

Calibration Techniques for Power-efficient Residue Amplifiers in Pipelined ADCs

Sehgal, R.K.

DOI

[10.4233/uuid:a6a5f8a8-e328-44d8-a265-dfe2b75a9bf8](https://doi.org/10.4233/uuid:a6a5f8a8-e328-44d8-a265-dfe2b75a9bf8)

Publication date

2021

Document Version

Final published version

Citation (APA)

Sehgal, R. K. (2021). *Calibration Techniques for Power-efficient Residue Amplifiers in Pipelined ADCs*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:a6a5f8a8-e328-44d8-a265-dfe2b75a9bf8>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

**Calibration Techniques for
Power-efficient Residue Amplifiers
in Pipelined ADCs**

Calibration Techniques for Power-efficient Residue Amplifiers in Pipelined ADCs

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology

by the authority of the Rector Magnificus Prof.dr.ir. T.H.J.J. van der Hagen
chair of the board of Doctorates
to be defended publicly on
Monday 5 July 2021 at 15:00 o'clock

by

Rohan SEHGAL

Master of Science in Electrical Engineering, Delft University of Technology,
the Netherlands
born in New Delhi, India

This dissertation has been approved by the promoters.

Composition of the doctoral committee:

Rector Magnificus	chairperson
Prof.dr.ir. K. Bult	Delft University of Technology, promotor
Prof.dr. K.A.A. Makinwa	Delft University of Technology, promotor

Independent members:

Prof.dr.ir. L.C.N. Vreede	Delft University of Technology
Prof.dr. P.J.A. Harpe	Eindhoven University of Technology
Prof.dr. N. Sun	Tsinghua University
Dr.ir. J. Craninckx	IMEC Belgium
Prof.dr.ir. W.A. Serdijn	Delft University of Technology, reserve member

Other members:

Dr.ir. F. van der Goes	Broadcom Netherlands
------------------------	----------------------



Keywords: CMOS, Pipelined Analog-to-Digital Converter, Discrete-time Amplifier, Digital Calibration, Linearization.

Printed by: Ipskamp, Enschede

Copyright © 2021 by R. Sehgal

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

"Whatever can go wrong, will go wrong"

Edward A. Murphy Jr.

"There are no wrong notes in jazz"

Miles Davis

Contents

1	Introduction	1
1.1	Pipelined ADCs.	2
1.2	Residue Amplifier.	3
1.3	Thesis Objective	5
1.4	Thesis Organization	6
	References.	7
2	Review of Pipelined ADCs	9
2.1	Pipelined ADC Architecture	9
2.1.1	Sub-blocks.	11
2.1.2	Sources of Error.	14
2.2	ADC Error Reduction Techniques	19
2.2.1	Digital Calibration	21
2.3	Conclusion	24
	References.	24
3	Analysis of Residue Amplifiers	27
3.1	Settling Behavior of Discrete-time amplifiers	28
3.1.1	Noise	30
3.1.2	Distortion	31
3.2	Discrete-time Amplifiers with Negative Feedback.	34
3.2.1	Gain Error	36
3.2.2	Noise	37
3.2.3	Distortion	40
3.3	Power Efficiency	45
3.3.1	SNR-limited case	46
3.3.2	SNR and Gain error-limited case	48
3.4	Conclusion	51
	References.	52

4	Split-ADC Calibration	53
4.1	Split-ADC Calibration Principle	54
4.1.1	Gain Error	55
4.1.2	Distortion	58
4.1.3	Mismatch between Split-ADCs.	60
4.1.4	Calibration Non-orthogonalities	61
4.2	Prototype Split-ADC Implementation	62
4.2.1	ADC Architecture	62
4.2.2	Op-amp	64
4.2.3	Comparator	65
4.2.4	Analog Input Interface	66
4.3	Measurement Results.	66
4.3.1	Calibration Performance	71
4.4	Conclusion	72
	References.	73
5	Residue Amplifier with Analog Error Correction	75
5.1	Analog Error Correction	76
5.2	Gain Error Correction	77
5.3	Distortion Correction.	79
5.3.1	Linearization Principle	81
5.3.2	Proposed Linearized Amplifier	84
5.3.3	Verification	87
5.4	MOS Implementation.	89
5.4.1	Linearization Parameters.	89
5.4.2	Even-order Distortion	93
5.5	Design Considerations	94
5.5.1	Higher-order Effects	95
5.5.2	Noise	95
5.5.3	Temperature Dependence	97
5.5.4	Common-mode Rejection.	98
5.6	Power Efficiency	98
5.7	Conclusion	100
	References.	102

6	A Split-ADC with a Linearized Integrating Amplifier	105
6.1	ADC Architecture	106
6.1.1	Calibration Architecture	107
6.1.2	"Split-over-time"-ADC	109
6.2	ADC Implementation.	111
6.2.1	Timing	111
6.2.2	Noise Budget	112
6.2.3	ADC Sub-blocks	115
6.3	Measurement Results.	120
6.4	Conclusion	125
	References.	126
7	Conclusions	129
7.1	Main Findings	129
7.2	Future Scope of Work	132
	References.	136
	APPENDICES	138
A	Appendix	139
A.1	Open-loop Amplifier Distortion Analysis	139
A.2	Closed-loop Amplifier Distortion Analysis.	141
	References.	146
B	Appendix	147
	Summary	151
	Samenvatting	157
	Acknowledgments	163
	List of Publications	167
	About the Author	169

1

Introduction

Major advances have been made in telecommunications over the last two decades, which have resulted in one of the most disruptive technologies of our times – the internet. The demand for internet has led to a proliferation of wireline and wireless network infrastructure around the globe, with new generations being adopted every few years. During the course of this doctoral research alone, we have seen the wireless networks evolve from 3G to 4G and now to 5G. This growth in connectivity across the world has transformed how economies, societies and geopolitics function. High bandwidth networks have allowed corporations to operate seamlessly across borders, made social media, e-commerce and cryptocurrencies ubiquitous and given countries more tools to interfere in each other’s democratic processes. For better or for worse, this explosion in networking and the internet has unalterably changed the world and a significant part of this has been fueled by innovation in wireline and wireless technologies.

Focusing on wireline transceivers in Ethernet physical layer interface (PHY) chips, the industry has transitioned from 10BaseT to 10GBaseT in the space of 20 years, with 25GBaseT now on the horizon. This >1000x improvement in throughput is also reflected in the spectral bandwidth, which has grown exponentially, as shown in Figure 1.1. To accomplish this, a lot of work has gone into improving each aspect of the Ethernet PHY, ranging from cables and magnetics to system-level innovation. However, the onus of driving this trend has largely been on the analog front-ends,

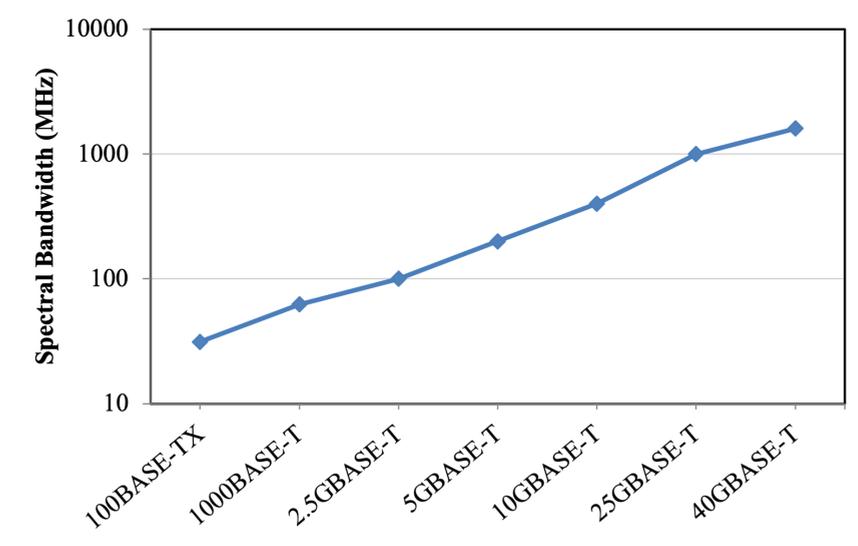


Figure 1.1: Ethernet standards with input spectral bandwidth specification

which have become faster and more precise while consuming less power, allowing the use of more complex modulation schemes to make the most out of the available channel bandwidth.

1.1. PIPELINED ADCs

Looking at the 10GBaseT PHY implementations, for example, nearly all the PHYs on the market use an ADC with $>9.5\text{b}$ ENOB, sampling at a rate of at least 800MS/s . To achieve such high sampling rates while maintaining a high resolution, primarily two techniques are employed – time-interleaving [1] and pipelining [2]. In time-interleaved ADCs, multiple ADC lanes are used in parallel to digitize the input signal. By sampling the signal sequentially, N number of ADC lanes can collectively sample the input signal at N times the rate of a single lane. While simplifying the design of the individual lane, time-interleaving imposes serious penalties on the overall ADC performance. Using multiple ADC lanes in parallel increases the load seen by the preceding buffer (or sample-and-hold) proportionally. In addition, any kind of mismatch between the ADC lanes – gain, timing and offset – generates spurious tones in the combined ADC output, severely limiting the overall ADC resolution. These two factors place a critical limitation on the number of lanes that can be interleaved to realize high-speed ADCs.

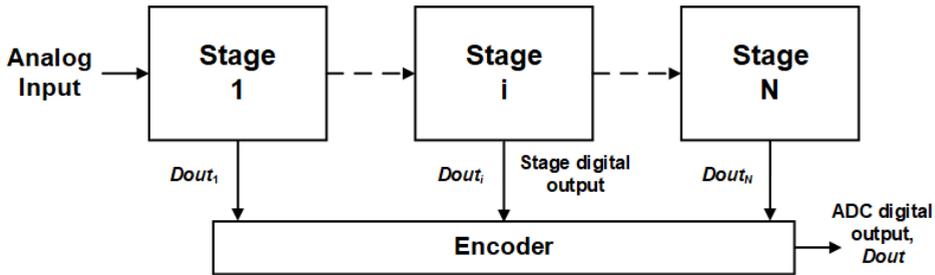


Figure 1.2: Pipelined ADC architecture

The other alternative, pipelining, does not suffer from either of the above drawbacks, making it a very attractive option. As shown in Figure 1.2, pipelining involves breaking the signal quantization process into multiple steps that can be implemented by multiple stages operating concurrently. Each stage now resolves a few bits, and passes on the remaining unquantized signal, or the residue, to the next stage. Considering that each stage is resolving only a few bits, the stages can be simplified and operated at a faster speed, allowing the ADC to achieve faster overall sampling rates at the cost of only increased data latency.

One of the key techniques that enable pipelining is residue amplification. In order to relax the noise requirements for later stages, amplifiers are used to boost the residue signal from the preceding stages. Since any amplifier non-ideality will corrupt the residue transferred to the following stages, the accuracy of the residue amplifier directly determines the accuracy of the ADC. Because of this, it often ends up being one of the most critical and power-hungry blocks in the signal chain.

1.2. RESIDUE AMPLIFIER

A lot of research has been undertaken on designing power efficient residue amplifiers, resulting in many innovative topologies. The common denominator in all these topologies, however, is the basic amplifier core, which can be described in its simplest form as four transistors charging a load capacitor in a fixed interval of time, as shown in Figure 1.3. These transistors are driven by a signal derived from the differential input and charge the capacitor to a differential output voltage with a certain amount of gain. The accuracy of this gain is usually limited by two constraints –

- **Settling time:** As most analog-to-digital converters operate on a sampled-and-held voltage input, the residue amplifier is usually a discrete-time circuit.

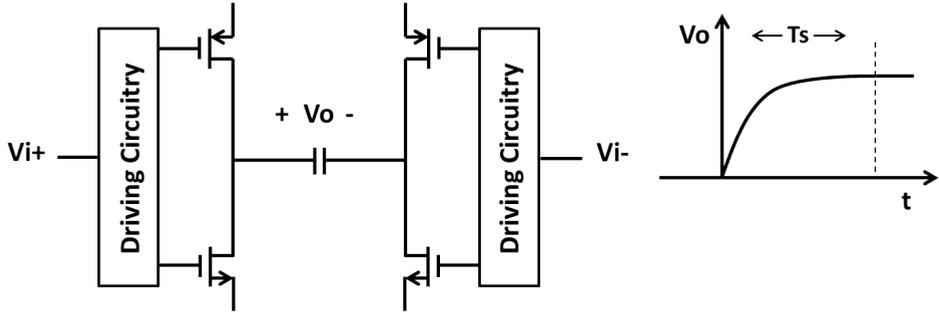


Figure 1.3: Core amplifier with circuit overhead and settling

Since the amplifier is normally reset prior to amplification, it takes a certain amount of time for the amplifier output voltage to reach its final value. If the amplifier is considered to be a single-pole system, its output voltage typically experiences an exponentially settling behavior. Hence, as the settling time for the amplifier is increased, the gap between its output voltage to its final steady-state desired value reduces. However, as most amplifiers are class-A, which burn a fixed amount of current during every instant of amplification, a longer settling time invariably results in higher power dissipation.

- **Circuit overhead:** Since the gain of the open-loop four-transistor core amplifier will vary over Process, Voltage and Temperature (PVT) variations beyond the desired accuracy requirements of a residue amplifier, some circuit overhead has to be used to improve the accuracy of the amplifier. Traditionally, this overhead has been in the form of high-gain amplifier stages and negative feedback, which, together, can provide a very accurate gain and low distortion. This, however, comes at the expense of significantly higher power dissipation.

The power efficiency of the residue amplifier will essentially be determined by how this core amplifier is utilized. Hence, we need to find solutions that allow us to use the core amplifier with the lowest possible settling and circuit overhead while achieving the desired gain accuracy.

There are several error sources in an amplifier that could degrade the accuracy of residue amplification in a pipelined ADC. The most significant of these are noise, linear gain error and distortion. Although all of these error sources limit the ADC dynamic range, they have significantly distinct characteristics and need to be tackled differently. Due to its stochastic nature, noise poses a fundamental limit on the ADC accuracy and, although certain choices can be made at architectural level to optimize

it, we have to scale up a design and its power dissipation in order to achieve a certain noise performance.

Absolute gain error and distortion, on the other hand, are more deterministic and can be suppressed in multiple ways. As mentioned previously, a classic approach is to use negative feedback with a very high loop-gain. However, this approach results in a very large circuit overhead and requires a lot of settling. Another more recent approach to achieve accurate residue amplification is to combine the amplifiers with some form of digital calibration. By fixing the errors in digital post-processing, the specifications on the residue amplifiers can be significantly relaxed. Considering that device scaling has made digital processing available at an increasingly lower footprint, this exchange of analog complexity for digital processing generally results in a lower overall ADC area and power dissipation. As a result, a lot of work has been published on utilizing digital calibration to fix the linear gain error and distortion arising from the residue amplifier, and achieving a much improved ADC figure of merit (FoM).

While digital calibration does shift some of the circuit overhead from the amplifier, it has certain limitations. The digital calibration can be broken into two parts – error *detection* and *correction*. While the error detection logic can be run at a slow rate, the error correction part of the calibration has to run at the full ADC speed. For high speed and high resolution correction logic, the area and power consumption can be significant.

Another major constraint of digital calibration is its finite resolution. The accuracy of digital correction is limited by the backend resolution. For a perfect calibration of distortion in an ADC stage’s residue signal, a precise estimate of the residue is required from its backend. And any imperfection in the backend will have a detrimental impact on the calibration accuracy. Furthermore, a residue amplifier with a significant amount of gain error and distortion will usually lead to a loss in the ADC dynamic range, something that cannot be restored through calibration.

1.3. THESIS OBJECTIVE

As discussed above, while digital calibration is an excellent way of realizing accurate residue amplifiers, it does suffer from certain limitations. This means that digitally assisted residue amplifiers can only go so far in improving the residue amplifier power efficiency, and we need to investigate more ways to further push the envelope.

If we shift the error correction to the analog domain, while keeping the error

detection in digital, we draw upon the best of both worlds. By using error estimation in digital along with a deterministic calibration method with fast convergence, the digital logic can be duty cycled to render its power consumption negligible. And more importantly, with the help of analog correction, we fix errors at the source, achieving a better calibration accuracy and a higher ADC dynamic range. Also, with analog-domain linearization, we can use much more power efficient amplifier topologies at much larger input swings which would not be possible solely with digital post-processing due to their high inherent nonlinearity.

In this thesis, we try to find the optimum design for digitally assisted residue amplifiers, by adopting a two-pronged approach –

1. Using very simple amplifier structures closely resembling the basic core amplifier along with low settling accuracy, with the help of linearization in the analog domain, and
2. a deterministic calibration architecture which allows us to calibrate linear gain errors and distortion in background while achieving a fast convergence

1.4. THESIS ORGANIZATION

The rest of the thesis is organized as follows. Chapter 2 provides an overview of pipelined ADCs and their main ADC sub-blocks. We also discuss the major sources of error in these sub-blocks and their impact on the ADC performance, along with a brief introduction to digital calibration and how it can be used to relax the design restrictions on these sub-blocks.

In Chapter 3, the settling behavior of discrete-time amplifiers is analyzed in detail. By looking at how noise and distortion vary in a discrete-time amplifier, we can identify design strategies to optimize their power consumption for achieving the best possible dynamic range.

Chapter 4 discusses the split-ADC calibration architecture, and how it can be used to calibrate gain and non-linearity errors emanating from the residue amplifier. A prototype split-ADC was implemented and its measurement results are presented to demonstrate the effectiveness and convergence speed of the calibration algorithm.

In Chapter 5, we discuss the principles of analog correction of amplifier gain and distortion, and some amplifier topologies based on analog linearization. The proposed amplifier topology is implemented as part of a custom split-ADC, whose architecture and circuit details are presented in Chapter 6, along with the ADC measurement results.

Finally, Chapter 7 concludes the thesis, summarizes the original contributions of this dissertation and outlines future scope of work.

REFERENCES

- [1] W. C. Black and D. A. Hodges, "Time interleaved converter arrays," *IEEE J. Solid-State Circuits*, vol. 15, pp. 1022–1029, Dec. 1980.
- [2] S. H. Lewis and P. R. Gray, "A pipelined 5MHz 9b ADC," in *ISSCC Dig. of Tech. Papers*, 1987.

2

Review of Pipelined ADCs

This chapter presents an overview of the pipelined ADC architecture. We will briefly review the operation of a pipelined ADC and of all its major sub-blocks, and the effect of their nonidealities on the ADC resolution.

One of the attributes that truly distinguishes pipelined ADCs from other ADC architectures is residue amplification, and the residue amplifier often occupies a significant portion of an ADC's error and power budget. Over the last two decades, several techniques have been proposed to relax the design requirements on the residue amplifier and simplify their design. We will review some of the prominent error correction techniques published in the literature in recent years, with a major focus on digital calibration.

2.1. PIPELINED ADC ARCHITECTURE

In a pipelined ADC, the quantization of the sampled input signal is performed over a cascade of multiple stages, as shown in Figure 2.1. Each stage consists of a coarse ADC which digitizes the input signal, $V_{\text{in},i}$, based on the ADC references, $\pm V_{\text{ref}}$, with a certain resolution, N_i . The coarse estimate of the signal, $D_{\text{out},i}$, is then subtracted from the input signal with the help of a sub-DAC and the remaining portion of the input signal, the residual quantization error, e_{qi} , is passed on to the next stage for further quantization. This residual quantization error can essentially

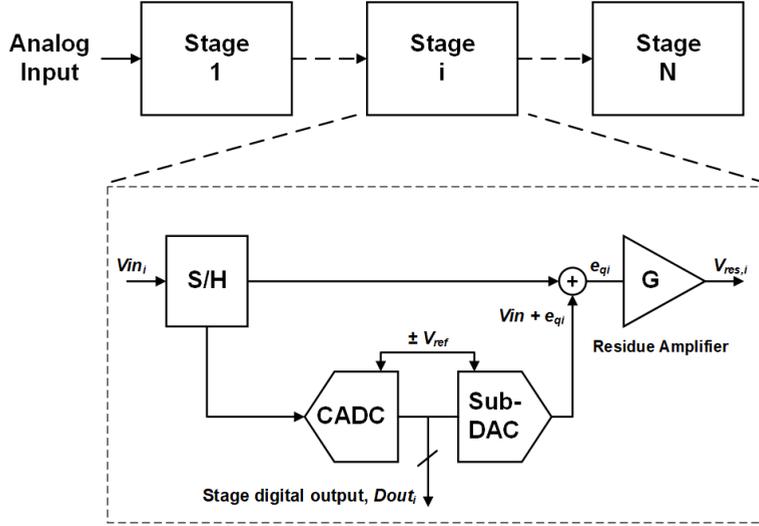


Figure 2.1: Pipelined ADC architecture overview

be expressed as –

$$e_{qi} = V_{in,i} - D_{out,i} \left(\frac{V_{ref}}{2^{N_i}} \right) \quad (2.1)$$

where $D_{out,i}(V_{ref}/2^{N_i})$ is the coarse ADC's output mapped onto the ADC reference voltage by the sub-DAC. Once the subsequent stage has sampled and started quantizing this residual signal, the previous stage is free and can be used to convert the next input signal sample. By operating the stages concurrently, the sampling speed of the overall ADC can be the same as that of each stage, while the overall ADC resolution equals the total number of bits resolved by all the stages. This allows pipelined ADCs to operate at high speeds without incurring limits on resolution or exponential growth in area, albeit at the expense of additional latency.

As the signal gets quantized through the stages, the number of bits left to be resolved reduces. However, the magnitude of the residual quantization error passed on from one stage to the next also decreases. This means that each stage needs to be designed with the same dynamic range. To relax the resolution required from later stages, some gain, $G_{a,i}$, can be added within each stage to scale up the residual signal to the full-scale range. The resulting amplifier output, the stage residue, $V_{res,i}$, can be expressed as –

$$V_{res,i} = G_{a,i} e_{qi} = G_{a,i} \left(V_{in,i} - D_{out,i} \left(\frac{V_{ref}}{2^{N_i}} \right) \right) \quad (2.2)$$

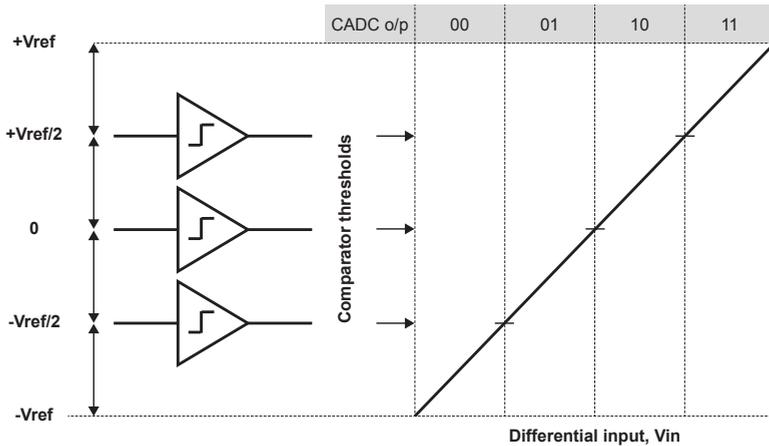


Figure 2.2: 2-bit Flash-type coarse ADC

Equation (2.2) basically represents the transfer function of the pipeline ADC stage, and shows how the stage residue voltage that is passed on to the subsequent stages for further quantization is scaled up to the full-scale input range. Due to this, the LSB size for the subsequent pipeline ADC backend also increases by the same gain. This lowers the resolution requirements for the latter stages, allowing them to be scaled down to optimize the overall ADC power dissipation.

2.1.1. SUB-BLOCKS

Each stage in a pipelined ADC (excluding the backend) consists of three primary building blocks – coarse ADC, sub-DAC and the residue amplifier. In this subsection, we will briefly review the operation of these three blocks.

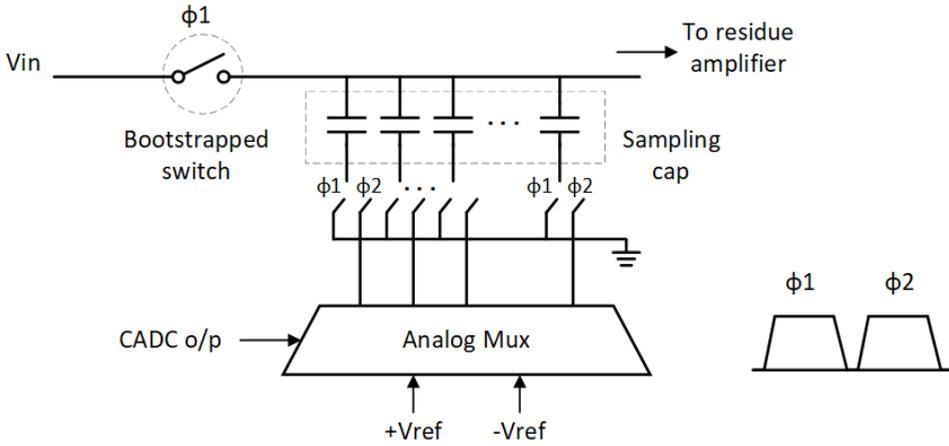
- a. **Coarse ADC:** The coarse ADC (CADC) samples and performs quantization on the input signal. Since the digitization in a pipelined ADC is spread over multiple stages, the coarse ADC is normally a low-resolution quantizer. Traditionally, the coarse ADCs used in the pipelined ADCs have been implemented using the flash architecture [1]. This is because for low resolution and high-speed, flash ADCs can be a good choice due to their simplicity. Figure 2.2 shows the transfer characteristic for a flash-type coarse ADC designed for a 2b/stage pipelined ADC.

Recently, SAR ADCs are increasingly being used as coarse ADCs due to their lower circuit overhead and power dissipation [2, 3]. This is because in advanced

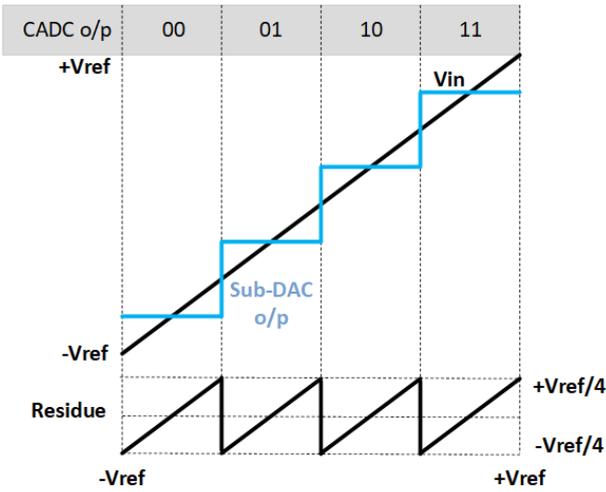
technology nodes, SAR ADCs have been shown to be fast enough to be used for pipelining in high-speed ADCs. And since they require less analog circuitry than flash ADCs, SAR ADCs are generally more power efficient. These pipelined SAR ADCs have demonstrated excellent figures of merit (FoM) and show a lot of promise for high-speed medium-resolution applications in finer technology nodes. However, in the course of this thesis, the flash architecture was preferred for implementing the CADC due to its lower complexity and ease of design.

- b. **Sub-DAC:** In order to generate the residue signal for subsequent stages, the CADC's output, $D_{\text{out},i}$, needs to be subtracted from the stage input signal. The sub-DAC is primarily used to generate the analog equivalent of the CADC's digital output, which can then be subtracted from the input signal. The most common sub-DAC topology employs a capacitive DAC, which can also be used to sample and hold the input signal, as shown in Figure 2.3(a). The cap-DAC is operated in two phases – during the track phase, it tracks and samples the input signal on the cap-DAC unit capacitors. In the hold phase, the unit capacitors of the DAC are connected to the positive and negative reference voltages based on the coarse ADC's estimate. Through charge redistribution, the residual quantization error, as shown in Figure 2.3(b), is generated on the top-plate of the cap-DAC.

The cap-DAC size is instrumental in determining the performance specifications of all the circuits around it – top- and bottom-plate switches, reference buffer and the residue amplifier – and the ADC stage itself. Since the cap-DAC is used for sampling the input signal, its total capacitance is determined by the kT/C noise limit and the ADC noise budget reserved for it. Apart from the noise budget, the entire sampling network also has to meet the bandwidth and linearity requirements. The top-plate sampling switch experiences the largest signal swing and is often the most challenging switch to design. Usually, some form of clock boosting has to be used to meet its linearity requirements. The bottom-plate switches are relatively simpler and are designed primarily to have low parasitic capacitance. These switches can be simplified by using reference attenuation. By switching only a part of the sampling capacitance to the reference voltages, the ADC reference voltages are effectively attenuated. This allows the use of larger voltages, such as the supply rails, as the two ADC reference voltages, making it easier to drive the switches.



(a)



(b)

Figure 2.3: Capacitive sub-DAC architecture along with its output and residue voltage

c. **Residue Amplifier:** As the signal passes through successive stages, the residue gets progressively smaller. In order to simplify the dynamic range requirements for the later stages, pipelined ADCs usually include some form of inter-stage residue amplification. To maintain the same signal swing for all the stages, ideally, the residue amplifier gain, $G_{a,i}$, should be equal to 2^{N_i} , where N_i is the number of bits resolved by that stage's coarse ADC. The stage residue w.r.t input signal is shown in Figure 2.4 for a 2-bit pipelined stage,

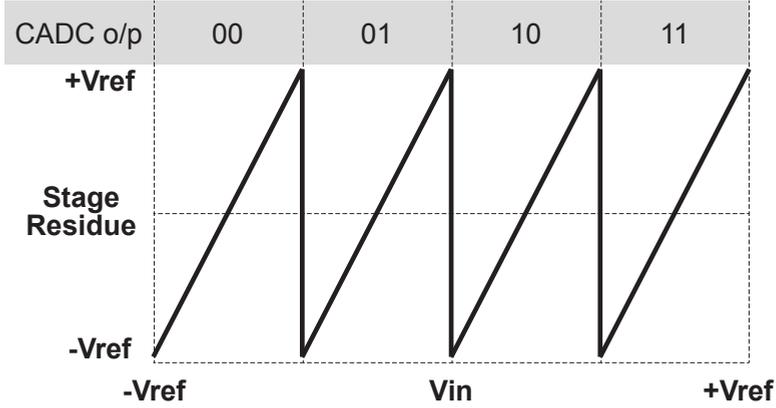


Figure 2.4: Residue amplifier output w.r.t input signal, V_{in}

with a gain of $4x$.

This stage gain is eventually taken into account when the stage output bits are combined in the digital encoder. The digital inverse of the ADC stage transfer function in (2.2) can be formulated to calculate the digital equivalent of the stage input, $D_{in,i}$, by using the digital encoder gain, $G_{d,i}$, of the i -th ADC stage along with the backend ADC code as –

$$D_{in,i} = D_{out,i} + \frac{D_{res,i}}{G_{d,i}} \quad (2.3)$$

where $D_{res,i}$ is the backend ADC code for the i -th stage and is essentially the digital representation of its output residue voltage calculated in (2.2). If we ignore the quantization error in $D_{res,i}$, then Equations (2.2) and (2.3) can be combined as –

$$D_{in,i} = D_{out,i} + \frac{G_{a,i}}{G_{d,i}} \left(V_{in,i} - D_{out,i} \left(\frac{V_{ref}}{2^{N_i}} \right) \right) \quad (2.4)$$

As seen from (2.4), by accurately matching the digital gain in the encoder with the analog gain in the ADC stage, the ADC stage digital outputs can be combined to produce the ideal digitized input.

2.1.2. SOURCES OF ERROR

As shown earlier in Figure 2.2, the comparator thresholds in the CADC are placed to divide the ADC input range into four equal subranges. Based on the subrange the

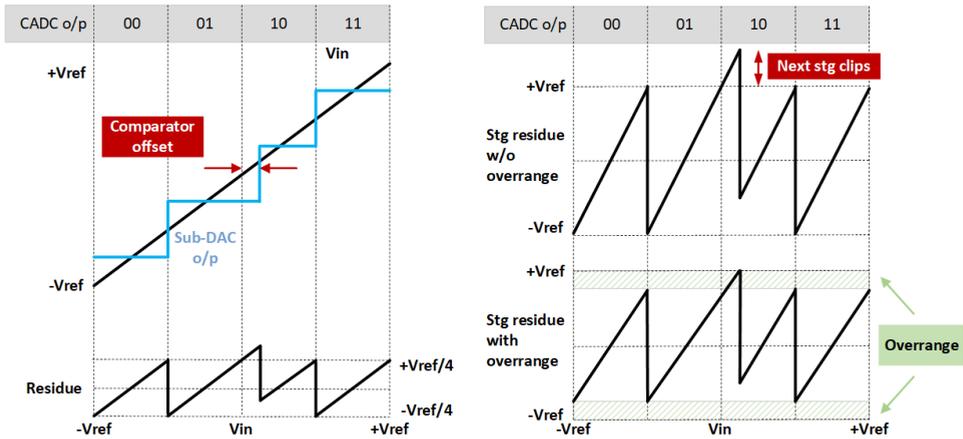


Figure 2.5: The effect of comparator offset in a CADC on a pipeline ADC stage with and without overrange

input signal is present in, the corresponding sub-DAC voltage level is selected. This sub-DAC output is subtracted from the sampled input and the resulting residue is amplified by the residue amplifier to generate the residue signal for the next stage.

Since each block serves a different purpose in a pipeline stage, any non-idealities in these sub-blocks have a distinct impact on the ADC performance. In this subsection, we will discuss some of the major sources of errors in these three pipelined ADC sub-blocks and how they can cause errors in the overall ADC output code.

- a. **Coarse ADC:** Since the generation of the residue relies on the CADC's estimate of the input signal, any inaccuracies in the comparator thresholds will affect the stage residue signal. The primary source of error in CADC is often the comparator offset, which shifts the ADC subranges, as illustrated in Figure 2.5. As this shift in the ADC subrange is not taken into account by the sub-DAC or the residue amplifier, this can push the output of the residue amplifier outside the input range of the next stage, while saturating the amplifier itself. This can lead to missing codes in the ADC output.

In order to relax the comparator offset requirements, some redundancy can be added to the subranges of the pipelined stage in the form of overrange. This can be done by reducing the residue amplifier gain a bit to create some spare headroom for CADC offset. In the example shown in Figure 2.5, the amplifier gain for the 2-bit stage is reduced from its ideal value ($4\times$) in order to create the overrange. A lower gain ensures that, in the presence of any offset in the

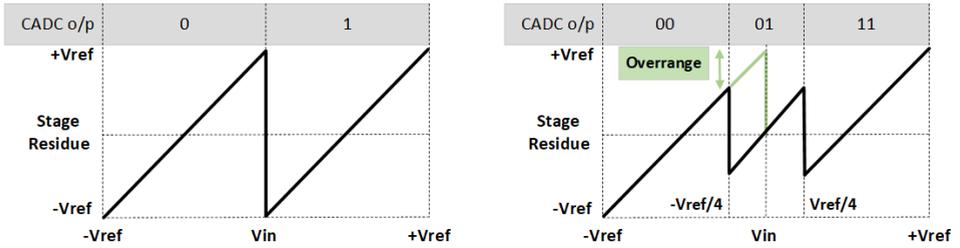


Figure 2.6: Comparison of stage residue signal in 1-bit vs 1.5-bit/stage

CADC, the stage residue will stay within the range of the next pipeline stage.

Another way of creating overrange is by adding extra comparators in the CADC to reduce the nominal input of the residue amplifier. One of the most popular architectures imbibing this method is the 1.5-bit/stage pipelined ADC [4]. In this architecture, an extra comparator is added to a nominally 1-bit ADC stage, with a residue amplifier gain of $2x$. By strategically placing the comparator thresholds, an overrange of $V_{ref}/2$ (and a redundancy of 0.5bit) is created, as shown in Figure 2.6. This allows for the comparators to have an input-referred offset of up to $\pm V_{ref}/4$ without saturating the next stage.

It should be noted that adding redundancy does not correct the errors arising due to the comparator offset, but only makes the ADC sufficiently tolerant of these errors. While the presence of overrange significantly relaxes the comparator requirements, it does incur a power penalty. Since the ADC resolution is a function of the overall gain seen in the pipeline chain, adding redundancy by reducing stage gain also reduces the ADC resolution. Similarly, increasing the resolution of the CADC without changing the stage gain will have no effect on the ADC resolution. However, the circuit overhead introduced by the addition of overrange is significantly smaller compared to the relaxation achieved in comparator design, making it a widely used technique to optimize ADC design.

- b. **Sub-DAC:** As the stage residue comprises of the sampled input and the sub-DAC output, any nonlinearity in the sub-DAC output will directly show up in the stage residue. This DAC nonlinearity could originate due to, for instance, mismatch between the unit capacitors in the capacitive-DAC. The effect of sub-DAC nonlinearity is illustrated in Figure 2.7 which compares the ADC output transfer curve of an ideal ADC with that of an ADC with a nonlinear

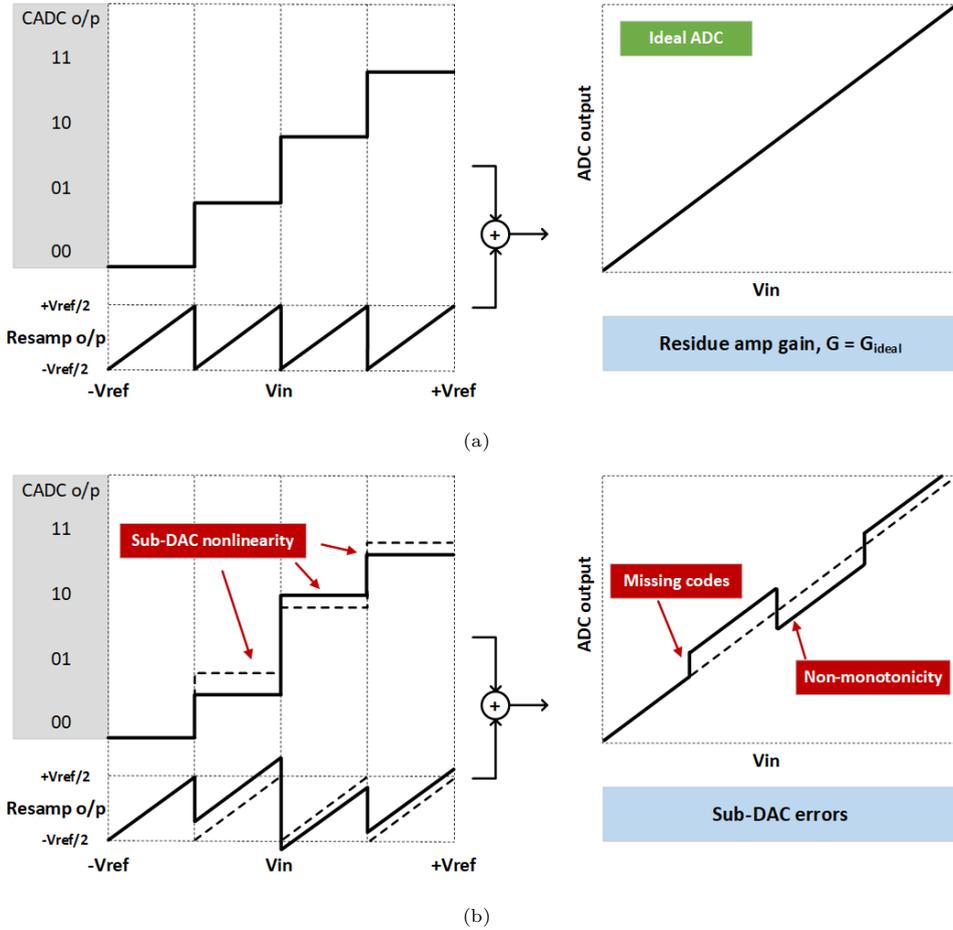


Figure 2.7: ADC output transfer for (a) an ideal ADC and (b) an ADC with nonlinear sub-DAC

sub-DAC. The ADC output curve is obtained by combining the CADC output of the pipelined ADC stage with the stage residue. In the ideal ADC, as shown in Figure 2.7(a), the sub-DAC output levels are perfectly linear. As a result, the ADC subranges align perfectly after encoding the stage residue with the stage CADC bits. Any nonlinearity in the sub-DAC output levels will shift the subranges of the stage residues as shown in Figure 2.7(b). This misalignment amongst the ADC subranges creates abrupt jumps in the ADC output code as the input signal transitions between the ADC subranges. These jumps can cause significant harmonic distortion and potentially missing codes and non-monotonicity in the ADC transfer characteristic.

In order to prevent any missing codes or non-monotonicity, the abrupt jumps caused by these misalignments need to be lower than $\pm \frac{1}{2}$ LSB of the ADC. This requirement can be used to define the maximum error allowed in the sub-DAC of the i -th stage of a pipeline ADC, δ_i , as [5] -

$$|\delta_i| \leq \frac{FS}{2^{1+r_i}} \quad (2.5)$$

where FS is the full-scale range of the ADC and r_i is the combined resolution of i -th stage and the later stages.

- c. **Residue Amplifier:** The residue amplifier scales up the residue of the current stage to the full-scale range of the next stage's quantizer. Since it is directly in the signal path, its non-idealities have a pivotal effect on the quality of the ADC output. The most significant amplifier non-idealities are offset, gain error, distortion. The residue amplifier's offset will have two effects. The ADC output code will contain the cumulative offset from the residue amplifiers in all the stages. Since a lot of applications do not place any restrictions on the ADC offset, the cumulative offset is normally not an issue. The main drawback of the amplifier offset is that it shifts the stage residue, hence, consuming part of the stage overrange.

The residue amplifier gain has much more stringent requirements. It has to match the radix used in the encoder to combine the bits from the ADC stages. Any deviation in the amplifier gain from that radix value creates a misalignment amongst the subranges at the CADC transitions, leading to errors in the ADC output code, as illustrated in Figure 2.8, which shows the discontinuities in the ADC output transfer function arising from gain error and amplifier nonlinearity. The specification for the residue amplifier gain error can again be derived by using the same method as the sub-DAC, resulting in

$$|\sigma_i| \leq \frac{1}{2^{1+r_{i+1}}} \quad (2.6)$$

where σ_i is the relative gain error in the i -th stage and r_{i+1} is the combined ADC resolution remaining after the i -th stage.

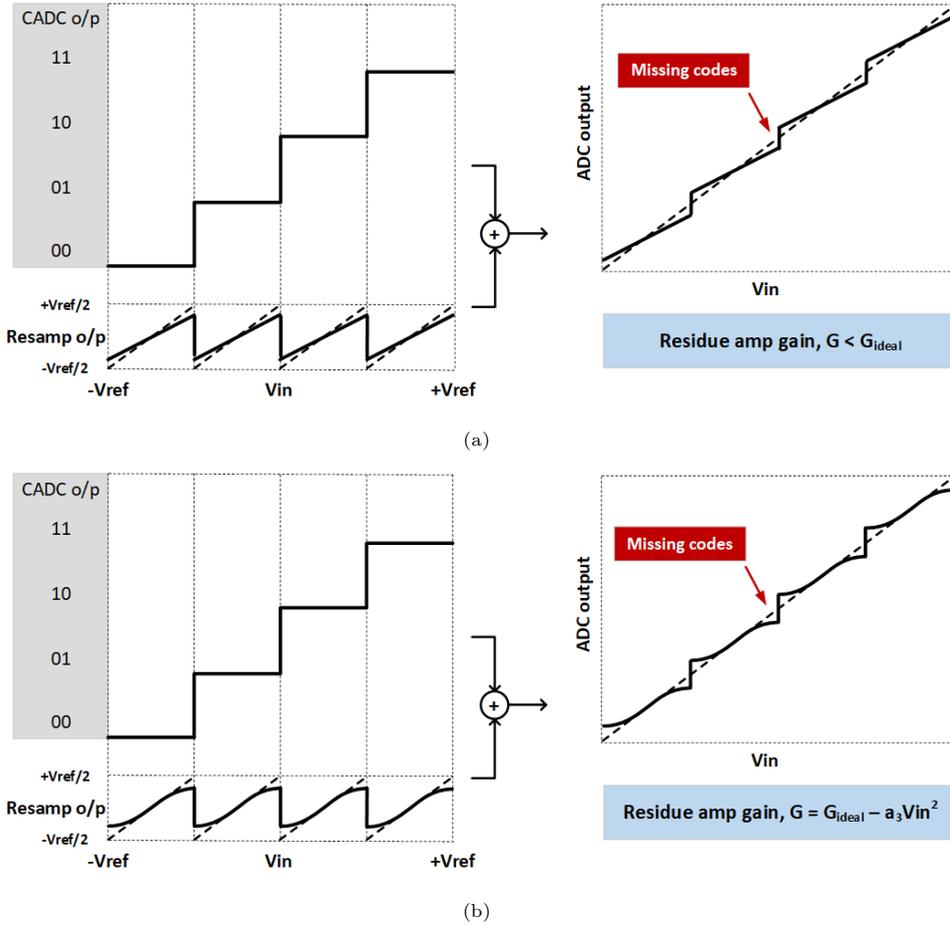


Figure 2.8: Effect of amplifier (a) gain error and (b) nonlinearity on ADC output code

2.2. ADC ERROR REDUCTION TECHNIQUES

Considering the accuracy requirements on the residue amplifier discussed in the previous section, using a conventional high-gain high-bandwidth closed-loop amplifier topology for a high-speed high-resolution pipelined ADC could be very expensive in terms of area, complexity and power consumption. In this section, we will focus on alternative, more efficient strategies for reducing the ADC errors due to amplifier non-idealities.

As the residue amplifier is one of the most power-hungry blocks inside the pipelined ADC, it has been the subject of extensive research over the last three decades. Several techniques have been proposed to relax the requirements from

residue amplifiers. These can broadly be divided into three groups. The first involves the use of analog techniques that allow the amplifier to achieve the required gain and linearity performance without using multiple high-gain stages. Some of the examples include gain-boosting [6], correlated double sampling (CDS) [7], correlated level shifting (CLS) [8], virtual ground reference buffers (VGRB)[9], etc. When compared to a conventional high gain high linearity opamp, these techniques often result in significant power savings and a much simpler amplifier. However, all these techniques incur certain penalties. Gain-boosting is an extremely effective technique to increase the DC gain of the amplifier. However, power still needs to be spent on increasing the amplifier bandwidth to make the most of the higher DC gain. Both CDS and CLS introduce an additional phase of amplification in order to sample the error signal from the first amplification (called the estimation) phase. In the VGRB technique, the reference buffers used to drive the cap-DAC in the amplification phase are referenced to the virtual ground of the closed-loop amplifier instead of the system ground. This effectively bootstraps the entire capacitance of the cap-DAC, removing it from the feedback network of the opamp and theoretically resulting in a feedback factor of 1. However, this improvement in the feedback factor is seriously limited by the gain of the reference buffer and the additional parasitics added by the reference buffers.

In a conventional opamp-based charge amplifier, the charge transfer from the sampling capacitor to the load capacitor settles exponentially and it concludes when the opamp's virtual ground signal reaches zero. Many designs have been published which have tried replacing these conventional op-amps with alternative circuits. These circuits, like zero-crossing based circuits (ZCBC) [10–12], ring amplifiers [13], pulse bucket brigade [14], charge pumps [15], etc, exhibit higher power efficiency because, instead of exponential settling, they rely on slew-based charging. For instance, in the ZCBC technique, instead of an amplifier, a current source is used to charge the load capacitor, with a high-precision comparator monitoring the zero-crossing point. This approach allows the amplification to be essentially open-loop. ADCs incorporated with these alternative topologies have demonstrated excellent power efficiency. However, their inherent accuracy is limited and, without the help of any of the previously mentioned accuracy-enhancing analog techniques or digital processing, these topologies are only suitable for low- to medium-resolution applications.

