

DOCTORAL THESIS

Design of a Highly Efficient Analog-to-Digital Converter for Wi-Fi 7 Devices

Vahur Kampus

TALLINN UNIVERSITY OF TECHNOLOGY
DOCTORAL THESIS
46/2025

Design of a Highly Efficient Analog-to-Digital Converter for Wi-Fi 7 Devices

VAHUR KAMPUS



TALLINN UNIVERSITY OF TECHNOLOGY
School of Information Technologies
Thomas Johann Seebeck Department of Electronics

This dissertation was accepted for the defence of the degree of Doctor of Philosophy in Electronics and Telecommunications on 06/06/2025

Supervisor: Professor Emeritus Toomas Rang
Thomas Johann Seebeck Department of Electronics
Tallinn University of Technology
Tallinn, Estonia

Co-supervisor: Head of the Department Doctor Alar Kuusik
Thomas Johann Seebeck Department of Electronics
Tallinn University of Technology
Tallinn, Estonia

Opponents: Dean of School Professor Jussi Ryyänen
School of Electrical Engineering
Aalto University
Espoo, Finland

Professor Márta Rencz
Department of Electron Devices
University of Technology and Economics Budapest
Hungary

Defence of the thesis: 27/06/2025, Tallinn

Declaration:

Hereby I declare that this doctoral thesis is my original development and achievement, submitted for the doctoral degree at Tallinn University of Technology and has not been submitted for any other academic degree.

Vahur Kampus

signature

Copyright: Vahur Kampus, 2025
ISSN 2585-6898 (publication)
ISBN 978-9916-80-326-4 (publication)
ISSN 2585-6901 (PDF)
ISBN 978-9916-80-327-1 (PDF)
DOI <https://doi.org/10.23658/taltech.46/2025>
Printed by Koopia Niini & Rauam

Kampus, V. (2025). *Design of a Highly Efficient Analog-to-Digital Converter for Wi-Fi 7 Devices* [TalTech Press]. <https://doi.org/10.23658/taltech.46/2025>

TALLINNA TEHNIKAÜLIKOOL
DOKTORITÖÖ
46/2025

Pidevaja Sigma-Delta muunduri disain Wi-Fi 7 seadmetele

VAHUR KAMPUS



Contents

List of Publications	8
List of Patents.....	9
List of Abbreviations	10
Introduction and Motivation.....	13
Problem Statement	15
Statement of Novelty and Author’s Contribution	16
Thesis Outline.....	17
1 IEEE 802.11 Standards (Wi-Fi)	18
2 Wi-Fi Receiver Chain (RX Chain).....	20
3 Analog to Digital Converters	24
3.1 Flash ADCs.....	24
3.2 Pipeline ADCs	25
3.3 Sigma-Delta ADCs	26
3.4 SAR ADCs.....	26
3.5 Other ADCs	27
4 Wi-Fi 7 and State of the Art Converters	28
4.1 RX Chain	28
4.2 Design Considerations for the Optimal ADC Type of Wi-Fi 7	29
4.3 Deriving the Specifications for the ADC	32
5 Sigma-Delta Conversion	35
6 Design of the CTΣΔADC	38
6.1 Initial Draft of the CTΣΔADC Specifications.....	38
6.2 Coefficient Development	40
6.3 Top-level Design of the ADC.....	46
6.4 Design of the Loopback Filter	47
6.4.1 Design Considerations for the Filter	47
6.4.2 The Design of Integrators	48
6.4.3 Design of Integrator1.....	50
6.4.4 Design of Integrator2.....	58
6.4.5 Design of Integrator3.....	62
6.4.6 Design of the Resonator	65
6.4.7 Summary of the Designed Filter	66
6.5 Design of the Quantizer	67
6.5.1 Function of the Quantizer	67
6.5.2 Comparator	68
6.5.3 Reference Generation and Pre-Amplifier	69
6.5.4 Latch and Buffers	69

