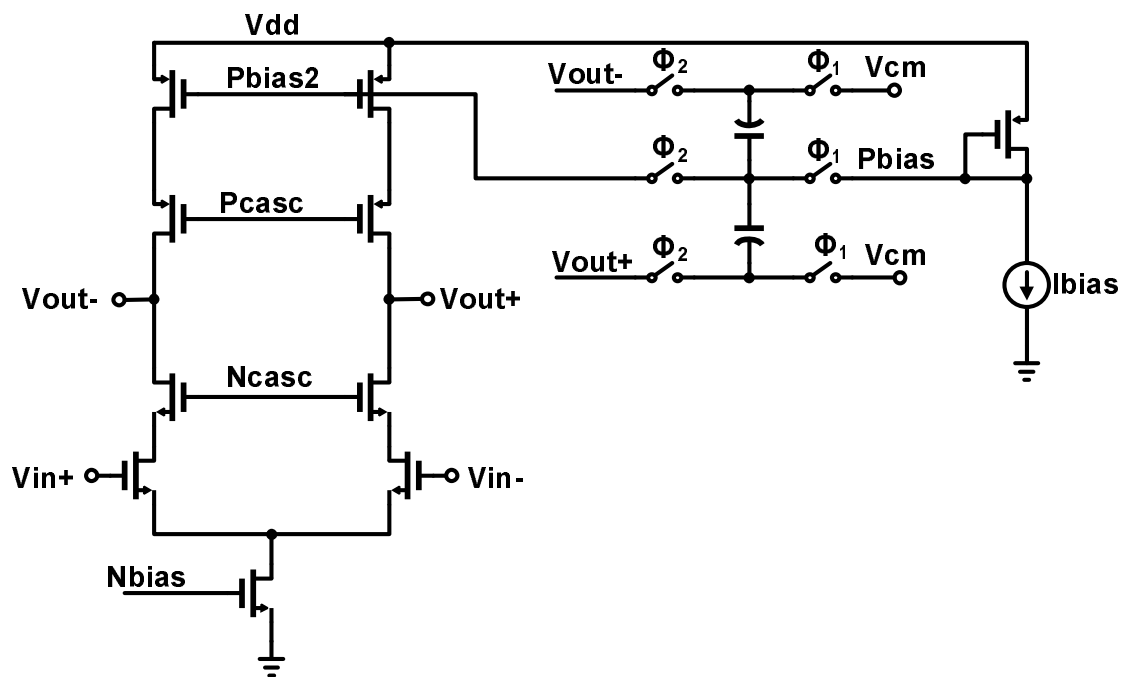


Figure 3.9: First Stage CT ADC Opamp with switched-capacitor common mode feedback

Table 3.2: Performance Summary

Resolution	11 bits
Conversion Rate	26MSPS
Input Range	3.6V _{ptp}
SNDR	61.13dB @ $F_{in}=1\text{MHz}$
SNR	62.76dB @ $F_{in}=1\text{MHz}$
SFDR	67dB @ $F_{in}=1\text{MHz}$
ENOB	9.85 bits
3dB Input Bandwidth	10MHz
Anti-alias attenuation @ $F_s=26\text{MHz}$	14dB
Analog Power	21.4mW
Digital Power	5.3mW
Supply voltage	1.8V
Die area	$1.4 \times 1.36 \text{ mm}^2$
Technology	0.18 μm CMOS

results in the bandwidth of interest. The overall ADC frequency response is shown in Fig. 3.11. It is clear that there is roughly 40dBs/decade filter roll-off and it is also clear that the sinc filtering is a clear component of the frequency response with nulls at multiples of 60MHz. This yields 14dB of filter attenuation at 26MHz. This filtering can be optimized further by increasing the value of the window with respect to the clock period. This would reduce the value of the first null frequency and yield more out of band attenuation.