The last group of techniques involves the use of digital calibration. In this approach, the residue amplifiers are allowed to possess insufficient gain and accuracy,

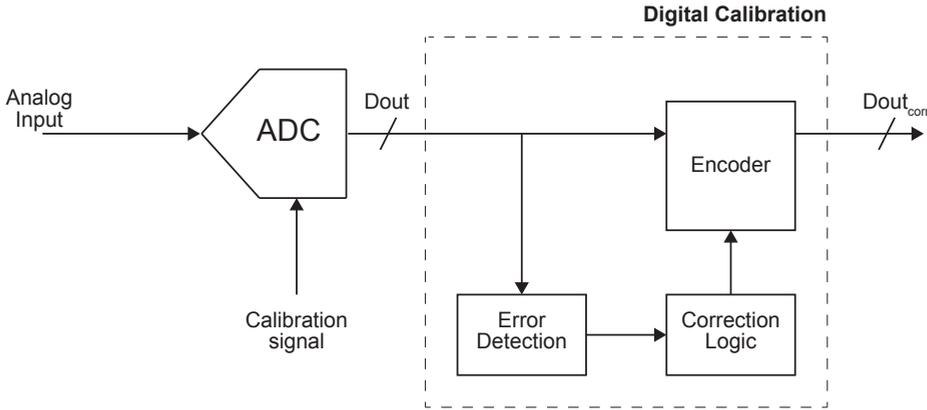


Figure 2.9: Digital calibration principle

and the resulting errors are detected and calibrated in digital post-processing. Since they are no longer required to achieve full gain accuracy, the residue amplifiers can be designed with simpler and more power-efficient topologies. Although implementing digital calibration entails its own set of complexities, a significant part of the calibration hardware is often not in the signal path, allowing it to operate at a much lower rate. As Moore’s law enables cheaper and lower power digital processing, such digitally-assisted amplifiers have quickly become the most widely adopted solution for residue amplification. In this section, we will cover some of the basics of digital calibration of residue amplifier gain and nonlinearity.

2.2.1. DIGITAL CALIBRATION

Digital calibration basically involves the estimation of the error(s) in the residue amplifier by analyzing the effect of the corresponding non-ideality on the ADC’s digital output. By observing certain signatures generated in the ADC output code by the respective error(s), the magnitude of the amplifier’s corresponding non-idealities can be estimated. As shown in Figure 2.9, this error estimation is typically carried out on the ADC output by a digital state machine, and often involves a calibration signal. And based on the estimated error signal, a correction term is applied to counteract the corresponding amplifier non-ideality. This is nominally also performed in the digital backend while encoding the ADC output bits, by adjusting the radix (gain error) or adding an input-dependent correction term (gain nonlinearity).

The nature of the error estimation process plays an important role in deciding the speed, complexity and the efficacy of the calibration algorithm. ADC errors can

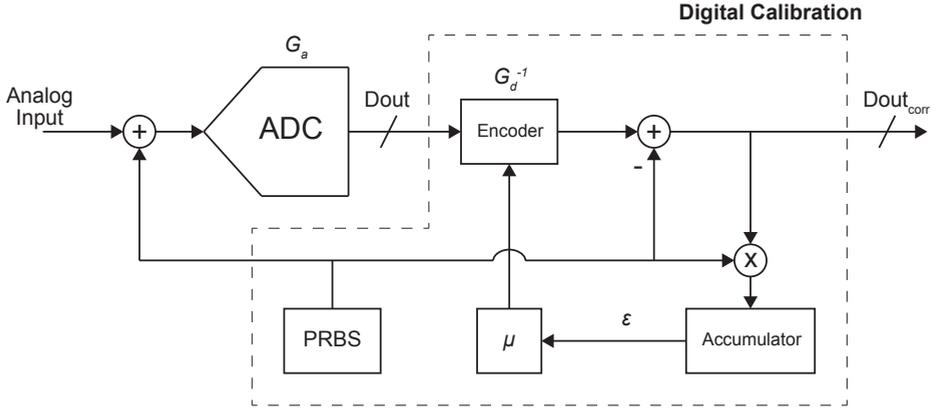


Figure 2.10: Correlation-based background calibration technique with PRBS signal

be estimated in two different modes – foreground and background. In the foreground mode [16–18], the ADC is taken offline by switching its input to a known calibration test signal, normally during startup. By analyzing its response to that test signal, the errors can be quickly determined. Since the ADC input is known, the amount of averaging required to suppress the interference of any other signal on the error estimation process is very low. Such deterministic methods typically achieve very fast convergence, while providing good calibration accuracy and low calibration logic power.

While foreground mode calibration methods can provide quick and accurate estimates of ADC errors, they cannot be used after startup. This limits their suitability in applications where it is necessary to track and calibrate the ADC errors during normal operation. For instance, in applications that involve large temperature and supply variations during the operation of the chip, it might be difficult to ensure the ADC resolution with just a foreground calibration at startup. For these applications, it is necessary to continuously monitor the ADC errors in background, while it is processing the input signal. Some approaches attempt to partially solve this by using a queue-based approach [19, 20], in which, during the normal operation of the ADC, some of the samples are reserved for calibration. During these samples, the ADC is no longer processing the input signal. While this approach helps the calibration engine to continuously monitor the PVT conditions, it comes at the cost of ADC throughput.

Using a background calibration algorithm [21–23] makes it possible to continuously track the ADC errors without sacrificing the ADC throughput. However, since

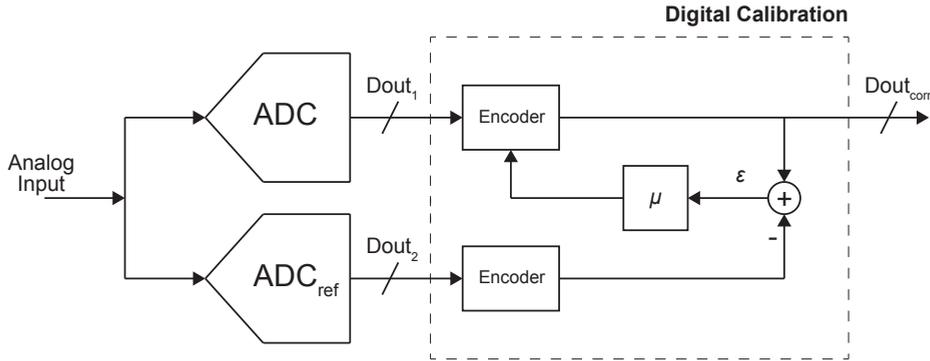


Figure 2.11: Deterministic background calibration technique with a parallel ADC path

the ADC is digitizing the normal input signal, it becomes a lot more complicated to decorrelate and extract the error signal from the ADC output. One of the ways to do that is by injecting a pseudo-random based binary sequence (PRBS) (+1,-1) as the calibration signal, as shown in Figure 2.10. This PRBS signal can be injected into the CADC or the sub-DAC of the pipeline stage, or directly into the input signal. In the digital backend, the digital equivalent of the PRBS signal is subtracted from the encoded ADC output. The resulting digital output contains the digitized ADC input and an error signal that is proportional to the difference between the analog and the digital radix. The error signal can be extracted by correlating this digital output with the PRBS signal through multiplication with the PRBS code, and then averaging the resulting product. This error signal can subsequently be used, with a certain gain, to adjust the radix in the encoder. To reliably separate the error signal from the input, a sufficiently random PRBS signal is required along with a lot of averaging, especially for amplifier nonlinearity calibration (empirically shown to be 2^N samples, where N bits is the desired ADC resolution [22]). Although these correlation-based approaches do not add any major additional analog hardware, they require a lot of ADC samples to perform a reasonably accurate calibration.

Some deterministic background calibration approaches have also been published in the literature [24–28]. These usually employ an auxiliary path such as a slow, accurate reference ADC which operates in parallel [24–26], as illustrated in Figure 2.11. The output of the ADC under calibration is compared to the reference ADC, and the difference between the two outputs provides the error signal. Since the linearity of the reference ADC limits the accuracy of the main ADC, the reference ADC is often also calibrated. These approaches, like the foreground methods, can converge quickly while providing an accurate estimate of the error signal. However,

they trade off calibration speed with higher power dissipation and area arising from the additional reference ADC. Some variants of this approach have tried to mitigate the additional hardware penalty by utilizing the ADCs in both the parallel paths for digitizing the input signal. Split-ADC technique [27, 28] is one of the primary examples of this and will be covered in more detail in the upcoming chapters.

2.3. CONCLUSION

In this chapter, one of the most popular ADC architectures – the pipelined ADC – has been reviewed. The key building blocks of a pipelined ADC - CADC, sub-DAC and the residue amplifier - were discussed in detail, along with the effect of their non-idealities on the ADC performance. Finally, some of the prominent techniques used to mitigate the effect of these circuit non-idealities were reviewed, with special emphasis on digital calibration.

REFERENCES

- [1] S. H. Lewis and P. R. Gray, “A pipelined 5MHz 9b ADC,” in *ISSCC Dig. of Tech. Papers*, 1987.
- [2] S. Louwsma, M. V. A. van Tuijl, and B. Nauta, “A 1.35 GS/s 10 b 175 mW time-interleaved A/D converter in 0.13 μ m cmos,” *IEEE J. Solid-State Circuits*, vol. 43, pp. 778–786, Apr. 2008.
- [3] C. Lee and M. Flynn, “A 12 b 50 MS/s 3.5 mW SAR-assisted 2-stage pipeline ADC,” pp. 239–240, 2010.
- [4] S. H. Lewis and et al., “A 10-b 20-Msample/s analog-to-digital converter,” *IEEE J. Solid-State Circuits*, vol. 27, pp. 351–358, Mar. 1992.
- [5] S. H. Lewis, “Optimizing the stage resolution in pipelined, multistage, analog-to-digital converters for video-rate applications,” *IEEE Trans. Circuits Syst.-II*, vol. 39, pp. 516–523, Aug. 1992.
- [6] K. Bult and G. Geelen, “A fast-settling CMOS op amp for SC circuits with 90-dB DC gain,” *IEEE J. Solid-State Circuits*, vol. 25, pp. 1379–1384, Dec. 1990.
- [7] K. Nagaraj and et al., “Switched-capacitor circuits with reduced sensitivity to amplifier gain,” *IEEE Trans. Circuits Syst.*, vol. 34, pp. 571–574, May 1987.
- [8] B. R. Gregoire and U. K. Moon, “An over-60dB true rail-to-rail performance using correlated level shifting and an opamp with only 30dB loop gain,” *IEEE J. Solid-State Circuits*, vol. 43, pp. 2620–2630, Dec. 2008.
- [9] H. H. Boo, D. S. Boning, and U. K. Moon, “A 12b 250MS/s pipelined ADC with virtual ground reference buffers,” *IEEE J. Solid-State Circuits*, vol. 50, pp. 2912–

- 2921, Dec. 2015.
- [10] T. Sepke and et al., "Comparator-based switched-capacitor circuits for scaled CMOS technologies," in *ISSCC Dig. of Tech. Papers*, pp. 220–221, 2006.
- [11] L. Brooks and H. S. Lee, "A zero-crossing-based 8b 200MS/s pipelined ADC," in *ISSCC Dig. of Tech. Papers*, pp. 460–461, 2007.
- [12] S. K. Shin and et al., "A 12bit 200MS/s zero-crossing-based pipelined ADC with early sub-ADC decision and output residue background calibration," *IEEE J. Solid-State Circuits*, vol. 49, pp. 1366–1382, June 2014.
- [13] B. Hershberg and et al., "Ring amplifiers for switched capacitor circuits," *IEEE J. Solid-State Circuits*, vol. 49, p. 2928–2942, Dec. 2012.
- [14] N. Dolev, M. Kramer, and B. Murmann, "A 12-bit, 200-MS/s, 11.5-mW pipeline ADC using a pulsed bucket brigade front-end," pp. 98–99, 2008.
- [15] I. Ahmed, J. Mulder, and D. A. Johns, "A low-power capacitive charge pump based pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 45, pp. 1016–1027, May 2010.
- [16] A. N. Karanicolas, H. S. Lee, and K. L. Bacrania, "A 15-b 1-Msample/s digitally self-calibrated pipeline ADC," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1207–1215, Dec. 1993.
- [17] I. E. Opris, L. D. Lewicki, and B. C. Wong, "A single-ended 12-bit 20 Msample/s self-calibrating pipeline A/D converter," *IEEE J. Solid-State Circuits*, vol. 33, pp. 1207–1215, Dec. 1998.
- [18] A. Verma and B. Razavi, "A 10-bit 500-MS/s 55-mW CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 44, pp. 3039–3050, Nov. 2009.
- [19] C. R. Grace, P. J. Hurst, and S. H. Lewis, "A 12-bit 80-MSample/s pipelined ADC with bootstrapped digital calibration," *IEEE J. Solid-State Circuits*, vol. 40, pp. 3039–3050, May 2005.
- [20] B. D. Sahoo and B. Razavi, "A 12-bit 200-MHz CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 44, pp. 2366–2380, Sep. 2009.
- [21] J. Ming and S. H. Lewis, "An 8b 80Msample/s pipelined ADC with background calibration," in *ISSCC Dig. of Tech. Papers*, pp. 42–43, 2000.
- [22] B. Murmann and B. E. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, p. 2040 – 2050, Dec. 2003.
- [23] A. Panigada and I. Galton, "A 130 mw 100 MS/s pipelined ADC with 69dB SNDR enabled by digital harmonic distortion correction," *IEEE J. Solid-State Circuits*, vol. 44, p. 3314 – 3328, Dec. 2009.
- [24] K. Dyer and et al., "Analog background calibration of a 10-b 40 MS/s parallel pipelined ADC," in *ISSCC Dig. of Tech. Papers*, pp. 142–143, 1998.

-
- [25] X. Wang, P. Hurst, and S. Lewis, "A 12-bit 20-Msample/s pipelined analog-to-digital converter with nested digital background calibration," *IEEE J. Solid-State Circuits*, vol. 39, pp. 1799–1808, Nov. 2004.
- [26] H. Wang and et al., "Nested digital background calibration of a 12-bit pipelined ADC without an input SHA," *IEEE J. Solid-State Circuits*, vol. 44, pp. 2780–2789, Oct. 2009.
- [27] J. Li and U.-K. Moon, "Background calibration techniques for multi-stage pipelined ADCs with digital redundancy," *IEEE Trans. Circuits Syst.-II*, vol. 50, pp. 531–538, Sep. 2003.
- [28] J. McNeill and et al., "Split ADC" architecture for deterministic digital background calibration of a 16-bit 1-MS/s ADC," *IEEE J. Solid-State Circuits*, vol. 40, pp. 2437 – 2445, Dec. 2005.

3

Analysis of Residue Amplifiers

One of the key steps in the digitization of an analog signal is sampling. This imposition of a time stamp on the input signal makes an ADC an inherently discrete-time system. As a consequence, all the major ADC sub-blocks, including the residue amplifiers, operate in a discrete-time environment with their operation synchronized to a periodic clock.

Discrete-time amplifiers are characterized by the fact that their operation is divided into different phases. For any given input sample, the amplification typically starts after a complete reset and the output voltage develops across a memory element, such as a capacitor. As a result, most of the metrics of a discrete-time amplifier -gain, SNR and THD- are dynamic and transient in nature.

In the past, residue amplifiers were allowed to settle completely and reach a steady-state, making their noise and distortion performance analogous to a continuous-time amplifier. However, as ADC speeds have increased, there has been a concerted focus on lowering their power consumption by moving to residue amplifiers that only settle partially. Due to the advent of low-cost digital processing, digital calibration has become a popular approach, whereby the inaccuracies of the residue amplifiers are detected and corrected for with the help of digital post-processing of the ADC output. While digital calibration can reduce the effect of insufficient gain and linearity in the residue amplifier on the ADC output linearity, it can not compensate

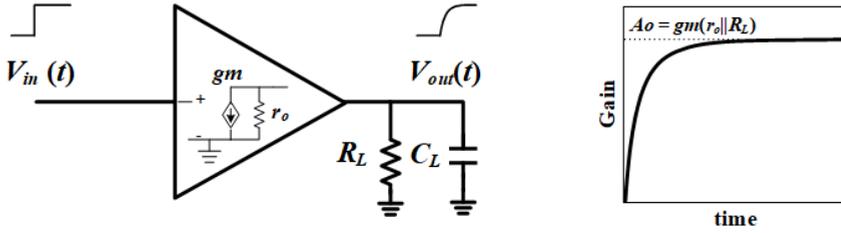


Figure 3.1: Operation of a discrete-time amplifier

for the loss of dynamic range. Hence, in order to optimize the power dissipation of the residue amplifiers, they need to be designed to achieve a certain performance before digital calibration is used to fill in the gap. However, due to their discrete-time behavior, the dynamic nature of a residue amplifier's performance can make it difficult to identify the optimum point of operation for calibration-assisted residue amplifiers. Hence, it is important to analyze the discrete-time amplifiers used for residue amplification in time-domain.

In this chapter, we will develop time-domain expressions for the most important design metrics of a residue amplifier - gain, noise, and distortion. We will consider the behavior of residue amplifiers in open- and closed-loop configurations, and identify ways to optimize their power efficiency in the presence of digital calibration.

3.1. SETTLING BEHAVIOR OF DISCRETE-TIME AMPLIFIERS

Residue amplifiers can be implemented either as open- or closed-loop configurations. Each configuration comes with certain advantages and drawbacks. Due to their simplicity, we will first analyze the behavior of discrete-time open-loop amplifiers.

Figure 3.1 shows a single-pole open-loop amplifier with an input transconductance of g_m and an output resistance of r_o (or admittance of g_o), yielding an intrinsic dc gain, A_o , of $g_m r_o$ (or g_m/g_o). As the intrinsic gain of the amplifier can be fairly large, it is often loaded by an additional load resistor, R_L , to achieve the desired gain. When driving this load resistor, along with a load capacitor of C_L , the amplifier output voltage response, $V_{out}(t)$, to a voltage input, $V_{in}(t)$, will settle with finite speed towards a steady-state gain of $g_m R_{out}$, where R_{out} is $r_o || R_L$. This output voltage response can be calculated by using the following differential equation

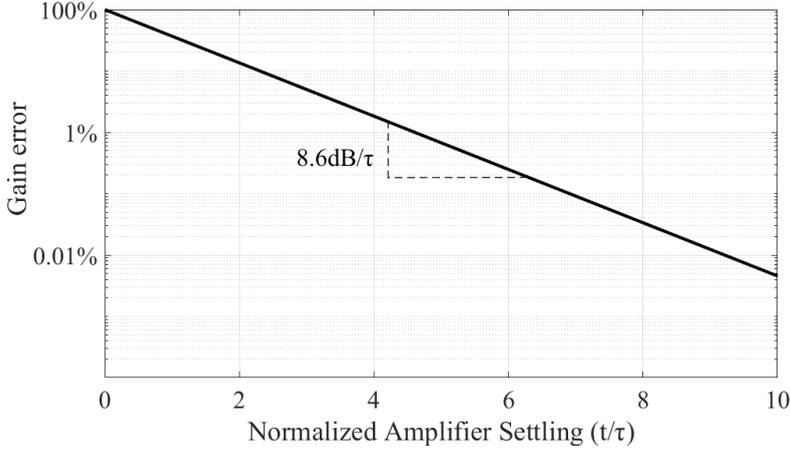


Figure 3.2: Gain error vs amplifier settling accuracy

$$I_x(t) = g_m V_{in}(t) = C_L \frac{dV_{out}(t)}{dt} + g_{out} V_{out}(t) \quad (3.1)$$

where $I_x(t)$ is the output current produced by the input transconductance, g_m , and g_{out} is the output admittance ($1/R_{out}$). Since a discrete-time amplifier operates on a sampled-and-held input, we can consider the input signal to be a step function of amplitude V_{in} . Hence, with the help of Laplace transform, equation (3.1) can be used to calculate $V_{out}(t)$ as –

$$V_{out}(t) = A(1 - e^{-t/\tau})V_{in}, \quad t \geq 0 \quad (3.2)$$

where A is the steady-state gain ($g_m R_{out}$) and τ is the time constant ($R_{out} C_L$) that defines the settling speed of the amplifier. Equation (3.2) describes a purely first-order exponential settling, resulting from the single-pole nature of the amplifier. Considering A to be the desired gain, the amplifier gain error over time can be derived as –

$$\varepsilon_g(t) = A - A(1 - e^{-t/\tau})V_{in} = Ae^{-t/\tau}, \quad t \geq 0 \quad (3.3)$$

The gain error is plotted in Figure 3.2 with respect to amplifier settling, which is expressed as t/τ . It can be seen that the gain error decreases at the rate of $20\log_{10}(e)$ or 8.6dB for every unit τ of settling. This particular rate of decrease

in the amplifier gain error also indicates first-order settling and can be used to characterize the settling response of an amplifier.

From (3.3), we can conclude that the gain error for a discrete-time amplifier is a function of time. This means that the amplifier makes a smaller gain error if it is allowed a longer amplification time. This transient behavior can also be seen in the amplifier's noise and linearity performance, and its power dissipation. Hence, in order to properly understand the design trade-offs, we will analyze the open-loop amplifier's noise and linearity performance w.r.t the amplifier settling in the following sub-sections.

3.1.1. NOISE

To calculate the noise power of the discrete-time amplifier shown in Figure 3.1, we need to consider the noise power spectral density from all the noise sources in the amplifier. To simplify the analysis, let us assume that the only noise sources are the input devices and the load resistor. Ignoring the contribution of flicker noise, the output-referred noise current power spectral density can be expressed as –

$$\overline{i_{o,n}^2(t)} = kT \left(g_m + \frac{1}{R_L} \right) \quad (3.4)$$

Using the Langevin form of the stochastic differential equation for the amplifier in Figure 3.1, the output-referred noise power can be calculated as shown in [1] –

$$\overline{v_{o,n}^2(t)} = \frac{kT}{C_L} (1 + A) (1 - e^{-2t/\tau}) + \frac{kT}{C_L} e^{-2t/\tau} \quad (3.5)$$

Since the load capacitor is assumed to undergo a complete reset before the amplification begins, this results in the capacitor sampling a noise power of kT/C_L at $t=0$. Based on the output noise power and the signal gain derived previously, the amplifier SNR can be expressed as –

$$\text{SNR}(t) = \frac{V_{out}^2(t)}{v_{o,n}^2(t)} = \frac{\{A(1 - e^{-t/\tau})\}^2}{\frac{kT}{C_L} (1 + A) (1 - e^{-2t/\tau}) + \frac{kT}{C_L} e^{-2t/\tau}} \quad (3.6)$$

Figure 3.3 shows the output-referred SNR versus amplifier settling. It can be seen that the amplifier is able to achieve more than 90% of its final SNR in the first three time-constants of settling.

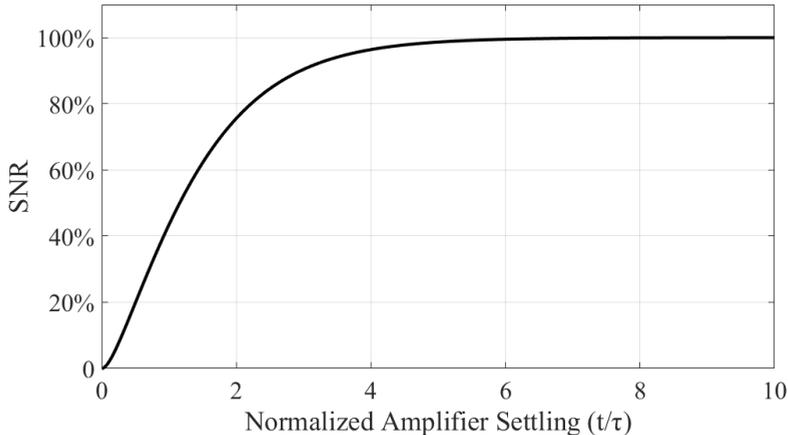


Figure 3.3: Settling behavior of SNR of an open-loop discrete-time amplifier

3.1.2. DISTORTION

Along with noise and gain error, amplifier distortion is amongst the most significant sources of error in residue amplification. Any nonlinear mechanism inside the amplifier could generate distortion components which will corrupt the linearity of the amplifier output. Since nonlinearity in the residue signal results in unwanted spurs in the ADC's digital output, any amplifier distortion will have a detrimental impact on the ADC's resolution. This makes it one of the most important design constraints for the amplifier.

Just like gain error and noise, amplifier distortion also shows a transient behavior in discrete-time amplifiers. Although we only care about the amplifier's linearity at the moment its output is sampled, it is important to know how the distortion components within the amplifier vary with time in order to optimize the amplifier power efficiency. In this sub-section, we will analyze the settling behavior of the distortion components arising from the input transconductance and the output impedance of an open-loop amplifier.

A non-linear single-stage amplifier modeled with input and output distortion sources is shown in Figure 3.4. For ease of calculation, the analysis is performed with admittances instead of impedances, and we ignore the effect of R_L . The input transconductance is modeled with second- and third-order distortion coefficients, g_{i2} and g_{i3} , respectively. Hence, the distortion components arising at the output due to the nonlinear g_m can be calculated by replacing the g_m in (3.1) with the nonlinear g_m characteristic (detailed derivation presented in Appendix A). As a result, the

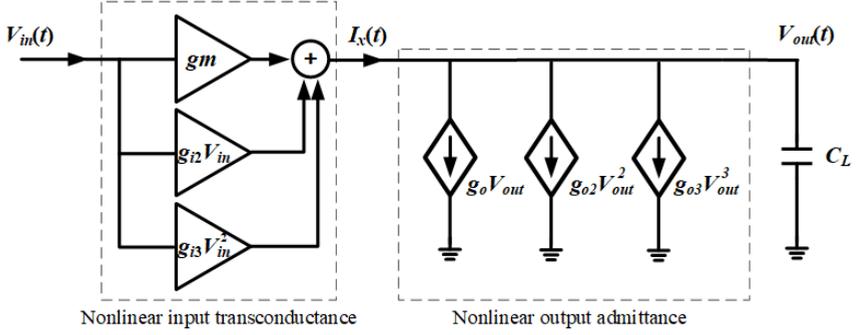


Figure 3.4: Amplifier distortion model

second-order input distortion component can be expressed as –

$$V_{out,in}^{(2)}(t) = \frac{g_{i2}}{g_o} (1 - e^{-t/\tau}) V_{in}^2, \quad t \geq 0 \quad (3.7)$$

Similarly, the third-order distortion component can be calculated as –

$$V_{out,in}^{(3)}(t) = \frac{g_{i3}}{g_o} (1 - e^{-t/\tau}) V_{in}^3, \quad t \geq 0 \quad (3.8)$$

Apart from the input transconductance, a voltage-dependent output impedance is also a major source of amplifier distortion. Depending on the voltage headroom available to the transistors, a relatively large signal swing at the output may end up modulating the amplifier output impedance, usually compressing the overall amplifier gain.

Since the output distortion depends on the magnitude of output voltage itself, its settling behavior is different from the input distortion. We can calculate it by again assuming the output impedance to have second- and third-order non-linearity as shown in Figure 3.4. If the amplifier output admittance can be written as –

$$g_{out} = g_o + g_{o2} V_{out}(t) + g_{o3} V_{out}^2(t) \quad (3.9)$$

then using a similar analysis as for the input distortion (presented in Appendix A), the second- and third-order output distortion components can be calculated as –

$$V_{out,out}^{(2)}(t) = \frac{g_{o2} g_m^2}{g_o^3} \left\{ 1 - 2 \left(\frac{t}{\tau} \right) e^{-t/\tau} - e^{-2t/\tau} \right\} V_{in}^2, \quad t \geq 0 \quad (3.10)$$

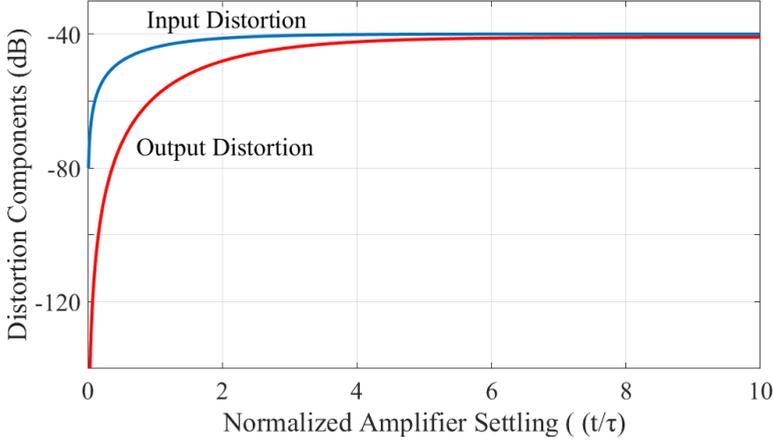


Figure 3.5: Settling behavior of amplifier distortion components

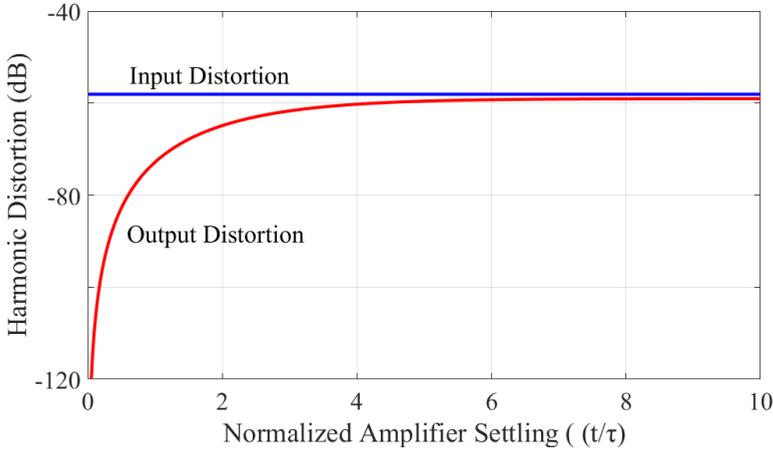


Figure 3.6: Settling behavior of harmonic distortion originating from the amplifier input transconductance and output impedance

$$\begin{aligned}
 V_{out,out}^{(3)}(t) = & \frac{g_{o3}g_m^3}{g_o^4} \left\{ (1 - e^{-t/\tau}) + \frac{3}{2}e^{-t/\tau} (1 - e^{-2t/\tau}) \right. \\
 & \left. + 3e^{-t/\tau} (1 - e^{-t/\tau}) - 3\left(\frac{t}{\tau}\right)e^{-t/\tau} \right\} V_{in}^3, \quad t \geq 0
 \end{aligned} \tag{3.11}$$

Figure 3.5 shows the settling behavior of the 2nd-order distortion components arising from input and output nonlinear mechanisms with a magnitude (normalized

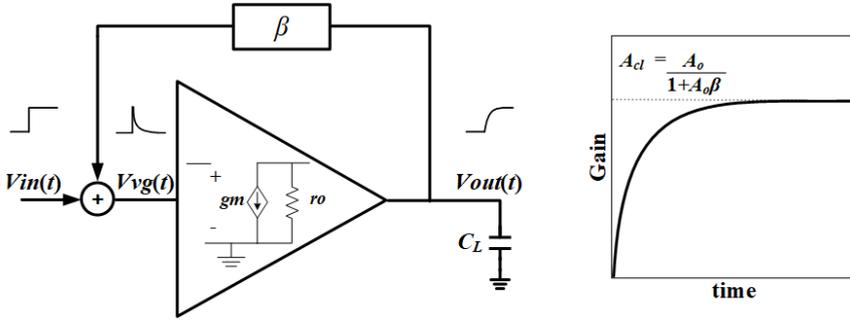


Figure 3.7: Amplifier in negative feedback with input, virtual ground and output signals

to the fundamental component) of approximately 0.1%. It can be seen that the input distortion component settles in the same way as the fundamental output signal, while the output distortion component settles a bit slower. This is also illustrated in Figure 3.6 where the two distortion components are plotted after being normalized with the fundamental output signal.

3.2. DISCRETE-TIME AMPLIFIERS WITH NEGATIVE FEEDBACK

As discussed in the previous section, the gain of an open-loop amplifier is directly dependent on the input transconductance and the output impedance of the amplifier. This makes the amplifier gain highly sensitive to the amplifier's bias point and its environment, and any kind of PVT variation could have a significant effect on the amplifier gain. Although this sensitivity towards PVT variations could be mitigated by applying design techniques such as a constant- g_m bias, their effectiveness is limited, precluding their use in high-resolution applications.

To make the gain more independent of PVT variations, amplifiers are often used in a negative-feedback configuration, which involves feeding the amplifier output back to its input through a feedback network, as shown in Figure 3.7. For amplifiers with high inherent gain, the output signal is driven in order to completely nullify any signal at the input nodes, creating a virtual ground at its input nodes. As a result, the overall closed-loop gain is completely determined by the feedback factor, β , which is a ratio of passive components, thus, reducing its sensitivity towards process variation, temperature drift, and other environmental changes.

For an amplifier employed in negative-feedback using a passive feedback network

with a gain of β , as shown in Figure 3.7, the overall output voltage $V_{out}(t)$ can be calculated by using the following series of equations –

$$\begin{aligned} V_{vg}(t) &= V_{in}(t) - \beta V_{out}(t) \\ I_x(t) &= g_m V_{vg}(t) = C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) \end{aligned} \quad (3.12)$$

where $V_{vg}(t)$ is the voltage signal at the virtual ground inputs of the amplifier, $I_x(t)$ is the output current produced by the input transconductance, g_m and g_o is the open-loop output transconductance of the amplifier, resulting in an open-loop gain, A_o , of g_m/g_o . If the input signal is again considered to be a step function of amplitude V_{in} , then, with the help of Laplace transform, (3.12) can be used to calculate $V_{out}(t)$ as –

$$V_{out}(t) = \frac{A_o}{1 + A_o\beta} (1 - e^{-t/\tau_f}) V_{in}, \quad t \geq 0 \quad (3.13)$$

where τ_f is the loop time-constant and can be expressed as –

$$\tau_f = \frac{\tau}{1 + A_o\beta} \quad (3.14)$$

For an extremely large loop gain ($A_o\beta \gg 1$), the amplifier closed-loop gain settles to simply $1/\beta$, making it independent of the amplifier's input g_m and output impedance. It is also evident from (3.14) that the closed-loop time-constant gets reduced by the loop-gain, thus increasing the 3-dB bandwidth of the amplifier. Apart from the output voltage, the virtual ground signal, $V_{vg}(t)$, is also an important indicator of settling in a closed-loop amplifier. By combining equations (3.12) and (3.13), $V_{vg}(t)$ can be expressed as –

$$V_{vg}(t) = V_{in} - \beta V_{out}(t) = \frac{1 + A_o\beta e^{-t/\tau_f}}{1 + A_o\beta} V_{in}, \quad t \geq 0 \quad (3.15)$$

From (3.15)), it can be seen that the virtual ground signal decays with time at the rate of 8.6dB per τ of settling, before settling to a residual value of $V_{in}/(1+A_o\beta)$. Just like in the open-loop case, this rate of decay is an indicator of a purely first-order exponential settling, resulting from the single-pole nature of the amplifier.

Monitoring the virtual ground signal provides key insights into how the amplifier behaves in a negative-feedback loop. As evident in (3.15)), the amplifier effectively tries to suppress any signal present at the virtual ground nodes, and while the limit of suppression is determined by the loop gain, the rate of suppression is defined by the closed-loop time constant.

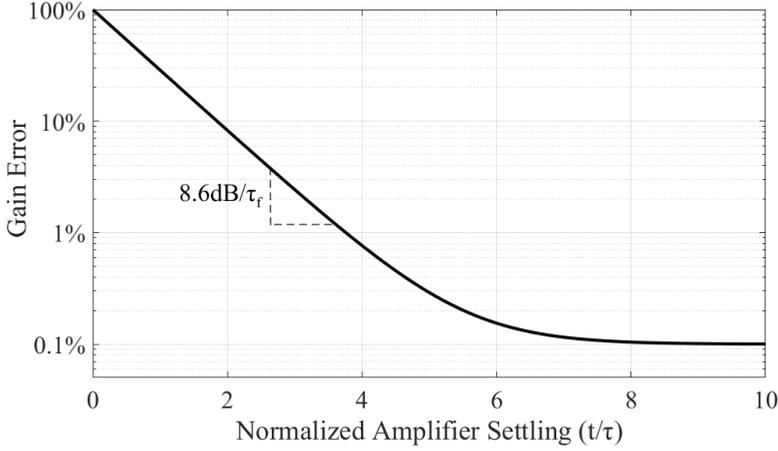


Figure 3.8: Settling behavior of gain error in closed-loop amplifier

3.2.1. GAIN ERROR

Considering the desired gain of the closed-loop amplifier shown in Figure 3.7 to be $1/\beta$, the gain error over time can be calculated using (3.15) as -

$$\begin{aligned} \text{Gain error, } \varepsilon_g(t) &= \frac{1}{\beta} - \frac{A_o}{1 + A_o\beta} (1 - e^{-t/\tau_f}) \\ &= \frac{1 + A_o\beta e^{-t/\tau_f}}{1 + A_o\beta} \frac{1}{\beta}, \quad t \geq 0 \end{aligned} \quad (3.16)$$

The gain error expression above consists of a static and a dynamic part. The static part represents the residual steady-state gain error due to the limited loop-gain. The dynamic part of the gain error describes the initial part of the amplification when the amplifier is still settling at a rate of $8.6\text{dB}/\tau_f$, as shown in Figure 3.8.

Using (3.16), and (2.2) from chapter 2 that describes the resolution required in the ADC stage, the loop gain and bandwidth requirements for a closed-loop residue amplifier of an ADC stage can be expressed as -

$$\text{Relative gain error} = \frac{1 + A_o\beta e^{-t/\tau_f}}{1 + A_o\beta} \leq \frac{1}{2^{1+r_{i+1}}} \quad (3.17)$$

where r_{i+1} is the combined ADC resolution remaining after the i -th stage. Hence, for a fully settled amplifier, the minimum required DC loop gain can be calculated as -

$$\frac{1}{1 + A_o\beta} \leq \frac{1}{2^{1+r_{i+1}}} \quad (3.18)$$

or,

$$A_o \geq \frac{2^{1+r_{i+1}}}{\beta} \quad (3.19)$$

To put that into perspective, for the first stage of a 12-bit pipelined ADC with 1-bit per stage architecture, the residue amplifier (with a gain of 2x) needs to achieve a relative accuracy of at least $1/2^{12}$. In order to accomplish this without any calibration, the amplifier needs to possess an open-loop DC gain of more than 8192 or 78dB.

Along with sufficient DC gain, the amplifier also needs to possess a certain bandwidth in order to achieve the gain with the required accuracy within the specified amplification period. This bandwidth requirement can also be calculated from (3.17), by assuming the op-amp to have sufficiently high DC gain ($A_o\beta \gg 1$). In that case, for an amplification period of t_s , (3.17) can be simplified to –

$$e^{-t_s/\tau_f} \leq \frac{1}{2^{1+r_{i+1}}} \quad (3.20)$$

or,

$$t_s \geq (1 + r_{i+1})\tau_f \ln 2 \quad (3.21)$$

Assuming t_s to be half of the ADC clock period, the 3-dB bandwidth of the closed-loop amplifier can be written as –

$$f_{3dB} \geq \frac{(1 + r_{i+1})f_s \ln 2}{\pi} \quad (3.22)$$

or,

$$f_u \geq \frac{(1 + r_{i+1})f_s \ln 2}{\pi\beta} \quad (3.23)$$

where f_s is the ADC sampling frequency and f_u is the unity-gain frequency of the op-amp and is expressed as βf_{3dB} . Based on the above expression, it can be seen that if the 12-bit pipelined ADC in the previous design example was to be operated at a speed of 200MS/s, then the amplifier in the first stage will need a unity-gain bandwidth of nearly 1GHz to completely settle to the desired gain with the required accuracy.

3.2.2. NOISE

In order to analyze the noise performance of the closed-loop amplifier, we consider an inverting amplifier with a capacitive feedback network, as shown in Figure 3.9, with all the capacitors being completely reset before the beginning of the amplification period. When an input step of V_{in} is applied at $t = 0$, it propagates through the capacitive network and generates spikes at the virtual ground input and the amplifier output nodes. These voltage spikes can be expressed as –

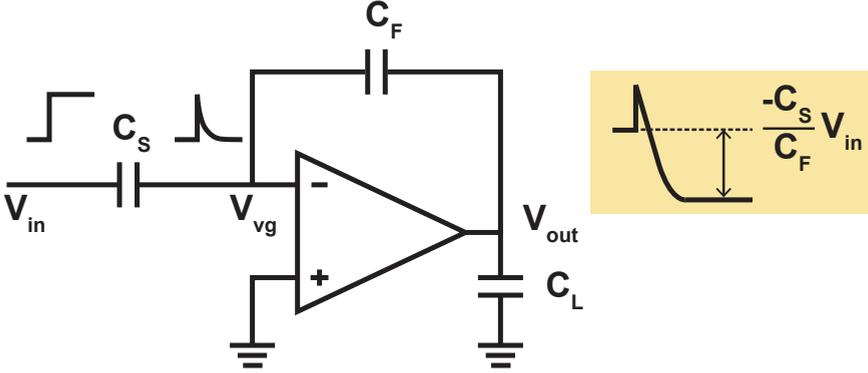


Figure 3.9: Closed-loop charge amplifier

$$\begin{aligned}
 V_{vg}(t = 0^+) &= V_{in} \frac{C_S(C_L + C_F)}{C_S C_F + C_S C_L + C_F C_L} \\
 V_{out}(t = 0^+) &= V_{in} \frac{C_S C_F}{C_S C_F + C_S C_L + C_F C_L}
 \end{aligned} \tag{3.24}$$

After the initial kick, the amplifier begins to drive the output in order to reduce the virtual ground signal to zero through negative feedback. As the amplifier output is settling, the transient closed-loop gain of the inverting amplifier can be expressed by taking a ratio of the amplifier output with the input step. Assuming the DC loop gain to be sufficiently large, the overall closed-loop gain, $A_{cl}(t)$, can be derived as –

$$\begin{aligned}
 A_{cl}(t) &= \frac{C_S C_F}{C_S C_F + C_S C_L + C_F C_L} e^{-t/\tau_f} - \frac{C_S}{C_F} (1 - e^{-t/\tau_f}) \\
 &= \frac{-C_S}{C_F} \left(1 - \frac{(C_L + C_F)(C_S + C_F)}{C_S C_F + C_S C_L + C_F C_L} e^{-t/\tau_f} \right)
 \end{aligned} \tag{3.25}$$

where τ_f is the time constant of the closed-loop amplifier. The first part in the equation represents the steady-state gain, $-C_S/C_F$, and the second half represents the gain error due to limited settling in the amplifier. Neglecting any other sources of noise in the amplifier, the noise current density from the input transistor is given by $4kTg_m$, with g_m being the input transconductance. If the noise bandwidth is limited only by C_L , then the output-referred noise power at time t can be calculated using the analysis in [2] –

$$\overline{v_{o,n}^2(t)} = \frac{kT}{C_L} \frac{A_o}{1 + A_o \beta} (1 - e^{-2t/\tau_f}) + \frac{kT}{C_L} e^{-2t/\tau_f} \tag{3.26}$$

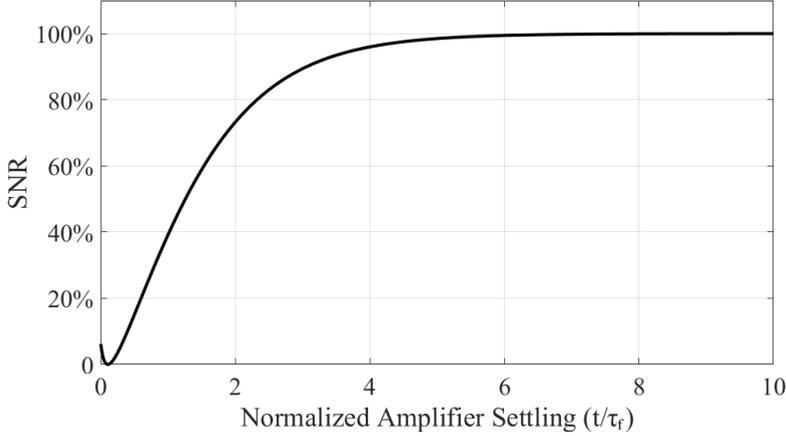


Figure 3.10: Settling behavior of SNR in closed-loop amplifier

where A_o is the open-loop gain of the amplifier, and β is the feedback factor, given by $C_F/(C_S + C_F)$. Just like in the case of an open-loop amplifier, the second term in (3.26) results from the noise power of load capacitor voltage at $t = 0$, which, if assumed to be perfectly reset, is given by kT/C_L . Considering $A_o\beta \gg 1$, we obtain the following expression -

$$\begin{aligned} \overline{v_{o,n}^2(t)} &\approx \frac{kT}{C_L} \frac{1}{\beta} \left(1 - (1 - \beta)e^{-2t/\tau_f}\right) \\ &\approx \frac{kT(C_S + C_F)}{C_F C_L} \left(1 - \frac{C_S}{C_S + C_F} e^{-2t/\tau_f}\right) \end{aligned} \quad (3.27)$$

The output signal power is proportional to the square of closed-loop gain. Therefore, using (3.25) and (3.27), the overall SNR at the output of the amplifier at any moment in time, t , can be expressed as -

$$\text{SNR}(t) = \frac{v_o^2(t)}{v_{o,n}^2(t)} = \frac{\left\{ \frac{C_S}{C_F} \left(1 - \frac{(C_L + C_F)(C_S + C_F)}{C_S C_F + C_S C_L + C_F C_L} e^{-t/\tau_f}\right) \right\}^2}{\frac{kT(C_S + C_F)}{C_F C_L} \left(1 - \frac{C_S}{C_S + C_F} e^{-2t/\tau_f}\right)} \quad (3.28)$$

Figure 3.10 shows the settling behavior of SNR versus the amplifier settling, as described by (3.28), for a closed-loop amplifier with a gain of $2x$ ($C_S = 2C_F = 2C_L$). As the amplifier output always crosses zero in its buildup from the initial feedforward kick to its final value, the output power initially reduces to zero before increasing exponentially. On the other hand, the noise power at the output sees a relatively small change during the amplifier settling. As a result, after the initial dip, the

SNR builds up roughly the same way as the output signal power and the amplifier achieves nearly 90% of its steady-state SNR in the first three time-constants.

3.2.3. DISTORTION

It was shown earlier that the application of feedback desensitized the overall amplifier gain to the amplifier's intrinsic gain by a factor of $1 + A_o\beta$. This means that any input-dependent variation in the closed-loop gain due to nonlinear mechanisms in the amplifier also gets suppressed by the loop-gain [3, 4]. This suppression of amplifier distortion is one of the key advantages of negative feedback, and has made it one of the mainstays of amplifier design techniques for high-precision applications.