6.5.6 Signal Propagation Challenges.....	71
6.5.7 Optimization of the Quantizer.....	73
6.5.8 Quantizer Calibration.....	74
6.5.9 Switching between 5-bit/6-bit Modes.....	75
6.5.10 Simulation Results.....	75
6.6 Design of the Feedback DAC.....	77
6.6.1 Design Considerations for the Feedback DAC.....	77
6.6.2 NRZ and RTZ Coding Schemes.....	78
6.6.3 Linearization of the Feedback DAC.....	79
6.6.4 Design of the feedback DAC.....	80
6.6.5 Reference Generation.....	81
6.6.6 DAC Calibration.....	82
6.6.7 Unit Cell Operations.....	84
6.6.8 Logic for DAC.....	85
6.6.9 6-bit/5-bit Modes.....	86
6.6.10 Thermal to Binary.....	86
6.6.11 Simulation Summary.....	87
6.7 Design of the Biasing.....	89
6.8 Design of the Master Clocking.....	91
6.9 Design of the Calibration Clocking.....	92
6.10 Design of the ELD DAC.....	93
6.11 Design of the Overload Recovery.....	95
6.12 Design of the Overload Masking.....	96
6.13 Design of the Calibration Machine.....	96
6.13.1 Calibration Machine Functionality.....	96
6.13.2 Calibration Impact on the QT.....	98
6.14 Design of the Digital Interface.....	98
6.15 Design of the Power Supply.....	99
6.15.1 The Need for On-Chip Power Supplies.....	99
6.15.2 Options of the LDO.....	101
6.15.3 LDO Topology.....	102
6.15.4 Error Amplifier Topology.....	104
6.15.5 Stability of the LDO.....	105
6.15.6 Efficiency of the LDO.....	107
6.15.7 Noise of the LDO.....	107
6.15.8 PSRR of the LDO.....	109
6.15.9 Extra Features of the LDO.....	110
6.16 Design of the Active Capacitor.....	111
6.17 Design of the VCO.....	112

6.17.1 Background of VCO-based Time Encoding ADCs	112
6.17.2 Functionality of VCO-based ADCs	113
6.17.3 VCO Design	114
6.17.4 Integration of the VCO-mode	115
6.17.5 Leslie & Singh Filter	116
6.17.6 Usage of the VCO-mode	117
6.18 Design of the Programmable Gain Controller	117
7 Optimization and Validation of the Design	121
7.1 Optimization of the Coefficients	121
7.2 Validation of the $\Sigma\Delta$ ADC	121
8 Layout and Verification	125
8.1 Layout of the $\Sigma\Delta$ ADC	125
8.2 Verification Methodology	126
8.3 Verification of the $\Sigma\Delta$ ADC	126
9 Testing and Measurements of the ADC	130
9.1. Measurements of the Standalone $\Sigma\Delta$ ADC Module	130
9.2. Measurements of the RX Chain	131
9.3. Further Nuggets from Measurements	133
10 Conclusion	135
10.1 Summary	135
10.2 Future Research and Wi-Fi 8	137
10.3 Experimental Features	138
10.4 Final Words	139
References	140
Acknowledgements	143
Abstract	144
Kokkuvõte	145
Appendix	147
Publication I	147
Publication II	151
Publication III	157
Publication IV	163
Publication V	169
Publication VI	177
Publication VII	185
Curriculum Vitae	203
Elulookirjeldus	204

List of Publications

Parts of the work of this thesis is based on the following publications. Copies of these publications can be found in the appendix of the thesis.