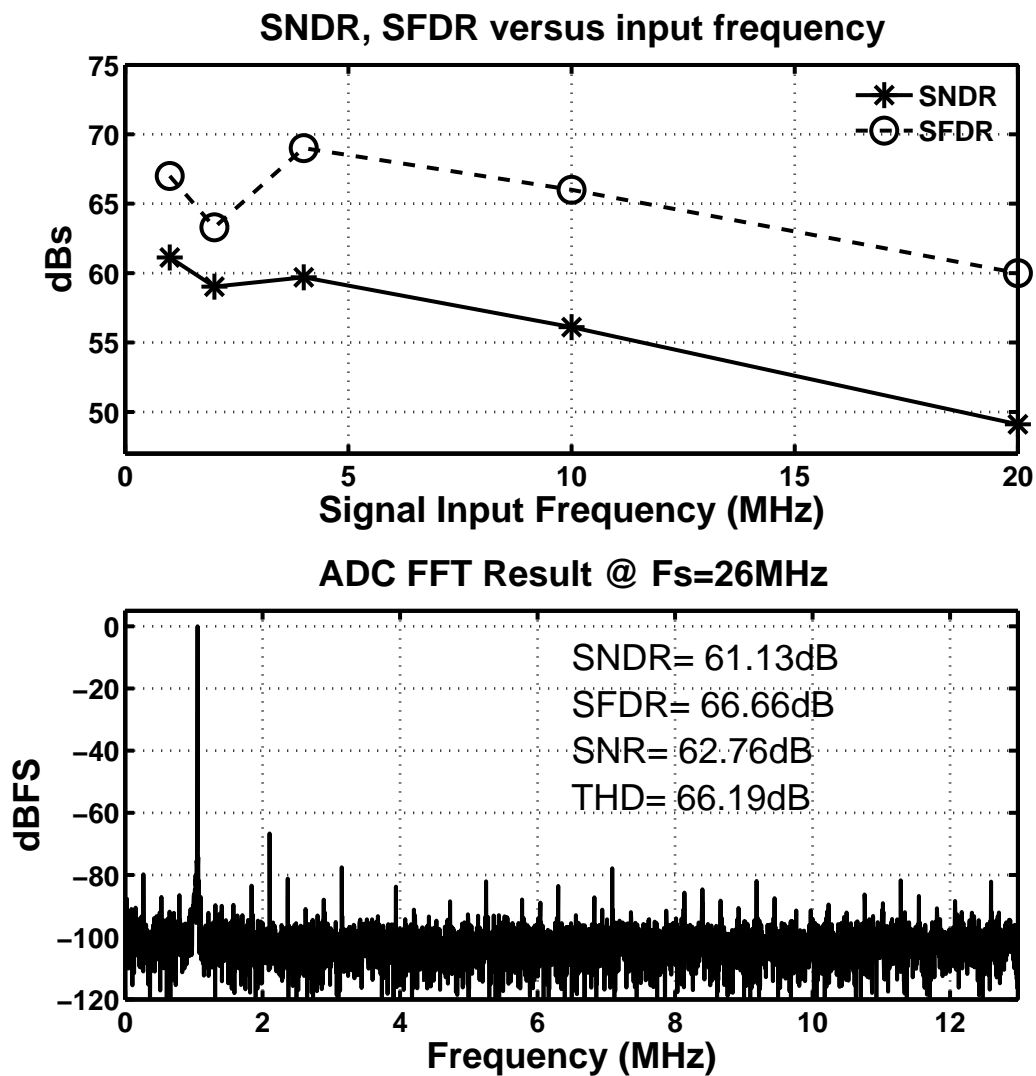


Figure 3.10: SNDR and SFDR over input frequency at 26MSPS

ADC Magnitude Response versus Input Frequency

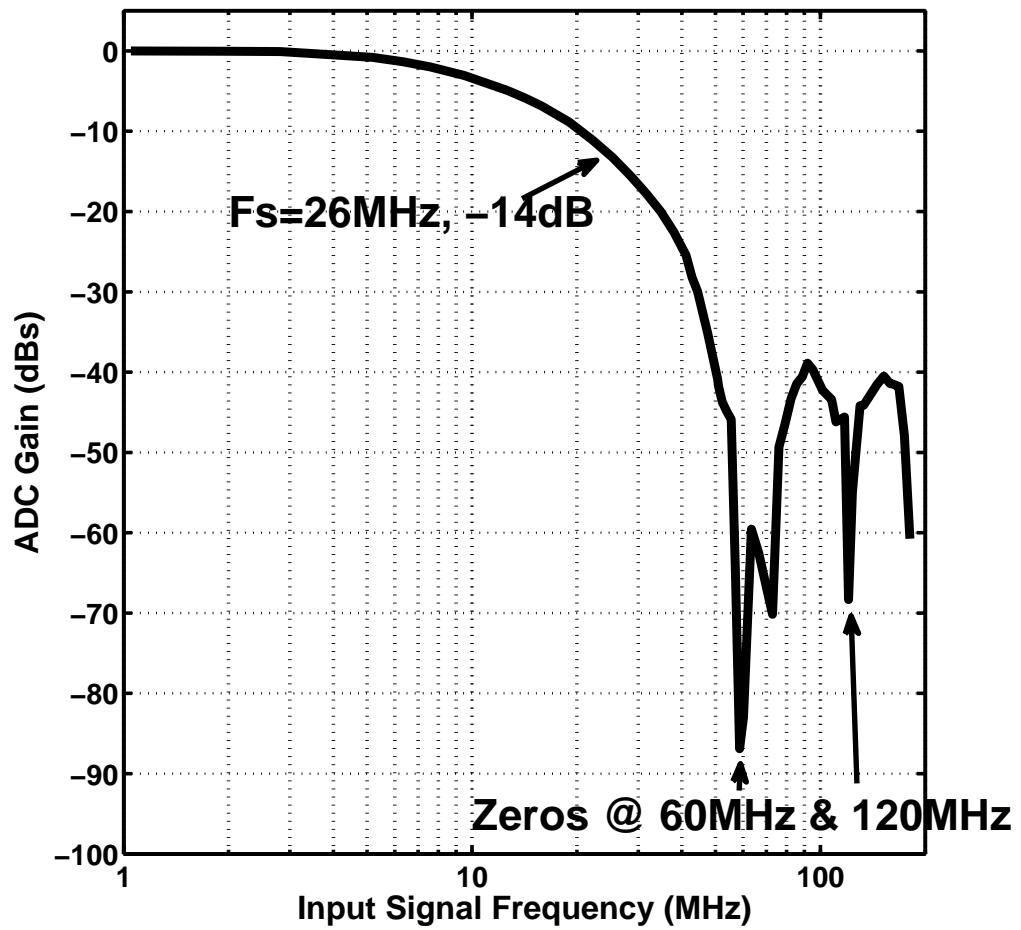


Figure 3.11: Measured Frequency Response of the pipeline ADC

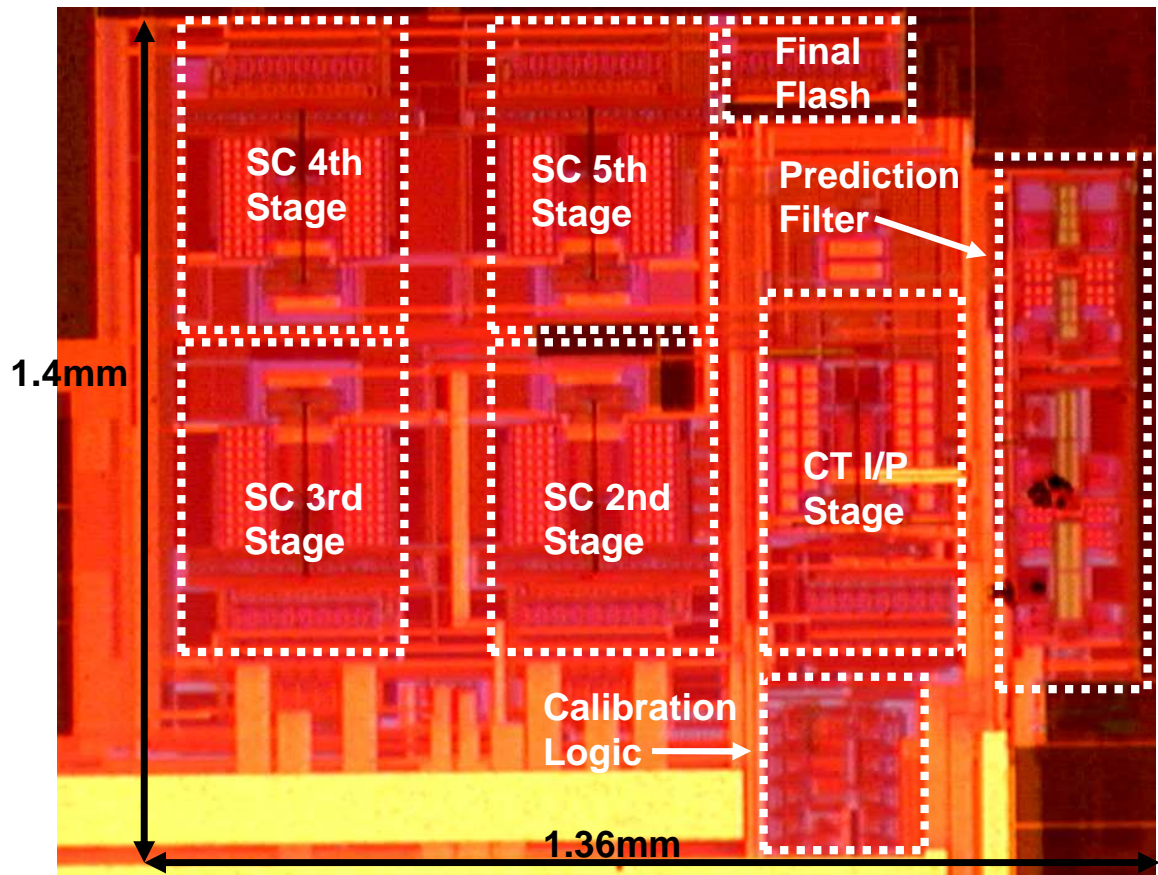


Figure 3.12: Photo of the die

Power Consumption (Total=21.4mW)




Figure 3.13: Power Breakdown of the prototype ADC

3.7 Conclusions

The concept of a continuous time input pipeline ADC with inherent anti-alias filtering has been verified. The elimination of the anti-aliasing filter for some applications will reduce system cost. The proposed CT input stage also offers a more favorable trade-off between noise, distortion and amplifier power when compared to switched-capacitor methods. This architecture also results in a lower area due to back-end capacitor scaling. Consequently switch Ron requirements are very much eased and sampling distortion is reduced. The benign resistive input load easily allows a rail to rail input signal to be coupled into the ADC. This continuous-time input architecture has been enabled by the introduction of low precision signal prediction within the ADC first stage.

ACKNOWLEDGMENT

The authors would like to thank National Semiconductor for supporting this work and providing chip fabrication. Much thanks is due to Professor Gabor Temes and Nima Maghari for useful discussions and layout help.