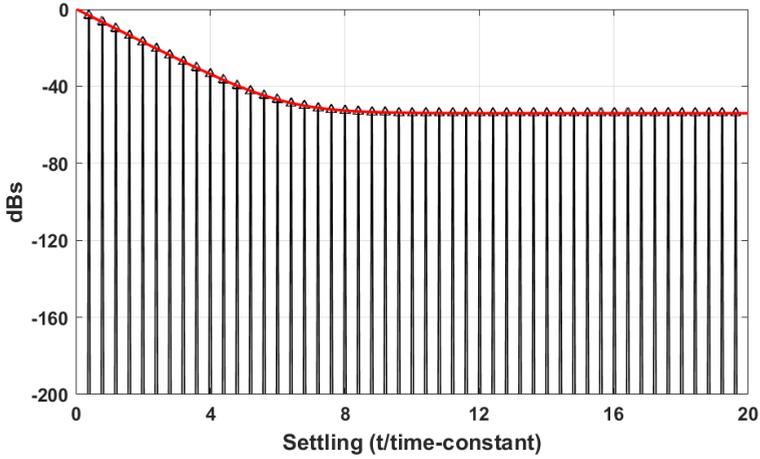
This mitigating effect of negative feedback on harmonic distortion has been documented extensively in the literature [3–6]. However, most of these approaches evaluate the distortion in frequency-domain. The frequency-domain approach essentially studies the amplifier in its steady-state, neglecting the way it behaves during the initial settling phase. For designing and optimizing discrete-time amplifiers for high-speed applications, however, it is important to understand how the effect of negative feedback evolves during the amplifier settling.

In this sub-section, we will extend the analysis of distortion components previously presented for open-loop amplifiers to study how negative-feedback affects the distortion components in discrete-time amplifiers. Along with mathematical expressions (with a detailed analysis in Appendix A), an intuitive description of the settling behavior of distortion components is presented below.

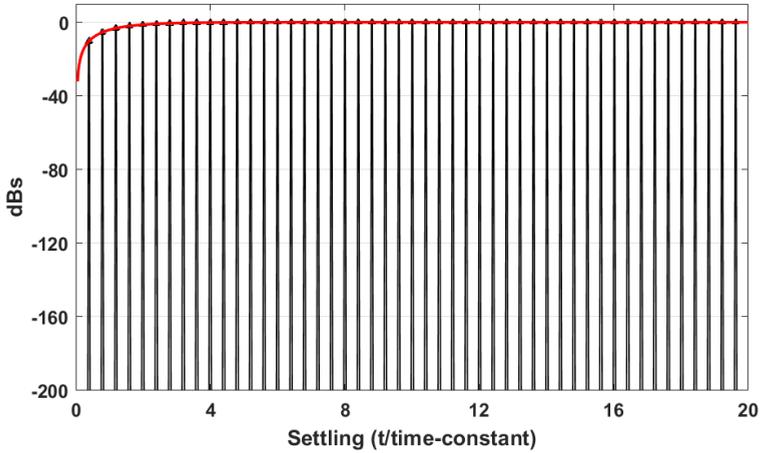
A. *Intuitive Explanation*

The transient behavior of a feedback amplifier's distortion components can be understood by considering the response of a negative-feedback system to a short voltage pulse. It is well established that any spurious signal in a negative-feedback loop gets suppressed by the overall loop gain [3]. However, the rate of suppression, as calculated previously, is determined by the 3-dB bandwidth or the time constant of the feedback loop. Once the pulse disappears, the amplifier's response to the pulse also starts to fade. However, this process is not infinitely fast, and, for a single-pole amplifier, occurs at the rate of 8.6dB per τ_f .

The above scenario can be extended to analyze the amplifier distortion by breaking down the signal responsible for generating the distortion into a series of pulses, as illustrated in Figure 3.11, and considering the response of the distortion mechanism and the negative-feedback loop to this series of pulses. This



(a)

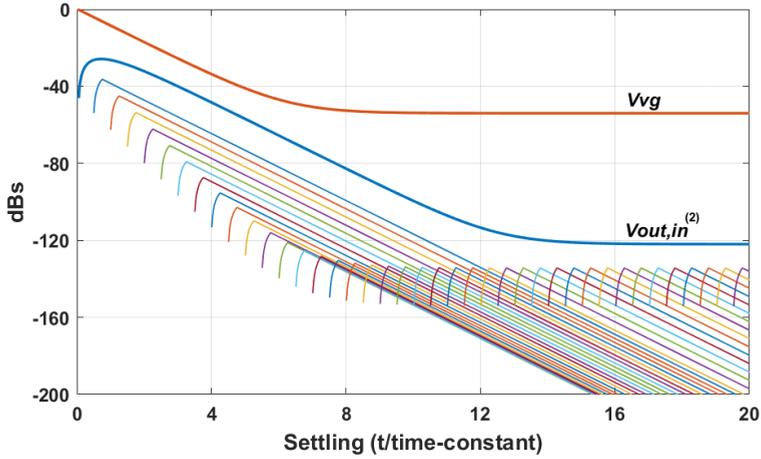


(b)

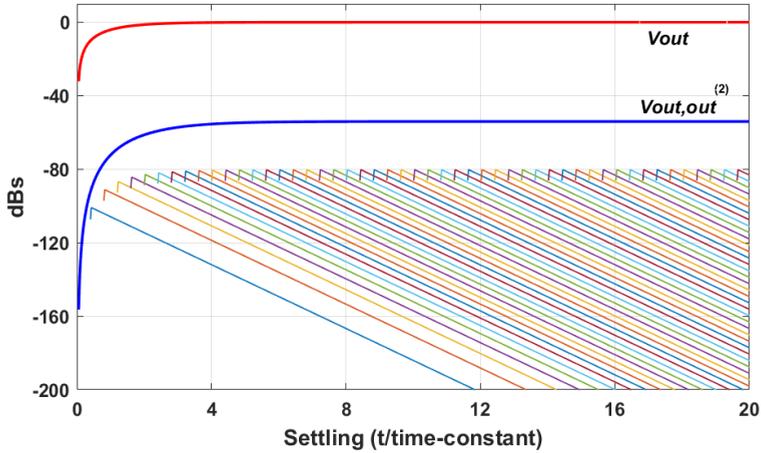
Figure 3.11: Amplifier (a) virtual ground signal and (b) output signal decomposed into individual pulses

process is basically like time-domain convolution, and the overall distortion at a certain moment in time will be the sum of all the distortion components triggered by these individual pulses while being suppressed by the feedback loop.

In the case of distortion originating from the input transconductance, the distortion mechanism is triggered by the virtual ground signal. A pulse at the



(a)



(b)

Figure 3.12: Second-order distortion generated in the amplifier by the (a) virtual ground signal and (b) output signal pulses

virtual ground node will generate a distortion voltage, which appears at the output at a finite speed due to the output pole. Once the input pulse disappears, the corresponding distortion voltage at the output also starts decaying at the rate defined by the loop time constant. While this distortion component is fading, the distortion component resulting from the next input pulse starts appearing at the output. As shown in Figure 3.12a, every individual input

pulse in the series of pulses will generate a distortion component that builds up and decays at the rate defined by the loop bandwidth.

Since negative-feedback essentially works towards eliminating the virtual ground signal, the magnitude of the input pulses decreases at the same rate. Hence, the magnitude of the distortion components generated by these pulses goes down by a factor of two in dBs in case of second-order distortion, and a factor of three for third-order distortion. However, as seen from Figure 3.12a, the overall sum of all distortion components at any moment is dominated initially by the first pulse and, as the amplifier approaches steady state, by the most recent pulse. The point of transition can be considered as the settling point for the input distortion.

The other significant non-ideal element limiting the amplifier linearity is the output impedance. The distortion arising from the nonlinear output impedance of the amplifier can be treated in a similar way as input distortion, as shown in Figure 3.12b. The output signal can be decomposed into a series of pulses, with each pulse creating a distortion component. At the beginning of the amplification period, the output voltage is close to zero and hence, and so are the corresponding distortion components. However, as the output voltage starts building up, the magnitude of the output distortion pulses also increases based on the loop time-constant and the order of distortion. These output distortion components keep increasing in magnitude before settling in a similar fashion as the amplifier output voltage. As evident from Figure 3.12b, the overall sum of all these distortion components at any instant in time is dominated by the most recent pulse.

B. *Mathematical Analysis*

A nonlinear negative-feedback amplifier can be analyzed by using the nonlinear amplifier model shown in Figure 3.4 in the closed-loop configuration illustrated in Figure 3.7. Starting with the nonlinear input transconductance, (3.12) can be rewritten to reflect the g_m nonlinearity in the following way -

$$\begin{aligned} V_{vg}(t) &= V_i(t) - \beta V_{out}(t) \\ I_x(t) &= g_m V_{vg}(t) + g_{i2} V_{vg}^2(t) + g_{i3} V_{vg}^3(t) \\ I_x(t) &= C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) \end{aligned} \quad (3.29)$$

It should be noted from (3.29) that there is no memory element present in the relationship between $I_x(t)$ and $V_{vg}(t)$. That is because the only memory

present in the circuit is associated with the load capacitance. By solving (3.29) with the help of some simplifications (Appendix A), the second-order input distortion component, $V_{out,in}^{(2)}(t)$, can be calculated as –

$$V_{out,in}^{(2)}(t) = \frac{g_{i2}}{g_o(1+A_o\beta)^3} V_{in}^2 \left((1 - e^{-t/\tau_f}) + (A_o\beta)^2 e^{-t/\tau_f} (1 - e^{-t/\tau_f}) + 2A_o\beta e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \quad (3.30)$$

This represents the transient behavior of the second-order distortion component from the nonlinear input g_m in response to a step input at $t = 0$. After the initial settling, the distortion component settles towards a steady-state value of $\frac{g_{i2}}{(g_o(1+A_o\beta))^3}$, which is the same as that derived in [4]. The third-order input distortion component, $V_{out,in}^{(3)}(t)$, can be similarly be expressed as –

$$V_{out,in}^{(3)}(t) = \frac{g_{i3}}{g_o(1+A_o\beta)^4} V_{in}^3 \left((1 - e^{-t/\tau_f}) + 3(A_o\beta)^2 e^{-t/\tau_f} (1 - e^{-t/\tau_f}) + \frac{(A_o\beta)^3}{2} e^{-t/\tau_f} (1 - e^{-2t/\tau_f}) + 3A_o\beta e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \quad (3.31)$$

It can be seen from (3.31) that the third-order distortion component, $V_{out,in}^{(3)}(t)$, also shows a similar settling behavior as $V_{out,in}^{(2)}(t)$, with both components being suppressed by the loop at the rate of 8.6dB/ τ_f as posited before.

The effect of a non-linear output impedance can be modeled by including the second- and third-order distortion coefficients, g_{o2} and g_{o3} , respectively in equation (3.12). The relationship between $I_x(t)$ and $V_{out}(t)$ can then be rewritten as –

$$I_x(t) = g_o V_{out}(t) + g_{o2} V_{out}^2(t) + g_{o3} V_{out}^3(t) + C_L \frac{V_{out}(t)}{dt} \quad (3.32)$$

The solution for (3.32) can be simplified in a similar way as the input distortion (Appendix A), resulting in the following expression for the second-order output distortion component –

$$V_{out,out}^{(2)}(t) = \frac{g_{o2}}{g_o} \frac{A_o^2}{(1+A_o\beta)^3} V_{in}^2 \left((1 - e^{-t/\tau_f}) + e^{-t/\tau_f} (1 - e^{-t/\tau_f}) - 2e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \quad (3.33)$$

The third-order distortion component in the output voltage, $V_{out,out}^{(3)}(t)$, can

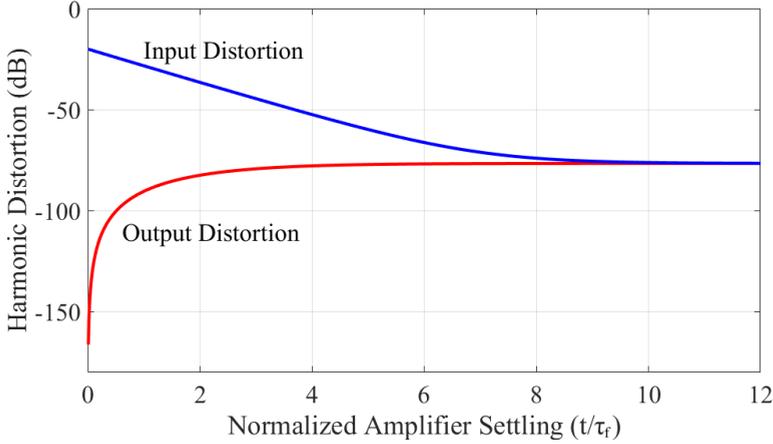


Figure 3.13: Settling behavior of harmonic distortion in closed-loop amplifiers

similarly be calculated as –

$$V_{out,out}^{(3)}(t) = \frac{g_{o3}}{g_o} \frac{A_o^3}{(1 + A_o\beta)^4} V_{in}^3 \left((1 - e^{-t/\tau_f}) + 3e^{-t/\tau_f} (1 - e^{-t/\tau_f}) - \frac{1}{2} e^{-t/\tau_f} (1 - e^{-2t/\tau_f}) - 3e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \quad (3.34)$$

It can be seen from (3.33) and (3.34) that the two distortion components, $V_{out,out}^{(2)}(t)$ and $V_{out,out}^{(3)}(t)$, follow the settling behavior of the output voltage. Figure 3.13 compares the settling behavior of second-order input and output distortion components after being normalized with the fundamental output signal. The amplifier distortion coefficients are chosen such that the steady-state values of the input and output distortion components are equal. While the distortion originating from the output settles fairly quickly towards its final value, the input distortion requires a lot more settling to reach its final value. This is primarily due to the negative-feedback suppressing the signal at the virtual-ground nodes, which lowers the amount of distortion originating from the amplifier input nonlinearity.

3.3. POWER EFFICIENCY

In the analyses in the previous sections, we studied the different metrics for a discrete-time amplifier, both in open- and closed-loop configuration, and how they

vary during the amplification period. Conventional design practices would dictate that all the amplifier metrics - gain, SNR and distortion - should be allowed to settle to their final values before the end of the amplification time. Many of these metrics, as indicated in the previous sections, settle exponentially with time, which means that the improvement in amplifier resolution saturates for higher rates of amplifier settling. The power dissipation, on the other hand, in a typical class-A amplifier that burns a fixed amount of current during the entire amplification period, increases linearly with amplifier settling.

In order to design energy-efficient pipelined A/D converters, it is important to optimize the way power is spent in the residue amplifiers. In this section, we will investigate how to trade amplifier power dissipation efficiently to achieve the desired resolution. Considering that digital calibration provides us the tools to relax some of the accuracy requirements from a residue amplifier, by examining the settling behavior of the amplifier characteristics analyzed in the previous sections, we will try to determine the optimum amount of amplifier settling in residue amplifiers. We will consider two design cases, in which the amplifier has to achieve (a) a fixed SNR and (b) a fixed SNR and gain.

3.3.1. SNR-LIMITED CASE

For studying power efficiency, we choose to consider class-A amplifiers for two reasons, namely because they are one of the most commonly used amplifier topologies and their power consumption is simpler to analyze. In class-A amplifiers, the current consumption remains the same during the entire amplification period. Hence, for a fixed amplification time t_s , the nominal power dissipation is directly proportional to the amount of settling allowed for the amplifier.

Since noise is an inherently non-deterministic error source, it imposes a fundamental limit on amplifier accuracy that can not be improved via calibration. Hence, in our analysis, we will look at how the amplifier power dissipation varies for different values of settling accuracy while achieving a fixed SNR. Apart from the initial signal feedthrough, SNR settles almost identically in open- and closed-loop amplifiers. As seen in both Figs. 3.3 and 3.10, when the settling is reduced, the SNR drops due to the loss of signal. If we consider only the SNR to be fixed by design specifications, this loss in SNR needs to be compensated by reducing the noise by sizing up the entire amplifier and capacitive network. Using the inverse square-law relationship between noise and power consumption as defined by Schreier FoM [7] for noise-limited cases, the resulting power dissipation, $P_{comp}(t_s/\tau_f)$, compensated

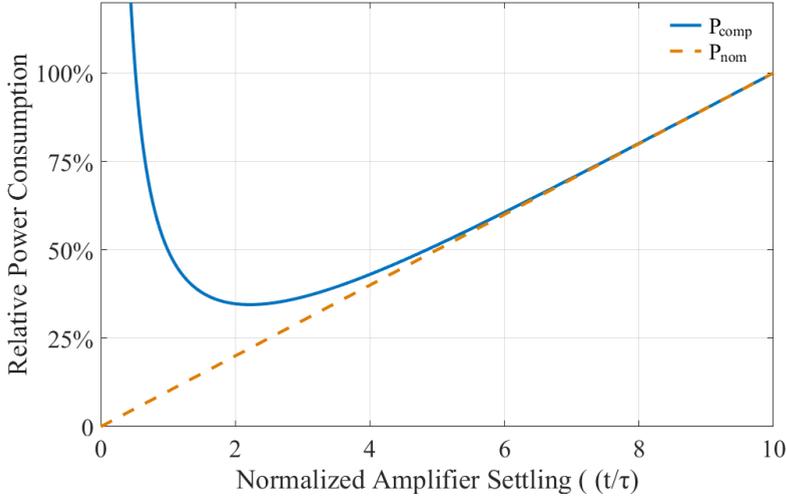


Figure 3.14: Power consumption versus settling for noise-limited amplifier

for a certain amount of settling can be expressed as –

$$P_{comp}(t_s/\tau_f) = P_{nom}(t_s/\tau_f) \left(\frac{\text{SNR}_{\text{final}}}{\text{SNR}(t_s)} \right)^2 \quad (3.35)$$

where P_{nom} is the nominal power of a class-A amplifier designed to achieve the specified SNR after complete settling (assumed to be 10τ settling in this analysis). This compensated power shoots up exponentially for low values of settling, as shown in Figure 3.14, due to very low amplifier gain. For higher rates of settling, the SNR saturates while the power dissipation continues to rise. Between these two extremes, there exists an optimum value of settling close to 2τ , where the power consumption is minimal. Around this value of settling, the loss in SNR is acceptable (as seen in Figure 3.10), making it more power-efficient to stop the amplifier from settling any further.

The other error sources limiting the amplifier resolution – gain error and distortion – benefit from amplifier settling in a different manner. While gain error reduces with amplifier settling in both open- and closed-loop amplifiers, only the input distortion in closed-loop amplifiers shows any improvement due to higher settling accuracy. In addition, for both of these error sources, the rate of settling is much slower than that for SNR. As noise is the only non-deterministic process limiting the amplifier resolution, power should only be spent towards achieving the desired SNR. Hence, to improve the amplifier’s power efficiency, gain and distortion

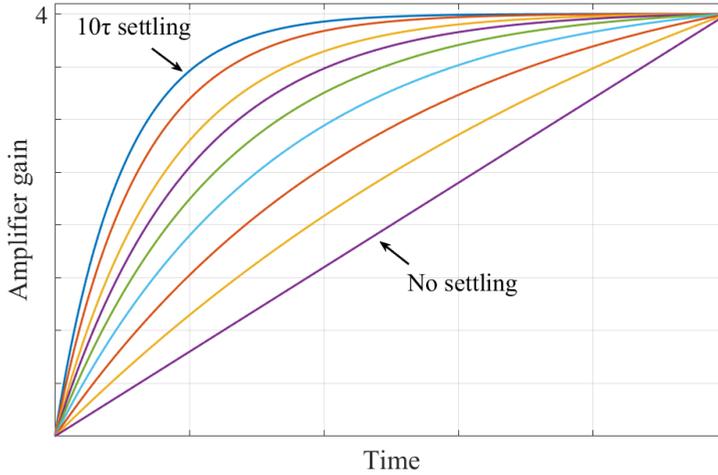


Figure 3.15: Amplifier gain compensated for different amounts of amplifier settling

calibration must be used to relax the settling requirements for achieving the desired gain accuracy to the optimum settling point identified in Figure 3.14.

3.3.2. SNR AND GAIN ERROR-LIMITED CASE

In the analysis shown in the previous sub-section, to maintain a constant SNR, the loss of signal gain is compensated purely by burning more power. We could also add another constraint in the analysis by also keeping the signal gain achieved at the end of the amplification period constant even with lower amplifier settling, as illustrated in Figure 3.15. This can be realized by increasing the amplifier's steady-state gain by using a smaller feedback factor in closed-loop amplifiers or using a larger load resistor in open-loop amplifiers.

As the settling starts to approach zero, the amplifier behavior approaches that of an integrator. As illustrated in Figure 3.16, as opposed to an exponentially settling amplifier, in an integrating amplifier, the gain increases linearly with time. For an ideal integration, the integrating amplifier needs to possess infinite DC gain and an infinitesimally small bandwidth (i.e. $\tau \rightarrow \infty$). While the intrinsic gain requirement is hypothetically infinite, the integrating amplifier trades it off with power consumption. The power efficiency of an integrating amplifier can be quantitatively analyzed by comparing the g_m required to achieve the same gain and SNR. The gain of an

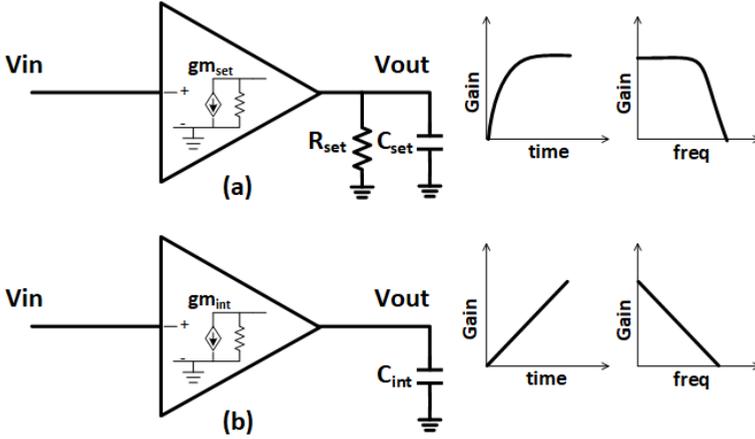


Figure 3.16: Settling vs integrating amplifier

integrating amplifier, A_{int} , can be expressed as –

$$A_{int} = \frac{g_{m,int} T_a}{C_{int}} \quad (3.36)$$

where $g_{m,int}$ and C_{int} are the integrating amplifier's transconductance and load capacitance, respectively, and T_a is the amplification time.

For the same amplification period, the gain of a conventional settling amplifier, A_{set} , shown in Figure 3.16(a), can be expressed as –

$$A_{set} = g_{m,set} R_{set} (1 - e^{-T_a/R_{set} C_{set}}) \quad (3.37)$$

with $g_{m,set}$, R_{set} and C_{set} being the settling amplifier's transconductance, output impedance, and load capacitance, respectively. If a settling of $N\tau$ is assumed for this amplifier, then (3.37) can be rewritten as –

$$A_{set} = g_{m,set} R_{set} (1 - e^{-N}) \quad (3.38)$$

Comparing the two amplifiers for the same achieved gain, (3.36) and (3.38) can be equated, resulting in a transconductance ratio of –

$$\frac{g_{m,int}}{g_{m,set}} = \frac{C_{int} R_{set}}{T_a} (1 - e^{-N}) = \frac{1}{N} \frac{C_{int}}{C_{set}} (1 - e^{-N}) \quad (3.39)$$

The ratio of C_{int} and C_{set} can be found by comparing the input-referred (single-ended) noise powers in these two cases as calculated in the previous section and [8],

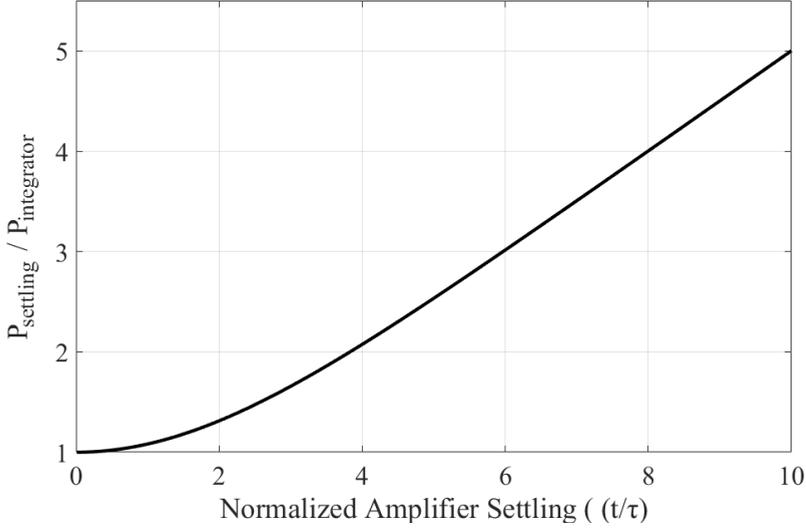


Figure 3.17: Normalized power dissipation of settling amplifier vs settling accuracy

respectively –

$$v_{n,int}^2 = \frac{4kT}{g_{m,int}} \frac{1}{2T_a} = \frac{2kT}{A_{int}C_{int}} \quad (3.40)$$

$$v_{n,set}^2 = \frac{kT}{A_{set}C_{set}} \left(\frac{1 - e^{-2N}}{1 - e^{-N}} \right)$$

where $1/2T_a$ is the noise bandwidth of the integrating amplifier. Since the two amplifiers are required to achieve the same SNR, the ratio of the required load capacitors in the two amplifiers can be derived by equating the expressions in (3.40), resulting in –

$$\frac{C_{int}}{C_{set}} = 2 \left(\frac{1 - e^{-N}}{1 - e^{-2N}} \right) \quad (3.41)$$

Using (3.39) and (3.41), the ratio between the two g_m 's can be expressed as –

$$\frac{g_{m,int}}{g_{m,set}} = \frac{N}{2} \frac{(1 - e^{-2N})}{(1 - e^{-N})^2} \quad (3.42)$$

Based on the g_m ratio in (3.42), the power consumption of a settling amplifier normalized to an integrating amplifier is shown in Figure 3.17 for different values of amplifier settling. Although the power dissipation of an integrator is always lower than that of a settling amplifier, this advantage is relatively low for settling $< 2\tau$. However, for higher values of settling, the settling amplifier consumes $N/2$ times the power of an integrating amplifier for the same gain and SNR.

3.4. CONCLUSION

One of the distinct characteristics of residue amplifiers is their discrete-time operation. In this chapter, the transient behavior of discrete-time amplifiers and their most important design metrics – gain, noise, and distortion- were studied in depth. Time-domain expressions were developed for amplifiers in both open- and closed-loop configuration to describe their settling behavior in detail. In order to maximize the amplifier power efficiency, optimum settling points for two cases of residue amplifiers were identified based on the settling behavior of their SNR. As for gain error and distortion, digital calibration was postulated as a power-efficient alternative for easing the burden of achieving the full gain accuracy from the residue amplifiers and allowing them to operate at their optimum settling point.

REFERENCES

- [1] E. Iroaga and B. Murmann, "A 12-bit 75MS/s pipelined ADC using incomplete settling," *IEEE J. Solid-State Circuits*, vol. 42, pp. 630–637, Apr. 2007.
- [2] R. Navid, T. Lee, and R. Dutton, "Minimum achievable phase noise of RC oscillators," *IEEE J. Solid-State Circuits*, vol. 40, pp. 630–637, Mar. 2005.
- [3] P. Gray and R. Meyer, *Analysis and Design of Analog Integrated Circuits*. Wiley, 1993.
- [4] W. Sansen, "Distortion in elementary transistor circuits," *IEEE Trans. Circuits Syst.-II*, vol. 46, pp. 315–325, Mar. 1999.
- [5] G. Palumbo and S. Pennisi, "High-frequency harmonic distortion in feedback amplifiers: analysis and application," *IEEE Trans. Circuits Syst.-I*, vol. 50, pp. 328–340, Mar. 2003.
- [6] W. Sansen, H. Quijing, and K. Halonen, "Transient analysis of charge transfer in SC filters – gain error and distortion," *IEEE J. Solid-State Circuits*, vol. 22, pp. 268 – 276, Apr. 1987.
- [7] R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*. Wiley-IEEE Press, Nov 2004.
- [8] T. Sepke and et al., "Noise analysis for comparator-based circuits," *IEEE Trans. Circuits Syst.-I*, vol. 56, pp. 541–553, Mar. 2009.

4

Split-ADC Calibration¹

In the previous chapter, the trade-off between an amplifier's settling accuracy and linearity was discussed. It was shown that the amplifier power efficiency can be significantly improved by employing incomplete settling in residue amplifiers along with digital calibration. By correcting gain and distortion errors through digital calibration, amplifier settling can be optimized to achieve the lowest possible power dissipation for a given SNR.

For pipelined ADCs with low-resolution front-ends, such as 1.5b/stage, it is beneficial to apply digital calibration to the first few stages. However, background calibration for distortion in multiple stages is nontrivial and most of the proposed techniques are either foreground [1], skip-and-fill [2], or queue-based [3, 4], and involve taking the ADC offline for a certain period of time. A few background techniques have also been published [5–9], but all of them are statistical in nature. In [5], pseudo-random calibration sequences are injected through the stage-DACs, while [6–8] rely on random selection between multiple residue-modes. [9] also uses a pseudo-random sequence to apply a digitally dithered reference step to the MDAC. When operating in continuous background mode, statistical approaches require a large number of cycles to estimate the effect of ADC non-idealities on the calibration signal with sufficient accuracy. This is primarily because the calibration signal is

¹This chapter is derived from publication: R. Sehgal, F. van der Goes and K. Bult, "A 12-b 53-mW 195MS/s pipeline ADC with 82dB SFDR using split-ADC calibration" *IEEE J. Solid-State Circuits*, vol. 50, no. 7, pp. 1592–1603, Jul. 2015.

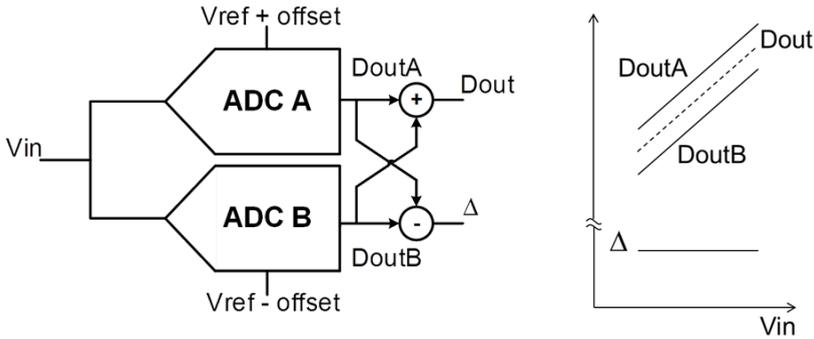


Figure 4.1: Split-ADC calibration principle

usually a low amplitude signal riding on a large uncorrelated signal. As a result, averaging over a substantially large number of samples is required to obtain sufficient accuracy, hence, driving up the calibration times for high-resolution ADCs. Most applications cannot tolerate long calibration cycles, and considering that this must be done with automatic test equipment (ATE) during production testing, a calibration time ranging in the hundreds of milliseconds is prohibitively expensive.

This chapter describes a fully deterministic continuous background calibration technique for residue amplifier distortion correction. Based on the split-ADC technique [10, 11], the calibration is shown to converge rapidly in less than 70,000 cycles. Its efficacy is demonstrated with the help of a prototype pipelined ADC with split-ADC calibration.

4.1. SPLIT-ADC CALIBRATION PRINCIPLE

In the split-ADC calibration architecture [10, 11], shown in Figure 4.1, the ADC is split into two identical halves, with each half-ADC digitizing the same input signal. By introducing a small offset (which could be a pseudo random (PR)-sequence [10] or a DC signal [11]) between them, the two half-ADCs are forced to go through different conversion trajectories. The two ADC outputs can be averaged to get the overall digital output. And by taking the difference, the input signal is effectively cancelled and the resulting digitized offset can be used to detect any deterministic ADC non-idealities. Since the offset was separated from the input signal without the need for any lengthy decorrelation procedure, the split-ADC calibration converges much faster than other full background calibration methods. Rather than using a

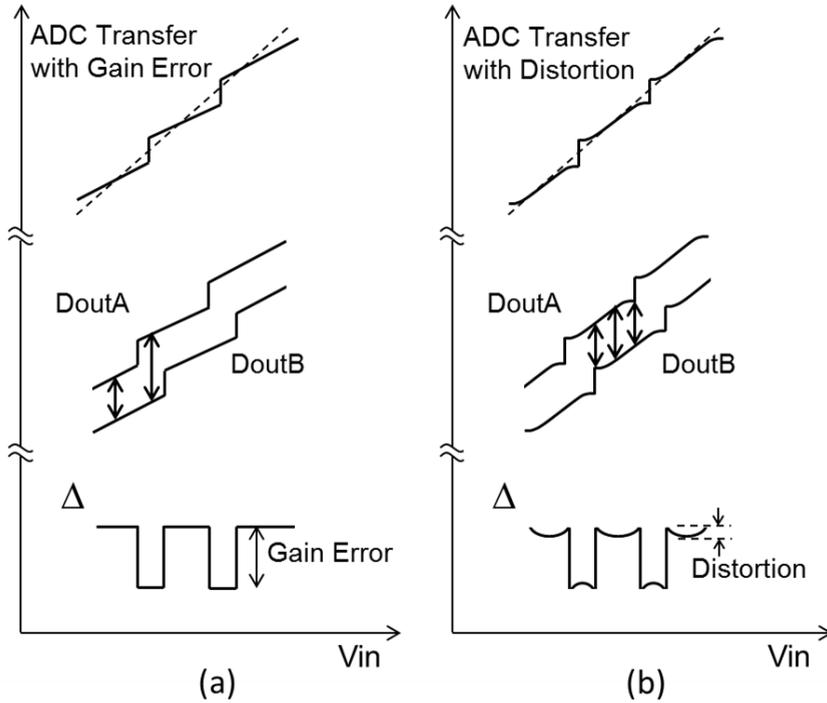


Figure 4.2: Split-ADC with (a) Gain Error and (b) Distortion

PR-sequence [10], a DC offset is used in this work for its ease of generation and effectiveness for nonlinearity calibration.

Since the analog portion is split into two halves, the noise power in each half-ADC increases by a factor of two. However, upon averaging the digital outputs of these ADCs, the uncorrelated noise is halved again. Hence, with respect to noise and analog power, the split-ADC technique has a negligible circuit overhead.

4.1.1.1. GAIN ERROR

Figure 4.2 shows the output transfer characteristic of a 1.5b/stg pipeline ADC with a gain error. This architecture consists of three ADC sub-ranges that are aligned when the gain of the inter-stage residue amplifier equals the digital gain ($2x$). However, in the presence of gain error, the slope of the output curve deviates from its ideal value, resulting in a misalignment of these sub-ranges. As highlighted in previous chapters, this leads to discontinuities in the ADC output every time it transitions from one sub-range to another.

The magnitude of the discontinuities between subranges in the ADC transfer function can be calculated by using (2.2) and (2.3). Considering the effect of gain error in the first stage of a 1.5-bit/stg ADC, the ADC stage transfer function can be written as –

$$V_{res,1} = G_{a,1} \left(V_{in} - D_{out,1} \left(\frac{V_{ref}}{2} \right) \right) \quad (4.1)$$

The effect of gain error on the ADC output can be analyzed by examining the reconstructed ADC input. By combining the stage residue with the stage digital output, the reconstructed ADC input, $V_{in,dig}$, can be expressed as –

$$V_{in,dig} = D_{out,1} \left(\frac{V_{ref}}{2} \right) + \frac{V_{res,1}}{G_{d,1}} = D_{out,1} \left(\frac{V_{ref}}{2} \right) + \frac{G_{a,1}}{G_{d,1}} \left(V_{in} - D_{out,1} \left(\frac{V_{ref}}{2} \right) \right) \quad (4.2)$$

To analyze the discontinuity between two ADC subranges, we need to consider an input that lies between those two subranges. For instance, in a 1.5bit/stg architecture, for an input voltage of $V_{ref}/4$, the ADC output occurs right between the 2nd and 3rd subrange. Due to the proximity of this input to the CADC thresholds, the CADC may generate two different digital outputs, resulting in completely opposite stage residue voltages. However, when the digital bits are combined in the encoder, the two trajectories should still result in the same overall ADC code. This can be verified by calculating the reconstructed ADC input for these two trajectories –

For CADC = 11, $D_{out,1} = 1$

$$V_{in,dig11} = \frac{V_{ref}}{2} + \frac{G_{a,1}}{G_{d,1}} \left(\frac{V_{ref}}{4} - \left(\frac{V_{ref}}{2} \right) \right) = \frac{V_{ref}}{2} - \frac{G_{a,1}}{G_{d,1}} \left(\frac{V_{ref}}{4} \right) \quad (4.3)$$

For CADC = 10, $D_{out,1} = 0$

$$V_{in,dig10} = 0 + \frac{G_{a,1}}{G_{d,1}} \left(\frac{V_{ref}}{4} - 0 \left(\frac{V_{ref}}{2} \right) \right) = \frac{G_{a,1}}{G_{d,1}} \left(\frac{V_{ref}}{4} \right) \quad (4.4)$$

As seen from above equations, for an ideal analog stage gain, these two trajectories lead to an identical reconstructed ADC input. However, the presence of any gain error will show up as difference between the ADC codes obtained from these two trajectories, resulting in discontinuities in the ADC transfer function. This discontinuity can be calculated as –

$$V_{in,dig11} - V_{in,dig10} = \frac{V_{ref}}{2} - \frac{G_{a,1}}{G_{d,1}} \left(\frac{V_{ref}}{4} \right) - \frac{G_{a,1}}{G_{d,1}} \left(\frac{V_{ref}}{4} \right) = \frac{V_{ref}}{2} \left(1 - \frac{G_{a,1}}{G_{d,1}} \right) \quad (4.5)$$

As seen from the above equation, this discontinuity is directly proportional to the gain error and can be eliminated by accurately matching the analog and digital gains.

In the split-ADC technique, this discontinuity is captured by introducing an offset between the signal paths of two identical ADCs and observing the difference between the two split-ADC outputs when they transition between subranges. For a differential offset, $2V_{off}$, the reconstructed ADC inputs for the two split-ADCs A and B can be expressed as –

$$V_{in,digA} = D_{out,1A} \left(\frac{V_{ref}}{2} \right) + \frac{G_{a,1}}{G_{d,1}} \left(V_{in} + V_{off} - D_{out,1A} \left(\frac{V_{ref}}{2} \right) \right) \quad (4.6a)$$

$$V_{in,digB} = D_{out,1B} \left(\frac{V_{ref}}{2} \right) + \frac{G_{a,1}}{G_{d,1}} \left(V_{in} - V_{off} - D_{out,1B} \left(\frac{V_{ref}}{2} \right) \right) \quad (4.6b)$$

To capture the subrange transition, we consider two input voltages, V_{inX} and V_{inY} , that ensure that the output of split-ADC B is in its 2nd and 3rd subranges, respectively, while split-ADC A operates in the third subrange for both inputs due to the offset. The reconstructed input arising from these two cases can be expressed as –

For $V_{in} = V_{inX}$, $D_{out,1A} = 1$ and $D_{out,1B} = 0$

$$V_{inX,digA} = \frac{V_{ref}}{2} + \frac{G_{a,1}}{G_{d,1}} \left(V_{inX} + V_{off} - 1 \left(\frac{V_{ref}}{2} \right) \right) \quad (4.7a)$$

$$V_{inX,digB} = \frac{G_{a,1}}{G_{d,1}} (V_{inX} - V_{off}) \quad (4.7b)$$

For $V_{in} = V_{inY}$, $D_{out,1A} = 1$ and $D_{out,1B} = 1$

$$V_{inY,digA} = \frac{V_{ref}}{2} + \frac{G_{a,1}}{G_{d,1}} \left(V_{inY} + V_{off} - 1 \left(\frac{V_{ref}}{2} \right) \right) \quad (4.8a)$$

$$V_{inY,digB} = \frac{V_{ref}}{2} + \frac{G_{a,1}}{G_{d,1}} \left(V_{inY} - V_{off} - 1 \left(\frac{V_{ref}}{2} \right) \right) \quad (4.8b)$$

Since the two split-ADCs transition between sub-ranges at different inputs due to the added offset, the discontinuities in their ADC transfer function appear at different inputs. This creates discontinuities or “steps” in the difference between the two split-ADC outputs, as illustrated in Figure 4.2(a), which can be calculated by observing the difference between the two split-ADCs –

$$\Delta = (V_{inY,digA} - V_{inY,digB}) - (V_{inX,digA} - V_{inX,digB}) = \frac{V_{ref}}{2} \left(1 - \frac{G_{a,1}}{G_{d,1}} \right) \quad (4.9)$$

It can be seen from (4.9) that the step, Δ , in the split-ADC difference is directly proportional to the gain error while being completely independent of the input signal. This signal-independent measure of the gain error is obtained without any lengthy decorrelation and, hence, allows the split-ADC calibration to converge relatively quickly while operating in full background mode.

The operation in (4.9) divides the difference signal between the two split-ADCs into different bins using the CADC values of ADC A and B from the stage under calibration. The bin values can then be averaged to reduce the impact of noise, and the resulting delta between these bin averages relates to the gain error estimate, which is used to update the digital gains. Eventually, the calibration loop is driven to achieve a difference signal that is a perfectly straight line without any discontinuities.

4.1.2. DISTORTION

The distortion in a class-A amplifier is almost always compressive in nature and leads to gain compression for larger input signals. Ignoring the fifth- and higher-order harmonics, the output voltage of a differential residue amplifier for an input signal $V_{in,RA}$, can be expressed as –

$$V_{res} = G_a V_{in,RA} - G_3 V_{in,RA}^3 \quad (4.10)$$

where G_a and G_3 are the linear gain and third-order distortion coefficients of the amplifier, respectively. By again considering the differential offset, $2V_{off}$, introduced between the two split-ADCs, the stage residues for ADCs A and B can be written as –

$$V_{resA} = G_a (V_{in,RA} + V_{off}) - G_3 (V_{in,RA} + V_{off})^3 \quad (4.11a)$$

$$V_{resB} = G_a (V_{in,RA} - V_{off}) - G_3 (V_{in,RA} - V_{off})^3 \quad (4.11b)$$

If the backend is accurate enough, then the two backend codes, D_{bA} and D_{bB} , can be used as estimates for the stage 1 residues. Hence, the difference of the backend codes from the two ADCs results in –

$$D_{bA} - D_{bB} = 2 \left(G_a V_{off} + G_3 V_{off}^3 + 3G_3 V_{off} V_{in,RA}^2 \right) \quad (4.12)$$

The difference between the two backend codes in (4.12) consists of a constant and a quadratic term that is proportional to the third-order distortion coefficient, as seen in Figure 4.2(b). The curvature of this quadratic component can be used to estimate the distortion present in the residue amplifier. The magnitude of this curvature can

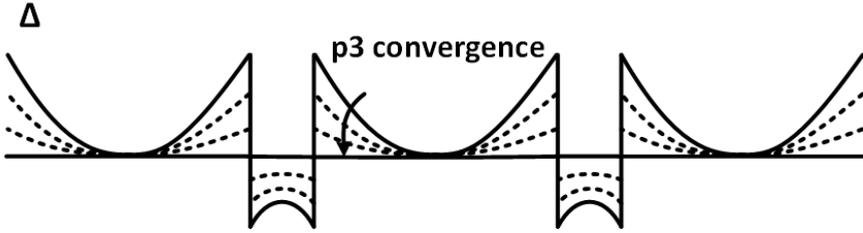


Figure 4.3: Effect of distortion correction on difference signal

be measured by dividing the difference signal into smaller bins with the help of the CADC bits from the following stage. By taking the difference between different bin averages, an error signal is generated which is proportional to the curvature of the difference signal. The error signal drives a parameter p_3 which is used to generate a look-up table (LUT) of the corrected residue signals –

$$\begin{aligned} D_{bA,corr} &= D_{bA} + p_3 D_{bA}^3 \\ D_{bB,corr} &= D_{bB} + p_3 D_{bB}^3 \end{aligned} \quad (4.13)$$

While the cubic correction term is an approximation compared to the actual inverse as calculated in [6], it is a lot simpler to implement and shows comparable accuracy for weak nonlinearities. The corrected residue signals are then scaled by the stage gains and added to the bits from stage 1 to get the overall digital output. By controlling p_3 , the calibration loop eventually drives the difference signal towards a straight line as shown in Figure 4.3.

One of the factors affecting the curvature of the difference signal in (4.12) is the offset between the split-ADCs. A larger offset makes the calibration loop more sensitive to the third-order distortion and hence resulting in faster convergence. In this work, an offset equivalent to 5.5% of the ADC full-scale was used, occupying 0.5dB of the ADC dynamic range. To efficiently utilize the headroom, the calibration offset can be reduced once convergence is reached.

The correction terms in (4.13) are based on the assumption that the backend code accurately represents the stage 1 residue. However, in this work, the offset was added between the two residues in every stage of the ADC. The presence of any random offsets added by the MDACs in the backend limits the accuracy of the distortion calibration. This limitation is experienced by any calibration approach that corrects distortion using digital post-processing. Both gain and distortion calibration loops are extended to multiple stages. As both of these loops rely on a relatively accurate

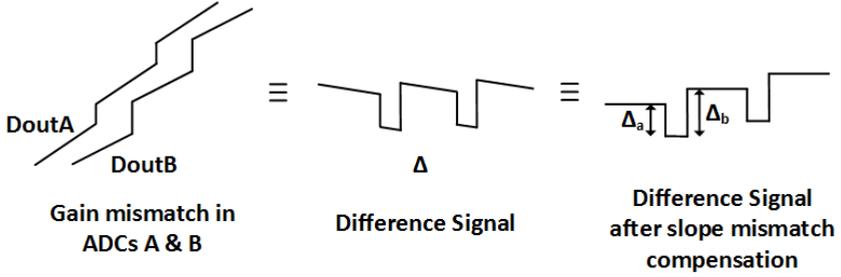


Figure 4.4: Gain mismatch in split-ADC calibration

backend, the calibration is first applied to the backend stages. While the bin sizes reduce for the latter stages, the number of bins goes up proportionally. Thus the calibration speed remains roughly the same for each stage in the pipeline chain.

4.1.3. MISMATCH BETWEEN SPLIT-ADCs

Due to random mismatch, the two half-ADCs will have different input-referred offsets. This makes the overall offset between the ADCs different from what is added through the ADC references. However, for gain calibration, the exact value of the offset between the digital outputs is not important, as the error signal is generated by taking the difference of the bin averages. As a result, the random offset between the two half-ADCs doesn't have any impact on the gain calibration.

The spread in the capacitive DACs and residue amplifiers also creates a gain mismatch between the two split-ADCs. Consequently, the input signal is not entirely cancelled by subtracting the split-ADC outputs, leading to a residual slope in the difference signal. This slope is equal to the cumulative gain mismatch of all the stages in the ADCs. The non-zero slope makes the difference signal input-dependent to a certain extent, making it difficult to calculate reliable bin averages. Hence, the overall gain of ADC A needs to be scaled in order to make the two output transfer curves parallel to each other. This scaling coefficient, α , can be calculated from the ratio of the slopes of the two fine ADC outputs [12].

While α calculates the cumulative mismatch, the gain mismatch for each stage needs to be estimated separately for an effective multi-stage gain calibration. This is done by treating the two discontinuities in the difference signal separately. As seen in Figure 4.4, Δ_a is created by ADC A transition, while Δ_b is created by ADC B. Δ_a and Δ_b are averaged separately to estimate the gain errors in each individual stage in ADC A and B, respectively.