- I Michael Fulde, Gernot Babin, Christoph Duller, Simon Gruenberger, Harun Habibovic, **Vahur Kampus**, Gerhard Knoblinger, Christian Krassnitzer, Franz Kuttner, Edwin Thaller, Davide Ponton, Andreas Santner, “Innovative RFDAC concepts for digital multi-mode transmitter in cellular applications” 9th Conference on Ph.D. Research in Microelectronics and Electronics (PRIME), Jun. 2013 - *Co-author. Used in the active capacitor design as described in thesis Chapter 6.16.*
- II **Vahur Kampus**, Toomas Rang, “A smart capless voltage regulator for very high bandwidth A/D and D/A converters in a standard 28nm CMOS process”, 15th Biennial Baltic Electronics Conference, Oct. 2016 – *Main author. Used for the supply design as described in thesis Chapter 6.15.*
- III Michael Fulde, Alexander Belitzer, Zdravko Boos, Michael Bruennert, Jonas Fritzin, Hans Geltinger, Marcus Groinig, Daniel Gruber, Simon Grünberger, Thomas Hartig, **Vahur Kampus**, Boris Kapfelsberger, Franz Kuttner, Stephan Leuschner, Thomas Maletz, Andreas Menkhoff, Jose Moreira, Alan Paussa, Davide Ponton, Harald Pretl, Daniel Sira, Ulrich Steinacker, Nenad Stevanovic, “A Digital Multimode Polar Transmitter Supporting 40MHz LTE Carrier Aggregation in 28nm CMOS”, IEEE International Solid-State Circuits Conference, Feb. 2017 – *Co-author. Responsible for supply concept. Used in the active capacitor design as described in thesis Chapter 6.16.*
- IV **Vahur Kampus**, Toomas Rang, Daniel Knaller, Christian Fleischhacker, Markus Korak, Jozef Kiss “A fully differential, 200MHz, programmable gain, level-shifting, hybrid amplifier/power combiner/test buffer, using pre-distortion for enhanced linearity”, 14th Conference on Ph.D. Research in Microelectronics and Electronics, Aug. 2018 – *Main author. Used in the PGA design as described in thesis Chapter 6.18.*
- V **Vahur Kampus**, Martin Trojer, Robert Teschner, “Unleashing the full power of feed-forward opamps: a 200MHz, fully differential, conditionally stable, 36dB gain PGA, using a four-stage multi-path 2.5V amplifier with double feed-forward compensation”, IEEE Nordic Circuits and Systems Conference, Oct. 2018 – *Main author. Used for the filter design as described in thesis Chapter 6.4.*
- VI **Vahur Kampus**, Robert Teschner, Ulrich Gaier, Thomas Linder, Gerhard Nössing, Martin Trojer, “Five-Stage, Power Efficient, Dual Rail, 100MHz, 10dB Programmable Gain Amplifier with Down-Stepping Functions in 28nm CMOS”, IEEE International Symposium on Circuits and Systems, May 2019 – *Main author. Used for the filter design as described in thesis Chapter 6.4.*
- VII Leidy Mabel Alvero-Gonzalez, Eric Gutierrez, Luis Hernandez, **Vahur Kampus**, Viktor Medina, Susana Paton, “Ring-Oscillator with Multiple Transconductors for Linear Analog-to-Digital Conversion”, MDPI Electronics Journal, Jun. 2021 – *Co-author. Responsible for the enablement of the VCO low power mode and overall integration into the ADC. Used in the design as described in thesis Chapter 6.17.*

List of Patents

- I Antonio Di Giandomenico, **Vahur Kampus**, Sergio Walter, Alexander Kahl, Steffen Trautmann, “Multi-purpose receiver chain for WiFi applications” US Patent number: 11115066, Filed: Dec. 2016 (Granted: Sep. 2021)
- II Antonio Di Giandomenico, **Vahur Kampus**, 2018, “SIGMA-DELTA-WANDLER” European Patent application number: 16205911.7, Filed: Dec. 2016
- III Antonio Di Giandomenico, **Vahur Kampus**, Sergio Walter, Alexander Kahl, Steffen Trautmann, “Multi-purpose receiver chain for wifi applications” US Patent number: 11716102, Filed: Sep. 2021 (Granted: Aug. 2023)

List of Abbreviations

ADC	Analog-to-Digital Converter/Conversion
Adj	Adjacent
AGC	Automatic Gain Control
Amp	Amplifier
Alt1	Alternate 1
Alt2	Alternate 2
AP	Access Point
APU	Accelerated Processing Unit
AR	Augmented Reality
BB	Baseband
BER	Bit Error Rate
BOM	Bill of Materials
BPSK	Binary Phase-shift keying
BW	Bandwidth
CM	Common Mode
CPK	Process Capability Index
CT	Continuous Time
DAC	Digital-to-Analog Converter/Conversion
DEM	Dynamic Element Matching
DFM	Design for Manufacturability
DNL	Differential Non-Linearity
DPD	Digital Pre-Distortion
DR	Dynamic Range
DT	Discrete Time
EHT	Extremely High Throughput
ELD	Excess Loop Delay
EUVL	Extreme Ultraviolet Lithography
EVM	Error Vector Magnitude
FB	Feed-Back
FF	Feed-Forward
FFT	Fast Fourier Transform
FOM	Figure of Merit
GBW	Gain Bandwidth
gm	Transconductance
GSM	Global System for Mobile
HD2	Second Harmonic Distortion
HD3	Third Harmonic Distortion
Ids	Drain-Source Current
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
ISI	Inter Symbol Interference
I/Q	In-Phase/Quadrature
LAN	Local Area Network
LDO	Low Dropout Regulator
LNA	Low Noise Amplifier
LO	Local Oscillator