Chapter 4 – Conclusions

4.1 What in a nutshell has been done so far?

This thesis has investigated the principle of CT input pipeline ADCs. Two CT input ADCs are presented. The first ADC presented in Chapter 2 processes the first stage residue in a CT fashion without filtering being applied to the input signal. Such an approach aims to show that by processing the first stage residue in a CT fashion the switched-capacitor impulse sampling function can be moved to the second stage input. The second ADC presented in Chapter 3 is an improvement on the first ADC of Chapter 2 in that filtering of the input signal is introduced in the first pipeline stage. Simple RC filtering and first order analog sinc filtering are employed. This allows the possibility that anti-alias filtering may be eased in complexity or eliminated. This of course depends greatly on the end application. Specifically it depends on the anti-alias attenuation required at $F_s - F_b$. Further improvements can be made to the second ADC to yield:-

- More aggressive anti-alias filtering
- Lower power operation
- Better SNDR performance
- A lower complexity prediction filter

These specific silicon improvements will be discussed next.

4.2 Specific Silicon Improvements

One limitation of the ADC described in Chapter 3 is the limited signal attenuation achieved at the frequency $F_s - F_b$. The main reason for this is that we initially just wanted to prove the concept. The out of band filtering can be improved by looking at the two filtering methods separately. Firstly the $3dB$ frequency of the RC filter can be reduced. This is easily achieved but it comes at the cost of signal droop at the passband edge. Secondly the sinc filtering can be optimized. The first thing to recognize is that the sampling clock used for the windowed integrator was a 50:50 duty cycle clock. This created the first sinc filter null to be at $2.F_s$. Such sinc filtering is far from optimum. For example if it were possible to window sample for the full period this would yield a first sinc null at F_s which would yield significantly improved anti-aliasing attenuation. However it not possible to window sample for the full period due to need to reset the amplifier for some minimum period of time. Using a $25MHz$ clock, setting the duty cycle to 88:12 and setting the RC $3dB$ to $7MHz$ yields a $3dB$ frequency of $6MHz$. This also yields anti-alias attenuation of $29dB$ at F_s and attenuation of $-17dB$ at $19MHz$. This frequency response is shown in Fig. 4.1. Yet another research avenue would be to increase the order of the filtering. Specifically any filtering techniques that can improve the transition band of the anti-alias filter and provide more aggressive filtering at $F_s - F_b$ would be useful. Specifically higher order sinc filtering might be an avenue