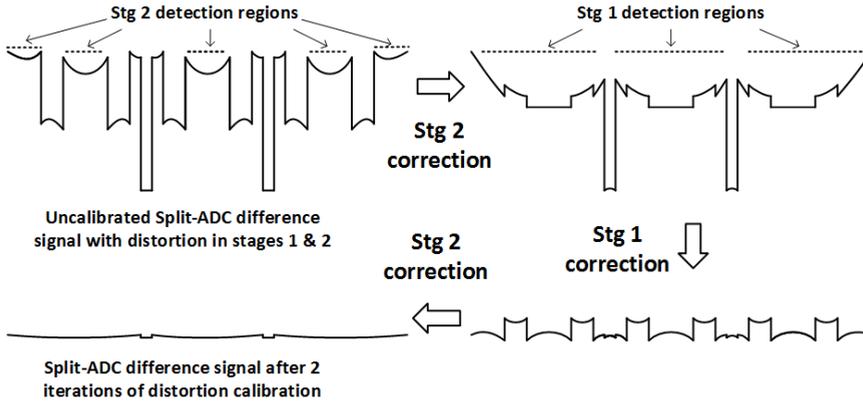


Figure 4.5: Non-orthogonalities between distortion calibration loops of stages 1 and 2

The mismatch between the third-order distortion coefficients of the two split-ADCs leads to a fraction of the signal and the third-order distortion component leaking into the difference signal. The former contributes towards the slope of the difference signal and hence, is treated as a part of the gain mismatch and taken into account by α . The presence of the residual third-order component in the difference signal, on the other hand, alters its quadratic nature and as a result, interferes with the estimation of the third-order distortion. However, due to its relatively small magnitude, this residual third-order component does not have a significant effect on the distortion estimation.

4.1.4. CALIBRATION NON-ORTHOGONALITIES

Due to the multi-stage gain and distortion calibration loops running in parallel, certain non-orthogonalities were observed between these loops. There were mainly two kinds of interactions affecting the calibration: between gain and distortion calibration and distortion calibration of different stages:

- a. As seen in Figure 4.2, the nonlinear gain compression created by distortion produces similar jumps in the difference signal as a linear gain error. Hence, for a given stage, the distortion calibration is run before the gain error calibration loop. But a gain estimate is also required to calculate the stage residue and for accurate compensation of the backend offset. This issue is resolved by sequentially running the two calibration loops iteratively in the background.
- b. Any given ADC output sample includes the distortion introduced by all the

stages and all of this distortion contributes towards the curvature of the difference signal. A correction loop for any stage will, thus, attempt to correct for the nonlinearity of the other stages as well. This issue can be resolved by running the stage 2 calibration loop followed by stage 1 loop iteratively in the background, as illustrated in Figure 4.5. In this work, two iterations were found to be sufficient for the correction coefficients to converge for both the stages.

4.2. PROTOTYPE SPLIT-ADC IMPLEMENTATION

In order to test the efficacy of the split-ADC calibration technique, a prototype pipelined split-ADC was built as a test vehicle. The goal was to see how much we could push the settling accuracy of residue amplifiers and still attain good linearity with the help of calibration. For this, we made the residue amplifier bias highly programmable with a large range.

4.2.1. ADC ARCHITECTURE

Figure 4.6 shows the complete ADC architecture. There are nine 1.5b stages, chosen for their simplicity and a 3-bit fine ADC (FADC). Two extra bits were resolved in the FADC to suppress quantization noise and improve calibration accuracy. ADC stages 2 and 3 were scaled by a factor of two to save power. To split the ADC into two identical halves, the amplifiers and the capacitance in each half-ADC were sized down by a factor of two. The amplifiers share the same bias and symmetry was maintained between the two half-ADCs to minimize mismatch. The CADCs and other digital blocks, being close to minimum size, were not scaled. Although this led to an increase in digital complexity, the overall area and power overhead was relatively small.

Gain calibration is used for stages 1-6, while the nonlinearity calibration is applied only to the first two stages. An offset was introduced differentially in the two half-ADCs via the externally generated ADC reference voltages. This method was preferred over the use of extra residue modes [11] in order to have greater flexibility.

Figure 4.7 shows the pipeline front-end stage architecture. The front-end was designed without a dedicated sample-and-hold amplifier (SHA) and the input networks for stage 1 MDAC and CADC were matched to reduce aperture errors. The sampling network and timing scheme that were used are similar to those proposed in [13], with a shortened sampling phase for both MDAC and CADC in stage 1 and

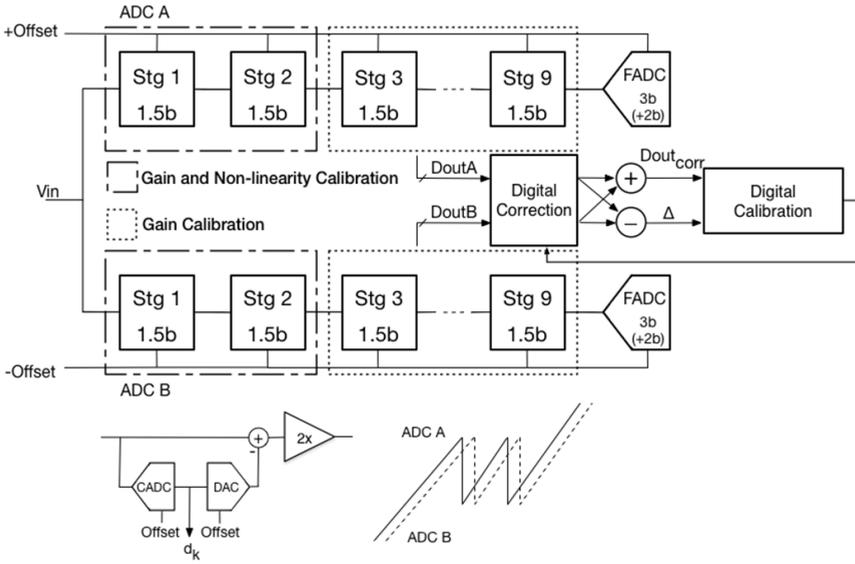


Figure 4.6: Prototype pipelined ADC architecture with split-ADC calibration

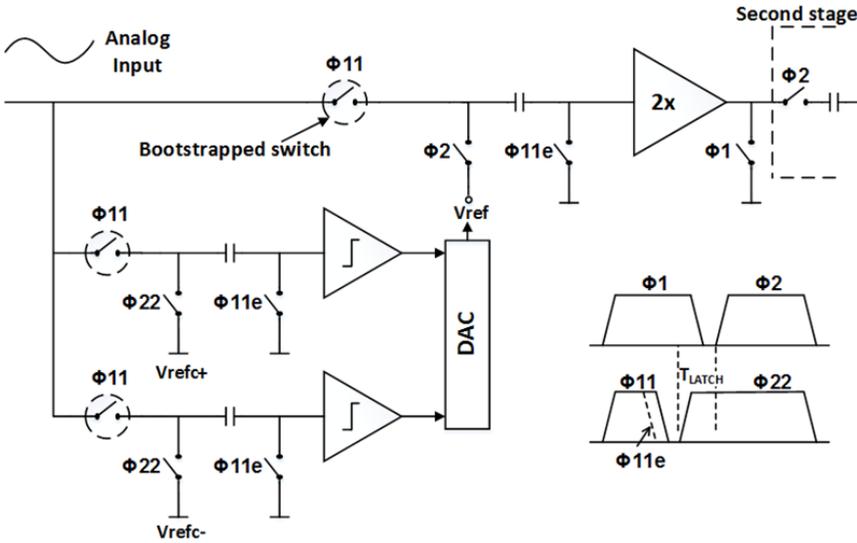


Figure 4.7: SHA-less front-end stage

CADCs in stages 2-9. A flip-around topology was used for the MDACs to achieve a higher feedback factor, with a total sampling capacitance of 4.4pF in the front-end for each split-ADC.

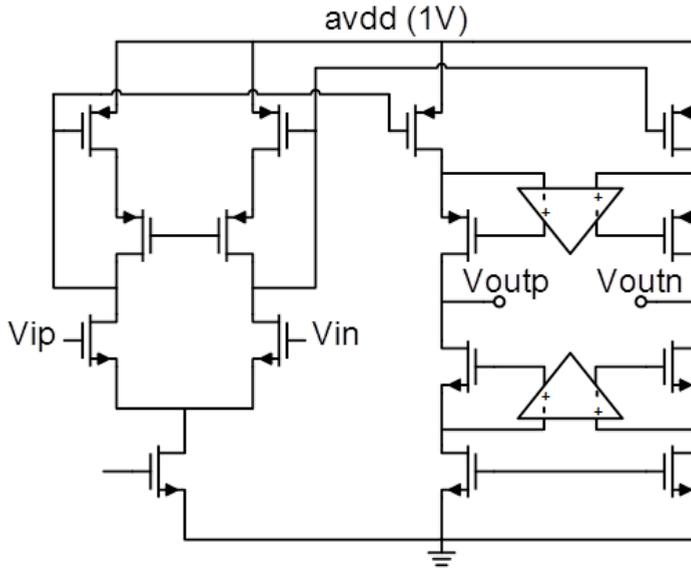


Figure 4.8: Opamp topology

4.2.2. OP-AMP

A key challenge of this work was the design of a residue amplifier suitable for testing the efficacy of gain and distortion calibration. The objective was to design the ADC with sufficiently accurate residue amplifiers (with high DC gain and settling) as a starting point and then tune the settling accuracy of the residue amplifiers to reduce their power consumption, while using calibration to correct for the resulting errors. This required an op-amp with a high DC gain and a wide range of programmability in its bandwidth.

A single-stage current mirror op-amp with cascodes, shown in Figure 4.8, was used as a residue amplifier because it provides a slightly higher DC gain than a folded-cascode topology for the same output swing. The amplifier's loop-gain was further boosted with the help of telescopic op-amps with level shifters at the input, achieving an overall DC gain of 65dB.

The current mirror ratio was set to unity to avoid excess noise and distortion from the current sources. Due to the low supply voltage (1V), the overdrive voltages of the current sources were optimized to achieve maximum output swing rather than minimum noise. A highly programmable master bias reference was used to vary the amplifier bias point to sweep its bandwidth over a wide range to observe the effect of limited settling on ADC performance in detail.

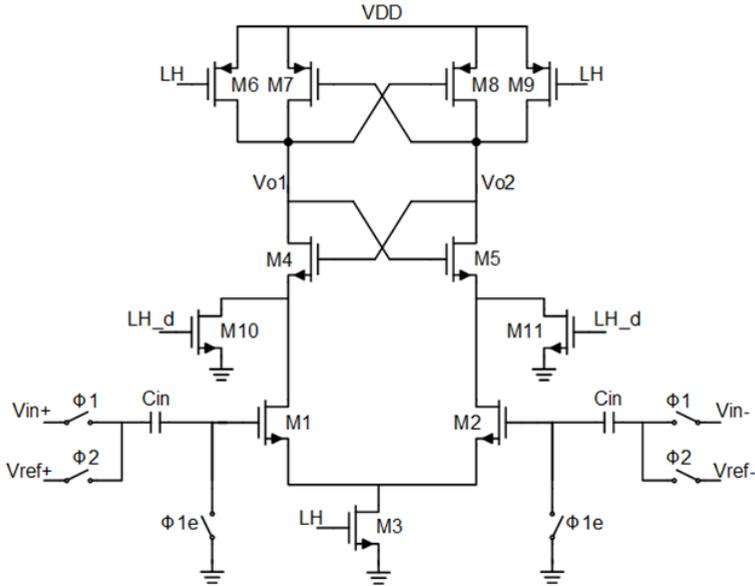


Figure 4.9: Comparator

4.2.3. COMPARATOR

The comparator used in the CADCs of the pipeline stages is shown in Figure 4.9. It is based on the dynamic switched-cap comparator [14], where the input is amplified by the differential pair, M1/2, and then inserted into the latch. This design was modified by adding two switches, M10 and M11, driven by a slightly delayed latch enable signal, LH_d, at the bottom of the latch. The delay allows the differential input to propagate through the differential pair and reach the latch. As it begins to regenerate the amplified differential input, the switches connect the bottom nodes of the latch to ground, giving it maximum headroom. This reduces the comparator time constant by 20%, improving the bit error rate significantly. The delay in LH_d should be sufficiently long to minimize the effect on the offset and was set at 50ps in this work.

As the speed requirements are relaxed for the 5b FADC, the switches were removed to simplify the latch. Since no over-range is implemented in the FADC, the comparator offset directly affects the ADC INL. Hence, differential pairs with resistive loads were used as preamplifiers with a gain of 6x to reduce the input-referred latch offset. Auto-zeroing was applied by storing the preamplifier offset at the input to further limit the overall comparator offset [15].

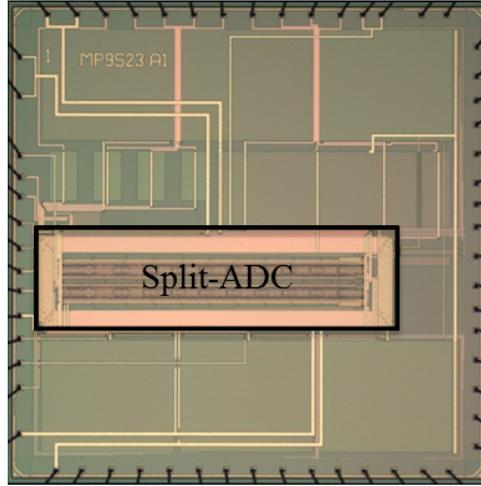


Figure 4.10: Chip micrograph

4.2.4. ANALOG INPUT INTERFACE

The input signals and the reference voltages being sampled by the ADC are driven directly from external pins. Due to this, the sampling glitch coming from the sampling network couples with the bond-wire inductance, creating ringing at their resonant frequency. This can significantly degrade both the input and reference signals. To avoid this ringing, on-chip low-pass RC networks with large MOS-caps of the order of 1nF were connected to all the references. However, for the input signals, to preserve bandwidth, a shunt cap identical to the first stage sampling cap was employed. An on-chip series metal resistor was also used to damp out the ringing. This, along with differential signaling, reduced the amount of spurious content on the input signals [16].

4.3. MEASUREMENT RESULTS

The ADC was implemented in 40nm 1V digital CMOS process and occupied an area of roughly 0.81mm^2 , including the bias, clock tree, output digital interface, and input test interface. This excludes the on-chip supply and reference decoupling caps. A die micrograph is shown in Figure 4.10. The digital outputs of each split-ADC were subsampled by the output digital interface before being taken off-chip. The output bits were then imported into MATLAB where the calibration sequence was performed. Measurements were carried out on three die samples and the overall

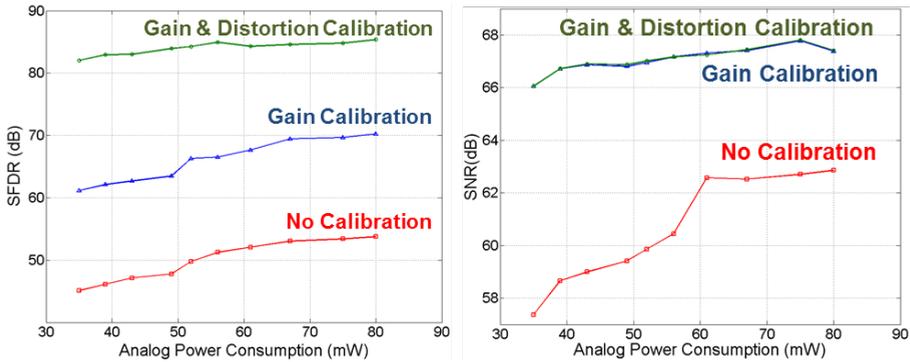


Figure 4.11: SFDR and SNR vs analog power sweep ($f_s = 195\text{MHz}$, $f_{in} = 96\text{MHz}$)

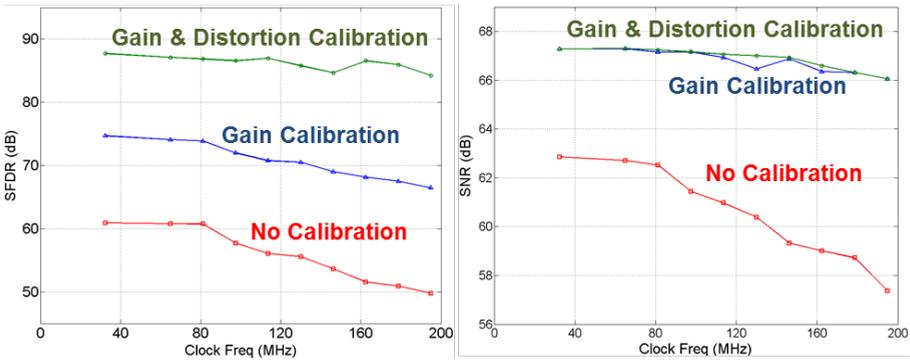


Figure 4.12: SFDR and SNR vs clock frequency sweep (Analog Power = 35mW, $f_{in} = 1\text{MHz}$)

performance was found to be fairly consistent over all samples.

The ADC was initially tested at a low clock frequency with a high current bias setting to ensure complete settling in the residue amplifiers. Then the settling accuracy was decreased by (a) lowering the bias currents (and hence, the ADC power consumption) and (b) raising the clock frequency, as seen in Figures 4.11 and 4.12. Starting with a settling accuracy of roughly 65dB (limited by DC loop-gain) in all stages, which is equivalent to 7.5τ , the settling was reduced to roughly 2τ (30%) in the first two stages and 3.5τ (50%) in the remaining stages. While the amplifier power consumption was effectively lowered by 60%, the overall SNR dropped by only 1dB. Hence, the residue amplifiers in all the stages were pushed close to the optimal settling point shown in Figure 3.6, while maintaining roughly the same SNDR using calibration.

The variation in the digital gain estimates and distortion correction coefficient

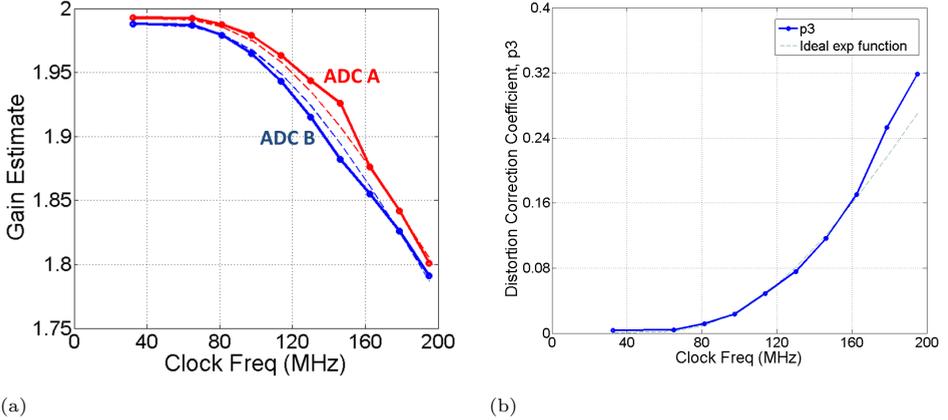


Figure 4.13: (a) Stage 1 gain estimates and (b) distortion coefficient vs clock frequency (Analog Power = 35mW, $f_{in} = 1\text{MHz}$)

for stage 1 in the two half-ADCs is shown in Figure 4.13 (along with the first-order theoretical predictions), and also points towards the amount of reduction in the amplifier settling accuracy versus clock frequency. Considering the amplifier to be close to settling to its full accuracy at low clock frequency ($f_s=32.5\text{MHz}$), the settling at $f_s=195\text{MHz}$ can be estimated to be roughly 2.3τ . As p_3 can be used as a measure of distortion in stage 1, Figure 4.13 implies an exponential increase in distortion as the settling accuracy is reduced, as predicted in Figure 3.6. The deviation observed at high clock frequencies can be explained by the increasingly significant effect of higher-order harmonics ($\geq 5^{\text{th}}$) and the non-overlap time between the track and amplification phases.

Figure 4.14 and 4.15 show the 4096-point FFT and INL of the ADC before and after calibration at $f_s = 195\text{MS/s}$ (subsampled by a ratio of 192) and an input signal frequency of 96MHz. The ADC shows an overall worst-case SNR (calculated by removing first 251 harmonics from the spectrum) of 66dB, an SFDR of 82dB and a maximum INL of 1.4 LSB at 195MS/s. All the results are presented for a sinusoidal input with an amplitude of 85% of the ADC full-scale, $1V_{\text{ppd}}$, to include a margin for the offset added in the half-ADCs.

With the help of gain and nonlinearity calibration, the amplifier power consumption was reduced from 64mW to 22.4mW. This includes the power dissipated in the preamplifiers used in the fine ADCs, which is estimated to be roughly 6mW. The bias consumed 12.6mW, a significant portion of which is attributed to the local bias within every stage and scales linearly with the amplifier power. The power drawn

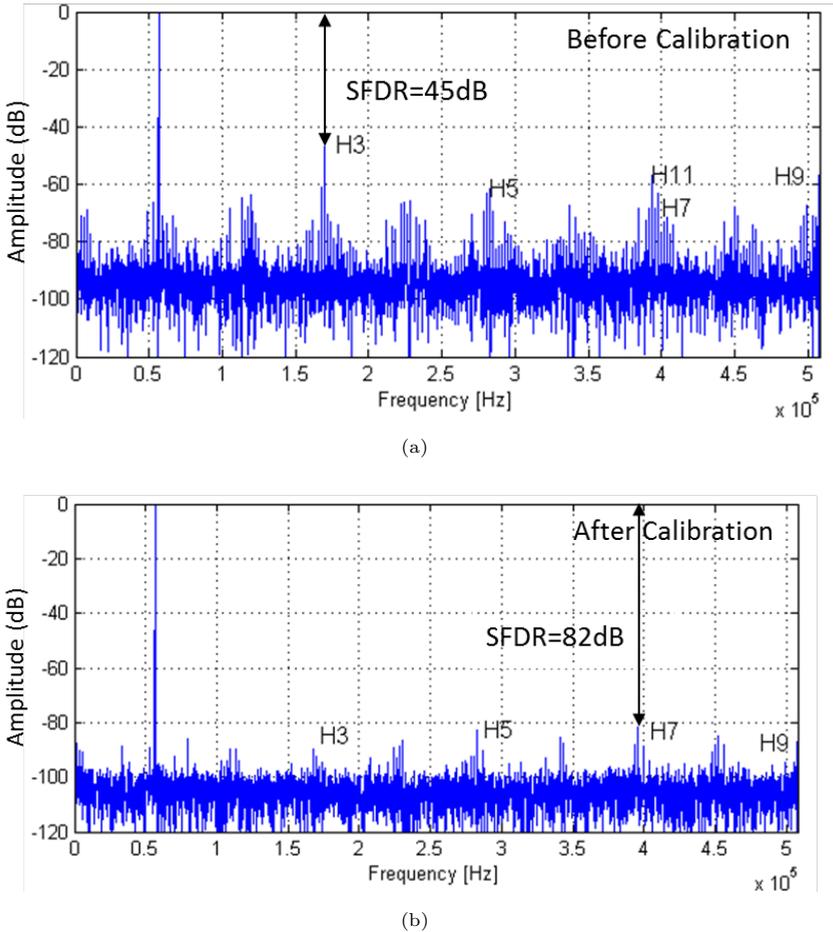


Figure 4.14: FFT of 192x sub-sampled output (a) before and (b) after calibration (Analog Power = 35mW, $f_s = 195\text{MHz}$, $f_{in} = 96\text{MHz}$)

by the clock tree and other miscellaneous digital blocks was 18mW. The clock tree forms a major part of the power consumption as it was optimized for operation up to 500MHz. As the comparators, encoders and other digital blocks were not scaled down, splitting the ADC created an estimated overhead of roughly 4mW. The entire ADC consumed 53mW, exhibiting a Schreier FoM, a more accurate measure for noise-limited ADCs [17], of 157.5dB. Table 4.1 summarizes and compares the performance of the ADC with other state-of-art pipeline ADCs utilizing closed-loop class-A amplifiers with distortion calibration. The power consumption of on-chip reference buffers and calibration logic has been excluded from all the ADCs to pro-

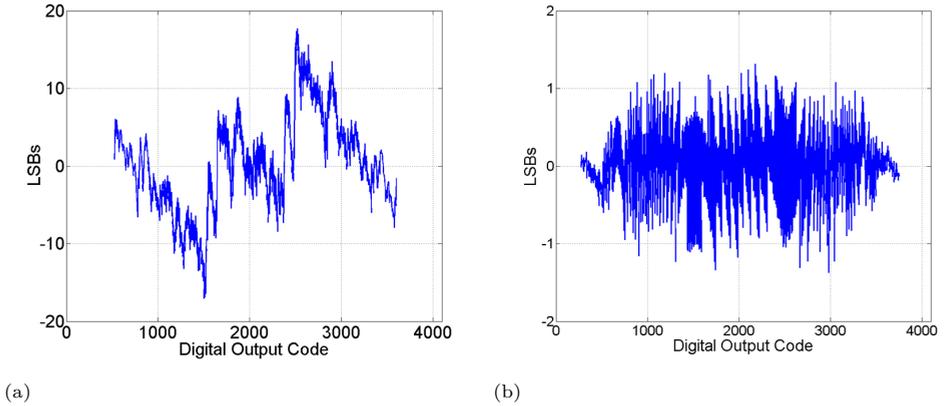


Figure 4.15: ADC INL (a) before and (b) after calibration (Analog Power = 35mW, $f_s = 195\text{MHz}$)

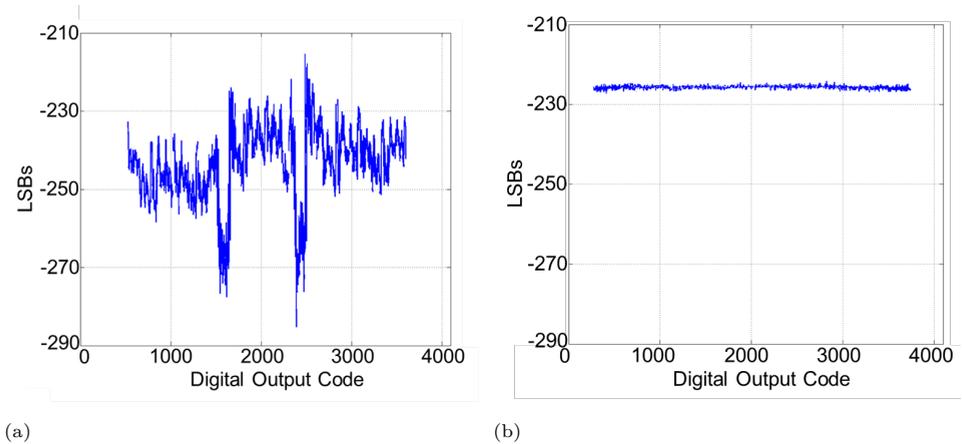


Figure 4.16: Split-ADC difference signal (a) before and (b) after calibration (Analog Power = 35mW, $f_s = 195\text{MHz}$)

vide a fair comparison. It should be noted that this ADC was designed with residue amplifiers with a high DC gain in order to achieve good performance even without calibration. Hence, to fully benefit from the split-ADC amplifier gain and distortion calibration, a simpler op-amp topology with a lower loop-gain or an open-loop amplifier can be chosen to design the residue amplifiers to optimize the ADC power efficiency.

	[2]	[3]	[5]	[7]	[8]	[9]	This work
Technology	90nm	0.25 μ m	90nm	90nm	0.18 μ m	0.18 μ m	40nm
Cal. Method	Skip-and-fill	Queue-based	Background RNG cal signal	Dual mode	Dual mode	Adaptive analog cancellation	Split-ADC
Supply (V)	1.2	2.5	1.2/1	1.2	1.8	1.6	1
ADC FS (V_{ppd})	1.2	1	1.5	1.6	2	2	1
F_s (MHz)	200	80	100	100	20	60	195
SNDR (dB)	59.4	72.6	69.8	70	60	73.3	64.8
SFDR (dB)	N/A	84.5	85	80	76	84	82
Power (mW)	186	340	113	200	2.9	58	53
Schreier FoM (dB)	146.7	153.2	156.2	154	154.4	160.4	157.4
No. of cal cycles	2×10^8	N/A	2×10^9	N/A	10^5	$3.2 \times 10^{4*}$	7×10^4

*Power-on calibration with DC input

Table 4.1: Performance Comparison

4.3.1. CALIBRATION PERFORMANCE

The calibration algorithm starts with a split-ADC difference signal as seen in Figure 4.16 and drives it towards a straight line by adjusting the digital gains and the nonlinearity LUTs. Each iteration of gain and nonlinearity calibration converges in roughly 10^4 samples, with Figure 4.17 showing the stage 1 coefficients settling during the first iteration. The overall ADC calibration, including the gain calibration for the first six stages, reached convergence within roughly 70,000 samples.

Since the calibration converges in less than 0.35ms (for $f_s = 195$ MS/s), the error estimation can be run at a much lower frequency, making its power consumption negligible. The nonlinearity LUTs consist of adders that are truncated to 6-bit precision words to reduce complexity. Due to the high-resolution multipliers running at the ADC sampling rate, gain correction is estimated to consume the majority of

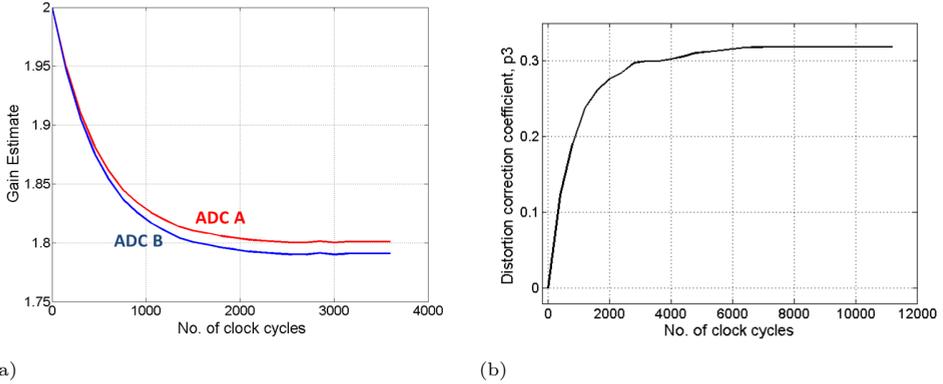


Figure 4.17: Settling of (a) gain error and (b) distortion coefficients for stage 1

the power associated with the calibration logic. The total power dissipation of the multipliers for all six stages running at 195MHz in 40nm CMOS was estimated to be 0.92mW. Hence, the correction logic is estimated to consume less than 2mW and not make a significant impact on the overall ADC power consumption.

4.4. CONCLUSION

In discrete-time closed-loop class-A amplifiers, a settling accuracy between $2-3\tau$ was shown in chapter 3 to achieve the best trade-off between SNR and power efficiency. In this chapter, a 12-bit pipelined ADC with split-ADC calibration architecture was presented, which enables the residue amplifiers to operate close to their optimum settling point. The split-ADC architecture was used to implement multi-stage gain and distortion calibration, and was shown to achieve fast convergence while dealing with non-orthogonalities between different calibration loops. Fabricated in 40nm 1V digital CMOS, the prototype ADC demonstrated an ENOB of 10.5b up to 195MS/s with a power dissipation of 53mW. While working continuously in background, the split-ADC calibration improved the ADC SFDR by 37dB within 70,000 samples. With the help of calibration, the power dissipation in the residue amplifiers was slashed by 60%, efficiently trading-off analog power for digital-post processing.

REFERENCES

- [1] A. Verma and B. Razavi, "A 10-bit 500-MS/s 55-mW CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 44, pp. 3039–3050, Nov. 2009.
- [2] B. D. Sahoo and B. Razavi, "A 12-bit 200-MHz CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 44, pp. 2366–2380, Sep. 2009.
- [3] C. R. Grace, P. J. Hurst, and S. H. Lewis, "A 12-bit 80-MSample/s pipelined ADC with bootstrapped digital calibration," *IEEE J. Solid-State Circuits*, vol. 40, pp. 3039–3050, May 2005.
- [4] J. Kim and B. Murmann, "A 12-b, 30-MS/s, 2.95-mW pipelined ADC using single-stage class-AB amplifiers and deterministic background calibration," *IEEE J. Solid-State Circuits*, vol. 47, pp. 2141–2151, Sep. 2012.
- [5] A. Panigada and I. Galton, "A 130 mw 100 MS/s pipelined ADC with 69dB SNDR enabled by digital harmonic distortion correction," *IEEE J. Solid-State Circuits*, vol. 44, p. 3314 – 3328, Dec. 2009.
- [6] B. Murmann and B. E. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, p. 2040 – 2050, Dec. 2003.
- [7] H. van de Vel, B. Butler, H. van der Ploeg, M. Vertregt, G. Geelen, and E. Paulus, "A 1.2-V 250-mW 14-b 100-MS/s digitally calibrated pipeline ADC in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, pp. 1047–1056, Apr. 2009.
- [8] N. Sun, H. Lee, and D. Ham, "A 2.9mW 11b 20MS/s pipelined ADC with dual-mode-based digital background calibration," in *ESSCIRC*, pp. 265–268, 2012.
- [9] Y. Miyahara, M. Sano, K. Koyama, T. Suzuki, K. Hamashita, and B.-P. Song, "A 14b 60MS/s pipelined ADC adaptively cancelling opamp gain and nonlinearity," *IEEE J. Solid-State Circuits*, vol. 49, pp. 416–425, Feb. 2014.
- [10] J. Li and U.-K. Moon, "Background calibration techniques for multi-stage pipelined ADCs with digital redundancy," *IEEE Trans. Circuits Syst.-II*, vol. 50, pp. 531–538, Sep. 2003.
- [11] J. McNeill and et al., "Split ADC" architecture for deterministic digital background calibration of a 16-bit 1-MS/s ADC," *IEEE J. Solid-State Circuits*, vol. 40, pp. 2437 – 2445, Dec. 2005.
- [12] H. Adel, M. Louerat, and M. Sabut, "Fast split background calibration for pipelined ADCs enabled by slope mismatch averaging technique," *Electronic Letters*, vol. 48, pp. 318–320, Mar. 2012.
- [13] D.-Y. Chang, "Design techniques for a pipelined ADC without using a front-end sample-and-hold amplifier," *IEEE Trans. Circuits Syst.-I*, vol. 51, p. 2123 – 2132, Nov. 2004.

-
- [14] M. Waltari and K. Halonen, "1-V 9-bit pipelined switched-opamp ADC," *IEEE J. Solid-State Circuits*, vol. 36, pp. 129–134, Jan. 2001.
 - [15] B. Razavi and B. Wooley, "Design techniques for high-speed, high-resolution comparators," *IEEE J. Solid-State Circuits*, vol. 27, p. 1916–1926, Dec. 1992.
 - [16] Analog Devices, *High Speed ADC Analog Input Interface Considerations*. Application Engineering Notebook, MT-228, 2012.
 - [17] R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*. Wiley-IEEE Press, Nov 2004.

5

Residue Amplifier with Analog Error Correction

In previous chapters, the motivation for designing residue amplifiers with low circuit overhead and reduced settling accuracy was established. However, despite achieving a better power efficiency, the usefulness of such designs is limited by their low gain accuracy and linearity. An effective approach to utilize residue amplifiers with low accuracy has been to combine them with digital calibration which, as demonstrated in Chapter 4, can be used to compensate for the ADC errors resulting from amplifier gain and distortion errors, leading to a significant reduction in residue amplifier power consumption. Due to the rising ubiquity of digital processing, most calibration approaches for residue amplifier non-idealities, including the one presented in Chapter 4, perform both error *detection* and *correction* in the digital domain through post-processing [1–3]. In this chapter, an alternative analog-domain approach for correcting amplifier gain and distortion errors is presented. This allows the amplifier non-idealities to be corrected at their source, thus achieving a much better amplifier performance without employing excessive circuit complexity or overhead.

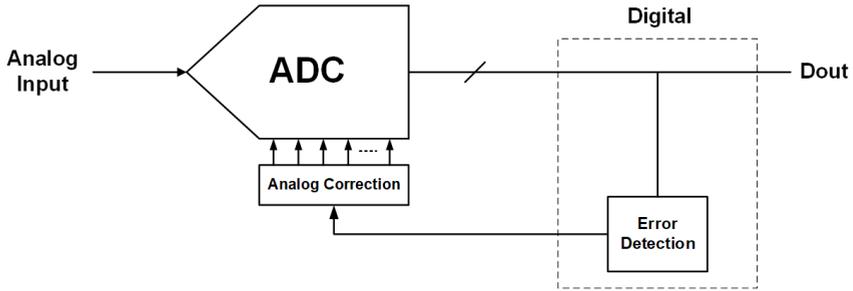


Figure 5.1: ADC calibration with analog error correction

5.1. ANALOG ERROR CORRECTION

Figure 5.1 shows the principle of ADC calibration using error correction in the analog-domain. Analog error correction typically involves deploying a tunable element in the circuit under calibration that can be used to correct its non-idealities. Analog correction has been used for offset calibration in comparators [4], curvature-compensation in bandgap references [5] and to correct the timing skew between time-interleaved ADC lanes [6]. There are several advantages of performing error correction in the analog domain. Firstly, although digital correction can correct deterministic errors, as discussed in Chapter 4, it cannot compensate for the lost dynamic range in the ADC output code. For instance, any reduction in amplifier gain will result in a lower ADC full-scale, which will reduce ADC SNR despite gain error calibration in post-processing. With the help of analog correction, however, the errors can be fixed directly at the source, restoring the ADC full-scale swing and resulting in a better ADC SNDR.

Secondly, any digital error correction that relies on the ADC backend code, such as distortion correction, is fundamentally limited by the accuracy of the backend. So, although the effect of noise in the ADC backend on amplifier distortion *detection* can be reduced through averaging, the error *correction* will be limited by the inaccuracies in the ADC backend code due to noise and distortion sources in the backend. Analog error correction, on the other hand, does not suffer from this limitation.

Finally, in many cases, it is much simpler to counteract analog non-idealities by using analog mechanisms rather than by using digital post-processing. For certain error sources, it is relatively simple to generate and apply a compensating error source in the analog-domain. For instance, to calibrate the offset of a comparator, it is much easier to create an opposite imbalance inside the comparator that is con-

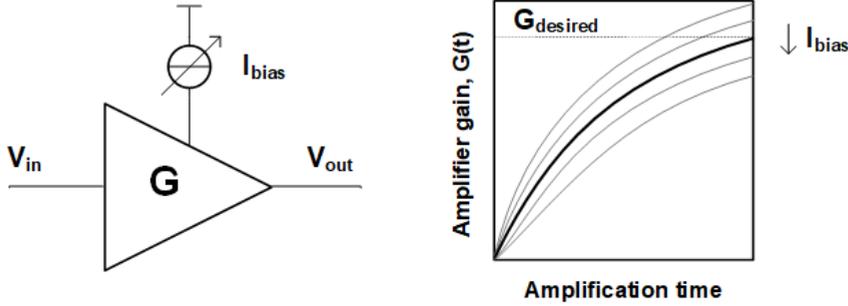


Figure 5.2: Amplifier gain error correction using incomplete settling

trolled by a calibration loop. Similarly, for amplifier distortion, it is fairly simple and area-efficient to generate a distortion of the opposite phase in the analog-domain. Hence, when compared to gain and distortion correction in the digital domain which require dedicated digital logic running at the full ADC clock speed to process and correct ADC errors, analog error correction methods can offer significantly more power and area efficient solutions. As the focus of this thesis is on residue amplifiers, in the next few sections, we will present ways of correcting amplifier gain and distortion errors in the analog-domain.

5.2. GAIN ERROR CORRECTION

Traditionally, amplifiers with a very high open-loop gain and settling accuracy are needed to achieve the desired gain accuracy in the first few stages of a high-resolution pipelined ADC. Since this can be very costly in terms of area and power dissipation, gain error calibration has become a popular and widely used approach. Depending on the amplifier topology, several approaches can be applied to correct the gain in the analog domain. One approach, suitable for both open- and closed-loop amplifiers, corrects the gain error by leveraging incomplete amplifier settling [7, 8], as shown in Figure 5.2. By using an amplifier with incomplete settling, its gain at the end of a fixed amplification time can be easily controlled by tuning the bias current. The range and resolution of the gain error correction depend on the programmability of the bias current and the amount of amplifier settling. The lower the amplifier settling, the more sensitive the amplifier gain will be to the bias current. This tunability is maximized in the case of integrating amplifiers, as the gain becomes, to first order, directly proportional to the bias current.

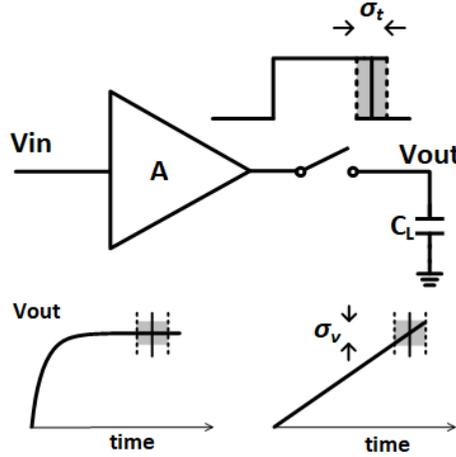


Figure 5.3: Effect of clock jitter in settling vs integrating amplifier

While the simplicity of this incomplete settling approach makes it very attractive for gain error correction, it also suffers from certain drawbacks. Incomplete settling compromises the amplifier settling accuracy, which makes the amplifier gain less well-defined. For closed-loop amplifiers, less settling can result in higher distortion, as discussed in Chapter 3. Low amplifier settling also makes the amplifier gain more susceptible to clock jitter. For an exponentially settling amplifier, the amplifier output approaches its final value well before the sampling point, as shown in Figure 5.3, making the sampled output voltage insensitive to any jitter in the sampling clock. At the other extreme is the case of an integrating amplifier, whose output never settles. As discussed in Chapter 3, the integrating amplifier's gain increases linearly with time, causing any uncertainty in the sampling moment to directly appear as noise on the output voltage sampled by the succeeding stage. Since clock jitter has a multiplicative effect on an integrating amplifier's output voltage, the resulting noise is directly proportional to the slope of the output and can be expressed as –

$$\text{Output jitter-induced noise voltage, } \sigma_v = V_{in} \frac{g_m}{C_L} \sigma_t \quad (5.1)$$

where σ_t is the RMS jitter in the sampling clock. Ignoring any other non-idealities, the SNR at the integrating amplifier's output resulting purely from clock jitter can

be calculated as –

$$\text{SNR}_{\text{jitter}} = \frac{(A_{\text{int}}V_{\text{in}})^2}{\left(V_{\text{in}} \frac{g_m}{C_L} \sigma_t\right)^2} = \left(\frac{T_a}{\sigma_t}\right)^2 \quad (5.2)$$

It can be inferred from (5.2) that as the amplification period decreases, the jitter-induced noise will become more significant. This implies that for higher ADC clock speeds, it can be beneficial to use interleaving to minimize the impact of clock jitter on the residue amplifier SNR.

Another drawback of amplifiers with incomplete settling is that their gain is no longer strictly a ratio of passive devices, and can vary significantly over the entire PVT range. While background calibration can help control the gain, the amplifier should also be paired with an on-chip constant gm biasing circuit to reduce the gain variation and, hence, relax the required gain correction range and increase robustness over a wide range of PVT conditions.

In order to maintain the benefits of negative feedback and complete settling, the amplifier gain error can also be corrected by adjusting the feedback factor with the help of a capacitive array [9]. This allows the use of op-amps with low DC gain. Although the attenuating effect of negative-feedback on distortion varies when the feedback factor is adjusted to fix the gain, the overall effect was shown to be minor. The main drawback of this scheme is the additional area required to implement the capacitive array to tune the amplifier's closed-loop gain, which could be significant even for medium-resolution gain correction. This limits their application to low- and medium-resolution ADCs.

5.3. DISTORTION CORRECTION

As demonstrated in Chapter 4, distortion correction in the digital domain can be performed by adding a cubic term that is equal and opposite to the third-order distortion detected in the residue amplifier. The motivation for performing this correction in the analog domain is based on the fact that it is more power-efficient to generate this compensating cubic term in the analog domain.

The idea of using two opposing distortion mechanisms as a linearization technique has been employed for cancelling the HD3 component in several types of amplifiers, ranging from RF low noise amplifiers (LNAs) [10, 11], programmable gain amplifiers (PGAs) [12], transconductors for g_m -C filters [13] and ADC residue amplifiers [14, 15]. These techniques can broadly be classified into two categories, shown in Figure 5.4, based on how the secondary distortion component is generated

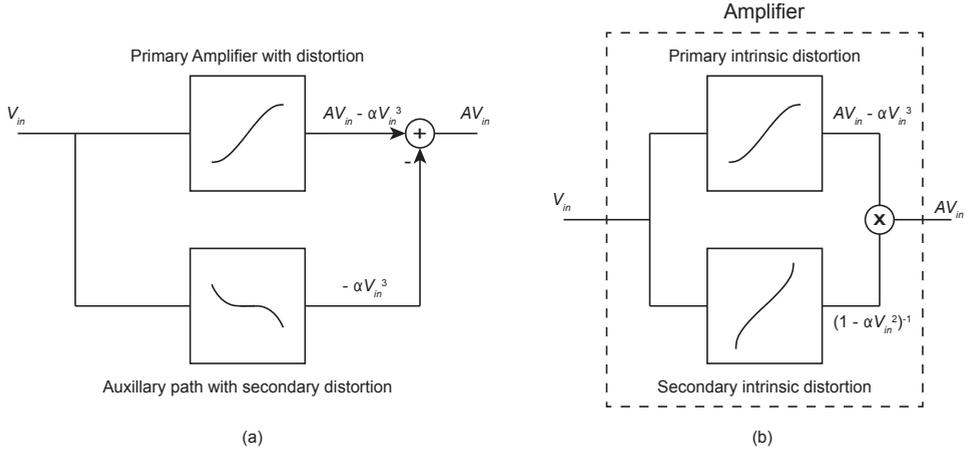


Figure 5.4: Analog-domain cancellation of amplifier distortion through (a) summation with an externally generated secondary distortion, or (b) multiplication with an internally generated secondary distortion

and used to cancel the primary intrinsic amplifier distortion. The first category [10, 13] is based on a summation of the amplifier output containing the primary amplifier distortion with an externally generated secondary distortion of the opposite phase. Both [10] and [13] use a feedforward approach, where an auxiliary path, consisting of a cubing circuit with a certain coefficient, is used to generate a term that is equal to the HD3 in the main signal path. By subtracting the outputs of the two paths, the HD3 component in the overall amplifier output can be cancelled. Since the competing distortion term is generated by a separate auxiliary path, it can be made easily tunable without disturbing the main path. This feature makes feedforward linearization an excellent candidate for working together with digital calibration. However, in most cases, the auxiliary path also generates a linear term that reduces the gain of the primary amplifier (26% gain reduction reported in [13]), while adding excess noise, resulting in significant degradation in the amplifier SNR.

The techniques in the other category utilize a secondary intrinsic distortion mechanism within the amplifier to cancel the primary distortion term. Since both the primary and secondary distortion mechanisms arise from the same amplifying circuit, it often results in a multiplication of these two distortion components. By making their influence on the amplifier gain equal and opposite, the amplifier gain can be linearized. An important advantage of this approach is that because there is no dedicated parallel path to generate the secondary distortion, the effect on the amplifier SNR is minimal. Several options for actuating this secondary distortion

mechanism have been identified in the literature. In [11, 16], an IM2 term generated by a squaring circuit is injected into the tail current of an input differential pair to cancel the HD3 term. [12] and [14] use amplifier stages with opposite distortion profiles in parallel and series, respectively, to generate an overall linear gain, while [15] utilizes the rate of change in the output common-mode of a dynamic amplifier to counterbalance the compressing HD3 in the amplifier output voltage. All these approaches promise good power efficiency. However, since the competing secondary distortion profile is now arising from within the amplifier itself, it can be difficult to implement it with a well-defined tuning knob, making it challenging to control the secondary distortion mechanism with a calibration loop. This is evident in all the designs mentioned above, where none of the approaches attempt to tune the competing distortion mechanism, resulting in only a moderately successful cancellation of the amplifier distortion (<10b SFDR).

For application in high-resolution ADCs, the holy grail of analog distortion correction is an approach that offers good tunability and control of the competing distortion mechanism without adding any excess noise and power dissipation. In the following sub-sections, we will introduce an amplifier topology that exhibits the potential for both these qualities.

5.3.1. LINEARIZATION PRINCIPLE¹

In Chapter 3, we studied the effect of negative feedback on a discrete-time amplifier's distortion components in the time-domain. In order to simplify the mathematical analysis, the effect of second-order distortion coefficient on the third-order distortion component was neglected. This effect can be calculated by analyzing the distortion of the negative-feedback amplifier shown in Figure 5.5 in steady-state. Its open-loop gain, A_{OL} , for an input voltage, V_a , can be expressed as –

$$A_{OL} = a_1 + a_2 V_a + a_3 V_a^2 \quad (5.3)$$

where a_1 , a_2 and a_3 are the first-, second- and third-order coefficients, respectively. When this amplifier is used in negative feedback with a feedback factor, f , its overall closed-loop gain, A_{CL} , can similarly be expressed as –

$$A_{CL} = b_1 + b_2 V_{in} + b_3 V_{in}^2 \quad (5.4)$$

¹This sub-section is partially derived from publication: Md. Akter, R. Sehgal and K. Bult, "A resistive degeneration technique for linearizing open-loop amplifiers" *IEEE Trans. Circuits and Systems II*, vol. 67, no. 11, pp. 2322–2326, Nov. 2020.

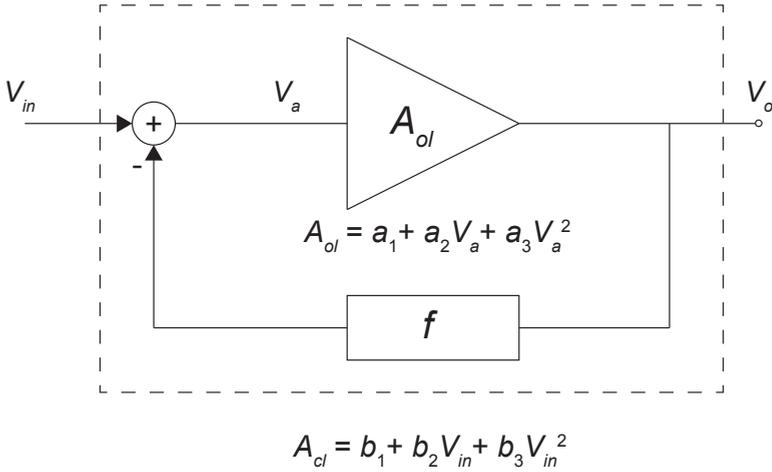


Figure 5.5: Closed-loop amplifier model with distortion

where b_1 , b_2 and b_3 are the coefficients defining the dependence of A_{CL} on the input voltage, V_{in} . Using (5.3) and (5.4), the overall third-order distortion component of the closed-loop amplifier can be calculated as shown in [17]–

$$b_3 = \frac{a_3(1 + a_1 f) - 2f a_2^2}{(1 + a_1 f)^5} \quad (5.5)$$

It can be seen that the third-order distortion coefficient for the closed-loop amplifier consists of two components. The first component originates from the open-loop third-order coefficient a_3 while the second term is a result of the second-order interaction around the feedback loop. This second term is only generated as a consequence of feedback and the presence of second-order distortion, and is always compressing in nature. Based on (5.5), the overall third-order distortion can be completely cancelled if the following condition is achieved –

$$a_3 = \frac{2f a_2^2}{(1 + a_1 f)} \quad (5.6)$$

It should be noted that to achieve this cancellation we need the third-order coefficient, a_3 , to have the same polarity as a_2^2 . In other words, the third-order coefficient needs to be expanding.

An example of this linearization principle is presented in [18] which analyses a common-emitter amplifier with emitter degeneration, as shown in Figure 5.6. A BJT has an exponential I - V transfer characteristic, and for a certain base-emitter

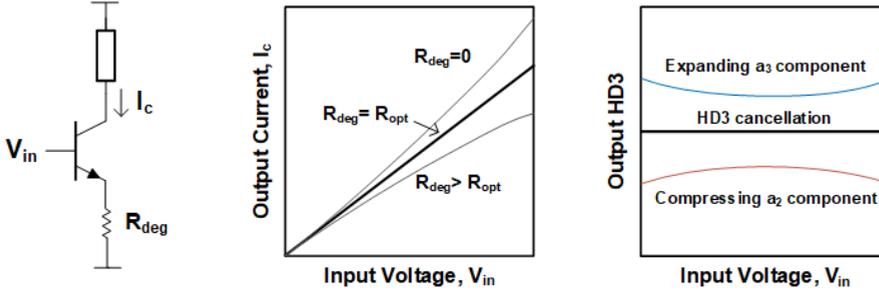


Figure 5.6: Linearization technique for HD3 cancellation through emitter degeneration

voltage, v_{be} , its collector current can be expressed as –

$$i_C = I_{CS} e^{v_{be}/V_T} \quad (5.7)$$

where I_{CS} is the collector saturation current and V_T is the thermal voltage and is equal to kT/q . Due to this exponential nature, the output collector current will contain both even- and odd-order components, all expanding in nature.

The degeneration resistor, R_{deg} , provides local negative feedback for the BJT common-emitter amplifier and controls the collector current by keeping v_{be} in check, with a loop-gain of $g_m R_{deg}$. While emitter degeneration is commonly used with a large loop-gain in order to make the g_m mostly dependent on $1/R_{deg}$, it is possible to cancel the HD3 component by picking a specific value of R_{deg} based on (5.6). By using the Taylor series expansion of (5.7) ($a_2 = a_1/2$; $a_3 = a_1/6$; $a_1 f = g_m R_{deg}$), the condition in (5.6) can be rewritten as –

$$g_m R_{deg} = \frac{1}{2} \text{ or, } R_{deg} = \frac{1}{2g_m} \quad (5.8)$$

Considering that the g_m of a BJT is equal to i_c/V_T , the voltage drop across the degeneration resistor, $V_{R_{deg}}$, can be calculated as –

$$V_{R_{deg}} = R_{deg} I_C = V_T/2 \quad (5.9)$$

The HD3 cancellation condition given by (5.9) can also be realized using MOSFETs in a common-source amplifier with source degeneration provided that the input MOS transistor is operating in deep weak-inversion. The MOS transistor g_m in weak inversion can be expressed as –

$$g_m = \frac{I_D}{nV_T} \quad (5.10)$$

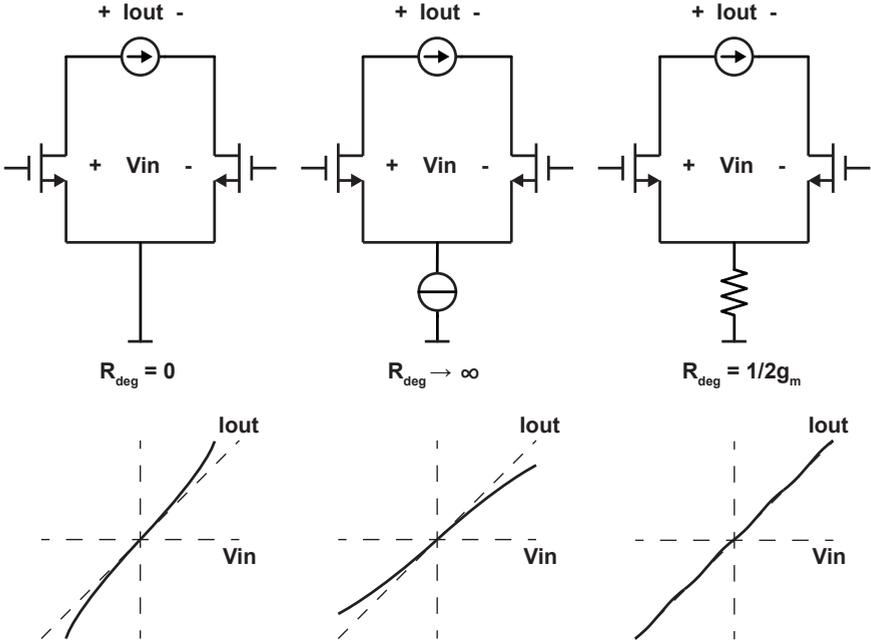


Figure 5.7: Amplifier linearization using common-mode degeneration resistance

where I_D is the drain bias current and n is the weak-inversion slope factor. Using the above expression, the condition for HD3 cancellation can then be calculated as

$$R_{deg} = \frac{1}{2ng_m}; \text{ and } V_{Rdeg} = V_T/2 \quad (5.11)$$

Considering the above equation, the optimum g_m degeneration factor ($1 + g_m R_{deg}$) for HD3 cancellation is $1 + 1/2n$. n is a technology-dependent parameter and is approximately 1.4, resulting in a g_m degeneration of nearly 35%.

5.3.2. PROPOSED LINEARIZED AMPLIFIER²

The previously discussed linearization principle can be extended to design a differential amplifier with third-order distortion cancellation, as shown in Figure 5.7, where the local negative feedback is now provided by a tail degeneration resistor, R_{deg} . An intuitive way of looking at this amplifier is by considering its operation across two extreme conditions, namely (1) $R_{deg} = 0$ and (2) $R_{deg} \rightarrow \infty$.

²This sub-section is partially derived from publication: R. Sehgal, F. van der Goes and K. Bult, "A 13-mW 64-dB SNDR 280MS/s pipelined ADC using linearized integrating amplifiers" *IEEE Journal of Solid-State Circuits*, vol. 53, no. 7, pp. 1878–1888, Jul. 2018.

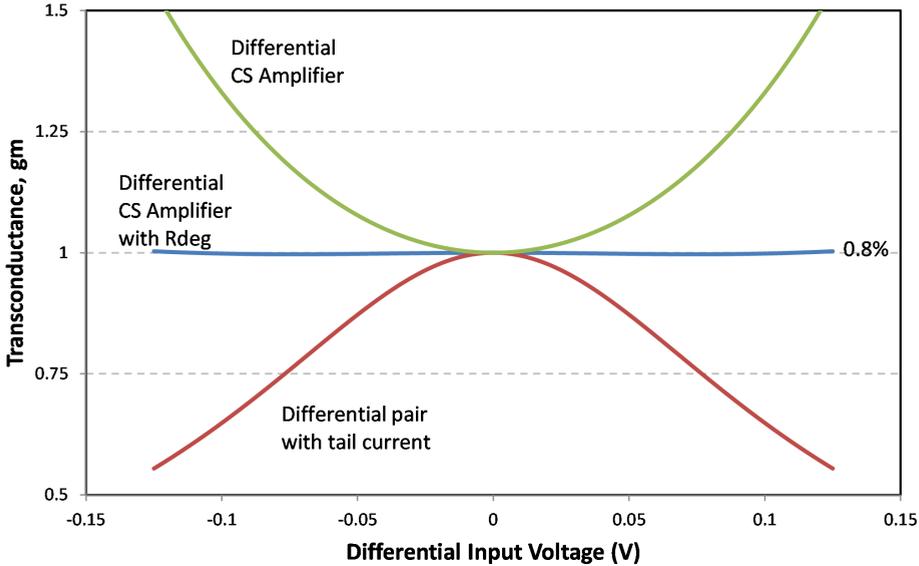


Figure 5.8: Transconductance of the proposed amplifier with CMD linearization compared with a conventional differential common-source amplifier and a differential pair

The first case represents a differential common source amplifier. If the input transistors are biased in deep weak inversion (WI), a zero source resistance means that there is no restriction on the total current being drawn by the two transistors. As a result, due to the exponential nature of the input transistors, this amplifier displays an expanding $V - I$ characteristic. In the second case, the infinite tail resistance exhibited by an ideal tail current source would restrict any change in total current. Due to the limitation on the total current, the amplifier would shift to a compressing $V - I$ characteristic. The linearization principle relies on the idea that between these two opposing distortion paradigms, an intermediate tail resistance can be identified for which the amplifier would exhibit an optimally linear input transconductance.

The value of this optimum tail resistance can be determined by analytically comparing this circuit with the one in Figure 5.6. Considering that the tail resistance only sees the sum of the two differential currents, it essentially acts as common-mode degeneration (CMD). This means that, effectively, only the even-order terms in the amplifier output current are degenerated. Hence, just like the example shown in Figure 5.6, the second-order term will again generate the compressing third-order distortion component through negative feedback, which can be used to cancel the

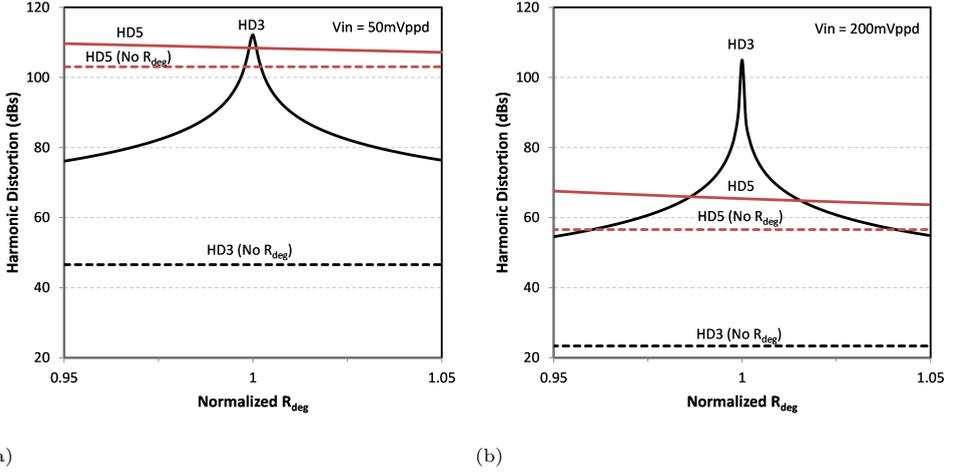


Figure 5.9: Amplifier nonlinearities versus their normalized sweep parameters for (a) 50mV_{ppd} and (b) $200\text{mV}_{\text{ppd}}$ input signals.

primary expanding third-order component. Although, due to the common-mode nature of the degeneration, neither the expanding or compressing third-order components will be suppressed by the loop-gain, the cancellation of the two components will again occur at the conditions defined in (5.4) and (5.9). This is illustrated in Figure 5.8, where the g_m of the amplifier with a CMD resistor of $R_{deg} = 1/2ng_m$ is compared to the g_m of a differential CS amplifier and a differential pair with an ideal tail current. By applying the CMD technique, the g_m variation across a differential input signal range of 125mV reduces from nearly 50% to less than 1%. This residual g_m variation can be attributed to the presence of higher-order distortion components.

A major advantage of using CMD is that it only degenerates the even-order terms. Due to this, the first-order gm of the amplifier remains largely unaffected. This is also seen in Figure 5.7 where the small-signal g_m 's of all three amplifiers are identical. Hence, using the CMD technique avoids the 35% reduction in the input g_m at the HD3 cancellation point.

Figure 5.9 shows the third- and fifth-order harmonic distortion components of a differential CS amplifier, with and without CMD, for two different input swings, as the R_{deg} is swept over a range of $\pm 5\%$ from its optimal value. It can be seen that the CMD linearization technique helps in suppressing the amplifier HD3 by more than 30dB across the entire range of R_{deg} for both 50mV_{ppd} and $200\text{mV}_{\text{ppd}}$ input swing. It should be noted that for 50mV_{ppd} input swing, HD5 is very small and the overall

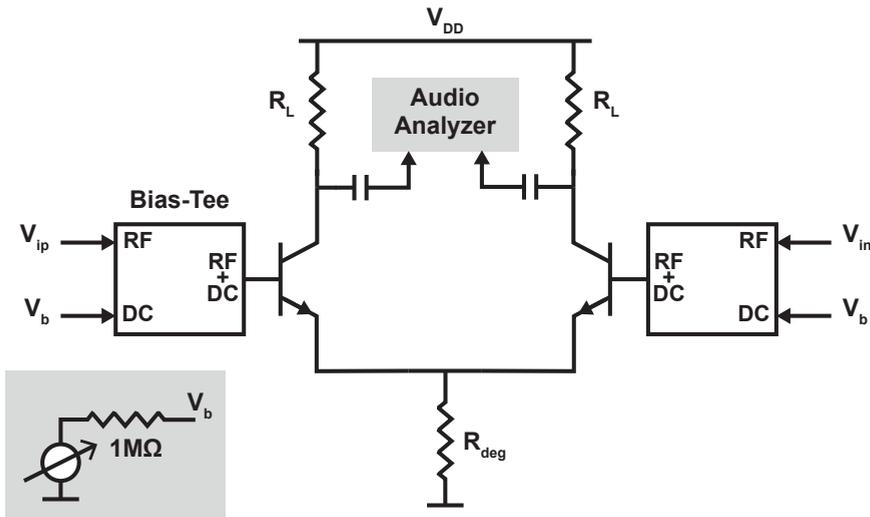


Figure 5.10: Measurement setup for verification of the CMD linearization principle

distortion is largely dominated by HD3. However, for larger output swings, HD5 starts to dominate the overall distortion performance of the CMD-based amplifier. In later sections, we will explore how the CMD technique can also be utilized for also cancelling HD5.

5.3.3. VERIFICATION

As an initial step towards the verification of the CMD linearization principle, the proposed amplifier was implemented on a stripboard using discrete BJTs and resistors, as shown in Figure 5.10. An audio analyzer, APx555, was used to assess the frequency spectrum of the amplifier output, along with an input amplitude of 50mV_{ppd} and an amplifier gain of 6x. Instead of adjusting the degeneration resistor, R_{deg} , the bias voltage V_b was tuned to drive the amplifier towards the HD3 cancellation point. The bias voltage was chosen as the sweep parameter because of the much finer tunability of the voltage source used for generating the bias voltage. The measured output spectra at the optimal linearity for 50mV_{ppd} input are shown in Figure 5.11. With the help of the CMD technique, the amplifier is able to achieve an excellent third-order harmonic distortion of 95dB. Apart from the several non-harmonic spurs in the output spectrum resulting from the rudimentary nature of the test setup, the amplifier output THD is dominated by the HD2 component that arises due to the mismatch between the discrete BJTs. In the next section, a

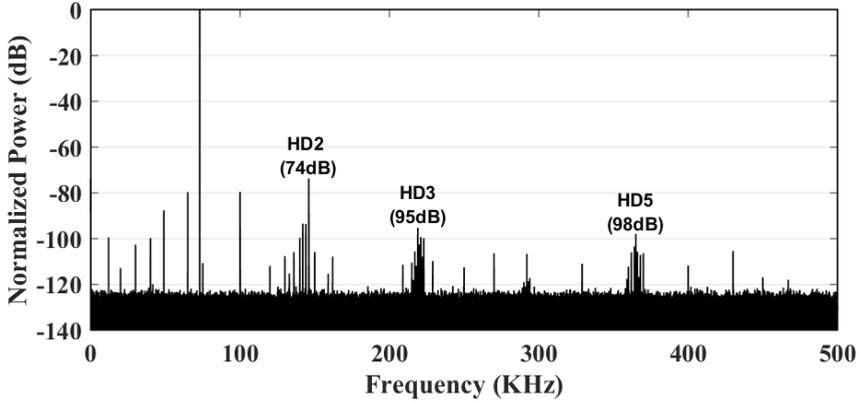


Figure 5.11: Amplifier output spectrum for optimal CMD resistance

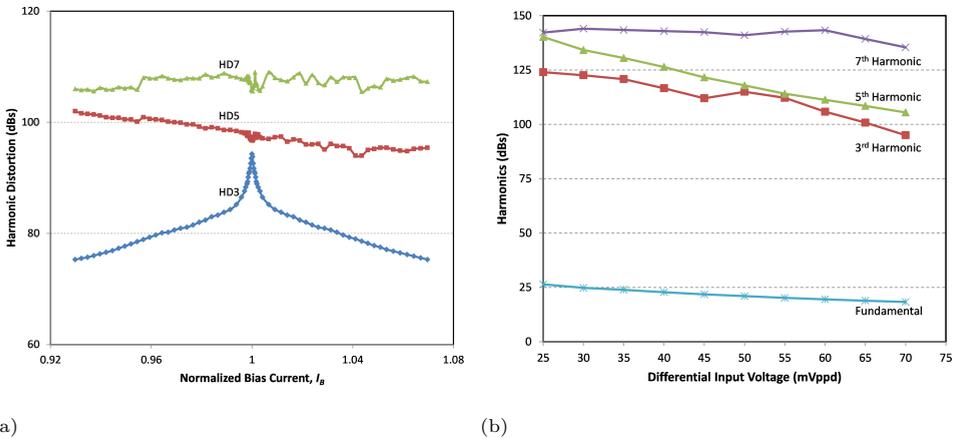


Figure 5.12: (a) Amplifier harmonic distortion vs bias current sweep (for 50mVpp-diff) and (b) Amplifier harmonic tones vs input voltage swing

modification to the proposed linearization technique will be presented that can be used to reduce the even-order distortion components.

The sensitivity of the CMD linearization technique is indicated in Figure 5.12, which shows the effect of variation in the amplifier biasing point and the input voltage swing on the measured distortion components of the CMD amplifier. As shown in Figure 5.12 (a), the amplifier exhibits a wide linear range and maintains better than -75dB HD3 even when the base current is varied by $\pm 5\%$ from the optimum value. Figure 5.12 (b) shows the linearity performance of the CMD amplifier over input amplitude sweep, with the HD3 cancellation optimized at an input

swing of 50mV_{ppd} . Over an input amplitude range of $25\text{-}70\text{mV}_{\text{ppd}}$, the amplifier exhibits better than -80dB HD3 without any re-calibration. These measurements, while performed over a basic amplifier setup on a stripboard with discrete components, demonstrate the validity of the linearization principle and its potential for implementation in integrated circuits.

5.4. MOS IMPLEMENTATION

As discussed in the previous subsections, since the linearization principle is based on an exponential V-I characteristic, it can be implemented in amplifiers consisting of MOS transistors biased in weak inversion. In this section, an all-MOS implementation of the amplifier, including the tunable tail degeneration resistor will be presented.

In order to achieve a stable HD3 cancellation, we need to achieve the condition described in (5.11) over a wide range of PVT conditions. In addition to that, R_{deg} needs to take into account any device mismatch and device & wiring parasitics, such as, the source series resistance, which can easily shift the cancellation point. Hence, while the cancellation will always improve the amplifier linearity, to ensure the amplifier operates at the optimum HD3 cancellation point, we need to supplement it with tunability and a calibration loop. In this section, we will discuss the details of the amplifier implementation and the several linearization parameters that were added to tune its secondary distortion mechanism.

5.4.1. LINEARIZATION PARAMETERS

As emphasized in the previous sections, for the amplifier to operate close to the optimum distortion cancellation point, the tail degeneration resistance needs to be adjustable by a calibration loop. Considering an implementation using polysilicon resistors, a programmable resistive array with even a moderate resolution can be very costly in terms of area. However, it can be easily implemented by a MOS transistor in the linear region, as seen in Figure 5.13. When biased in the strong-inversion triode region, the MOS drain-source resistance can be expressed as –

$$R_{out} = \frac{1}{\frac{W}{L} \mu_n C_{ox} (V_{GS} - V_{TH} - \frac{V_{DS}}{2})} \quad (5.12)$$

Since the voltage swing at the tail node is relatively low, the MOS resistor can operate with sufficient linearity to effectively degenerate the CS amplifier. And by

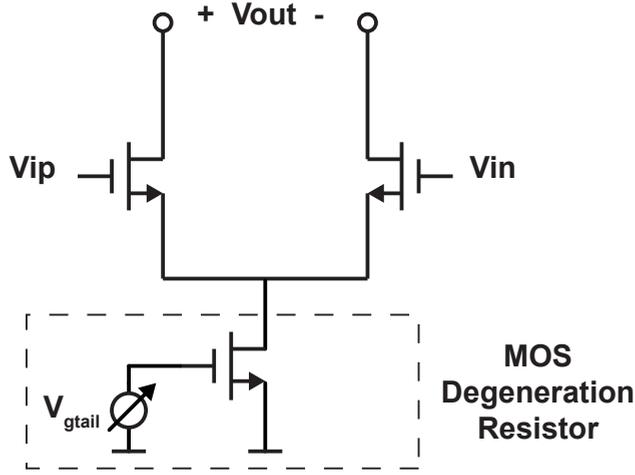


Figure 5.13: CMD amplifier with MOS as degeneration resistor

tuning the bias voltage of the tail NMOS, the linearity of the overall amplifier can be tuned with much finer steps.

So far, the focus has been on using a specific degeneration resistance value to cancel the HD3. The overall THD of the amplifier after linearization can be assessed by considering the most significant odd-order distortion components. As shown in Figure 5.9, for larger input swings ($\geq 200\text{mV}_{\text{ppd}}$), HD5 starts limiting the overall amplifier linearity after HD3 cancellation to $<75\text{dB}$. For applications requiring even higher linearity, it would be interesting to explore degeneration schemes that can be used for cancelling both HD3 and HD5, allowing the amplifier to achieve higher linearity ($>75\text{dB}$) without limiting its voltage swing.

Considering again a differentially driven input pair, with a differential input of V_{id} , consisting of MOS transistors biased in deep weak inversion, its differential output current can be perfectly linearized by ensuring that the sum of the currents in the two input transistors is equal to (Appendix B) –

$$I_{tail} = g_m V_{id} \coth\left(\frac{V_{id}}{2nV_T}\right) \quad (5.13)$$

where g_m is the desired differential transconductance of the input pair and I_{tail} is the total current flowing through the input pair. Equation (5.13) can be further simplified with the help of Taylor series expansion –

$$I_{tail} = g_m n V_T \left(1 + \frac{1}{3} \left(\frac{V_{id}}{2nV_T}\right)^2 - \frac{1}{45} \left(\frac{V_{id}}{2nV_T}\right)^4 + \dots\right) \quad (5.14)$$

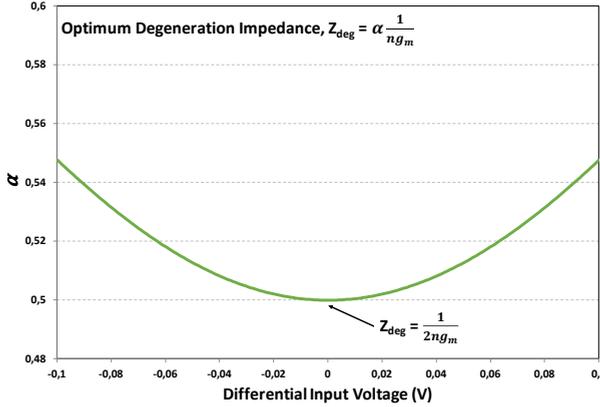


Figure 5.14: Optimum common-mode degeneration impedance for a differential weak-inversion CS amplifier

The common-mode degeneration impedance required to ensure the condition in (5.13) can be derived to be (Appendix B) –

$$Z_{deg} = \frac{1}{2ng_m} \left(\frac{\cosh(x) - \sinh(x)/x}{\cosh(x) - x/\sinh(x)} \right) \quad (5.15)$$

where x is equal to $V_{id}/2nV_T$. The degeneration resistance calculated in (5.15) has a certain dependence on the input signal, V_{id} . The variation of this optimum degeneration resistance over V_{id} is plotted in Figure 5.14. If only the third-order term is considered, using the Taylor series expansions for hyperbolic functions, the above expression can be simplified to the same value as calculated in equation (5.11), $1/2ng_m$.

Although it would be very difficult to realize the exact impedance function described in (5.15), the ideal value for I_{tail} in (5.14) can be approximated by two transistors biased in weak-inversion and connected in parallel, while being driven by the differential input V_{id} , as shown in Figure 5.15. The value of the tail current can be tuned by changing its bias voltage and, since the input signal is coupled to the tail transistors through a coupling capacitor, the extent to which the input signal modulates this degeneration current (or the effective tail resistance) can be adjusted by using a tunable attenuation capacitor to ground. By using a programmable capacitive array, the drive strength of the input signal towards the tail degeneration can be tuned. By selecting the right attenuation capacitor value, the tail degeneration characteristic can be optimized.

These two knobs, tail bias voltage and attenuation capacitor, establish how

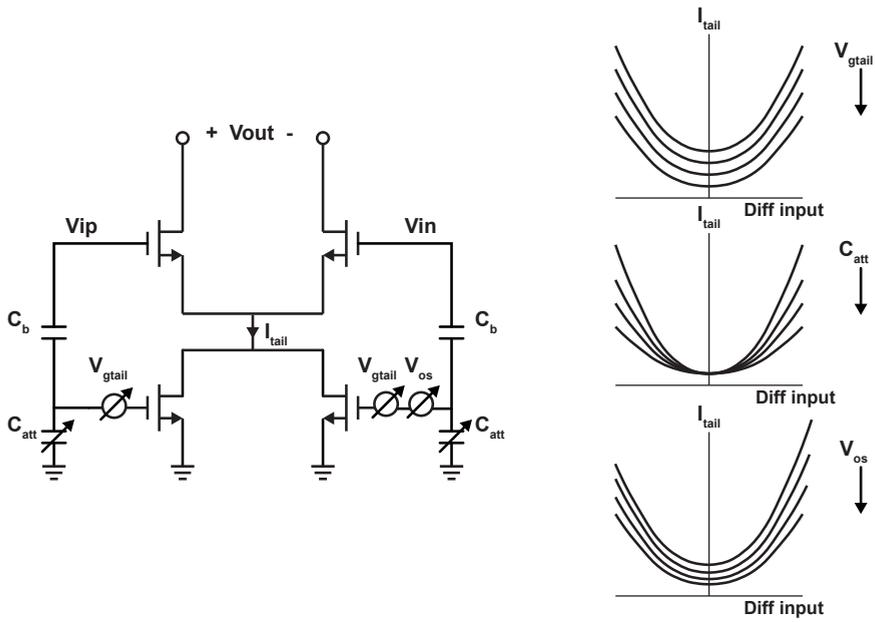


Figure 5.15: Proposed CMD amplifier with three linearization parameters

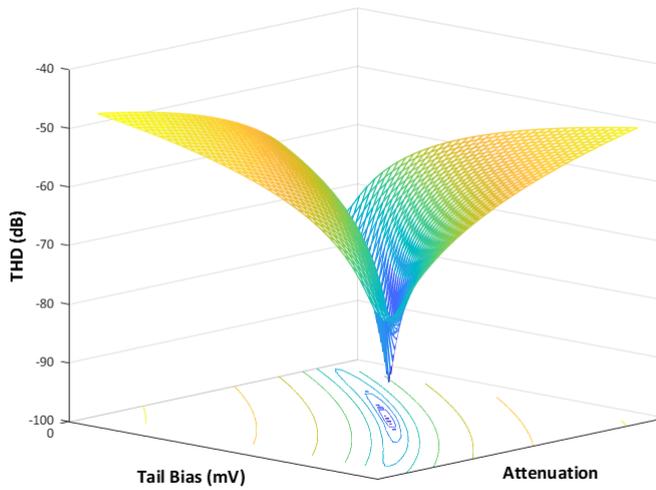


Figure 5.16: Simulated amplifier THD versus linearization parameters

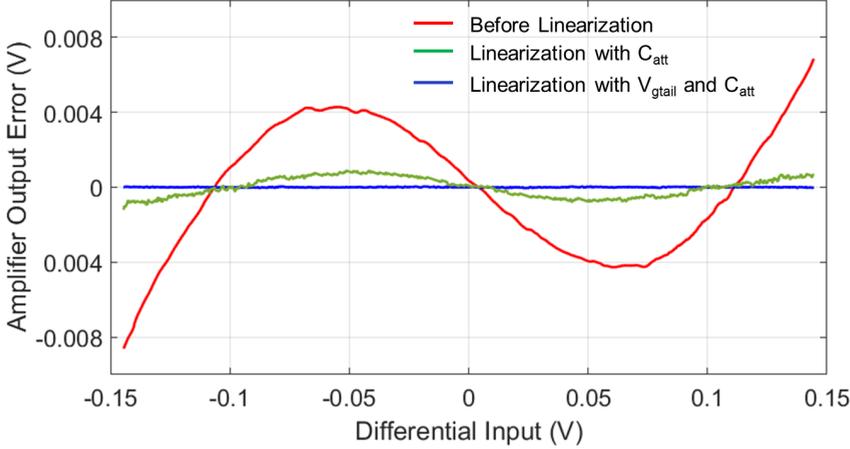


Figure 5.17: Amplifier output error before and after linearization

closely the tail current characteristic approximates the ideal function, hence having a significant effect on the linearity of the amplifier. This is reflected in Figure 5.16, which shows the simulated THD (considering the first five harmonics) of a 250MHz switched-cap amplifier using the proposed CMD amplifier, with an input swing of $300\text{mV}_{\text{ppd}}$ and a gain of 4x, as these two parameters are swept. It can be seen that over the entire search space of combinations of these two parameters, a unique set can be found for which the amplifier distortion exhibits a minimum, where the THD is $< -90\text{dB}$.

Considering that these two parameters will be tuned by a calibration loop, the uniqueness of this combination of parameters should make it easier for the calibration to locate this optimum linearity point. Figure 5.17 shows the deviation of the amplifier output from its ideal value. For an input swing of $300\text{mV}_{\text{ppd}}$, the error reduces from nearly -40dB to -55dB by sweeping only C_{att} and nearly -90dB after linearization with both C_{att} and V_{gtail} .

5.4.2. EVEN-ORDER DISTORTION

While these two tuning parameters are very effective against odd-order distortion, they only adjust the circuit symmetrically. Due to the high inherent distortion of the WI input pair, any offset or gm mismatch in the amplifier will create a significant even-order distortion. This is also reflected in the measurement results shown in Figure 5.11. In order to avoid the even-order distortion from limiting the

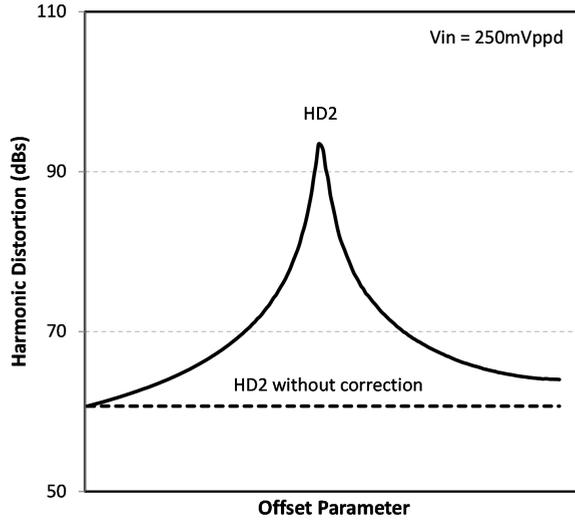


Figure 5.18: HD2 vs Tail degeneration offset parameter

overall amplifier linearity, the mismatch can be compensated for by adding an extra tunable offset, V_{os} , in the tail current source, as seen in Figure 5.15, with the help of an additional programmable current source in the bias diode of one of the tail transistors. This offset provides a knob to counteract the imbalance arising from the mismatch of the input transistors. Figure 5.18 shows the effect of the tail offset on the HD2 of the proposed amplifier with a 2% mismatch in its input devices. It should be noted that this tunable tail current offset does not cancel the inherent offset of the amplifier, but merely corrects the even-order distortion by counteracting the imbalance within the amplifier.

5.5. DESIGN CONSIDERATIONS

In the previous sections, an amplifier linearization principle using a tunable input-driven common-mode degeneration resistance was presented. Although the proposed amplifier was able to achieve excellent linearity without excessive circuit overhead, other practical factors could potentially affect its performance or power-efficiency. In this section, we will discuss some of these pertinent design aspects of the proposed amplifier topology.

5.5.1. HIGHER-ORDER EFFECTS

In Figure 5.12 (b), it is shown that as the input swing increases, the linearization starts becoming relatively less effective. This is because several higher-order effects become more consequential for larger input swings. The most significant higher-order mechanism that affects the MOS implementation of the CMD amplifier is the degree of channel inversion. For large input signals, one of the input transistors starts approaching moderate inversion. As a result, the $V - I$ characteristic of that input transistor is no longer strictly exponential and hence its distortion coefficients begin to significantly deviate from those calculated using the Taylor series expansion. Due to this, the conditions described in (5.11), (5.13) and (5.15) no longer hold. This particular dependency of the linearization on the input swing can be relaxed by biasing the input transistors in deep weak inversion so that they remain firmly in weak inversion across the entire input range.

Other significant factors that could limit the HD3 cancellation include the influence of the body effect on the threshold voltage and the weak-inversion slope factor, and the distortion arising from the nonlinear transistor output impedance. The weak-inversion slope factor, n , is a function of the depletion layer width, which is affected by the gate-bulk voltage and short-channel effects like DIBL, etc. This results in n , and hence the HD3 cancellation point, varying with gate and drain voltages. The impact of output impedance distortion also limits the overall linearity that can be achieved from the amplifier and becomes more prominent as the desired residue amplifier gain is increased. The effect of the output swing on the input transistor's output impedance can be somewhat reduced by adding cascode devices.

5.5.2. NOISE

Compared to a differential common-source amplifier, the proposed amplifier has two additional circuit elements that could affect the amplifier SNR –

- a. Tail degeneration resistors
- b. Coupling and attenuation capacitors

The noise originating from the tail degeneration is comparable to the noise introduced by the tail current source of a differential pair, and will have a multiplicative effect on the overall amplifier differential output noise. In case of small input signals, due to the common-mode nature of the tail degeneration, its noise current will split equally between the positive and negative branches, hence causing no effect

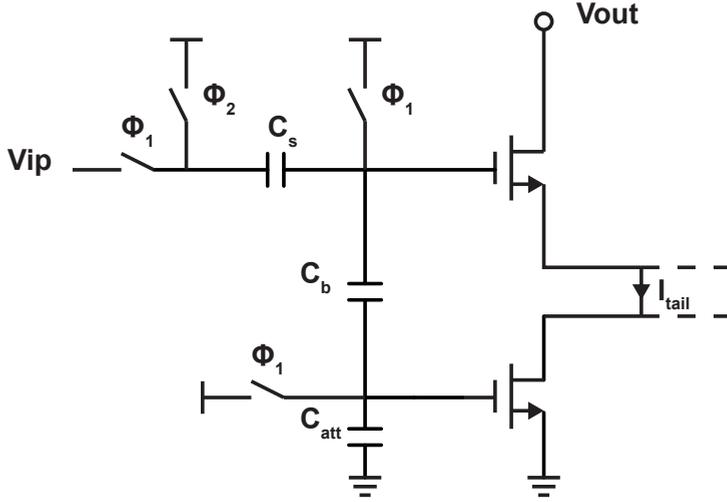


Figure 5.19: Amplifier half-circuit for noise analysis

on the differential output noise. However, for large input swings, one of the input transistors will be almost completely shut off, and the majority of the tail noise current will flow through the other input transistor. As a result, the full noise of the tail degeneration resistor will show up in the overall differential output noise.

The second component in the proposed amplifier that affects its noise performance is the set of capacitors, as shown in Fig. 5.19, used for setting the drive strength of the signal towards the tail degeneration. Firstly, due to the presence of C_{att} , the input signal voltage at the gate of the input transistor will be attenuated by a ratio of $C_s / (C_s + C_b || C_{att})$, where C_b is the coupling capacitance and C_s is the input sampling capacitance. As this signal loss will degrade the amplifier SNR, it is, hence, important to keep C_b and C_{att} small as compared to C_s .

Secondly, since the gates of the input and tail transistors are all reset to bias voltages before amplification, these capacitors experience a sampling moment at the end of the reset phase. Hence, as the amplifier is released from the reset phase, the bias voltages stored on these three capacitors will also be corrupted by the thermal noise power defined by kT/C . This is an approximation since the noise bandwidth of the network around each capacitor is not only defined by the respective capacitor, but also by the other two capacitors.

During the amplification phase, these three capacitors are connected together. As a result, the total thermal noise power at the gate of the input transistor can be roughly expressed as $kT / (C_s + C_b || C_{att})$. Consequently, the amplifier SNR gets

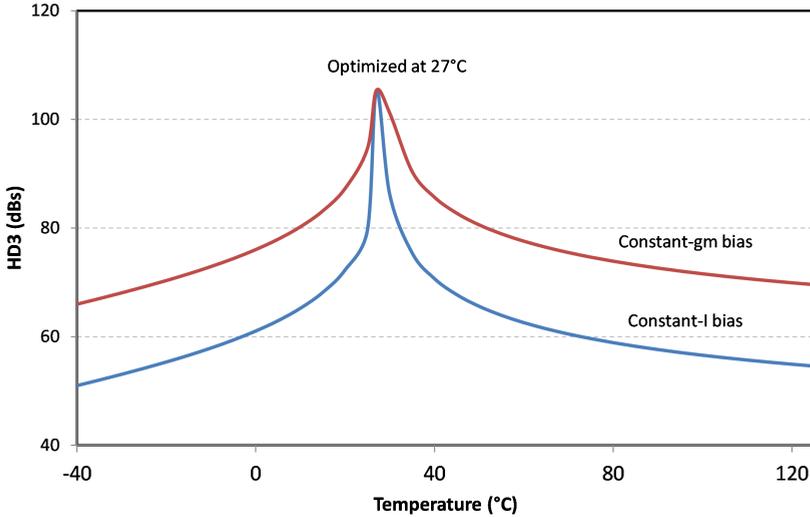


Figure 5.20: HD3 vs temperature for different bias schemes

lowered by $\sqrt{C_s/(C_s + C_b || C_{att})}$, which again mandates the use of small C_b and C_{att} capacitors.

5.5.3. TEMPERATURE DEPENDENCE

As seen from (11), the condition for HD3 cancellation is a function of temperature. This results in the HD3 varying strongly with temperature, especially when biased with a constant-current biasing scheme. There are multiple ways to alleviate this temperature dependence of the HD3 cancellation. One of them is by using a constant- g_m biasing scheme [19]. With the help of a constant- g_m biasing circuit, the input transistor can be biased at a g_m that is only determined by a resistor. If the temperature coefficient of the resistor in the biasing circuit can be matched to that used in the linearized amplifier, then the sensitivity of the linearization technique towards temperature variations can be significantly reduced. This is reflected in Fig. 5.20, which compares the variation in HD3 over temperature for the two different biasing schemes for the proposed amplifier with a poly resistor as CMD degeneration.

Another way to make the HD3 cancellation more stable over a wide temperature range is to use a background calibration loop. This calibration loop needs to have a good resolution and should be able to converge in a relatively short amount of time. Fast convergence ensures that the calibration loop can accurately track any

degradation in the HD3 due to micro-variations in temperature and adjust the tail degeneration accordingly.

5.5.4. COMMON-MODE REJECTION

As the input-driven tail degeneration transistors are effectively acting as an impedance with a relatively low magnitude, they do not offer much common-mode rejection. Since, for common-mode signals, the input g_m does get degenerated by the tail resistance, the common-mode rejection ratio (CMRR) will be equal to the degeneration ratio, $1 + g_m R_{deg}$. As the R_{deg} around the HD3 cancellation point is equal to $1/2ng_m$, the CMRR will be roughly 1.35 or 2.6dB. Although better than a differential common-source amplifier, a CMRR of 2.6dB will not be sufficient to ensure robustness against CM signals. As a result, other techniques will have to be employed to improve the CM rejection of the proposed amplifier topology, which will be discussed in Chapter 6.

5.6. POWER EFFICIENCY

Table 5.1 presents a qualitative comparison of the proposed amplifier with other recent residue amplifier topologies from literature. Both closed-loop [7, 20–22] and open-loop G_m - R residue based amplifiers [1, 23] rely on RC -based exponential settling, which improves the gain accuracy but at the cost of power dissipation.

The other topologies – zero-crossing based (ZCB) [24–26], ring amplifiers [27, 28] and dynamic amplifiers [29–31] – drive the output load with slew-based charging which maximizes their power efficiency. Zero-crossing based amplifiers typically consist of a current-source output-stage driving the load capacitor and a comparator that monitors the virtual ground nodes of the amplifier. Ring amplification is another interesting idea that features a (minimum) three-stage amplifier in negative-feedback. In contrast to the “classic” feedback approach, the ring amplifier does not employ any compensation, and exploits the resulting instability to achieve a fast operation. In order to ensure that the amplifier achieves the desired gain, a “dead-zone” is created in the third-stage, which switches OFF the transistors in this stage once the output voltage reaches its desired value and stops the ring amplifier from oscillating. Since both ZCB and ring amplifiers are used in negative-feedback, they have reported good linearity and gain accuracy while achieving much better power efficiency than the exponentially settling amplifiers. However, both these approaches require extra circuitry to drive the “core” amplifier. The ZCB approach relies on an

	Settling	Achievable Linearity	Power Overhead
Closed-loop amplifier [7,20-22]	Exponential	High	High. High DC loop gain and feedback network required to achieve accurate closed-loop gain. Inverter-based amplifier [7,21-22] simplify overhead with help of calibration
Open-loop Gm-R amplifier [1,23]	Exponential	Low, needs calibration	Medium. Requires moderate DC gain
Zero-Crossing Based [24-26]	Slew-based	Moderate, requires offset-compensation	Medium. Requires an accurate comparator
Ring Amplifier [27-28]	Slew-based	High	Requires 2 stages of inverters to drive the output stage and dead-zone control circuitry
Dynamic Amplifier [29-31]	Slew-based	Low. [29-31] resolve large number of bits before amplification	Negligible
Proposed Integrating Amplifier	Slew-based	High after linearization	Low. 1dB SNR drop due to input loss from driving tail current

Table 5.1: Comparison of recently published amplifier topologies

accurate comparator to achieve the desired gain, while the ring amplifier requires at least three gain-stages along with circuitry to control the dead-zone of the output stage. This circuit overhead essentially limits the power efficiency of these amplifier topologies.

Dynamic amplifiers [29–31] are inverter-based open-loop amplifiers which, due to their simplicity and slew-based settling, exhibit excellent power efficiency. However,

due to the lack of negative-feedback, their gain accuracy and distortion suffers. As a consequence, their applications have been mostly limited to pipelined-SAR ADCs, which resolve a large number of bits before residue amplification, hence relaxing their accuracy requirements.

The proposed linearization technique allows a simple differential common-source amplifier topology to achieve a linear gain for a wide input-range without relying on settling accuracy and negative feedback. As the amplifier is used in an open-loop configuration, it does not require a high DC gain, thus entailing a low circuit overhead. And, as the addition of tail degeneration has a very minor impact on the effective transconductance and noise, and consumes a very small voltage headroom ($< V_T/2$), the increase in power overhead is minimal. Since the linearization technique requires the transistors to be biased in the deep weak-inversion region, the amplifier power efficiency further benefits from an excellent g_m/I_d ratio. Although the transit frequency, f_t , lowers as the transistor is pushed deeper into the weak-inversion region, the transistor speeds are still sufficiently high in advanced technology nodes. All these factors help the proposed amplifier topology to potentially have one of the highest power efficiencies amongst contemporary residue amplifier topologies suitable for high-speed operation.