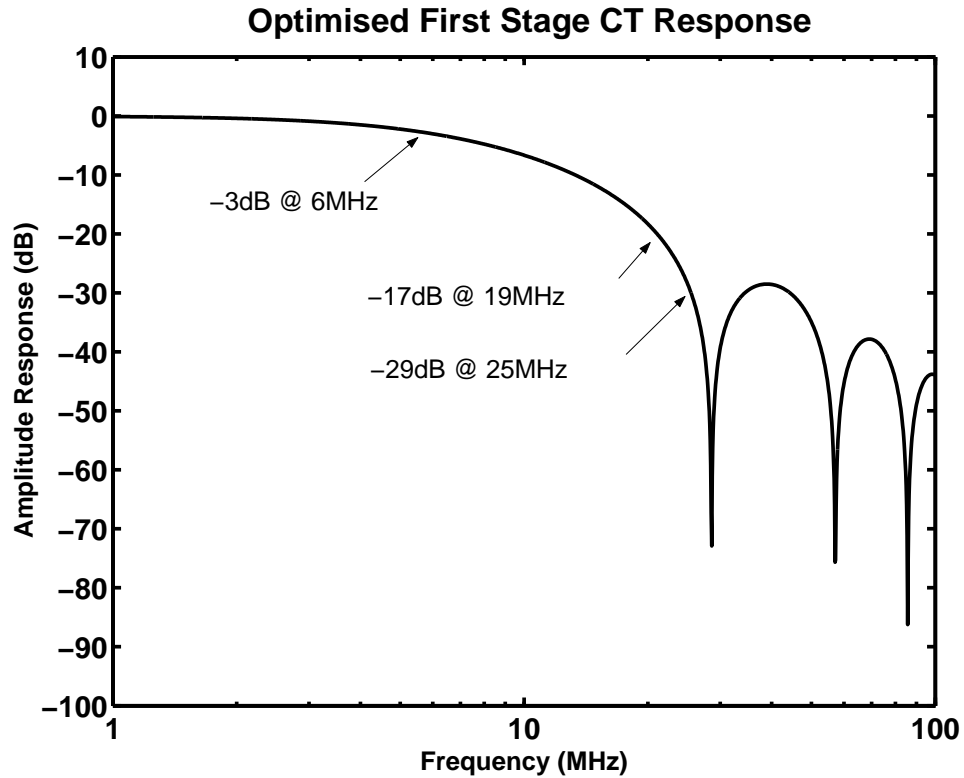


Figure 4.1: Optimized Frequency Response of the CT pipeline ADC

worth pursuing.

The power consumption can be reduced in this ADC. The main areas where power can be reduced is in the back-end ADC amplifiers. This can be done by improving the feedback factor in the back-end MDACs. Specifically separate sampling capacitors and reference capacitors were used in the back-end ADC switched-capacitor MDACs. The reason this was done was to ensure the charge delivered to the MDACs from the reference voltages was constant from cycle to cycle. This eases the requirements on the reference voltages distribution. However this approach reduced the switched-capacitor MDAC feedback factor significantly and

thus increased the power required to maintain the required settling time. By using the same capacitors to sample the stage input signal and inject the reference voltage the feedback factor can be improved significantly.

Experimental results showed that the SNDR of the second ADC described in Chapter 3 was limited by the back-end ADC. The back-end ADC had an ENOB of roughly 7.5 bits when tested on its own. When combined with the first 2.5 bit stage the full ADC performance came to about 9.8 bit ENOB. The reason that the back-end ADC had poor SNDR was due to the fact that the back-end MDACs had high noise gain. This was also due to the poor feedback factor described already.

It is felt that the CT ADC described in Chapter 3 could operate with a lower complexity prediction filter. The filter used was originally designed for a full half clock time advance at $50MHz$ ($10nS$) over $20MHz$ of bandwidth roughly. This corresponds to the ADC in Chapter 2. However the same filter was reused for the CT ADC in Chapter 3. Here a quarter clock time advance was required over $10MHz$ or so. Here the bandwidth requirements are eased. The question is whether this would lead to a much simpler prediction filter. This is worthy of further study.

4.3 Future Directions

It is felt that the CT input pipeline ADC architecture can be further optimized to yield a solution that is more robust, easy to test, lower power and higher performance. The objective is to resolve as many bits as possible in the CT first stage. Specifically if the prediction filter can be made as simple as possible then perhaps it

can be implemented using an over-sampled switched-capacitor filter. For example if it could be implemented with a biquad then perhaps switched capacitor circuitry might add value due to the fact that the filter would then be clock programmable. Yet a further improvement that is worthy of investigation is whether a digital filter could be used for the prediction filter in a case where the flash ADC and digital prediction were over-sampled. This scenario was investigated for ADC 1 during this thesis work and it was found that the coarse quantization in conjunction with the high out of band gain of the prediction filter cause the prediction filter output to consist of gross errors due to noise. It would be worthy of further investigation for ADC 2 described in Chapter 3.