5.7. CONCLUSION

As fully digital error correction can entail significant complexity and power for high-speed and high-resolution ADCs, a more efficient calibration loop can be built by performing the error correction in the analog domain. This chapter presented some analog domain methods for the correction of residue amplifier gain and distortion. Open-loop integrating amplifiers were considered as an amplifier topology with a simple analog gain-error correction mechanism. Due to their extremely low settling, the gain of the integrating amplifier is almost linearly proportional to the bias current. This makes the tuning of their gain very straightforward.

For amplifier distortion correction, a new topology was proposed that is based on the cancellation of the expanding third-order distortion component in the input g_m of the amplifier with a compressing distortion component generated due to resistive degeneration. By picking a specific degeneration resistance value, the third-order distortion in the input g_m of the amplifier can be perfectly linearized. The degeneration's impact on the amplifier input transconductance can be almost completely negated by using the degeneration in a common-mode fashion, while still improving

the g_m variation from 50% in a differential pair to $<1\%$. The linearization principle was further verified through measurements on a strip-board amplifier version built with discrete BJT components that showed an HD3 of -95dB for an input swing of 50mV_{ppd} .

The distortion cancellation was extended to 2nd and 5th order distortion components by using a pair of input-driven weak inversion MOS transistors as the common-mode tail degeneration. By tuning the bias voltage and the attenuation of the input signal driving the degeneration network, the proposed amplifier topology was able to achieve a THD of $<-90\text{dB}$ for an input swing of $300\text{mV}_{\text{ppd}}$ and a gain of 4x. By improving the amplifier THD by $>40\text{dB}$ without adding excessive circuit or power overhead, the proposed CMD linearization principle promises to be a very power-efficient method for analog distortion correction that is suitable for residue amplifiers when combined with a digital calibration loop.

REFERENCES

- [1] B. Murmann and B. E. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, p. 2040 – 2050, Dec. 2003.
- [2] H. van de Vel, B. Butler, H. van der Ploeg, M. Vertregt, G. Geelen, and E. Paulus, "A 1.2-V 250-mW 14-b 100-MS/s digitally calibrated pipeline ADC in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, pp. 1047–1056, Apr. 2009.
- [3] A. Panigada and I. Galton, "A 130 mw 100 MS/s pipelined ADC with 69dB SNDR enabled by digital harmonic distortion correction," *IEEE J. Solid-State Circuits*, vol. 44, p. 3314 – 3328, Dec. 2009.
- [4] M. Lee, W. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1591–1599, Nov. 2000.
- [5] L. Jin, H. Xing, D. Chen, and R. Gieger, "A self-calibrated bandgap voltage reference with 0.5ppm/°C temperature coefficient," in *IEEE ISCAS*, pp. 265–268, 2006.
- [6] C. Huang, C.-Y. Wang, and J.-T. Wu, "A CMOS 6-bit 16-GS/s time-interleaved ADC using digital background calibration techniques," *IEEE J. Solid-State Circuits*, vol. 46, pp. 848–858, Apr. 2011.
- [7] M. Akter, R. Sehgal, F. van der Goes, and K. Bult, "A 66-dB SNDR pipelined split-ADC in 40-nm CMOS using a class-AB residue amplifier," *IEEE J. Solid-State Circuits*, vol. 53, pp. 2939–2950, Oct. 2018.
- [8] B.-N. Fang and J.-T. Wu, "A 10-bit 300-MS/s pipelined ADC with digital calibration and digital bias generation," *IEEE J. Solid-State Circuits*, vol. 48, pp. 670–683, Mar. 2013.
- [9] C.-J. Tseng and et al., "A 10-b 320-MS/s stage-gain-error self-calibration pipeline ADC," *IEEE J. Solid-State Circuits*, vol. 47, pp. 1334–1343, Jun. 2012.
- [10] Y. Ding and R. Harjani, "An 18-dBm IIP3 LNA in 0.35 μ m CMOS," in *ISSCC Dig. of Tech. Papers*, pp. 162–163, 2001.
- [11] S. Lou and H. Luong, "A linearization technique for RF receiver front-end using second-order intermodulation injection," *IEEE J. Solid-State Circuits*, vol. 43, pp. 2404–2412, Nov. 2008.
- [12] T. Kim and B. Kim, "A 13-dB IIP3 improved low-power CMOS RF programmable gain amplifier using differential circuit transconductance linearization for various terrestrial mobile D-TV applications," *IEEE J. Solid-State Circuits*, vol. 41, pp. 945–953, Apr. 2006.
- [13] D. Kim, B. Kim, and S. Nam, "A transconductor and tunable gm-C high-pass filter linearization technique using feedforward gm3 cancelling," *IEEE Trans. Circuits Syst.-II*, vol. 62, pp. 1058–1062, Nov. 2015.

- [14] L. Yu, M. Miyahara, and A. Matsuzawa, "A 9-bit 1.8 GS/s 44mW pipelined ADC using linearized open-loop amplifiers," *IEEE J. Solid-State Circuits*, vol. 51, pp. 2210–2221, Oct. 2016.
- [15] L. Yu, M. Miyahara, and A. Matsuzawa, "A 9-bit 500MS/s 6.0mW dynamic pipelined ADC using time-domain linearized dynamic amplifiers," in *Proc. IEEE ASSCC*, pp. 65–68, 2016.
- [16] K. Bult and H. Wallinga, "A CMOS four-quadrant analog multiplier," *IEEE J. Solid-State Circuits*, vol. 21, pp. 430–435, Jun. 1986.
- [17] D. Pederson and K. Mayaram, *Analog Integrated Circuits for Communication: Principle, Simulation and Design*. Norwell, MA: Kluwer, 1991.
- [18] W. Sansen, "Distortion in elementary transistor circuits," *IEEE Trans. Circuits Syst.-II*, vol. 46, pp. 315–325, Mar. 1999.
- [19] J. Steininger, "A transconductor and tunable gm-C high-pass filter linearization technique using feedforward gm3 cancelling," *IEEE Circuits and Devices Mag.*, pp. 26–31, May. 1990.
- [20] R. Sehgal, F. van der Goes, and K. Bult, "A 12b 53mW 195MS/s pipelined ADC with 82dB SFDR using split-ADC calibration," *IEEE J. Solid-State Circuits*, vol. 50, pp. 1592–1603, Jul. 2015.
- [21] Y. Chae and G. Han, "Low voltage, low power, inverter-based switched-capacitor delta-sigma modulator," *IEEE J. Solid-State Circuits*, vol. 44, pp. 458–472, Feb. 2009.
- [22] J. Kim and B. Murmann, "A 12-b, 30-MS/s, 2.95-mW pipelined ADC using single-stage class-AB amplifiers and deterministic background calibration," *IEEE J. Solid-State Circuits*, vol. 47, pp. 2141–2151, Sep. 2012.
- [23] E. Iroaga and B. Murmann, "A 12-bit 75MS/s pipelined ADC using incomplete settling," *IEEE J. Solid-State Circuits*, vol. 42, pp. 630–637, Apr. 2007.
- [24] T. Sepke and et al., "Comparator-based switched-capacitor circuits for scaled CMOS technologies," in *ISSCC Dig. of Tech. Papers*, pp. 220–221, 2006.
- [25] L. Brooks and H. S. Lee, "A zero-crossing-based 8b 200MS/s pipelined ADC," in *ISSCC Dig. of Tech. Papers*, pp. 460–461, 2007.
- [26] S. K. Shin and et al., "A 12bit 200MS/s zero-crossing-based pipelined ADC with early sub-ADC decision and output residue background calibration," *IEEE J. Solid-State Circuits*, vol. 49, pp. 1366–1382, June 2014.
- [27] B. Hershberg and et al., "Ring amplifiers for switched capacitor circuits," *IEEE J. Solid-State Circuits*, vol. 49, p. 2928–2942, Dec. 2012.
- [28] Y. Lim and M. Flynn, "A 100MS/s, 10.5bit, 2.46mW comparator-less pipelined ADC using self-biased ring amplifiers," *IEEE J. Solid-State Circuits*, vol. 50, pp. 2331–2341, Oct. 2015.

-
- [29] B. Verbruggen and et al., “A 2.6mW 6 bit 2.2 GS/s fully dynamic pipelined ADC in 40nm digital CMOS,” *IEEE J. Solid-State Circuits*, vol. 45, pp. 2080–2090, Oct. 2010.
- [30] B. Verbruggen, M. Iriguchi, and J. Cranineckx, “A 1.7mW 11b 250 MS/s 2-times interleaved fully dynamic pipelined SAR ADC in 40nm digital CMOS,” *IEEE J. Solid-State Circuits*, vol. 47, pp. 2880–2887, Dec. 2012.
- [31] F. van der Goes and et al., “A 1.5mW 68dB SNDR 80MS/s 2x interleaved pipelined SAR ADC in 28nm CMOS,” *IEEE J. Solid-State Circuits*, vol. 49, pp. 2835–2845, Dec. 2014.

6

A Split-ADC with a Linearized Integrating Amplifier¹

In the previous chapter, an integrating residue amplifier with an analog linearization scheme was presented. This chapter discusses the prototype pipelined ADC that was built around the proposed amplifier to test the efficacy of the linearization principle. In order to achieve a stable performance over a wide PVT range, the linearization scheme needs to be driven by an error detection loop. This calibration loop should preferably be able to run in full background mode with fast convergence without introducing significant overhead in the ADC. In Chapter 4, the split-ADC calibration approach was discussed in detail and, owing to its deterministic nature, was demonstrated to be one of the fastest background calibration techniques reported in literature. For the prototype ADC, a modified split-ADC calibration architecture is presented which reduces the mismatch between the two half-ADCs. The proposed split-ADC architecture is used to detect gain error and harmonic distortion arising from the residue amplifiers and accordingly tune the

¹This chapter is derived from publication: R. Sehgal, F. van der Goes and K. Bult, "A 13-mW 64-dB SNDR 280MS/s pipelined ADC using linearized integrating amplifiers" *IEEE J. Solid-State Circuits*, vol. 53, no. 7, pp. 1878–1888, Jul. 2018.

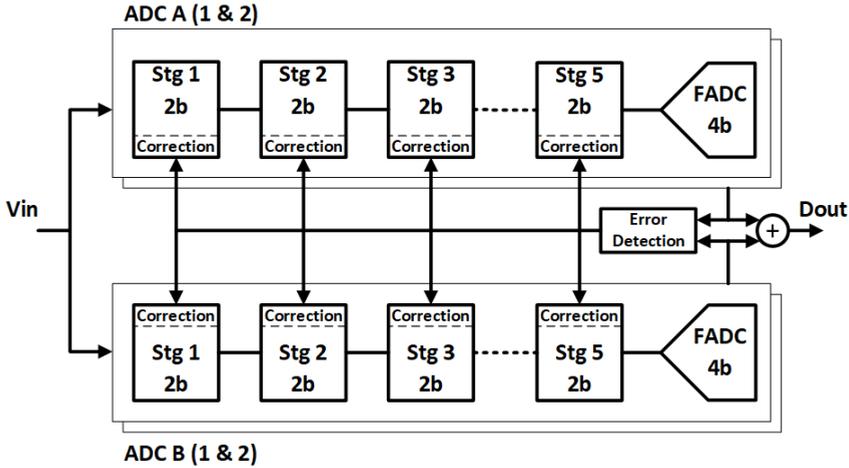


Figure 6.1: Prototype ADC architecture

corresponding knobs. Finally, circuit details of a silicon implementation of the prototype ADC are presented and the improvements achieved in the ADC linearity due to the linearization scheme are verified with the help of measurement results.

6.1. ADC ARCHITECTURE

In order to test the efficiency of the proposed integrating amplifier, it was incorporated in a two-lane 12-bit 280MS/s SHA-less pipelined ADC. As shown in Figure 6.1, the pipelined ADC consists of five 2-bit stages followed by a 4-bit fine ADC. Two additional bits were resolved in the fine ADC to reduce the impact of quantization noise to 14-bit level and improve the calibration accuracy. A resolution of 3-bit per stage was chosen with 1-bit overrange, as a compromise between ADC power efficiency and ease of design. Stages 2-5 are scaled by a factor of two with respect to stage 1 to save power. The residue amplifier gain and nonlinearity errors are detected in the digital domain with the help of the split-ADC calibration technique, while the correction of these errors is performed in the analog domain by tuning the bias and linearization parameters, respectively, of the stage residue amplifier. All the pipeline stages are calibrated for amplifier gain and nonlinearity errors. The residue amplifiers in the first two stages utilize all three linearization parameters (for cancelling HD3, HD5 and HD2). Due to lower accuracy requirements, stages 3-5 are implemented with only the tail bias (HD3) and offset (HD2) parameters

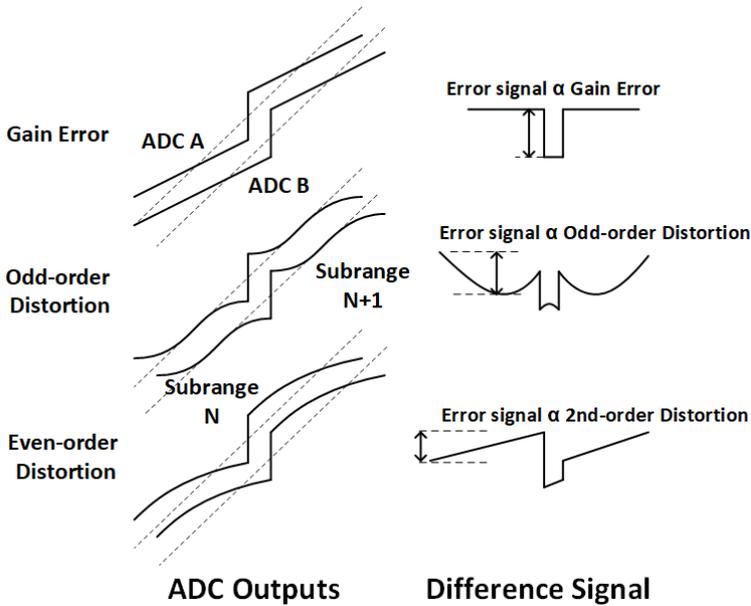


Figure 6.2: Detection of gain and distortion errors through Split-ADC technique

without the coupling cap between input and tail nodes.

6.1.1. CALIBRATION ARCHITECTURE

As discussed in Chapter 4, the split-ADC technique [1–3] involves splitting an ADC into two identical halves. These two half-ADCs digitize the same input sample but are forced to take different trajectories by injecting an offset between the references of the two half-ADCs. The outputs of the two half-ADCs can be averaged to obtain the overall offset-free digital output, while their difference can be used to detect any non-idealities in the half-ADCs. Each of the non-idealities – gain error, even-order, and odd-order distortion – have a unique effect on the difference signal, as seen in Figure 6.2, and that effect can be measured to estimate the magnitude of the respective non-ideality present in the residue amplifier.

Once the digitized outputs are generated, they are imported into MATLAB where the difference signal is generated and processed. Similar to the calibration algorithm presented in chapter 4, this difference signal is driven towards a straight line by tuning the ADC stage gains and linearization parameters. The gain errors are corrected in two steps - coarse manual adjustments of the on-chip programmable bias and off-chip automatic fine-tuning of the digital gain in MATLAB. In a future

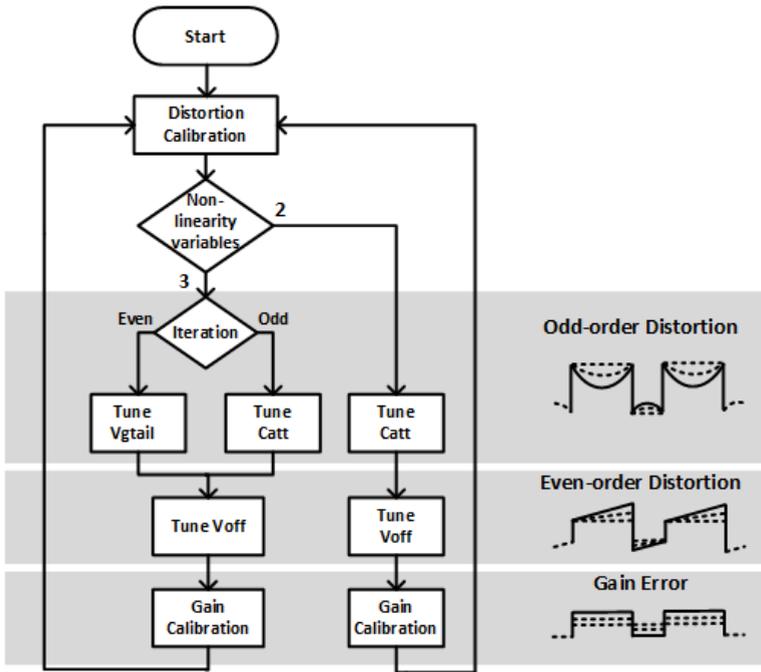


Figure 6.3: ADC stage calibration flowchart

prototype, the coarse tuning of the on-chip bias can be substituted by an on-chip constant gm-biasing circuit or a bias loop consisting of a replica amplifier, which makes the amplifier gain more immune to PVT variations. The fine adjustment of digital gain can also be replaced completely by a programmable amplifier bias with higher resolution.

Figure 6.3 illustrates the calibration procedure followed for stages 1-2. To test the efficacy of the linearization parameters, the nonlinearity calibration algorithm was designed with two different flows that swept (a) all three parameters - capacitive attenuation, C_{att} , tail bias, V_{gtail} , and tail offset, V_{os} , or (b) only C_{att} and V_{os} . While the former provides the desired control over the tail degeneration, the latter case was used as an additional test-mode. By sweeping all three parameters, the tail transistors can be tuned sufficiently to better approximate the optimum tail degeneration characteristic, leading to a much better linearization, as discussed in the previous chapter. Since both C_{att} and V_{gtail} affect the odd-order distortion, they are adjusted in separate iterations. The non-orthogonalities between nonlinearity and gain calibration are resolved by running them over multiple iterations as discussed in Chapter 4.

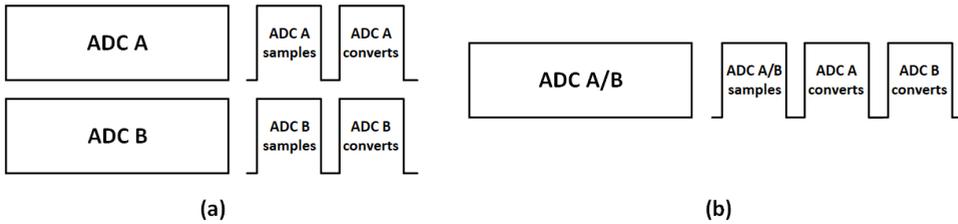


Figure 6.4: (a) Conventional Split-ADC versus the (b) “Split-over-time”-ADC concept

The artifacts in the difference signal shown in Figure 6.2 correspond to the difference of two completely identical half-ADCs. However, in reality, there will always be a certain mismatch between these two ADCs. Since this design utilizes open-loop residue amplifiers, this mismatch will be even more pronounced. There are several techniques that deal with the gain mismatch between the two half-ADCs [1, 2] with relatively low additional complexity. However, the mismatch between the distortion components in the residue amplifiers of the two half-ADCs introduces much more significant issues in the ADC calibration. For instance, the mismatch between the 3rd-order distortions in the two half-ADCs also results in an even-order distortion, which also produces a linear slope in the difference signal. Hence, this mismatch in the half-ADCs would directly interfere with the detection of even-order distortion coming from the amplifier. This makes it imperative to minimize the mismatch between the two half-ADCs.

6.1.2. "SPLIT-OVER-TIME"-ADC

The mismatch between the two half-ADCs can be minimized by using the same ADC for two consecutive conversions of the same input sample. The offset between the two half-ADCs can now be added over two separate clock phases. This “split-over-time”-ADC architecture is conceptually shown in Figure 6.4 and is similar to the perturbation-based calibration method presented in [4]. As compared to the conventional split-ADC approach, where the ADC is split physically into two separate half-ADCs, the “split-over-time” architecture performs the ADC split temporally. The main advantage of this approach is that it allows the two half-ADC operations to share the same residue amplifier, helping in minimizing the mismatch in the gain and distortion errors between the two half-ADCs. And while the “split-over-time” approach lowers the effective ADC sampling speed, it offsets that through area savings by sharing the major ADC sub-blocks. Hence, when compared to the conventional split-ADC technique, the “split-over-time” method achieves roughly

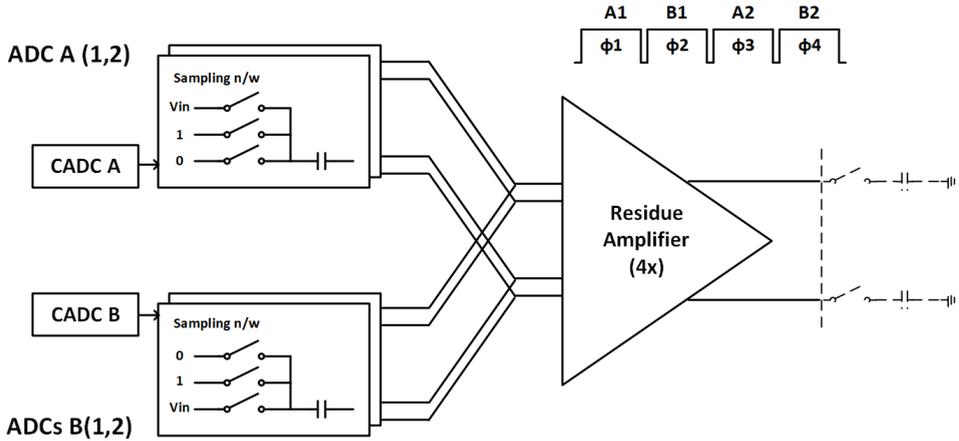


Figure 6.5: Split-ADC MDAC architecture with 2x interleaving

the same sampling speed, for a fixed ADC area and SNR.

Although it is possible to use entirely the same ADC to convert the input sample again after adding the calibration offset, there are certain challenges in sharing some of the ADC sub-blocks. As discussed in more detail later, charge sharing between the sampling capacitance and comparator/amplifier parasitics will create a significant gain mismatch between the two half-ADCs. Hence, in this design, while the residue amplifier is shared between the two half-ADCs, the rest of the ADC hardware – coarse ADCs (CADCs), capacitive DACs (CDACs) and digital logic – is physically split into two.

Figure 6.5 shows the 2x-time interleaved MDAC based on the “split-over-time”-ADC architecture. The capacitive networks belonging to the two half-ADCs, A and B, both connect to the same residue amplifier. As discussed in more detail in the next section, in order to better optimize ADC power and area, the residue amplifier is also shared by the ADC stages in the two lanes, 1 and 2. Since the gain/distortion mismatch between the half-ADCs is normally dominated by the device mismatch within the amplifier, by using the same residue amplifier, the gain and nonlinearity mismatch between the two half-ADCs is eliminated to a large extent. And, although the mismatch in the capacitive DACs will still create variations in the gain, they are expected to be small due to the much better inherent matching of metal capacitors. And by sharing the residue amplifier across the two ADC lanes, the calibration loop is also simplified by reducing the number of linearization parameters being driven by the calibration backend.

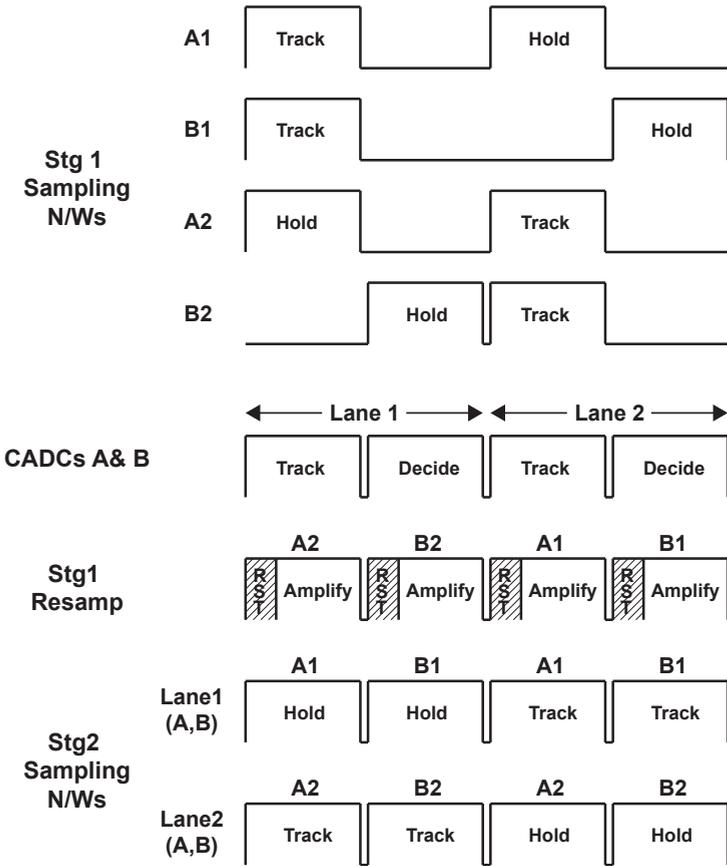


Figure 6.6: ADC timing scheme

6.2. ADC IMPLEMENTATION

In this section, the design specifics of the “split-over-time” pipelined ADC are discussed. An overview of the ADC timing and noise budget is presented followed by implementation details of some of the prominent ADC sub-blocks.

6.2.1. TIMING

The ADC timing scheme used to implement the offset addition over time is shown in Figure 6.6. In the “split-over-time” concept shown in Figure 6.4(b), the overall ADC is operated on a three-phase clock in which the residue amplifier is used for amplification in the latter two phases while being idle during the track phase. In order to avoid such idle periods for the residue amplifier, the ADC sampling speed is

doubled by using two time-interleaved lanes, 1 and 2, with a single residue amplifier shared between them. This also allows the ADC to be implemented with an even-phase (four-phase) clocking scheme instead of an odd-phase (three-phase) scheme, as shown in Figure 6.6. A reset pulse is used between the amplification periods for half-ADCs A and B to reduce the inter-symbol interference (ISI) between the two half-ADCs. Two interleaved sampling networks are used for each half-ADC stage while the CADCs, operating at twice the speed, are also shared between the two lanes. Sharing the CADCs and, more importantly, the residue amplifiers between the ADC lanes helps in significantly reducing the interleaving errors.

While the stage 1 sampling networks for half-ADCs A & B sample together, in order to share the residue amplifier, the residue of half-ADC B has to wait for one clock cycle. To ensure that this does not have any effect on the residue voltage, the bottom-plate switch of the sampling network was designed to have sufficiently low leakage. This asymmetrical clocking is necessary only in stage 1, as that extra clock cycle delay creates the required latency between split-ADCs A & B residues for the backend stages.

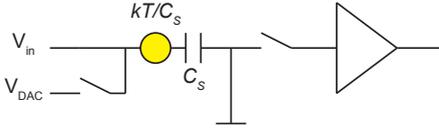
With the help of interleaving and the proposed timing scheme, the prototype ADC can operate at the same overall sampling frequency as the conventional split-ADC shown in chapter 4. Although using additional ADC lanes for interleaving often leads to an increase in area, the proposed timing scheme somewhat mitigates the area penalty by enabling the CADCs and the residue amplifiers to be shared between the two lanes. The four-phase timing scheme also allows the use of an entire clock-cycle for CADCs to make a decision. This significantly relaxes the speed requirements from the comparators. And as the proposed ADC timing scheme utilizes the residue amplifier in all four phases, it can maximize the amplifier power efficiency without shutting down the amplifier, avoiding the undesired transients in bias and common-mode voltages observed in dynamic amplifiers [5–7].

6.2.2. NOISE BUDGET

The ADC was designed to achieve an SNR of 11.5b for a full-scale input swing of $1.2V_{ppd}$. This translates into an input-referred differential rms noise voltage of $120\mu V$. Since the quantization noise is at 14b level, the overall input-referred noise is dominated by the thermal noise of the ADC sub-blocks.

Figure 6.7 describes shows the main noise sources of a pipeline ADC stage based on the proposed integrating amplifier. In this case, the input sampling bandwidth is assumed to be determined by the sampling switch and the stage MDAC sam-

Sampling Phase



Amplification Phase

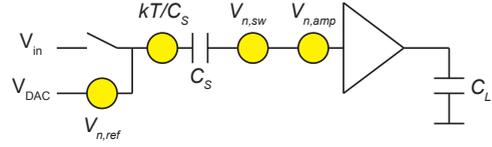


Figure 6.7: Noise analysis of the pipelined ADC stage

pling capacitance, C_S . The input-referred noise power of the pipeline stage can accordingly be calculated as -

$$V_{n,stage}^2 = \frac{kT}{C_S} + V_{n,amp}^2 + V_{n,sw}^2 + V_{n,ref}^2 \quad (6.1)$$

where $V_{n,amp}$ is the input-referred amplifier noise voltage, $V_{n,sw}$ is the noise voltage arising from the interleaving switch at the input of the residue amplifier and $V_{n,ref}$ is the noise contribution of the ADC reference also referred to the amplifier input. As the noise bandwidth of both $V_{n,amp}$ and $V_{n,sw}$ is determined by the amplifier noise bandwidth, assuming the amplifier to be an ideal integrator, the contribution from these two noise sources can be calculated as -

$$V_{n,amp}^2 + V_{n,sw}^2 = \left(\frac{kT}{g_m} + 4kTR_{sw} \right) \frac{1}{2T_{int}} \quad (6.2)$$

where T_{int} is the amplification (or integration) period, g_m is the input transconductance of the integrating amplifier and R_{sw} is the on-resistance of the interleaving switch. Considering that the gain of the integrating amplifier is a direct function of the integration period, it can be expressed as -

$$A = \frac{g_m}{C_L} T_{int} \quad (6.3)$$

where C_L is the load capacitance of the integrating amplifier. By using (6.2) and (6.3), the noise can be expressed as -

$$V_{n,amp}^2 + V_{n,sw}^2 = \frac{2kT}{AC_L} (1 + g_m R_{sw}) \quad (6.4)$$

It should be noted that the C_L in (6.4) does not include the sampling capacitance of the CADCs of that stage. If the residue amplifier has an inter-stage gain of 4x, then (6.4) can be rewritten as -

$$V_{n,amp}^2 + V_{n,sw}^2 = \frac{kT}{2C_L} (1 + g_m R_{sw}) \quad (6.5)$$

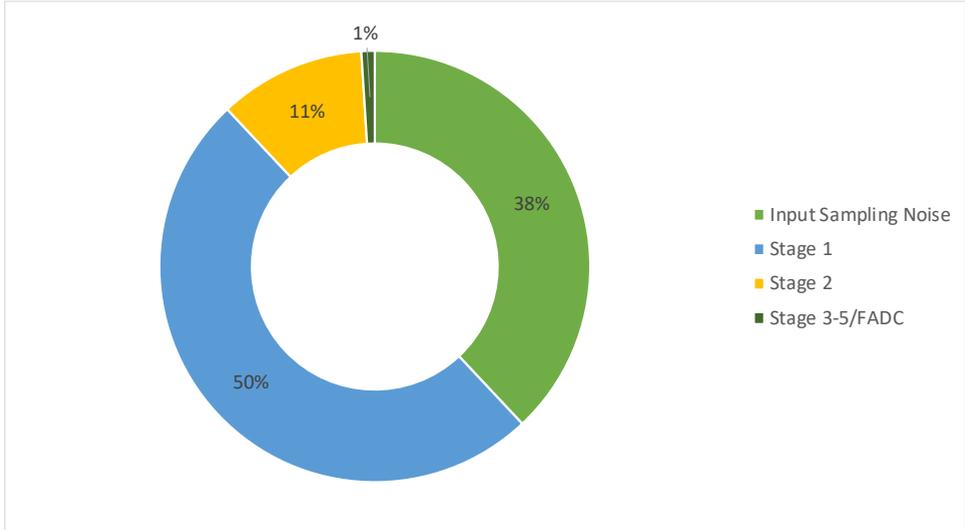


Figure 6.8: ADC noise budget

	MDAC Sampling Cap	Residue Amplifier Load Cap
Stage 1	1.5pF	700fF
Stage 2	510fF	320fF
Stage 3-5	140fF	320fF

Table 6.1: ADC stage sampling/load capacitances

Based on (6.1) and (6.5), the noise contribution from the sampling capacitance and the residue amplifier to the total ADC noise budget can be apportioned in every stage. The capacitance values used in each stage are listed in Table 6.1. It should be noted that in a pipelined ADC stage, the sampling capacitance of the MDAC effectively determines the SNR in the signal chain. Hence, the load capacitor, C_L , for a particular stage in (6.5) only includes the sampling capacitance of the MDAC and not that of the CADC in the subsequent stage. As the residue amplifier has to drive both the MDAC and CADC sampling capacitance, this leads to an overhead in power dissipation in pipelined ADC architectures that utilize flash-type CADCs.

Since the tracking switches at the output of the residue amplifier, shown in Figure 6.5, are effectively in series with a high-impedance current source, their noise contribution to the amplifier output noise is negligible. The noise of the tracking switches at the output of the residue amplifier is filtered by the amplifier noise bandwidth and attenuated by the amplifier gain. Hence, their contribution to the

stage noise budget is not very significant. The noise power on the load capacitor, C_L , is also affected by the initial condition imposed on it during the reset phase. Once sampled on the capacitor, this reset noise power will decay at an exponential rate depending on the time constant, τ_o , at that node and can be expressed as $(kT/C_L)e^{-2t/\tau_o}$ [8].

The distribution of the noise budget over the five ADC stages and the fine-ADC is illustrated in Figure 6.8. Since every pipelined stage introduces a gain of 4x, the noise contribution from the later stages diminishes and becomes relatively insignificant. Hence, the noise from the first stage residue amplifier along with the input driver/sampling network is allocated the majority of the noise budget, with the second and the third stages contributing most of the remaining noise power.

6.2.3. ADC SUB-BLOCKS

RESIDUE AMPLIFIER

Based on the proposed linearization technique, a complementary integration-based CMD amplifier was implemented with the differential input being applied to both NMOS and PMOS transistors, as shown in Figure 6.9. The sampling capacitance for PMOS and NMOS sides have been split in order to set their input common-modes independently. Cascode devices are used to boost the output impedance of the amplifier. The effect of the parasitic gate-drain capacitances of the input transistors on the charge stored on the input cap is cancelled with the help of cross-coupled neutrodyning caps [9] (not shown in the figure). While the tail degeneration bias, V_{gtail} , and the input attenuation cap, C_{att} , are implemented on both sides, the tail degeneration offset, V_{os} , is implemented only in the NMOS tail degeneration, as this is enough to correct for the entire amplifier imbalance.

Although the proposed amplifier has a tail degeneration, it does not provide a reliable common-mode (CM) control. Hence, two push-pull current sources, tied to the output and driven by a switched-cap CM feedback loop [10], as shown in Figure 6.10, are used to regulate the output CM level. These current sources are one-fourth the size of the input pair to limit their impact on total amplifier noise and power, and the cap, C_{cm} , is 7% of the total load cap for the first stage residue amplifier and 12% for the residue amplifiers in the remaining stages. The switches operating during the ϕ_1 phase in Figure 6.10 are sized carefully to reduce their loading on the amplifier output nodes and are switched during the reset phase to minimize the impact of non-linear charge-injection from the switches on the amplifier linearity.

Due to the susceptibility of the integration-based amplifier to memory effects,

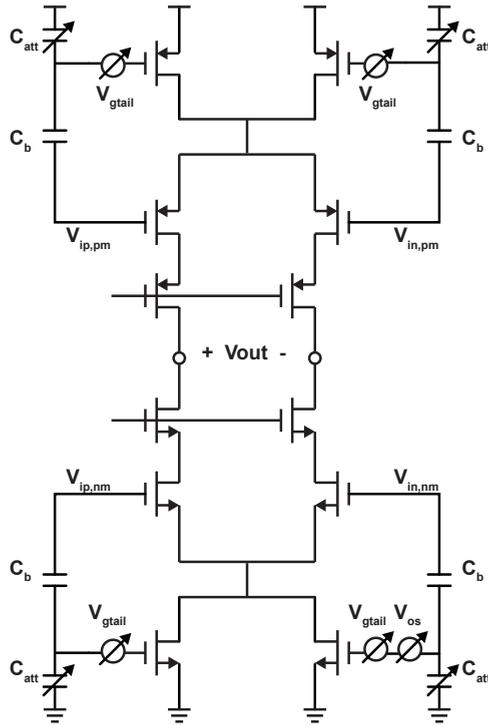


Figure 6.9: Proposed push-pull integrating amplifier with analog linearization

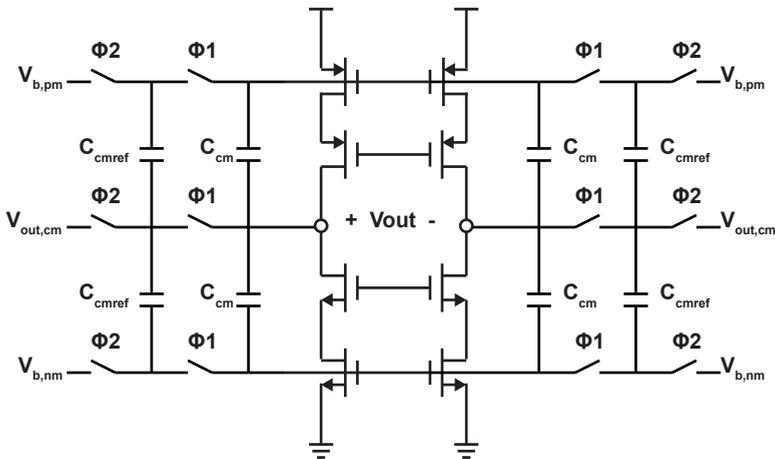


Figure 6.10: Switched-cap CMFB

all the capacitors and differential nodes inside the amplifier need to be reset using switches. This is performed during a reset pulse to purge all previous signal content

before the amplification phase begins. This reset is also used to establish the bias for all the differential pair and tail degeneration inputs. This allows the use of differential sampling [11] to mitigate the propagation of common-mode variations between stages.

Since the residue amplifiers are used in an open-loop configuration, there is no charge redistribution between the sampling and load capacitors. Hence, once the input signal has been sampled by the stage 1 sampling capacitor, that charge remains on that capacitor until it is reset. This means that, theoretically, the same sampling capacitor can be used again for the amplification phase of the other half-ADC, resulting in significant area savings. However, due to charge sharing between the sampling capacitor and the input parasitic capacitor of the amplifier, the input signal for the latter half-ADC will always experience an additional attenuation. This creates a deterministic gain mismatch between the half-ADCs that was ascertained through simulations to be as high as up to 5%. Hence, separate sampling capacitors were used for the two half-ADCs.

CALIBRATION DACS

In order to tune the tail degeneration bias and input attenuation, programmable current DACs and capacitive DACs were built-in along with each residue amplifier. Both the DACs are implemented with 7b resolution to have sufficient tuning resolution in the tail degeneration. A segmentation ratio of 4b thermometer + 3b binary was used to optimize the calibration logic without introducing excessive DNL in the DACs. The residue amplifiers in the first two stages utilize all three linearization parameters. Due to lower accuracy requirements, stages 3-5 are implemented with only the tail bias and offset parameters without the coupling cap between input and tail nodes. The coupling capacitance is only used in the first two stages and is 150fF and 70fF, respectively.

CAPACITIVE DAC

The capacitive DAC network used in the first-stage MDAC is shown in Figure 6.11. The sampling paths for the PMOS- and NMOS-side of the amplifier are entirely separated to allow independent biasing points. As the ADC was designed for a $1.2V_{ppd}$ swing with a common-mode of 0.5V (or half-supply), ideally, 0.8V and 0.2V are required as ADC references. However, in order to reduce the switch sizes, the MDACs and the CADCs are designed with 1V and ground as ADC reference voltages with reference attenuation capacitors used to reduce the effective reference voltage to 0.6V. These are part of the sampling capacitance, but unlike the 3b cap-DAC,

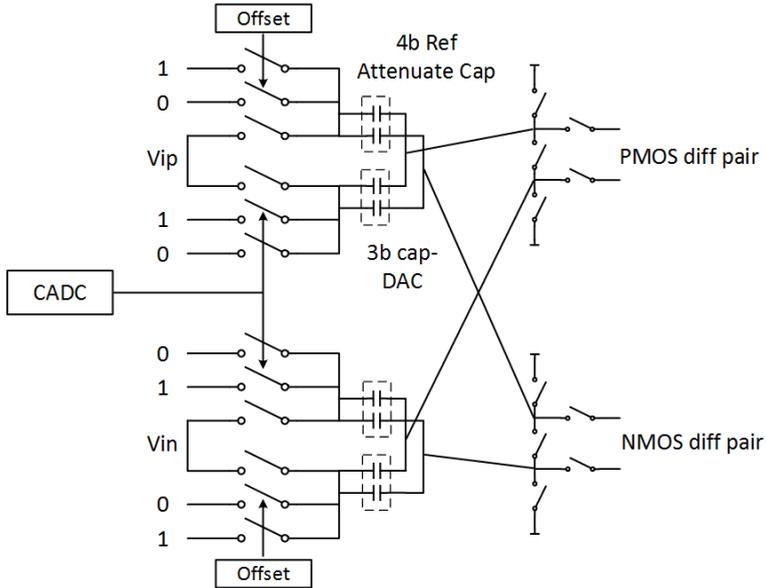


Figure 6.11: Stage 1 capacitive DAC

they are switched based on a static code. Typically, the attenuation capacitance is switched to the common-mode voltage by splitting it into two equal halves and connecting them to symmetrically to the positive and negative rails. To introduce the calibration offset required for split-ADC calibration, the reference attenuation capacitors are implemented as 4b thermometer cap-DACs, along with programmable connections to the two references. This offset cap-DAC is only included in the first pipeline stage as it is sufficient to introduce the calibration offset in only the first stage of the ADC.

Due to its high output impedance, the open-loop integrating amplifier has an inherently low bandwidth. This relaxes the bandwidth requirements of the sampling networks of pipeline stages 2-5 and the fine ADC. Hence, their sampling switches are no longer required to be low-ohmic and can be scaled down accordingly. This simplifies the switching and clock distribution in the sampling networks. As for the switch linearity requirements, they are most stringent in the first ADC stage and subsequently reduce by 2 bits after every stage. Hence, clock bootstrapping was only used in the first stage, while simple transmission gates were utilized for the remaining stages.

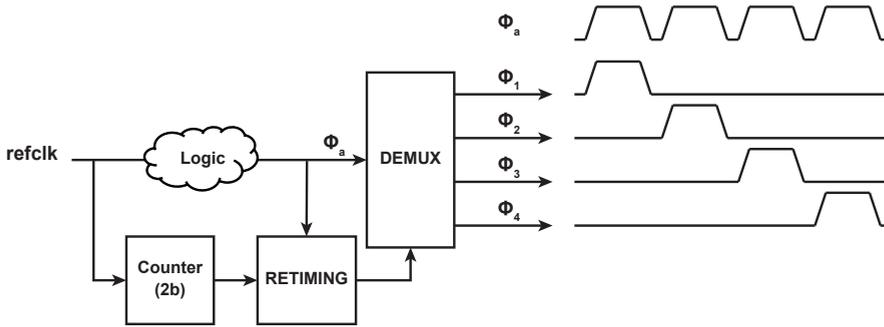


Figure 6.12: Four-phase clock generation for 2x time-interleaved “split-over-time”-ADC

CLOCK GENERATOR

As shown in Figure 6.6, the MDAC operates in a four-phase timing scheme, with the same stage residue amplifier being used in each of the four amplification phases, which correspond to the two half-ADCs across two lanes. Due to the integrating nature of the residue amplifier, any mismatch between these four phases will result in a gain mismatch either between the two half-ADCs or between the ADC lanes resulting in interleaving tones. Hence, a low mismatch clocking scheme was implemented to minimize the variations amongst the four clock phases. As illustrated in Figure 6.12, all the master clocks, running at four times the frequency of a single half-ADC, are generated from a reference clock signal with the help of programmable digital logic. Then, using a synchronized 2-bit counter output, these master clock signals are demultiplexed into 4 separate groups of clock phases. By keeping the majority of the clock generation logic common amongst the four clock paths, the 3σ Monte Carlo spread between the four amplification and track clock phases was limited to less than 0.8ps ($< 0.05\%$ of the clock period).

COARSE ADC

In the timing scheme illustrated in Figure 6.6, there is a latency of one extra clock cycle between the track phase and residue amplification. As a result, an entire clock cycle can be earmarked for the CADCs to make a decision. This eases the speed requirements of the comparators, allowing for their design to be optimized for low power. A simple strong-arm latch connected to the input through a differential pair, similar to the design shown in Chapter 4, was utilized in this ADC. As the speed requirements are relaxed, the extra NMOS switch, used previously to increase the latch speed, is not included to simplify the comparator design and lower their input-referred offset and power consumption.

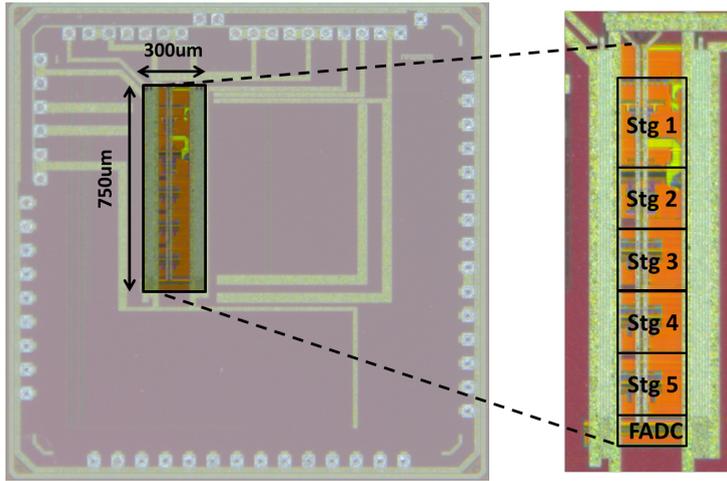


Figure 6.13: Chip micrograph

FINE ADC

A 4-bit flash architecture was used as the fine ADC (FADC). Instead of a resistive ladder, capacitive interpolation was used by splitting each comparator's sampling capacitance between the ADC references in a specific ratio to generate the references for the comparator array. Since there is no overrange present in the FADC, the comparator offset needs to be low to reduce its impact on the ADC DNL. This was accomplished by using autozeroed preamplifiers with input offset storage [12] to reduce the 3σ comparator offset from 42mV to 13mV.

6.3. MEASUREMENT RESULTS

The prototype 12-bit ADC was implemented in 28nm 1V CMOS process and occupied an overall area of 0.22mm^2 as shown in Figure 6.13. This includes the bias, clock tree, output digital interface, and an input digital test interface. The on-chip circuits for amplifier linearization parameters occupy around 2% of the total area. The reference voltages use off-chip decoupling caps and are provided to the ADC through separate pads. The ADC digital outputs are subsampled by the output digital interface and then imported into MATLAB for error detection. Based on the detected error signal, the calibration loop sets the configuration bits for the respective analog correction parameter using the input digital test interface. By running the calibration iteratively as described in Chapter 4, the overall ADC cali-

	Power Consumption
Clock Generation	6mW
ADC References	2mW
Bias Generation	1mW
Residue Amplifiers	0.4mW
Misc.	3.6mW
Total	13mW

Table 6.2: ADC power consumption

bration reaches convergence in less than 100K samples while running completely in the background.

The distribution of ADC power dissipation is shown in Table 6.2. Sampling at 280MS/s, the ADC dissipates an overall power of 13mW from a 1V supply. Out of this, 6mW is attributed to clocking circuits, 1mW for the reference bias generation and 2mW for the reference voltages. 4mW is drawn from the analog supply, which includes the stage amplifiers and their local bias, comparators, decoders, and other digital logic. Simulations indicate that the majority of that power is spent on logic, as the amplifiers are extremely low power and consume less than 0.4mW in total.

Since the amplifier nonlinearity in the prototype ADC is corrected in the analog domain, the error correction part of the calibration consumes negligible power. And due to the fast convergence speed of the split-ADC calibration, the error detection can be run at a lower clock speed to minimize its impact on the overall ADC power consumption.

Figure 6.14 shows the ADC output spectrum (after subsampling by a ratio of 129) before and after calibration for an input frequency close to 137.5MHz and an input swing of $1.1V_{ppd}$ (92% of ADC full scale). The remaining 8% headroom is reserved for the addition of calibration offset, implying an SNR reduction of 0.75dB. However, since the convergence speed of the split-ADC calibration is, to some extent, proportional to the size of the offset voltage, this calibration offset can be optimized and reduced once the calibration has converged.

As the first two pipeline stages were implemented with both V_{gtail} and C_{att} parameters, the nonlinearity calibration was performed in two steps to see the effectiveness of the two linearization parameters. After gain and nonlinearity calibration with C_{att} and V_{os} in the first two stages, the SNDR improved by more than 25dB. However, the overall ADC resolution is still limited by distortion, with an SFDR

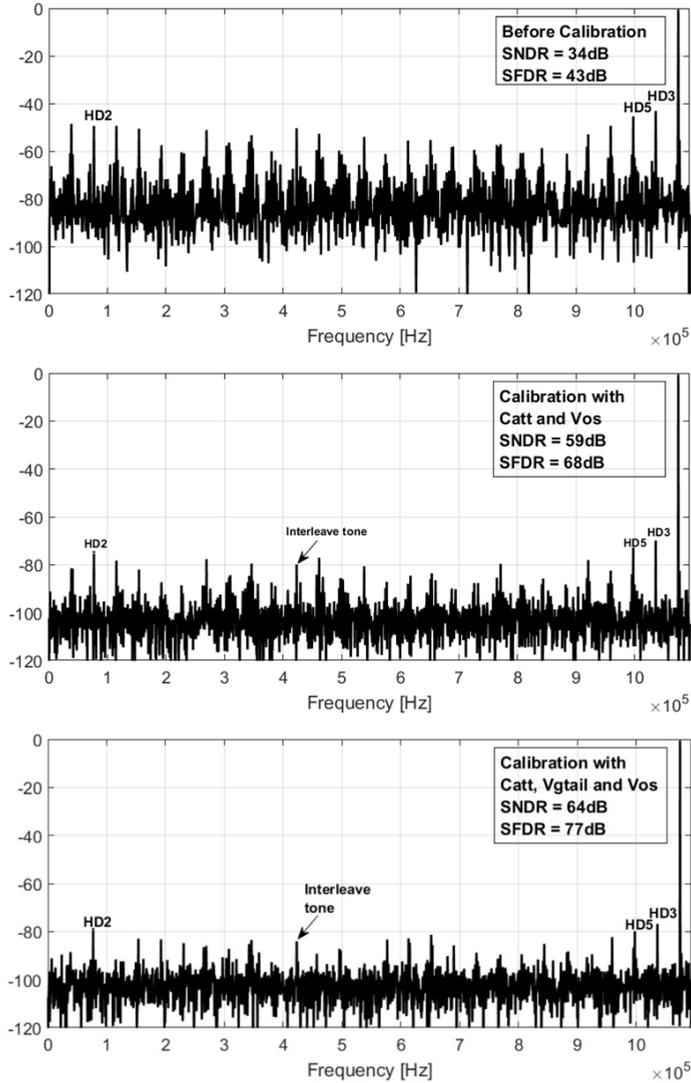


Figure 6.14: ADC spectrum before and after calibration

of 68dB. By using all three linearization parameters – V_{gtail} , C_{att} and V_{os} – all the harmonics are suppressed to about -80dB level, with HD2, 3 and 5 being the dominant tones. The two ADC lanes show good matching and the only significant interleaving tone appears to be due to lane gain mismatch. This could potentially be due to a mismatch between the sampling capacitance of the two lanes. With the help of calibration, the ADC shows an overall 77dB SFDR and 64dB SNDR.

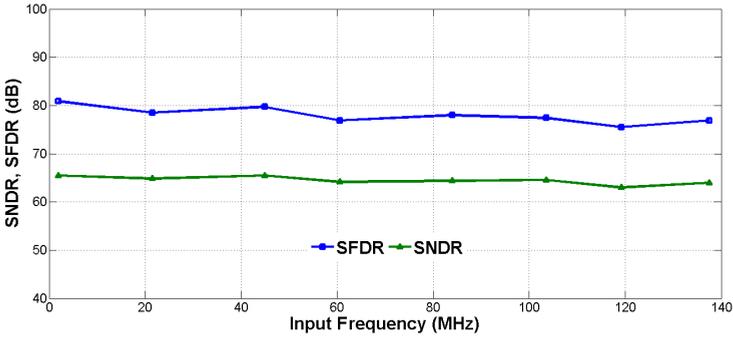


Figure 6.15: ADC SNDR and SFDR versus input signal frequency sweep

A similar performance is achieved by the ADC over four different die samples after calibration.

One of the non-idealities limiting the ADC linearity is the cap-DAC mismatch. Since the capacitive matching is normally much better than the amplifier distortion, the effect of the cap-DAC mismatch is visible only after calibration. Although the cap-DAC mismatch can also be calibrated using split-ADC calibration as shown in [13], this was not done in this work.

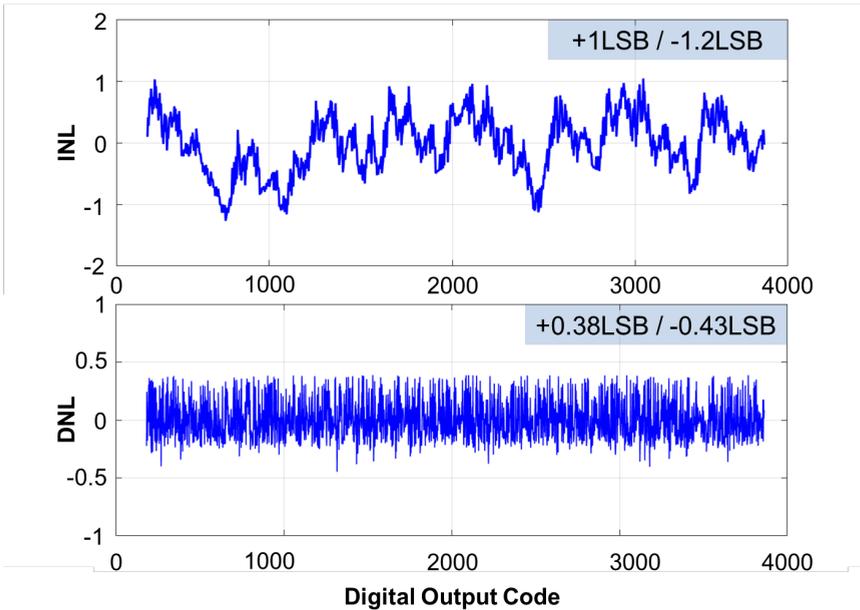


Figure 6.16: ADC INL and DNL after calibration

	Shin JSSC'14	Dolev VLSI '13	Oh VLSI'13	Lim JSSC'15	This work
Technology	55nm	65nm	130nm	65nm	28nm
Amplifier Topology	ZC- Based	Pulsed Bucket Brigade	Time- based charge pump	Ring am- plifier	Integrating CMD am- plifier
Supply(V)	1.1	1	1.3	1.2/0.75	1
ADC FS (V_{ppd})	2	-	2.4	-	1.2
F_S (MHz)	200	200	70	100	280
SNDR (dB) @Nyquist	63.2	57.6	62.6	56.6	64
SFDR(dB) @Nyquist	76	82*	81*	64.7	77
Power (mW)	30.7	11.5	6.38	2.5	13
Schreier FoM (dB)	158.3	157	161.5	159.6	164.3
Walden FoM (fJ/conv)	130	92.8	61.3	44.5	35.8

*at low input signal frequency

Table 6.3: ADC performance comparison

The ADC was tested at multiple input signal frequencies to check the effectiveness of amplifier linearization and exhibited a fairly stable performance of $>10.3b$ ENOB over the entire frequency range as shown in Figure 6.15. The ADC static performance is represented in the INL and DNL plots shown in Figure 6.16. At 12 bit level, the ADC exhibits a worst-case INL of $+1/-1.2$ LSBs and a DNL of $+0.38/-0.43$ LSBs.

Table 6.3 summarizes the performance of the prototype ADC. As evident from [14], the ADC power efficiency strongly depends on the chosen architecture. As the focus of this work is on amplifier performance and not on ADC performance per se, in order to conduct a fair evaluation of the proposed amplifier's performance, the prototype ADC is compared with other pipelined ADCs with similar architecture, sampling speed and resolution.

The prototype achieves a Walden FoM of $35.8fJ/conv$ and a Schreier FoM of $164.3dB$, the latter being a much more relevant measure for noise-limited ADCs.

As seen in Table 6.3, when compared to recent state-of-the-art pipelined ADCs with other power-efficient residue amplifier topologies [15–18], this work shows an improvement of at least 3dB in Schreier FoM, indicating a power efficiency advantage of at least 2x w.r.t to the other ADCs.

Although the above comparison of amplifier power efficiency is complicated by differences in technology, supply voltage and speed, one of the focal points of this prototype ADC was to verify the effectiveness of the proposed linearization technique. In that regard, when compared to other ADC designs with integration-based residue amplifiers [5–7], this work shows a similar or better SFDR for a much higher input signal swing and frequency with the help of the proposed linearization scheme.

6.4. CONCLUSION

This chapter discussed a prototype pipelined ADC employing an integration-based residue amplifier with the linearization technique proposed in Chapter 5. The amplifier gain error and non-linearity were detected in the background through the split-ADC calibration technique. In order to reduce the mismatch between the two half-ADCs, a modified split-ADC calibration architecture was implemented. By digitizing the input sample twice over two separate clock cycles, the calibration offset can be introduced over time and the same ADC hardware can be utilized for the two half-ADC operations. This “split-over-time” ADC allowed the residue amplifiers in the prototype ADC to be shared between the two half-ADCs. To further improve the power efficiency of the residue amplifiers, they are shared between two time-interleaved ADC lanes. Fabricated in 28-nm 1-V digital CMOS, the ADC achieved an ENOB of 10.34 bit at 280 MS/s with a power consumption of 13 mW, displaying an overall Schreier FoM of 164.3dB. The linearization scheme, along with the split-ADC calibration scheme operating in the background, enables the prototype ADC to achieve >77dB SFDR while displaying an improvement of at least 3dB in Schreier FoM over other comparable state-of-the-art pipeline ADC designs.

REFERENCES

- [1] J. Li and U.-K. Moon, "Background calibration techniques for multi-stage pipelined ADCs with digital redundancy," *IEEE Trans. Circuits Syst.-II*, vol. 50, pp. 531–538, Sep. 2003.
- [2] J. McNeill and et al., "'Split ADC' architecture for deterministic digital background calibration of a 16-bit 1-MS/s ADC," *IEEE J. Solid-State Circuits*, vol. 40, pp. 2437 – 2445, Dec. 2005.
- [3] R. Sehgal, F. van der Goes, and K. Bult, "A 12b 53mW 195MS/s pipelined ADC with 82dB SFDR using split-ADC calibration," *IEEE J. Solid-State Circuits*, vol. 50, pp. 1592–1603, Jul. 2015.
- [4] W. Liu, P. Huang, and Y. Chiu, "A 12-bit, 45-MS/s, 3-mW redundant successive-approximation-register analog-to-digital converter with digital calibration," *IEEE J. Solid-State Circuits*, vol. 46, p. 2661–2672, Nov. 2011.
- [5] B. Verbruggen, M. Iriguchi, and J. Craninckx, "A 1.7mW 11b 250 MS/s 2-times interleaved fully dynamic pipelined SAR ADC in 40nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 47, pp. 2880–2887, Dec. 2012.
- [6] B. Verbruggen, K. Deguchi, B. Malki, and J. Craninckx, "A 70 dB SNDR 200 MS/s 2.3 mW dynamic pipelined SAR ADC in 28 nm digital CMOS," pp. 1–2, 2014.
- [7] F. van der Goes and et al., "A 1.5mW 68dB SNDR 80MS/s 2x interleaved pipelined SAR ADC in 28nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, pp. 2835–2845, Dec. 2014.
- [8] T. Sepke and et al., "Noise analysis for comparator-based circuits," *IEEE Trans. Circuits Syst.-I*, vol. 56, pp. 541–553, Mar. 2009.
- [9] P. Gray and R. Meyer, *Analysis and Design of Analog Integrated Circuits*. Wiley, 1993.
- [10] O. Choksi and L. Carley, "Analysis of switched-capacitor common-mode feedback circuit," *IEEE Trans. Circuits Syst.-II*, vol. 50, pp. 906–917, Dec. 2003.
- [11] J. Kim and B. Murmann, "A 12-b, 30-MS/s, 2.95-mW pipelined ADC using single-stage class-AB amplifiers and deterministic background calibration," *IEEE J. Solid-State Circuits*, vol. 47, pp. 2141–2151, Sep. 2012.
- [12] B. Razavi and B. Wooley, "Design techniques for high-speed, high-resolution comparators," *IEEE J. Solid-State Circuits*, vol. 27, p. 1916–1926, Dec. 1992.
- [13] I. Ahmed and D. Johns, "An 11-bit 45 MS/s pipelined ADC with rapid calibration of DAC errors in a multi-bit pipeline stage," *IEEE J. Solid-State Circuits*, vol. 46, p. 1626–1637, Jul. 2008.
- [14] B. Murmann, "ADC Performance Survey 1997-2020." <http://web.stanford.edu/~murmman/adcsurvey.html>.

-
- [15] S. K. Shin and et al., “A 12bit 200MS/s zero-crossing-based pipelined ADC with early sub-ADC decision and output residue background calibration,” *IEEE J. Solid-State Circuits*, vol. 49, pp. 1366–1382, June 2014.
 - [16] N. Dolev, M. Kramer, and B. Murmann, “A 12-bit, 200-MS/s, 11.5-mW pipeline ADC using a pulsed bucket brigade front-end,” pp. 98–99, 2013.
 - [17] T. Oh and et al., “A 70MS/s 69.3dB SNDR 38.2fJ/conversion-step time-based pipelined ADC,” pp. 96–97, 2013.
 - [18] Y. Lim and M. Flynn, “A 100MS/s, 10.5bit, 2.46mW comparator-less pipelined ADC using self-biased ring amplifiers,” *IEEE J. Solid-State Circuits*, vol. 50, pp. 2331–2341, Oct. 2015.

7

Conclusions

Residue amplification plays a key role in determining the energy efficiency, area and performance of high-speed pipelined ADCs. At its core, a residue amplifier simply consists of four transistors that transfer a differential input voltage to a capacitive load with the desired gain. However, in order to ensure gain accuracy over PVT, the core amplifier has to be augmented with extra circuitry to achieve high settling accuracy, at the expense of area and power dissipation. In this dissertation, techniques to improve the power efficiency of residue amplifiers were investigated. The main objective of this research was to devise a high-resolution residue amplifier topology that relies on minimum circuit overhead and settling accuracy. This final chapter reviews the key ideas that were explored and implemented to achieve this objective, followed by some proposals for future exploration.

7.1. MAIN FINDINGS

The main contributions of this doctoral work are outlined as follows -

- The analysis of harmonic distortion in amplifiers has been so far limited to frequency-domain modeling under steady-state conditions. In Chapter 3, a mathematical analysis was developed to understand and describe the transient settling behavior of thermal noise and harmonic distortion in discrete-time amplifiers. By assessing the effect of incomplete settling accuracy on the

amplifier SNR and THD, an optimum settling point for a digitally-assisted residue amplifier was calculated.

- The split-ADC architecture [1, 2] has been previously demonstrated as a deterministic method for calibrating residue amplifier gain errors with fast convergence while operating completely in background. In Chapter 4, the split-ADC technique was extended to the calibration of harmonic distortion in residue amplifiers. With the help of the split-ADC technique, the amplifier gain and distortion errors of a prototype ADC were calibrated over multiple ADC stages through digital post-processing, resulting in an overall ADC SFDR of 82dB. The calibration allowed the power dissipation of the residue amplifiers to be reduced by 60% without any impact on the ADC resolution.
- To address the complexity and accuracy limitations of digital distortion correction, a new analog linearization scheme is proposed in Chapter 5. It exploits the fact that MOSFETs biased in weak inversion have an exponential V - I characteristic. It was shown that the distortion due to the expanding characteristic of the input transconductance of a differential common-source amplifier can be cancelled by using the compressive effect of resistive degeneration. By implementing the degeneration in a common-mode fashion, this distortion cancellation can be achieved without reducing the input transconductance. A simple stripboard implementation of a common-emitter amplifier with a common-mode degeneration (CMD) resistor was used to verify the linearization principle and exhibited an $HD_3 < -95\text{dB}$ for a 50mV_{ppd} input swing.
- Based on the CMD linearization principle, a new amplifier topology is proposed in Chapter 5 that allows the cancellation of HD_2 , HD_3 and HD_5 . By using input-driven transistors for common-mode degeneration, the input range of a differential CS amplifier was extended to $300\text{mV}_{\text{ppd}}$ for a THD of $< -90\text{dB}$. The degeneration network was implemented with tunable knobs that could be used to control the cancellation of the three harmonic distortion components.
- The conventional split-ADC architecture relies on splitting an ADC into two physical halves. In Chapter 6, a new split-ADC architecture is described that helps reduce the mismatch between the two half-ADCs. The “split-over-time” ADC architecture splits the ADC in time-domain by utilizing the same ADC to convert an input sample twice while introducing the calibration offset between

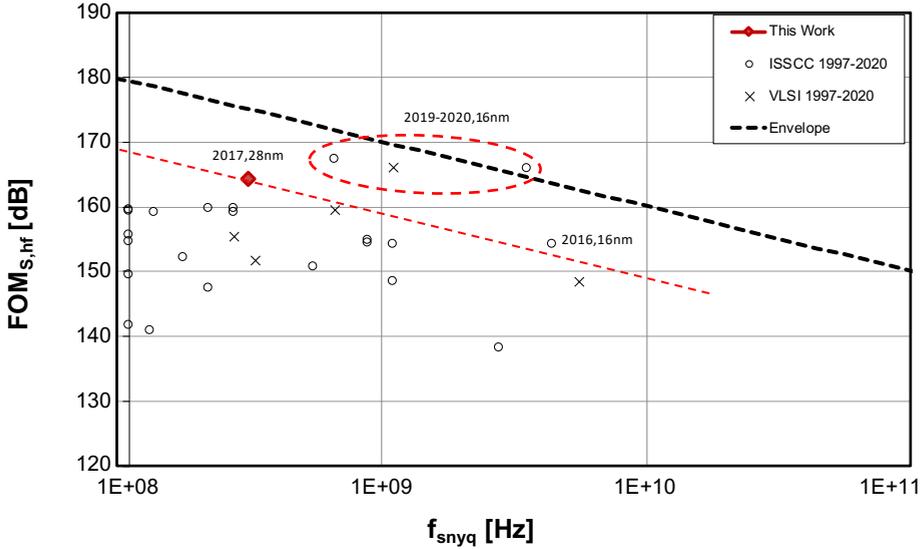


Figure 7.1: Schreier FoM versus ADC sampling frequency for state-of-the-art high-speed high resolution pipelined ADCs (ENOB > 9b and $f_s > 100\text{MS/s}$)

the two split-ADCs sequentially. This allows the two half-ADCs to share the residue amplifier, resulting in a significantly lower mismatch between the half-ADCs and mitigating the orthogonality between the calibration of even- and odd-order distortion components.

During the course of this research, two prototype split-ADCs were designed and measured to verify the ideas outlined above. The first split-ADC [3, 4] utilized a conventional opamp with high DC gain in negative feedback for residue amplification, along with a highly programmable bias to tune its settling accuracy. Fabricated in 40nm 1V digital CMOS, the ADC achieved an ENOB of 10.5b at a sampling frequency of 195MS/s, exhibiting a Schreier FoM of 157.4dB. The split-ADC calibration enabled the residue amplifiers to operate close to their optimum settling points shown to be between $2-3\tau$ for closed-loop amplifiers in Chapter 3) while reaching convergence in only 70K samples, making it one of the fastest background distortion calibration techniques reported in literature.

With a suitable candidate for distortion calibration identified, a second prototype ADC [5, 6] was built around the proposed open-loop CMD amplifier that promised a high linearity with an extremely low circuit overhead and settling accuracy ($< 0.5\tau$). Implemented with the “split-over-time” ADC architecture, the ADC was fabricated

in 28nm 1V CMOS, and demonstrated an ENOB of 10.34b at 280MS/s, displaying a Schreier FoM of 164.3dB. Compared to the previous prototype ADC with closed-loop residue amplifiers, the “split-over-time” ADC with open-loop CMD amplifier improved the Schreier FoM by 3.5x after adjusting for the impact of technology scaling [7]. Figure 7.1 shows a comparison of [5] with other state-of-the-art high-speed high-resolution pipelined ADCs with similar architecture (flash-type CADC) and speed & resolution (ENOB > 9b and $f_s > 100\text{MS/s}$). An envelope with a slope of -10dB/decade is included to capture the effect of sampling frequency on the Schreier FoM. It can be seen that the prototype ADC achieves one of the best FoMs amongst state-of-the-art designs¹ further highlighting the power efficiency of the proposed CMD amplifier.

The duration of this doctoral work presents an interesting opportunity to assess its significance and its adoption by the wider scientific community. Although the principle of analog linearization is widely used in broadband RF applications, it was, till a few years ago, rarely used in ADCs. Prior to the publication of [5] in 2017, the prevalent solution for distortion correction in residue amplifiers was through digital correction, with only 3 publications reporting analog linearization techniques. Since then, however, 7 amplifier designs [8–14] have been published in the last three years alone that feature some form of analog linearization, including capacitive degeneration [8] which relies on the same principle while using a capacitor as the degeneration impedance. This spurt in publications, along with the number of citations received by [5, 6] in a relatively short span of time (24 citations in 3 years) points to a growing interest in linearized open-loop dynamic amplifiers.

7.2. FUTURE SCOPE OF WORK

Due to limited design time, the residue amplifier presented in Chapter 6 was designed with a programmable constant-current bias. This resulted in a rather large gain variation (nearly $\pm 25\%$) over the entire PVT range. As the residue amplifier in the first ADC stage needed to have an accuracy of less than 0.1%, an entirely analog domain gain correction would have necessitated a 10-bit calibration DAC. To simplify the design, a combination of digital and analog correction was used for correcting the amplifier gain error. The calibration DACs were designed to cover the entire gain correction range and were used for coarse correction, while the digital

¹It should be noted that the four designs with a better FoM than [5] were all designed in 16nm CMOS, which is a FinFet process and cannot be easily compared with planar technologies.

stage gain in the encoder was finetuned to attain the desired gain accuracy. This two-step calibration can be simplified by reducing the gain variation range with the help of a constant- g_m or a replica-bias loop. This will relax the specifications for the calibration DACs, which can subsequently be designed with sufficient resolution to allow a completely analog-domain gain correction.

The key advantage of the “split-over time” architecture is that it allows sharing of sub-blocks between the two half-ADCs. One of the sub-blocks that was not shared in this work was the input sampling network. This was because the charge sharing between the input sampling cap and the residue amplifier input parasitics would cause additional attenuation in the residue voltage of the “latter” half-ADC, resulting in a significant gain mismatch between the two half-ADCs. However, some interesting advantages vis-à-vis noise arise by sharing the input sampling network. While the total area of the sampling network would remain roughly the same as the noise on the sampling cap would be correlated between the two half-ADCs and would not get suppressed through the averaging, the noise will be cancelled in the difference signal. As the input sampling noise is often one of the dominant noise sources in an ADC, the cancellation of this noise in the difference signal can allow a significantly faster convergence of the calibration loop.

As mentioned before, the analog-domain linearization approach used for the CMD amplifier has also been implemented by utilizing capacitive degeneration, albeit in a pseudo-differential way [8]. Using essentially the same principle, capacitive degeneration was shown to cancel all the odd-order harmonics in the input transconductance and achieve excellent linearity (THD < -100dB). However, due to the pseudo-differential nature of the capacitive degeneration, the input transconductance also gets degenerated, and decreases by 33% at the optimum linearization point. This drawback can be avoided by also implementing capacitive degeneration in a common-mode fashion.

In a complementary CMOS amplifier, both PMOS and NMOS sides will have a degeneration capacitor. The two capacitors can be combined, resulting in an amplifier operating purely from the charge on the degeneration capacitor during the amplification phase. This floating amplifier with a capacitor acting as a supply has been shown in [15] where it has been used as the preamplifier of a comparator. The attractive quality of this floating amplifier is its implicit common-mode control. This eliminates the power and noise overhead of adding a dedicated common-mode feedback loop and significantly improves the power efficiency of the residue amplifier. A variant of this topology suitable for residue amplification is shown in Figure 7.2.

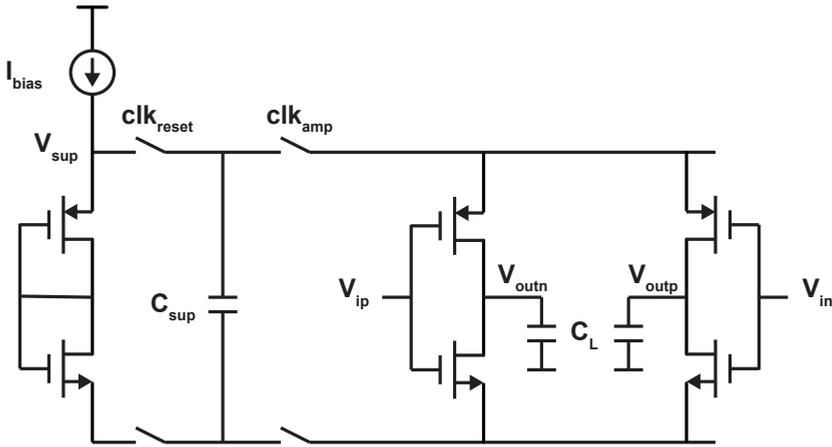


Figure 7.2: Floating residue amplifier with capacitive supply

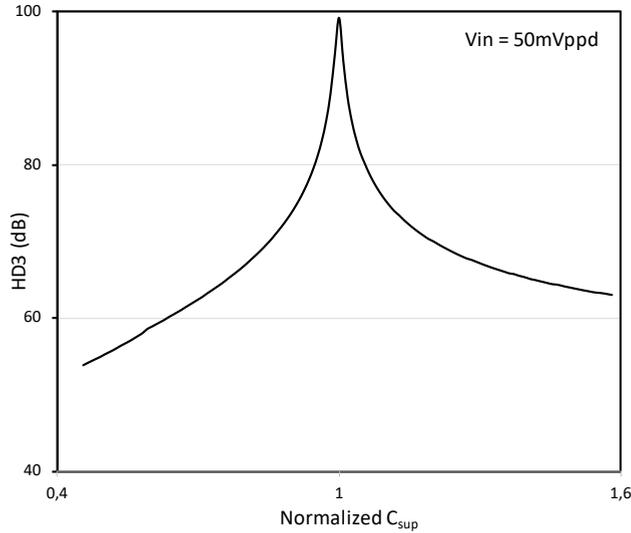


Figure 7.3: Floating residue amplifier HD3 versus C_{sup} sweep

The amplifier bias voltages are mapped on to the floating amplifier transistors via the supply capacitor, C_{sup} , which also acts as the common-mode degeneration. By biasing the amplifier in weak inversion region, its input transconductance can be linearized by tweaking the size of either the supply capacitor or the bias current, I_{bias} . The HD3 cancellation achieved by this amplifier is shown in Figure 7.3, which plots the effect of sweeping the supply capacitor, C_{sup} , on the amplifier HD3.

The floating amplifier achieves the same HD3 as the CMD amplifier presented in Chapter 5 (nearly -100dB) under similar design conditions (input swing = 50mV_{ppd} , gain = 4x).

REFERENCES

- [1] J. Li and U.-K. Moon, "Background calibration techniques for multi-stage pipelined ADCs with digital redundancy," *IEEE Trans. Circuits Syst.-II*, vol. 50, pp. 531–538, Sep. 2003.
- [2] J. McNeill and et al., "Split ADC" architecture for deterministic digital background calibration of a 16-bit 1-MS/s ADC," *IEEE J. Solid-State Circuits*, vol. 40, pp. 2437 – 2445, Dec. 2005.
- [3] R. Sehgal, F. van der Goes, and K. Bult, "A 12 b 53 mW 195 MS/s pipelined ADC with 82 dB SFDR using split-ADC calibration," in *Proc. IEEE ESSCIRC*, pp. 67–70, 2014.
- [4] R. Sehgal, F. van der Goes, and K. Bult, "A 12b 53mW 195MS/s pipelined ADC with 82dB SFDR using split-ADC calibration," *IEEE J. Solid-State Circuits*, vol. 50, pp. 1592–1603, Jul. 2015.
- [5] R. Sehgal, F. van der Goes, and K. Bult, "A 13-mW 64-dB SNDR 280MS/s pipelined ADC using linearized open-loop class-AB amplifiers," in *Proc. IEEE ESSCIRC*, pp. 131–134, 2017.
- [6] R. Sehgal, F. van der Goes, and K. Bult, "A 13-mW 64-dB SNDR 280MS/s pipelined ADC using linearized integrating amplifiers," *IEEE J. Solid-State Circuits*, vol. 53, p. 1878–1888, Jul. 2018.
- [7] K. Bult, "The Effect of Technology Scaling on Power Dissipation in Analog Circuits," in *Analog Circuit Design* (M.Steyeart, A. v. Roermund, and J. Huijsing, eds.), Springer, 2006.
- [8] M. Akter, R. Sehgal, F. van der Goes, and K. Bult, "A 66-dB SNDR pipelined split-ADC in 40-nm CMOS using a class-AB residue amplifier," *IEEE J. Solid-State Circuits*, vol. 53, pp. 2939–2950, Oct. 2018.
- [9] L. Luo and et al., "A capacitively-degenerated high-linearity dynamic amplifier using a real-time gain detection technique," in *IEEE ISCAS*, pp. 1–4, 2019.
- [10] W. Jiang and et al., "A temperature-stabilized single-channel 1-GS/s 60-dB SNDR SAR-assisted pipelined ADC with dynamic Gm-R based amplifier," *IEEE J. Solid-State Circuits*, vol. 55, pp. 322–332, Feb. 2020.
- [11] Y. Park and et al., "An 11-b 100-MS/s fully dynamic pipelined ADC using a high-linearity dynamic amplifier," *IEEE J. Solid-State Circuits*, vol. 55, p. 2468–2477, Sep. 2020.
- [12] Z. Zheng and et al., "A single-channel 5.5mW 3.3GS/s 6b fully dynamic pipelined ADC with post-amplification residue generation," in *ISSCC Dig. of Tech. Papers*, pp. 254–256, 2020.
- [13] Y. Kim and et al., "A 41 μ W 16MS/s 99.2dB-SFDR capacitively degenerated dynamic

- amplifier with nonlinear-slope-factor compensation,” in *ISSCC Dig. of Tech. Papers*, pp. 251–253, 2020.
- [14] S. Pan and K. Makinwa, “A $10\text{fJ}\cdot\text{K}^2$ wheatstone bridge temperature sensor with a tail-resistor-linearized OTA,” *IEEE J. Solid-State Circuits*, vol. 55, p. Early Access, Sep. 2020.
- [15] X. Tang and et al., “An energy-efficient comparator with dynamic floating inverter amplifier,” *IEEE J. Solid-State Circuits*, vol. 55, p. 1011–1022, Apr. 2020.

A

Appendix

In this Appendix, the mathematical analysis of distortion components in discrete-time amplifiers is presented.

A.1. OPEN-LOOP AMPLIFIER DISTORTION ANALYSIS

Considering the model for a nonlinear amplifier, shown in Figure A.1, used in an open-loop configuration. The first-order output voltage response of the amplifier can be expressed as –

$$I_x(t) = g_m V_{in}(t) = C_L \frac{dV_{out}(t)}{dt} + g_{out} V_{out}(t) \quad (\text{A.1})$$

If the input gm is assumed to have second- and third-order distortion, for a given input signal, $V_{in}(t)$, its output current can be expressed as –

$$I_x(t) = g_m V_{in}(t) + g_2 V_{in}^2(t) + g_3 V_{in}^3(t) \quad (\text{A.2})$$

where g_{i2} and g_{i3} are the second- and third-order distortion coefficients, respectively. In order to calculate the 2nd order output-referred distortion component, we begin with the assumption of $g_{i2} \neq 0$ and $g_{i3} = 0$. Rewriting equation (A.1) to include the second-order distortion coefficient –

$$g_m V_{in}(t) + g_{i2} V_{in}^2(t) = C_L \frac{dV_{out}(t)}{dt} + g_{out} V_{out}(t) \quad (\text{A.3})$$

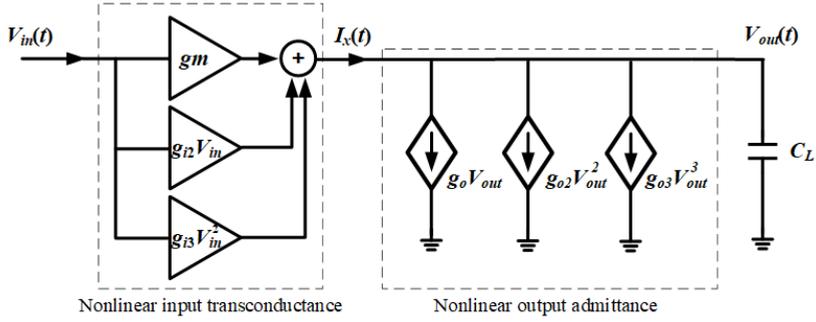


Figure A.1: Amplifier distortion model

Equation (A.3) can be simplified by transforming it to Laplace domain and rearranging the terms and bringing it back in time-domain, resulting in –

$$V_{out}(t) = A(1 - e^{-t/\tau})V_{in} + \frac{g_{i2}}{g_{out}} \frac{e^{-t/\tau}}{\tau} \otimes V_{in}^2(t) \quad t \geq 0 \quad (\text{A.4})$$

Hence, the second-order distortion component can be written as –

$$V_{out,in}^{(2)}(t) = \frac{g_{i2}}{g_{out}} (1 - e^{-t/\tau}) V_{in}^2 \quad t \geq 0 \quad (\text{A.5})$$

Similarly, the third-order distortion component can be calculated as –

$$V_{out,in}^{(3)}(t) = \frac{g_{i3}}{g_{out}} (1 - e^{-t/\tau}) V_{in}^3 \quad t \geq 0 \quad (\text{A.6})$$

Apart from the input distortion, the amplifier linearity is also often limited by the nonlinear output impedance. The behavior of this distortion depends entirely on the output voltage and can be calculated by assuming the output impedance to have second- and third-order non-linearity. For ease of calculation, the analysis is performed with admittances instead of impedances. If the amplifier output admittance can be written as –

$$g_{out} = g_o + g_{o2}V_{out}(t) + g_{o3}V_{out}^2(t) \quad (\text{A.7})$$

then the output current can be expressed as –

$$I_x(t) = g_m V_{in}(t) = C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) + g_{o2}V_{out}^2(t) + g_{o3}V_{out}^3(t) \quad (\text{A.8})$$

Solving the above equation for $V_{out}(t)$ is extremely complex due to its recursive nature. Hence, in order to simplify the analysis, we assume the distortion to be weak, which means we can approximate the distortion components by considering

them to be generated purely by the linear part of the output voltage. This results in -

$$I_x(t) = g_m V_{in}(t) = C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) + g_{o2} V_{out_lin}^2(t) + g_{o3} V_{out_lin}^3(t) \quad (\text{A.9})$$

where $V_{out_lin}(t)$ is the output voltage if g_{o2} and g_{o3} are zero, and is described by equation (3.2). Considering only the second-order term, we can use the same method as before to express it as -

$$V_{out,out}^{(2)}(t) = \frac{g_{o2}}{g_o} \frac{e^{-t/\tau}}{\tau} \otimes \left(\frac{g_m}{g_o} (1 - e^{-t/\tau}) \right)^2 \quad t \geq 0 \quad (\text{A.10})$$

which results in -

$$V_{out,out}^{(2)}(t) = \frac{g_{o2} g_m^2}{g_o^3} \left(1 - 2 \left(\frac{t}{\tau} \right) e^{-t/\tau} - e^{-2t/\tau} \right) \quad t \geq 0 \quad (\text{A.11})$$

Similarly, the third-order distortion component can be calculated as -

$$V_{out,out}^{(3)}(t) = \frac{g_{o3}}{g_o} \frac{e^{-t/\tau}}{\tau} \otimes \left(\frac{g_m}{g_o} (1 - e^{-t/\tau}) \right)^3 \quad t \geq 0 \quad (\text{A.12})$$

$$\begin{aligned} &= \frac{g_{o3} g_m^3}{g_o^4} \left((1 - e^{-t/\tau}) - \frac{e^{-t/\tau}}{2} (1 - e^{-2t/\tau}) \right. \\ &\quad \left. + 3e^{-t/\tau} (1 - e^{-t/\tau}) - 3e^{-t/\tau} \left(\frac{t}{\tau} \right) \right) \quad t \geq 0 \end{aligned} \quad (\text{A.13})$$

A.2. CLOSED-LOOP AMPLIFIER DISTORTION ANALYSIS

INPUT DISTORTION

A negative-feedback amplifier with a nonlinear input transconductance is modeled using a block diagram as shown in Figure A.2. The input g_m is assumed to have second- and third-order distortion, and for a given virtual ground signal, $V_{vg}(t)$, its output current can be expressed as -

$$I_x(t) = g_m V_{vg}(t) + g_{i2} V_{vg}^2(t) + g_{i3} V_{vg}^3(t) \quad (\text{A.14})$$

where g_{i2} and g_{i3} are the second- and third-order distortion coefficients, respectively. Based on (A.14), the output voltage of the amplifier, $V_{out}(t)$, can be calculated through the following series of equations -

$$\begin{aligned} V_{vg}(t) &= V_{in}(t) - \beta V_{out}(t) \\ I_x(t) &= g_m V_{vg}(t) + g_{i2} V_{vg}^2(t) + g_{i3} V_{vg}^3(t) \\ I_x(t) &= C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) \end{aligned} \quad (\text{A.15})$$

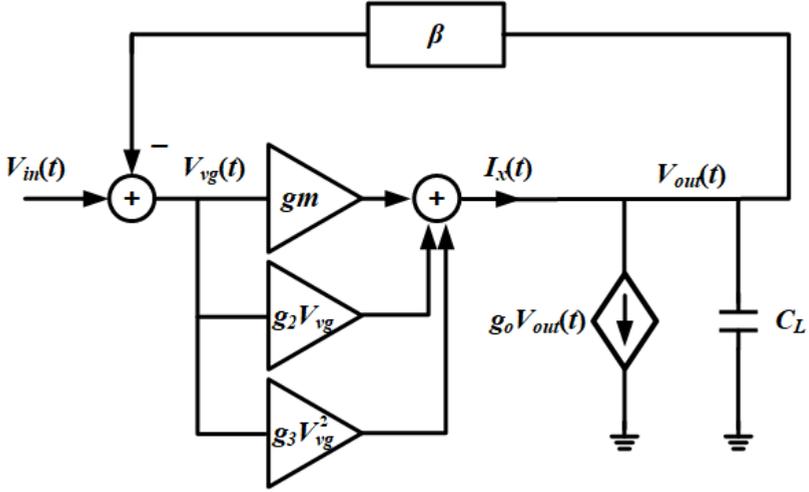


Figure A.2: Negative-feedback amplifier model with nonlinear input transconductance

It should be noted that, as seen from (A.14), there is no memory element present in the relationship between $I_x(t)$ and $V_{vg}(t)$. That is because the only memory present in the circuit is associated with the load capacitance, which is taken into account in (A.2).

Due to its recursive nature, solving (A.2) in order to calculate $V_{out}(t)$ requires very complex mathematical calculations. Hence, in order to simplify the analysis, we use the same simplifications to the model in Figure A.2 as used in the previous analysis. We assume that the transconductance is weakly nonlinear, i.e. $g_m \gg g_{i2}g_{i3}$. This entails that while $V_{vg}(t)$ does include the distortion components, they are relatively small. So their effect can be ignored by assuming that the input distortion is generated only by the linear part of $V_{vg}(t)$. By taking this simplification into account, (A.14) can be rewritten as –

$$I_x(t) = g_m V_{vg}(t) + g_{i2} V_{vg_lin}^2(t) + g_{i3} V_{vg_lin}^3(t) \quad (\text{A.16})$$

where $V_{vg_lin}(t)$ is the virtual ground signal generated for the same amplifier with no distortion (i.e. $g_{i2} = g_{i3} = 0$) and is the same as the expression in (3.15). Equations (A.2) and (A.16) can be combined as –

$$g_m V_{vg}(t) + g_{i2} V_{vg_lin}^2(t) + g_{i3} V_{vg_lin}^3(t) = C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) \quad (\text{A.17})$$

We first focus on the second-order distortion component by assuming $g_{i2} \neq 0$ and

$g_{i3} = 0$. This results in –

$$g_m V_{vg}(t) + g_{i2} V_{vg_lin}^2(t) = C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) \quad (\text{A.18})$$

$$\text{or, } g_m V_{in}(t) + g_{i2} V_{vg_lin}^2(t) = C_L \frac{dV_{out}(t)}{dt} + (g_o + g_m \beta) V_{out}(t) \quad (\text{A.19})$$

Hence, again assuming $V_{in}(t)$ to be a step function with an amplitude of V_{in} , $V_{out}(t)$ can be calculated by using the Laplace transform to rearrange (A.19) and converting it back to the transient domain –

$$V_{out}(t) = \frac{A_o}{1 + A_o \beta} (1 - e^{-t/\tau_f}) V_{in} + \frac{g_{i2}}{g_o (1 + A_o \beta)} V_{vg_lin}^2(t) \otimes \frac{e^{-t/\tau_f}}{\tau_f}, t \geq 0 \quad (\text{A.20})$$

where $A_o (= g_m/g_o)$ defines the open-loop gain of the amplifier and τ_f is the closed-loop time constant and is equal to $\tau/(1 + A_o \beta)$. The first term in (A.20) is the linear part of the amplifier output signal and is equal to the expression in (3.2), while the second term represents the component of the output voltage arising due to the second-order distortion in the input transconductance, and will be referred to as $V_{out,in}^{(2)}(t)$ –

$$V_{out,in}^{(2)}(t) = \frac{g_{i2}}{g_o (1 + A_o \beta)} V_{vg_lin}^2(t) \otimes \frac{e^{-t/\tau_f}}{\tau_f} \quad t \geq 0 \quad (\text{A.21})$$

where $V_{vg_lin}^2(t)$ is given by –

$$V_{vg_lin}^2(t) = \left(\frac{1 + A_o \beta e^{-t/\tau_f}}{1 + A_o \beta} V_{in} \right)^2 \quad (\text{A.22})$$

$$= \left(1 + (A_o \beta)^2 e^{-2t/\tau_f} + 2A_o \beta e^{-t/\tau_f} \right) \frac{V_{in}^2}{(1 + A_o \beta)^2}, \quad t \geq 0 \quad (\text{A.23})$$

By taking the convolution of $V_{vg_lin}^2(t)$ with the rest of the expression in (A.21), $V_{out,in}^{(2)}(t)$ can be calculated as –

$$V_{out,in}^{(2)}(t) = \frac{g_{i2}}{g_o (1 + A_o \beta)^3} V_{in}^2 \left((1 - e^{-t/\tau_f}) + (A_o \beta)^2 e^{-t/\tau_f} (1 - e^{-t/\tau_f}) + 2A_o \beta e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \quad (\text{A.24})$$

This represents the transient behavior of the second-order distortion component from the nonlinear input g_m in response to a step input at $t = 0$. After the initial settling, the distortion component settles towards a steady-state value of $\frac{g_{i2}/g_o}{(1 + A_o \beta)^3} V_{in}^2$

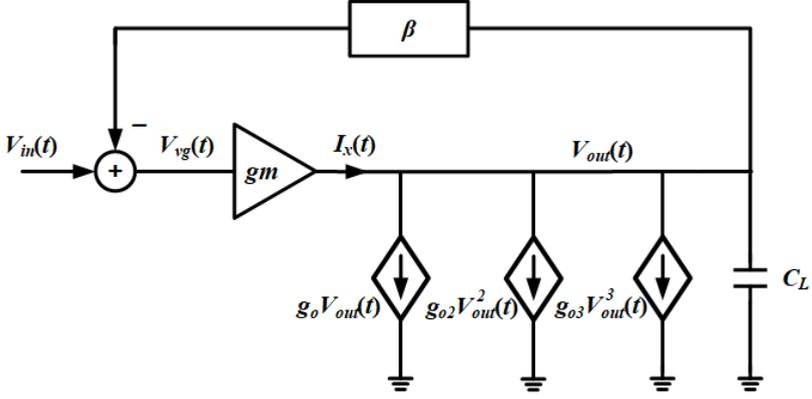


Figure A.3: Negative-feedback amplifier model with nonlinear output admittance

, which is the same as that derived in [1]. The third-order distortion component, $V_{out,in}^{(3)}(t)$, can be similarly be expressed as –

$$V_{out,in}^{(3)}(t) = \frac{g_{i3}}{g_o(1 + A_o\beta)} V_{vg_lin}^3(t) \otimes \frac{e^{-t/\tau_f}}{\tau_f} \quad t \geq 0 \quad (\text{A.25})$$

which can be solved as –

$$V_{out,in}^{(3)}(t) = \frac{g_{i3}}{g_o(1 + A_o\beta)^4} V_{in}^3 \left((1 - e^{-t/\tau_f}) + 3(A_o\beta)^2 e^{-t/\tau_f} (1 - e^{-t/\tau_f}) + \frac{1}{2}(A_o\beta)^3 e^{-t/\tau_f} (1 - e^{-2t/\tau_f}) + 3A_o\beta e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \quad (\text{A.26})$$

It can be seen from (A.26) that the third-order distortion component, $V_{out,in}^{(3)}(t)$, also shows a similar settling behavior as $V_{out,in}^{(2)}(t)$, with both components being suppressed by the loop at the rate of 8.6dB/ τ as posited in Chapter 3.