A key challenge that was not tackled for ADC 2 was the temperature variation of the interstage gain in the first stage. Specifically the fact that the interstage gain is given by W/RC means that the interstage gain is subject to temperature variation. Any calibration used needs to track the gain variation in accordance with how fast it changes.

Bibliography

- [1] S. Bardsley et al, "A 100dB SFDR 80MSPS 14-bit 0.35um BiCMOS Pipeline ADC", *IEEE J.Solid-State Circuits*, Vol.41,No.9, September 2006.
- [2] D. W. Cline and P. R. Gray, "A power optimized 13-b 5-Msamples/s pipelined analog-to-digital converter in 1.2um CMOS", *IEEE J.Solid-State Circuits*, Vol.31,No.3, Mar. 1996.
- [3] S. Lewis, "Optimising the stage resolution in Pipelined, multistage, Analog-to-Digital Converters for Video-Rate Applications" *IEEE TCAS 2*, Vol.39, No.8, August 1992.
- [4] D. Kelly et al, "A 3-V 340-mW 14-b 75-Msample/s CMOS ADC With 85-dB SFDR at Nyquist Input", *IEEE J.Solid-State Circuits*, Vol.36, No.12, December 2001.
- [5] R. Reeder, "AN742 Frequency Domain Response of Switched-Capacitor ADCs", *Analog Devices Application Note*.
- [6] D. Robertson, "Selecting Mixed-Signal Components for Digital Communications Systems, Part V, Aliases, images, and spurs", *Analog Dialogue*, 31-3, 1997.
- [7] H.W. Bode, "Network analysis and feedback amplifier design", 1945.
- [8] D.Y. Chang, "Design Techniques for a Pipelined ADC Without Using a Front-End Sample-and-Hold Amplifier" *IEEE TCAS 1* Vol.51, No.11, November 2004.
- [9] T. Kuyel.F. Tsay, "Optimal analog trim techniques for improving the linearity of pipeline ADCs" Proc. International Test Conference, pp.367-375 Oct 2000.
- [10] A. Baschiroto et al, "A 150 Msample/s 20 mW BiCMOS switched-capacitor biquad using precise gain op amps" *ISSCC Digest of Technical Papers*, pp. 212-213, Feb 1995.

- [11] V. Saari, M. Kaltiokallio, S. Lindfors, J. Ryynanen, K.A Halonen, "A 1.2V 240MHz CMOS Continuous-Time Low-Pass Filter for a UWB Radio Receiver" *ISSCC Digest of Technical Papers*, pp. 122-591, Feb 2007.
- [12] A. Baschiroto et al, "A 200-Ms/s 10-mW switched-capacitor filter in 0.5- μ m CMOS technology", *IEEE J.Solid-State Circuits*, Vol.35, No.8, Aug. 2000.
- [13] Singer, L., Mehr, I., "A 55-mW, 10-bit, 40-Msample/s Nyquist-rate CMOS ADC" *IEEE J.Solid-State Circuits*, Vol.35, No.3, March 2000.
- [14] Richard Schreier, Jos Silva, Jesper Steensgaard, Gabor C. Temes, "Design-Oriented Estimation of Thermal Noise in Switched-Capacitor Circuits", *IEEE TCAS 1* Vol.52, No.11, November 2005.
- [15] S. J. Orfanidis, "Introduction to Signal Processing", Prentice-Hall Inc, 1995.
- [16] N. Gray, T. Smith, "Improve Group Delay Response In The Anti-Aliasing Filter", National Semiconductor Application Note, *Electronic Design*, May, 1998.
- [17] R. Bagheri et al, "An 800-MHz6-GHz software-defined wireless receiver in 90-nm CMOS", *IEEE J.Solid-State Circuits*, Vol. 41, No.9, December 2006
- [18] D.Gubbins, B. Lee, P. Hanumolu, U. Moon, "A Continuous-time Input Pipeline ADC", *Proc. IEEE Custom Integrated Circuits Conf.*, Sep. 2008
- [19] A. Karanicolas, H. S. Lee, and K. Bacrania, "A 15-b 1MSample/s digitally self-calibrated pipeline ADC", *IEEE J.Solid-State Circuits*, Vol.28, No. 12, pp. 1207-1215, Dec. 1993.