OUTPUT DISTORTION

Apart from the input distortion, the amplifier linearity is also often limited by the nonlinear output impedance. The closed-loop amplifier with output distortion is modeled in Figure A.3, where g_{o2} and g_{o3} are the second- and third-order distortion coefficients of the output admittance. The block diagram can be expressed through the following equations –

$$\begin{aligned} V_{vg}(t) &= V_{in}(t) - \beta V_{out}(t) \\ I_x(t) &= g_m V_{vg}(t) \\ I_x(t) &= C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) + g_{o2} V_{out}^2(t) + g_{o3} V_{out}^3(t) \end{aligned} \quad (\text{A.27})$$

The solution for (A.27) can be simplified in a similar way as that for the input distortion, by assuming that only the linear part of $V_{out}(t)$ creates the distortion. This would lead to –

$$I_x(t) = C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) + g_{o2} V_{out_lin}^2(t) + g_{o3} V_{out_lin}^3(t) \quad (\text{A.28})$$

where $V_{out_lin}(t)$ represents the output voltage for $g_{o2} = g_{o3} = 0$, and yields the same expression as (3.13) –

$$V_{out_lin}(t) = \frac{A_o}{1 + A_o\beta} (1 - e^{-t/\tau_f}) V_{in} \quad t \geq 0 \quad (\text{A.29})$$

To calculate the second-order distortion, we follow a similar process as used for calculating the input distortion by assuming $g_{o2} \neq 0$ and $g_{o3} = 0$. By combining (A.27)-(A.29), we can describe the output voltage as –

$$g_m V_{vg}(t) = C_L \frac{dV_{out}(t)}{dt} + g_o V_{out}(t) + g_{o2} V_{out_lin}^2(t) \quad (\text{A.30})$$

$$\text{or, } g_m V_{in}(t) = C_L \frac{dV_{out}(t)}{dt} + (g_o + g_m\beta) V_{out}(t) + g_{o2} V_{out_lin}^2(t) \quad (\text{A.31})$$

By again assuming a step input, (A.31) can be simplified by converting it to Laplace domain, rearranging and then bringing it back in the transient domain, resulting in –

$$V_{out}(t) = \frac{A_o}{1 + A_o\beta} (1 - e^{-t/\tau_f}) V_{in} + \frac{g_{o2}}{g_o (1 + A_o\beta)} V_{out_lin}^2(t) \otimes \frac{e^{-t/\tau_f}}{\tau_f}, t \geq 0 \quad (\text{A.32})$$

The first term again represents the linear part of the output voltage. To derive the second-order distortion component, $V_{out,out}^{(2)}$, we consider the second term in (A.32) –

$$V_{out,out}^{(2)}(t) = \frac{g_{o2}}{g_o (1 + A_o\beta)} V_{out_lin}^2(t) \otimes \frac{e^{-t/\tau_f}}{\tau_f}, t \geq 0 \quad (\text{A.33})$$

where $V_{out_lin}^{(2)}(t)$ is given by –

$$V_{out_lin}^{(2)}(t) = \left(\frac{A_o}{1 + A_o\beta} (1 - e^{-t/\tau_f}) \right) V_{in}^2 \quad (\text{A.34})$$

$$= \left(\frac{A_o}{1 + A_o\beta} \right)^2 (1 + e^{-2t/\tau_f} - 2e^{-t/\tau_f}) V_{in}^2 \quad t \geq 0 \quad (\text{A.35})$$

By taking the convolution of $V_{out_lin}^2(t)$ with the rest of the equation in (A.33), $V_{out,out}^{(2)}(t)$ can be calculated as –

$$V_{out,out}^{(2)}(t) = \frac{(g_{o2}/g_o) A_o^2}{(1 + A_o\beta)^3} V_{in}^2 \left((1 - e^{-t/\tau_f}) + e^{-t/\tau_f} (1 - e^{-t/\tau_f}) - 2e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \quad (\text{A.36})$$

The third-order distortion component in the output voltage, $V_{out,out}^{(3)}(t)$, can similarly be calculated by solving for –

$$V_{out,out}^{(3)}(t) = \frac{(g_{o3}/g_o)}{(1 + A_o\beta)} V_{out_lin}^3(t) \otimes \frac{e^{-t/\tau_f}}{\tau_f}, \quad t \geq 0 \quad (\text{A.37})$$

$$\begin{aligned} &= \frac{(g_{o3}/g_o) A_o^3}{(1 + A_o\beta)^4} V_{in}^3 \left((1 - e^{-t/\tau_f}) + 3e^{-t/\tau_f} (1 - e^{-t/\tau_f}) \right. \\ &\quad \left. - \frac{1}{2} e^{-t/\tau_f} (1 - e^{-2t/\tau_f}) - 3e^{-t/\tau_f} \left(\frac{t}{\tau_f} \right) \right) \end{aligned} \quad (\text{A.38})$$

It can be seen from (A.36) and (A.38) that the two distortion components, $V_{out,out}^{(2)}(t)$ and $V_{out,out}^{(3)}(t)$, follow the settling behavior of the output voltage. The settling plots for both input and output distortion components are shown visually in 3.

REFERENCES

- [1] W. Sansen, “Distortion in elementary transistor circuits,” *IEEE Trans. Circuits Syst.-II*, vol. 46, pp. 315–325, Mar. 1999.

B

Appendix

In chapter 5, the linearization principle based on utilizing common-mode degeneration was presented. In this Appendix, the mathematical derivation for the optimum common-mode degeneration impedance required for linearization is presented.

Considering a differential common-source MOS amplifier with a common-mode degeneration impedance. For a MOS transistor biased in weak-inversion (with its bulk tied to ground), its drain current, I_D , can be expressed as –

$$I_D = I_S e^{\left(\frac{V_{Gb}-V_{th}}{nV_T}\right)} e^{\left(\frac{-V_{Sb}}{V_T}\right)} \quad (\text{B.1})$$

where V_{Gb} and V_{Sb} are the bias voltages at the gate and source terminals of the transistor, respectively. If a differential voltage, V_{id} , is applied across the MOS transistor gate terminals, then the differential and total tail current can be expressed as –

$$\begin{aligned} I_{diff} &= I_S e^{\left(\frac{V_{Gb}-V_{th}}{nV_T}\right)} e^{\left(\frac{-V_{Sb}+V_S}{V_T}\right)} \left(e^{\frac{V_{id}}{2nV_T}} - e^{\frac{-V_{id}}{2nV_T}} \right) \\ &= 2I_D e^{\frac{-V_S}{V_T}} \sinh\left(\frac{V_{id}}{2nV_T}\right) \end{aligned} \quad (\text{B.2})$$

$$\begin{aligned} I_{tail} &= I_S e^{\left(\frac{V_{Gb}-V_{th}}{nV_T}\right)} e^{\left(\frac{-V_{Sb}+V_S}{V_T}\right)} \left(e^{\frac{V_{id}}{2nV_T}} + e^{\frac{-V_{id}}{2nV_T}} \right) \\ &= 2I_D e^{\frac{-V_S}{V_T}} \cosh\left(\frac{V_{id}}{2nV_T}\right) \end{aligned} \quad (\text{B.3})$$

where V_S is the swing at the tail terminal of the amplifier for a differential input V_{id} . If the amplifier has a perfectly linear input transconductance, g_m , over the entire differential input range, then its differential current can also be expressed as -

$$I_{diff} = g_m V_{id} \quad (\text{B.4})$$

Using equations (B.2), (B.3) and (B.4), we can derive the following expression for I_{tail} -

$$\frac{I_{tail}}{I_{diff}} = \coth\left(\frac{V_{id}}{2nV_T}\right) \quad (\text{B.5})$$

$$\text{or, } I_{tail} = g_m V_{id} \coth\left(\frac{V_{id}}{2nV_T}\right) \quad (\text{B.6})$$

I_{tail} is the optimum total current (or tail current) flowing through the two input transistors that ensures a perfectly linear differential input transconductance. In order to achieve this tail current, the common-mode degeneration impedance can be calculated as -

$$Z_{deg} = \frac{dV_S}{dI_{tail}} = \frac{dV_S}{dV_{id}} \left(\frac{dI_{tail}}{dV_{id}}\right)^{-1} \quad (\text{B.7})$$

The two derivative terms in (B.7) can be calculated by differentiating (B.2) and (B.6) w.r.t V_{id} . The derivative of I_{tail} can be calculated as -

$$\frac{dI_{tail}}{dV_{id}} = g_m \coth\left(\frac{V_{id}}{2nV_T}\right) - \frac{g_m V_{id}}{2nV_T} \frac{1}{\sinh^2\left(\frac{V_{id}}{2nV_T}\right)} \quad (\text{B.8})$$

$$= \frac{g_m}{\sinh(x)} \left(\cosh(x) - \frac{x}{\sinh(x)} \right) \quad (\text{B.9})$$

where x is defined as $V_{id}/2nV_T$.

To calculate the derivative of V_S w.r.t V_{id} , we need to consider (B.2), which can be combined with equation (B.4) and be rewritten as -

$$g_m V_{id} = 2I_D e^{\frac{-V_S}{V_T}} \sinh\left(\frac{V_{id}}{2nV_T}\right) \quad (\text{B.10})$$

In weak-inversion, the g_m of a MOS transistor can be expressed as -

$$g_m = \frac{I_D}{nV_T} \quad (\text{B.11})$$

Using (B.11), equation (B.10) can be rewritten as -

$$V_{id} = 2nV_T e^{\frac{-V_S}{V_T}} \sinh\left(\frac{V_{id}}{2nV_T}\right) \quad (\text{B.12})$$

$$\text{or, } e^{\frac{V_S}{V_T}} = \frac{\sinh\left(\frac{V_{id}}{2nV_T}\right)}{\frac{V_{id}}{2nV_T}} = \frac{\sinh(x)}{x} \quad (\text{B.13})$$

Differentiating both sides w.r.t V_{id} –

$$\frac{dV_S}{dV_{id}} e^{\frac{V_S}{V_T}} = \frac{1}{2nV_T} \left(\frac{\cosh(x)}{x} - \frac{\sinh(x)}{x^2} \right) \quad (\text{B.14})$$

Using (B.13), equation (B.14) can be rewritten as –

$$\frac{dV_S}{dV_{id}} \frac{\sinh(x)}{x} = \frac{1}{2n} \left(\frac{\cosh(x)}{x} - \frac{\sinh(x)}{x^2} \right) \quad (\text{B.15})$$

$$\text{or, } \frac{dV_S}{dV_{id}} = \frac{1}{2n \sinh(x)} \left(\cosh(x) - \frac{\sinh(x)}{x} \right) \quad (\text{B.16})$$

Using (B.9) and (B.16), the optimum degeneration impedance can be expressed as –

$$Z_{deg} = \frac{\frac{1}{2n \sinh(x)} \left(\cosh(x) - \frac{\sinh(x)}{x} \right)}{\frac{g_m}{\sinh(x)} \left(\cosh(x) - \frac{x}{\sinh(x)} \right)} \quad (\text{B.17})$$

$$= \frac{1}{2ng_m} \left(\frac{\cosh(x) - \frac{\sinh(x)}{x}}{\cosh(x) - \frac{x}{\sinh(x)}} \right) \quad (\text{B.18})$$

Equation (B.18) can be simplified by using Taylor series expansion for the hyperbolic functions. To derive Z_{deg} for only HD3 cancellation, all higher than third-order terms in the Taylor series expansion can be ignored, resulting in –

$$Z_{deg} = \frac{1}{2ng_m} \quad (\text{B.19})$$

Summary

The past few decades have seen a sustained demand for higher bandwidth in communication networks. This push for higher bandwidths has been supported by increase in sampling rates of high-resolution data converters. Pipelined analog-to-digital converters (ADCs) have been a popular choice for high-speed ADCs, as they can significantly relax the speed-requirements by breaking down the quantization process in multiple concurrent steps.

One of the key ingredients for pipelining is residue amplification, and the design choices used for residue amplifier often determine the power efficiency of the entire ADC. The residue amplifier, in essence, consists of a four-transistor core that drives a capacitive load, and in order to attain gain accuracy, this core amplifier is augmented with circuit overhead and settling accuracy, both of which cause significant increase in power consumption. This dissertation investigates residue amplifiers that can achieve the desired gain accuracy and distortion performance while relying on minimum settling accuracy and circuit overhead.

Chapter 2 reviews the pipelined ADC architecture and the operation of its major sub-blocks – Coarse ADC (CADC), sub-digital-to-analog converter (DAC) and the residue amplifier. In a pipelined ADC, the quantization process is spread across multiple ADC stages. Each stage consists of a coarse ADC that generates a coarse digital estimate of the input signal. This coarse estimate is subtracted from the input signal with the help of a sub-DAC that generates an analog equivalent of the CADC output. After the subtraction, the remaining residue signal is amplified by the residue amplifier and passed on to the following stage.

The effect of non-idealities in these three sub-blocks on the ADC performance is also examined in Chapter 2. While combining the digital output bits from different ADC stages, the digital encoder needs to account for the gain of the residue amplifier. Any difference between the residue amplifier gain and the digital encoder gain will corrupt the ADC output. This difference in digital and analog gains can arise from either a fixed gain error or distortion in the residue amplifiers.

For high-resolution ADCs, the constraints on the residue amplifier gain error for the first few stages of the pipelined ADC can be very strict. Using classic amplifier

solutions with a high DC gain amplifier in negative feedback and high settling accuracy can be extremely costly in terms of power dissipation. Several techniques have been proposed to relax the accuracy requirements of the residue amplifier and are reviewed in Chapter 2. The most popular of these techniques is digital calibration. The magnitude of the amplifier gain error and distortion can be judged by injecting a known calibration signal into the residue amplifier and analyzing its digitized equivalent. Digital calibration can be used during start-up (foreground-mode) or in normal operation when the ADC is digitizing the input signal (background-mode). Background-mode calibration allows continuous monitoring of the amplifier errors and help relax the gain sensitivity of the amplifier. The main challenge of background digital calibration, however, is the separation of the calibration signal from input signal.

Due to their discrete-time nature, the key characteristics of residue amplifiers – gain, noise and distortion – are transient in nature. Chapter 3 presents a time-domain analysis of open- and closed-loop residue amplifiers and the settling behavior of their design metrics. For both open- and closed-loop amplifiers, the amplifier gain error was shown to settle at 8.6dB per unit τ of settling, where τ is the time-constant of the amplifier and is determined by the output impedance and the capacitive load of the amplifier. Both open- and closed-loop residue amplifiers were also shown to achieve roughly 90% of their final SNR in the first three time-constants of settling. And the suppression of distortion originating from the amplifier input in closed-loop amplifiers was also shown to occur at 8.6dB/ τ .

Based on the above analysis, the optimum settling points for two residue amplifier design cases using class-A amplifier topologies were identified. In the first case, a closed-loop amplifier with a fixed SNR was considered and was shown to achieve maximum power efficiency between $2-3\tau$ settling. For the second design case, both SNR and gain error of the residue amplifier were constrained to the desired specification over all values of settling accuracy. It was shown that for higher values of settling, N (where $N = t/\tau > 2$), the power dissipation of a settling amplifier is $N/2$ times higher than the power of an integrating amplifier ($t/\tau \approx 0$), for the same gain and SNR. While the above analysis optimizes the power dissipation of the two residue amplifier design cases for a certain SNR, the gain accuracy suffers in both cases due to the low settling accuracy. Hence in order to allow the residue amplifier to operate at the optimum settling accuracy points, they need to be supported by gain and distortion calibration, which allow them to achieve the desired gain accuracy while ideally imposing a low overhead on the ADC power and complexity.

Chapter 4 reviews the split-ADC calibration technique, which is a fully deterministic background calibration method that has been shown to achieve fast convergence with relatively low analog overhead. In this chapter, the split-ADC technique is extended for calibrating the residue amplifier distortion in multiple stages of a pipelined ADC. In the split-ADC calibration method, the ADC is split into two identical halves, with each half-ADC digitizing the same input signal. By introducing a calibration offset between them, the two half-ADCs are forced to go through different conversion trajectories. By taking the difference of the two digital outputs, the input signal can be effectively cancelled, leaving behind the digitized calibration offset that can be used to detect the ADC errors. Since the calibration offset can be separated from the input signal without any lengthy decorrelation, the split-ADC calibration can achieve fast convergence while operating completely in background.

The error detection for residue amplifier gain error and distortion using split-ADC technique can be implemented by studying their effect on the difference signal. The gain error introduces discontinuities in the difference signal whenever one of the half-ADCs transitions from one subrange to another. The size of these discontinuities is proportional to the magnitude of the gain error. Amplifier odd-order distortion injects a quadratic component in the difference signal, and the curvature of this quadratic component can be used to estimate the third-order distortion in the residue amplifier. The size of the error signal for both gain and distortion calibration is directly proportional to the calibration offset between the split-ADCs. This means that a larger calibration offset will result in a faster convergence, albeit at the cost of ADC dynamic range. In the prototype presented in Chapter 4, an offset equivalent to 5.5% of the ADC full-scale was added to limit the impact on the ADC dynamic range to 0.5dB.

A key consideration of split-ADC calibration is the effect of mismatch between the half-ADCs on the difference signal. Any gain mismatch between the two half-ADCs will introduce a fraction of the input signal into the difference signal and causes it to have a slope. This slope has to be compensated for by scaling one of the half-ADC outputs. Calibrating amplifier gain and distortion errors across multiple ADC stages also presents certain challenges. Since all calibration loops are essentially working off the same difference signal, there are non-orthogonalities between these calibration loops that can affect the convergence of these loops. These non-orthogonalities can be resolved by running the calibration loops iteratively.

In order to test the efficacy of the split-ADC calibration technique, a prototype 12-bit pipelined split-ADC was implemented. The prototype ADC consists of nine

1.5-bit stages followed by a 5-bit fine ADC. Gain calibration was used to correct the gain for the first 6 stages, while nonlinearity calibration was applied only to the first two stages. Both gain and distortion errors were corrected through digital post-processing.

The prototype ADC was designed with a residue amplifier with high DC gain and bandwidth as a sufficiently accurate starting point. The calibration was then used to relax the settling accuracy of the residue amplifier without compromising the ADC resolution. A single-stage current mirror opamp with gain-boosting was utilized to achieve a DC loop gain of 65dB. The prototype ADC was implemented in 40nm 1V digital CMOS process and occupied an area of 0.81mm^2 , with the digital post-processing implemented in MATLAB. With the help of calibration, the prototype ADC achieved 10.5b ENOB at a sampling rate of 195MS/s with a power dissipation of 53mW. While working continuously in background, the split-ADC calibration improved the ADC SFDR by 37dB within 70000 samples along with reducing the power dissipation of the residue amplifiers by 60%.

The prototype ADC presented in Chapter 4 relied on digital correction of gain and distortion errors, which face dynamic range and accuracy limitations while adding considerable overhead in the digital backend. Chapter 5 presents an amplifier topology that allows the correction of gain and distortion error directly at the source in the analog domain. Open-loop integrating amplifiers provide a simple way to control the gain through the bias current due to their extremely low settling. For distortion correction, a linearization principle based on the degeneration of a weak-inversion common-source amplifier was proposed. The degeneration impedance counteracts the expanding input g_m characteristic of a weak-inversion common-source amplifier, and the third-order component was shown to be completely cancelled for $R_{deg} = 1/2ng_m$.

The impact of degeneration on the amplifier input transconductance can be eliminated by using the degeneration in a common-mode fashion, while still improving the g_m variation from 50% in a differential pair to <1%. The common-mode degeneration (CMD) linearization principle was verified through measurements on a strip-board amplifier version built with discrete BJT components that showed an HD3 of -95dB for an input swing of 50mV_{ppd} . The distortion cancellation was extended to 2nd and 5th order distortion components by using a pair of input-driven weak inversion MOS transistors as the common-mode tail degeneration. By tuning the bias voltage and the attenuation of the input signal driving the degeneration network, the proposed amplifier topology was able to achieve a THD of <-90dB for

an input swing of $300\text{mV}_{\text{ppd}}$ and a gain of 4x.

Chapter 6 presents a prototype 12-bit pipelined split-ADC that was built around the proposed integration-based CMD amplifier to test the efficacy of the linearization principle. The amplifier gain error and non-linearity (even- and odd-order distortion) were detected in the background through the split-ADC calibration technique. In order to reduce the mismatch between the two half-ADCs, a modified split-ADC calibration architecture was implemented. By digitizing the input sample twice over two separate clock cycles, the calibration offset can be introduced over time and the same ADC hardware can be utilized for the two half-ADC operations. This “split-over-time” ADC allowed the residue amplifiers in the prototype ADC to be shared between the two half-ADCs. To further improve the power efficiency of the residue amplifiers, they were shared between two time-interleaved ADC lanes.

The 12-bit pipelined ADC consists of five 3-bit stages (with one bit as overrange) followed by a 4-bit fine ADC. The proposed CMD amplifier is implemented in a complementary fashion to improve power efficiency. Fabricated in 28-nm 1-V digital CMOS, the ADC occupied an area of 0.22mm^2 . The ADC output data was imported into MATLAB for error detection and the calibration loop was closed by driving the analog correction parameters in the residue amplifier degeneration. Sampling at 280MS/s , the ADC achieved an ENOB of 10.34b with a power consumption of 13mW, displaying an overall Schreier FoM of 164.3dB. The linearization scheme, along with the split-ADC calibration scheme operating in the background, enabled the prototype ADC to achieve $>77\text{dB}$ SFDR while displaying at least 3dB better Schreier FoM w.r.t other comparable state-of-the-art pipeline ADC designs.

Finally, Chapter 7 concludes the thesis by summarizing the original contributions and comparing the second prototype ADC with state-of-the-art pipelined ADCs with similar speed, resolution and architecture. The prototype ADC was shown to achieve one of the best Figure of Merits (FoMs) amongst state-of-the-art designs, highlighting the power efficiency of the proposed CMD amplifier. The amplifier distortion calibration using split-ADC technique was also seen to be one of the fastest background distortion calibration techniques reported in literature. Chapter 7 also presents some threads for future exploration, most prominent of which includes the using capacitive common-mode degeneration for linearizing the residue amplifier.

Samenvatting

De afgelopen decennia is er een aanhoudende vraag naar hogere bandbreedte in communicatienetwerken geweest. Dit streven naar hogere bandbreedtes werd ondersteund door een toename van de bemonsteringsfrequenties van dataconverters met hoge resolutie. Gepijplijnde analog-naar-digitaal-omzetters (ADCs) zijn een populaire keuze voor snelle ADCs, omdat ze de snelheidsvereisten aanzienlijk kunnen verminderen door het kwantisatieproces in meerdere gelijktijdige stappen op te splitsen.

Een van de belangrijkste ingrediënten voor pipelining is versterking van het residu, en de ontwerpkeuzes die worden gebruikt voor residuversterkers bepalen vaak de energie-efficiëntie van de gehele ADC. In de kern bestaat de residuversterker uit een kern met vier transistoren die een capacitieve belasting aandrijft, en om de versterkingsnauwkeurigheid te bereiken, wordt deze versterkerkern uitgebreid met overhead van andere componenten en settlingnauwkeurigheid, die beide een aanzienlijke toename van het stroomverbruik veroorzaken. Dit proefschrift onderzoekt residuversterkers die de gewenste versterkingsnauwkeurigheid en vervormingsprestaties kunnen bereiken, terwijl ze vertrouwen op minimale settlingnauwkeurigheid en overhead van andere componenten.

Hoofdstuk 2 bespreekt de gepijplijnde ADC-architectuur en de werking van de belangrijkste subblokken - Grove ADC (CADC), sub-DAC en de residuversterker. In een gepijplijnde ADC is het kwantisatieproces verspreid over meerdere ADC trappen. Elke trap bestaat uit een grove ADC die een grove digitale schatting van het ingangssignaal genereert. Deze grove schatting wordt van het ingangssignaal afgetrokken met behulp van een sub-DAC die een analog equivalent van de CADC-uitgang genereert. Na de aftrekking wordt het resterende residusignaal versterkt door de residuversterker en doorgegeven aan de volgende trap.

Het effect van niet-idealiteiten in deze drie subblokken op de ADC-prestaties wordt ook onderzocht in Hoofdstuk 2. Terwijl de digitale uitgangsbits van verschillende ADC-trappen worden gecombineerd, moet de digitale encoder rekening houden met de versterking van de residuversterker. Elk verschil tussen de werkelijke versterking van de residuversterker en de versterking van de digitale encoder zal de

ADC-uitgang vervormen. Dit verschil in digitale en analoge versterking kan het gevolg zijn van een vaste versterkingsfout of vervorming in de residuversterkers.

Voor ADC's met een hoge resolutie kunnen de eisen voor de versterkingsfout van de residuversterker voor de eerste paar trappen van de gepijplijnde ADC erg streng zijn. Het gebruik van klassieke versterkeroplossingen met een versterker met hoge DC-versterking in negatieve feedback en een hoge settlingnauwkeurigheid kan extreem kostbaar zijn in termen van vermogensverbruik. Er zijn verschillende technieken voorgesteld om de nauwkeurigheidseisen van de residuversterker te verminderen en deze worden besproken in Hoofdstuk 2. De meest populaire van deze technieken is digitale kalibratie. De grootte van de versterkingsfout en de vervorming van de versterker kunnen worden beoordeeld door een bekend kalibratiesignaal in de residuversterker te injecteren en het gedigitaliseerde equivalent ervan te analyseren. Digitale kalibratie kan worden gebruikt tijdens het opstarten (voorggrondmodus) of bij normaal bedrijf wanneer de ADC hetingangssignaal digitaliseert (achtergrondmodus). Kalibratie in de achtergrondmodus maakt continue bewaking van de versterkerfouten mogelijk en helpt de versterkingsgevoeligheid van de versterker te reduceren. De belangrijkste uitdaging van digitale achtergrond kalibratie is echter de scheiding van het kalibratiesignaal van hetingangssignaal.

Vanwege hun discrete-tijd karakter hebben de belangrijkste kenmerken van residuversterkers - versterking, ruis en vervorming - een transient karakter. Hoofdstuk 3 presenteert een analyse in het tijdsdomein van open- en gesloten-lus residuversterkers en het settlinggedrag van hun ontwerpstatistieken. Voor zowel open- als gesloten-lus versterkers bleek de versterkingsfout van de versterker te stabiliseren op 8,6 dB per eenheid τ settling, waarbij τ de tijdconstante van de versterker is en wordt bepaald door de uitgangsimpedantie en de capacatieve belasting van de versterker. Zowel residuversterkers met open als gesloten lus bleken ook ongeveer 90% van hun uiteindelijke SNR te bereiken in de eerste drie tijdconstanten van settling. En hoewel negatieve feedback in gesloten lus versterkers de vervorming onderdrukt die afkomstig is van de versterkeringang, bleek de onderdrukking ook op te treden met $8.6 \text{ dB}/\tau$.

Op basis van de bovenstaande analyse werden de optimale settlingpunten voor twee ontwerpgevallen van een residuversterker op basis van klasse-A-versterkertopologieën geïdentificeerd. In het eerste geval werd een gesloten lus versterker met een vaste SNR overwogen en bleek deze een maximale energie-efficiëntie te bereiken tussen $2-3\tau$ settling. Voor het tweede ontwerpgeval waren zowel de SNR als de versterkingsfout van de residuversterker beperkt tot de gewenste specificatie voor alle waarden

van de settlingnauwkeurigheid. Er werd aangetoond dat voor hogere waarden van settling, N (waarbij $N = t/\tau > 2$), het vermogensverbruik van een settelende versterker $N/2$ keer hoger is dan het vermogensverbruik van een integrerende versterker ($t/\tau \approx 0$), voor dezelfde versterking en SNR. Hoewel de bovenstaande analyse het vermogensverbruik van de twee ontwerpgevallen van de residuversterker optimaliseert voor een bepaalde SNR, lijdt de versterkingsnauwkeurigheid in beide gevallen vanwege de lage settlingnauwkeurigheid. Om de residuversterker in staat te stellen op de optimale settlingnauwkeurigheidspunten te werken, moeten ze dus worden ondersteund door versterkings- en vervormingskalibratie, waardoor ze de gewenste versterkingsnauwkeurigheid kunnen bereiken terwijl deze idealiter een lage overhead op het ADC-vermogen en de complexiteit hebben.

Hoofdstuk 4 bespreekt de split-ADC-kalibratietechniek, een volledig deterministische achtergrondkalibratiemethode waarvan is aangetoond dat deze een snelle convergentie bereikt met een relatief lage analoge overhead. In dit Hoofdstuk wordt de split-ADC-techniek uitgebreid voor het kalibreren van de residuversterkervorming in meerdere trappen van een gepijplijnde ADC. Bij de split-ADC-kalibratiemethode wordt de ADC in twee identieke helften gesplitst, waarbij elke halve ADC hetzelfdeingangssignaal digitaliseert. Door er een kalibratie-offset tussen te introduceren, worden de twee halve ADC's gedwongen verschillende conversietrajecten te doorlopen. Door het verschil van de twee digitale uitgangen te nemen, kan hetingangssignaal effectief worden geannuleerd, waarbij de gedigitaliseerde kalibratie-offset achterblijft, en die vervolgens kan worden gebruikt om de ADC-fouten te detecteren. Aangezien de kalibratieoffset kan worden gescheiden van hetingangssignaal zonder enige langdurige decorrelatie methode, kan de split-ADC-kalibratie een snelle convergentie bereiken terwijl deze volledig op de achtergrond werkt.

Het effect van een versterkingsfout en vervorming van de residuversterker wordt bestudeerd in Hoofdstuk 4. De versterkingsfout introduceert discontinuïteiten in het verschilsignaal wanneer een van de halve ADC's overgaat van het ene subbereik naar het andere. De grootte van deze discontinuïteiten is evenredig met de grootte van de versterkingsfout. De oneven-orde vervorming van de versterker injecteert een kwadratische component in het verschilsignaal, en de kromming van deze kwadratische component kan worden gebruikt om de derde-orde vervorming in de residuversterker te schatten. De grootte van het foutsignaal voor zowel versterkings- als vervormingskalibratie is recht evenredig met de kalibratieoffset tussen de gesplitste ADC's. Dit betekent dat een grotere kalibratieoffset zal resulteren in een snellere convergentie,

zij het ten koste van het dynamische bereik van ADC. In het prototype gepresenteerd in Hoofdstuk 4 werd een offset gelijk aan 5.5% van het ingangsbereik van de ADC toegevoegd om de impact op het dynamische bereik van de ADC te beperken tot 0.5dB.

Een belangrijke overweging van de split-ADC-kalibratie is het effect van mismatch tussen de halve ADC's op het verschilsignaal. Elke mismatch in de versterking tussen de twee halve ADC's introduceert een fractie van hetingangssignaal in het verschilsignaal en zorgt ervoor dat het een helling krijgt. Deze helling moet worden gecompenseerd door een van de half-ADC-uitgangen te schalen. Het kalibreren van versterking en vervormingsfouten over meerdere ADC-trappen brengt ook bepaalde uitdagingen met zich mee. Aangezien alle kalibratielussen in wezen op hetzelfde verschilsignaal werken, zijn er niet-orthogonaliteiten tussen deze kalibratielussen die de convergentie van deze lussen kunnen beïnvloeden. Deze niet-orthogonaliteiten kunnen worden opgelost door de kalibratielussen iteratief uit te voeren.

Om de doeltreffendheid van de split-ADC-kalibratietechniek te testen, werd een prototype 12-bits gepijplijnde split-ADC geïmplementeerd. Het prototype ADC bestaat uit negen 1.5-bits trappen gevolgd door een 5-bits fijne ADC. Versterkingskalibratie werd gebruikt om de versterking voor de eerste 6 trappen te corrigeren, terwijl niet-lineariteitskalibratie alleen werd toegepast op de eerste twee trappen. Zowel versterkings- als vervormingsfouten werden gecorrigeerd door middel van digitale nabewerking.

Als voldoende nauwkeurig startpunt is het prototype ADC ontworpen met een residuversterker met een hoge DC-versterking en bandbreedte. De kalibratie werd vervolgens gebruikt om de settlingnauwkeurigheid van de residuversterker te verminderen zonder de ADC-resolutie in gevaar te brengen. Een eentraps stroomspiegelgebaseerde operationele versterker met boosting van de versterking werd gebruikt om een DC-lusversterking van 65 dB te bereiken. Het prototype ADC werd geïmplementeerd in een 40nm 1V digitaal CMOS proces en besloeg een oppervlakte van 0.81 mm^2 , met de digitale nabewerking geïmplementeerd in MATLAB. Met behulp van kalibratie bereikte het prototype ADC 10.5b ENOB met een bemonsteringssnelheid van 195 MS/s bij een vermogensverbruik van 53 mW. Terwijl hij continu op de achtergrond werkte, verbeterde de split-ADC-kalibratie de ADC SFDR met 37dB binnen 70000 samples, samen met een vermindering van het vermogensverbruik van de residuversterkers met 60%.

Het prototype ADC dat in Hoofdstuk 4 wordt gepresenteerd, was gebaseerd op digitale correctie van versterkings- en vervormingsfouten, die beperkingen op het

gebied van dynamisch bereik en nauwkeurigheid hebben, terwijl ze een aanzienlijke overhead toevoegen aan de digitale backend. Hoofdstuk 5 presenteert een versterkertopologie die de correctie van versterkings- en vervormingsfouten direct bij de bron in het analoge domein mogelijk maakt. Integreerbare versterkers met open lus bieden een eenvoudige manier om de versterking door middel van de instelstroom te regelen vanwege hun extreem lage settling. Voor vervormingscorrectie werd een linearisatieprincipe voorgesteld op basis van de degeneratie van een common-source versterker in weak inversion. De degeneratie-impedantie gaat de uitdijende ingangsgm-karakteristiek van een common-source-versterker in weak inversion tegen, en de derde-ordecomponent bleek volledig te zijn geannuleerd voor $R_{deg} = 1/2ng_m$.

De impact van degeneratie op de ingangstransconductantie van de versterker kan worden geëlimineerd door de degeneratie op een common-mode manier toe te passen, terwijl de g_m -variatie nog steeds wordt verbeterd van 50% in een differentieel paar tot <1%. Het common-mode degeneratie (CMD) linearisatieprincipe werd geverifieerd door middel van metingen op een strip-board versterkerversie gebouwd met discrete BJT-componenten die een HD3 van -95dB vertoonden voor een ingangszwaai van 50mV_{ppd} . De vervormingsonderdrukking werd uitgebreid tot 2e en 5e orde vervormingscomponenten door een paar ingangsgestuurde MOS-transistoren in weak inversion te gebruiken als de common-mode staartdegeneratie. Door de instelspanning en de verzwakking van het ingangssignaal dat het degeneratienetwerk aanstuurt af te stemmen, kon de voorgestelde versterkertopologie een THD van <-90dB bereiken voor een ingangszwaai van 300mV_{ppd} en een versterking van 4x.

Hoofdstuk 6 presenteert een prototype 12-bit gepijplijnde split-ADC die is gebouwd rond de voorgestelde op integratie gebaseerde CMD-versterker om de doeltreffendheid van het linearisatieprincipe te testen. Derversterkingsfout van de versterker en niet-lineariteit (even- en oneven-ordevervorming) werden op de achtergrond gedetecteerd via de split-ADC-kalibratietechniek. Om de mismatch tussen de twee halve ADC's te verminderen, werd een aangepaste split-ADC-kalibratiearchitectuur geïmplementeerd. Door hetzelfde ingangsmuster twee keer te digitaliseren over twee afzonderlijke klokcycli, kan de kalibratieoffset in de loop van de tijd worden ingevoerd en kan dezelfde ADC-hardware worden gebruikt voor de twee halve ADC-bewerkingen. Door deze "split-over-tijd"ADC konden de residuversterkers in de prototype-ADC worden gedeeld tussen de twee halve ADC's. Om de energie-efficiëntie van de residuversterkers verder te verbeteren, werden ze gedeeld tussen twee in de tijd verweven ADC-banen.

De 12-bits ADC met pijplijn bestaat uit vijf 3-bits trappen (met één bit als

overbereik) gevolgd door een 4-bits fijne ADC. De voorgestelde CMD-versterker is op een complementaire manier geïmplementeerd om de energie-efficiëntie te verbeteren. Gefabriceerd in 28-nm 1-V digitale CMOS, besloeg de ADC een oppervlakte van 0.22 mm². De ADC-uitgangsdata werden geïmporteerd in MATLAB voor foutdetectie en de kalibratielus werd gesloten door de analoge correctieparameters aan te sturen in de degeneratie van de residuvertersterker. Met een bemonsteringsfrequentie van 280 MS/s behaalde de ADC een ENOB van 10.34b bij een stroomverbruik van 13mW, met een totale Schreier FoM van 164.3dB. Het linearisatieschema, samen met het split-ADC-kalibratieschema dat op de achtergrond werkte, stelde het prototype ADC in staat om > 77dB SFDR te bereiken, terwijl het ten minste 3dB betere Schreier FoM weergaf ten opzichte van andere vergelijkbare state-of-the-art pijplijn-ADC-ontwerpen.

Ten slotte besluit Hoofdstuk 7 het proefschrift door de oorspronkelijke bijdragen samen te vatten en het tweede prototype ADC te vergelijken met state-of-the-art pijplijn-ADC's met vergelijkbare snelheid, resolutie en architectuur. Het prototype ADC bleek een van de beste Figure of Merits (FoM's) te hebben onder de modernste ontwerpen, wat de energie-efficiëntie van de voorgestelde CMD-versterker benadrukt. De kalibratie van versterkervervorming met behulp van de split-ADC-techniek werd ook gezien als een van de snelste kalibratietechnieken voor achtergrondvervorming die in de literatuur wordt vermeld. Hoofdstuk 7 presenteert ook enkele aanknopingspunten voor toekomstig onderzoek, waarvan de meest prominente het gebruik van capacitieve common-mode degeneratie is voor het lineariseren van de residuvertersterker.

Acknowledgments

And finally, the most widely read section of the dissertation ☺

While the technical information in this thesis can be, at times, esoteric, there is something universal about wanting to be a part of a greater whole, and to be grateful for it. We all parse through the acknowledgements of a dissertation because we want to sample the candidate's life during those 5 years (or 11 years in this case), and all its pieces. Count the candidate's friends, develop an admiration towards his advisors and run through a very quantized version of his life and relationships. We do this because an opportunity to feel gratitude is something truly valuable, and looking back, there is a lot to be thankful for.

First and foremost, I would like to thank Klaas Bult, who has played a lot of roles in my life over the last decade – promotor, manager, mentor. His enthusiasm towards circuit design is infectious, and I have been in awe of him ever since that first lecture on analog CMOS design in 2008. When it comes to circuit design – he really is an artist, driven by intuition, mixing transistors like colors on a palette to create something new. He is also a virtuoso communicator and has the ability to distil an idea down to its clearest form. Every conference presentation had to go through several rigorous dry runs till it was absolutely right. And now his words echo in my head every time I make a slide. He is an excellent and patient teacher, and his humility sometimes makes us forget what a titan he is.

Another person who played an important role during the course of this PhD is Prof. Kofi Makinwa. Being “promotor-less” for most of my time as a doctoral candidate, I really appreciate his time and effort in taking up the formalities as my promotor and reviewing the manuscript. During the last two years, whenever I slowed down, it was him tightening the screws that pushed me towards completion.

Frank van der Goes was one of the most crucial pillars of my PhD. As my daily supervisor, he was the go-to guy for technical discussions and questions. Frank is a master at solving problems creatively on his feet, and I really admire how he can view engineering problems through both industry and research perspectives. It was an absolute privilege to be able to work with and learn from him during my time at Broadcom Netherlands.

I also want to thank Prof. Leo de Vreede, Prof. Pieter Harpe, Prof. Nan Sun, Dr. Jan Craninckx and Prof. Wouter Serdijn for taking the time to read the thesis and be a part of my defence committee. I really appreciated the discussions and the feedback I received on my work. A special mention to Prof. Wouter Serdijn, who drastically altered the trajectory of my life by inviting me to the Netherlands for an internship in 2007. He took a chance on me when not many would, and for that I will be eternally grateful.

Having spent nearly a third of my life working at Broadcom Netherlands, I got to develop some close friendships with my colleagues. I would like to thank all of them – Chris, Jan W., Jan M., Xiaodong, Qiongna, Stefano, Han, Zeng, Jiansong, Erol, Jeff and everyone else. They are not just talented engineers but good friends who took me in as an intern and made an engineer out of me. As the names above suggest, Broadcom Bunnik team is an international melting pot and the colorful lunch-time conversations about politics, food and politics of food will forever be etched in my memory. A special thanks to Jeff for synthesizing the fn digital interface for both chips, and Reinier & Ulli for their help with the layout. Doing a PhD in the industry can often be lonely and I would like to thank Shakil for sharing the pain of a very unique PhD journey.

I am grateful to my Ethernovia colleagues for being a part of this new adventure that we embarked on 2 years ago. We went through some challenging times and I feel very lucky to be a part of this incredibly talented and driven team.

Els is probably the most compassionate person I know. She consistently sought ways to make my life easier in Broadcom and Ethernovia, from helping me move houses to being always available for a chat and a hug. You are truly wonderful.

During the last ten years, I was lucky to meet some amazing people, and share some incredible memories, travels and conversations with them. Andrei, Natalia – you both have been my family in Utrecht, and the time we lived together were some of the best years of my life. You introduced me to good wines, weird music, snowboarding and Christmas traditions, and were always with me, whether I was watching horror movies or getting married. You two left an indelible mark on my life in ways nobody else could. A special thanks to Natalia for her photoshop wizardry on the thesis cover.

Brought together through a shared passion for gossip, Iniyavan has been my best friend and my work spouse for the last five years, as we worked together across two companies. Ini understands me better than most people and talking to him has always brought me peace. I'd be lost without him.

Living away from home was made easier by having a nucleus of close Indian friends to share dinners, festivals and cultural references with. Sumeet & Sharmishtha, Raj, Yash, Sridevi, Kajal, Vaibhav – thank you for your friendships, and for providing me refuge from the identity crises arising from straddling two extremely different cultures.

As an engineer, I had mostly lived in a tech bubble till I moved to Utrecht, where I had the fortune to meet people from very diverse backgrounds. Luisa, Vita, Joesry, Jort, Neele – your friendships enriched my life and transformed my perspective. Andrea, thank you for being part of some of the most surreal moments of my life. Mateja and Roman have been wonderful friends & neighbours and a constant source of wines and cakes. Chiara and Joanna – you were wonderful housemates, and my only sense of a PhD social life came through you.

A big part of my twenties was spent with my Lindy Hop fraternity, as I got consumed by the world of swing dancing. Lisa, Maartje, Mechteld, Jojanneke, Jackson, Ellen, Jolien, Malou, Praveen, Dolinde, Sarah, Janna, Heidi – the hours I spent with you dancing, practicing, travelling to festivals and talking about swing music were truly magical. Thank you for being fellow addicts.

My extended family – aunts, uncles and cousins – in India is a tight-knit group that consistently comes together to celebrate and share moments of joy and, lately, to pull together during times of intense grief. You are a part of my fabric and I hope to spend more time with you all. My parent-in-laws have been very encouraging during this unexpectedly long PhD, and I am grateful for their support.

My sister, Ashima, to paraphrase, “is a force of nature. At her core, all the strength of the universe.” Her generosity and spirit have always stood out, and she has a limitless capacity for empathy. In a life spanning across three different continents during the course of this PhD, rarely have I seen her take the easy way. She has been my moral compass and a guide, and one could not hope for a better sister.

My parents instilled in me a sense of fortitude and patience, traits that proved invaluable during this long and hard journey. They supported me as I moved away from home to pursue a path that has undoubtedly been hard on them, and for that I am forever grateful. I carry your lessons in my heart, and I hope this feat makes you proud.

My two kids – Yuki and Nilay – have been an abiding source of joy and happiness. And distraction. Yuki, who will never realize it, taught me to follow my nose and stretch often, and to always live in the moment. And being around Nilay is like a

continuous time-travel loop, as I venture from being a parent to a child myself.

Did I mention how long this PhD lasted? Creating tension, like a note held too long. And during all those years, Sarah has kept me going, holding down the fort with a baby, a dog and a husband often physically attached to his laptop. Every word in this thesis was carried on the back of her love and support. She is a supermom, and her sacrifices are reflected not just in this thesis but also in the health and happiness of our kids. Sarah, you are so much stronger than you realize. We have shared some of the most thrilling moments of our lives, and I look forward to embarking on new adventures with you.

List of Publications

Journal Papers

- **R. Sehgal**, F. van der Goes and K. Bult, "A 12b 53mW 195MS/s pipelined ADC with 82dB SFDR using Split-ADC calibration," *IEEE J. Solid-State Circuits*, vol. 50, no. 7, pp. 1592–1603, Jul. 2015.
- **R. Sehgal**, F. van der Goes and K. Bult, "A 13-mW 64-dB SNDR 280MS/s pipelined ADC using linearized integrating amplifiers," *IEEE J. Solid-State Circuits*, vol. 53, no. 7, pp. 1878–1888, Jul. 2018.
- M. Akter, **R. Sehgal**, F. van der Goes and K. Bult, "A 66-dB SNDR pipelined split-ADC in 40-nm CMOS using a class-AB residue amplifier," *IEEE J. Solid-State Circuits*, vol. 53, no. 10, pp. 94–97, Oct. 2019.
- M. Akter, **R. Sehgal** and K. Bult, "A Resistive degeneration technique for linearizing open-loop amplifiers," *IEEE Trans. Circuit and Systems II*, vol. 67, no.11, pp. 2322–2326, Nov. 2020.

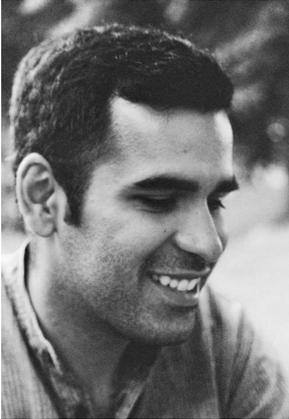
Conference Papers

- **R. Sehgal**, F. van der Goes and K. Bult, "A 12b 53mW 195MS/s pipelined ADC with 82dB SFDR using Split-ADC calibration," *Proc. IEEE ESSCIRC*, pp. 67–70, 2014.
- **R. Sehgal**, F. van der Goes and K. Bult, "A 13-mW 64-dB SNDR 280MS/s pipelined ADC using linearized open-loop class-AB amplifiers," *Proc. IEEE ESSCIRC*, pp. 131–134, 2017.
- M. Akter, **R. Sehgal**, F. van der Goes and K. Bult, "A 66-dB SNDR pipelined split-ADC using a class-AB residue amplifier with analog gain correction," *IEEE J. Solid-State Circuits*, vol. 53, no. 10, pp. 94–97, Oct. 2019.
- K. Bult, M. Akter and **R. Sehgal**, "High-efficiency residue amplifiers," *Proc. AACD*, Edinburgh, May 2018.

Book Chapters

- K. Bult, M. Akter and **R. Sehgal**, "High-efficiency residue amplifiers," *Low-Power Analog Techniques, Sensors for Mobile Devices, and Energy Efficient Amplifiers*, pp. 253-296, Springer, 2019.

About the Author



Rohan Sehgal received the Bachelor of Engineering degree in Electronics and Communications from Delhi University, India, in 2007, and the M.Sc. degree (*cum laude*) from TUDelft, the Netherlands, in 2010. From 2010 to 2015, he worked at Broadcom Netherlands as a doctoral candidate in collaboration with TUDelft.

He joined Broadcom Netherlands in 2016 as a Senior Design Engineer and was involved in analog and mixed-signal circuit design for wireline applications. Since 2019, he has been working at Ethernovia as a Senior Member of Technical Staff, developing high-speed analog front-ends for automotive applications.

He received the Best Thesis Award in 2007, and the NXP fellowship during 2008–2010.