MCS	Modulation Coding Scheme
MLO	Multi-Link Operation
MTPR	Missing Tone Power Ratio
NF	Noise Figure
NRZ	Non-return to Zero
NTF	Noise Transfer Function
OEM	Original Equipment Manufacturer
OFDM	Orthogonal Frequency-division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
OoB	Out-Of-Band
Opamp	Operational Amplifier
opa1	First Operational Amplifier
opa2	Second Operational Amplifier
opa3	Third Operational Amplifier
OSR	Over Sampling Ratio
PCB	Printed Circuit Board
PFM	Pulse-frequency Modulation
PGC	Programmable Gain Controller
PL	Phase Loss
PM	Phase Margin
PSD	Power Spectral Density
PSR	Power Supply Rejection
PTAT	Proportional to Absolute Temperature
PWM	Pulse-width Modulation
QAM	Quadrature Amplitude Modulation
QT	Quantization/Quantizer
RF	Radio Frequency
RTZ	Return to Zero
RX	Receiver
SAR	Successive Approximation Register
SDR	Signal to Distortion ratio
SNDR	Signal to Noise and Distortion ratio
SNR	Signal to Noise ratio
Spec	Specification
STA	Station
STF	Signal Transfer Ratio
SoC	System on Chip
TDM	Time Division Multiplexing
THD	Total Harmonic Distortion
TTM	Time to Market
TWT	Target Wake-up Time
TX	Transmitter
UGF	Unity Gain Frequency
UHF	Ultra High Frequency
VCO	Voltage Controlled Oscillator
VDD	Supply Voltage
VR	Virtual Reality
VSS	Ground Voltage

WPA	Wi-Fi Protected Access
$\Delta\Sigma$	Delta-Sigma
$\Sigma\Delta$	Sigma-Delta

Explanations of abbreviations used in the thesis.

Introduction and Motivation

Throughout history, communication methods have continuously evolved. From rudimentary smoke signals and early days of homing pigeons, going through the invention of postal services and paper mail to the days of telegraph and radio, mankind has relentlessly pursued faster and more information-dense communication channels. With the advent of the telecommunications age, data transmission has witnessed exponential growth. As we delve deeper into the digital era, the demand for high-speed, reliable, and efficient wireless radio communication only intensifies. Mankind has always sought ways to find ways to communicate faster and pack more information into the surrounding medium.

The year 2024 marks the introduction and year 2025 the roll out of the latest iteration of the Wi-Fi standard, IEEE 801.11be, also known as Extremely High Throughput (EHT) or Wi-Fi 7 [1]. Wi-Fi 7 delivers significant speed improvements over its predecessors, boasting theoretical maximum speeds of up to 46Gbps, far exceeding Wi-Fi 6 capabilities. This translates to faster downloads, seamless 8K video streaming, and efficient data transfers, ensuring smooth operation of high-bandwidth applications without bottlenecks [2].

Wi-Fi 7 leverages the full potential of the 6GHz band, offering wider channels ranging from 40MHz to 320MHz. The ability to utilize the newly available 6GHz band, alongside existing 2.4GHz and 5GHz bands, reduces congestion and interference, especially in densely populated areas. This capability is vital for environments like apartment complexes, stadiums, and large office buildings where numerous networks can overlap. Additionally, Wi-Fi 7 introduces Multi-Link Operation (MLO), allowing devices to operate on multiple frequency bands (2.4GHz, 5GHz, and 6GHz) simultaneously, significantly boosting network capacity and efficiency. This ensures optimal connectivity for everyone, even in homes and workplaces with numerous connected devices. With the potential for connecting a significantly higher number of devices to a single access point, Wi-Fi 7 paves the way for broader adoption of Internet of Things (IoT) [3].

Wi-Fi 7 prioritizes drastically reduced latency, a critical factor for applications like online gaming, virtual reality (VR) industrial automation, and augmented reality (AR) [4]. Lower latency translates to faster response times and a smoother user experience. This improvement is particularly beneficial for industries relying on real-time communication and remote collaboration, such as telemedicine and remote work setups. Wi-Fi 7 also incorporates advanced technologies like Orthogonal Frequency Division Multiple Access (OFDMA) and Target Wake Time (TWT). OFDMA enhances data transmission efficiency by allowing multiple devices to share a channel simultaneously, while TWT improves battery life for connected devices by scheduling communication times. These innovations contribute to a more reliable and energy-efficient networks.

Moving to Wi-Fi 7 can be seen as a strategic move to future-proof network infrastructure. As technology evolves and demands for higher performance increase, a Wi-Fi 7 network ensures preparedness to handle upcoming advancements. This forward-thinking approach minimizes the need for frequent upgrades and investments in the near future. On top, Wi-Fi 7 prioritizes data privacy and security with advanced encryption protocols, such as Wi-Fi Protected Access 3 (WPA3) and WPA4 and improved security measures to safeguard networks from potential threats.

Wi-Fi 7 represents a significant advancement in wireless networking technology, delivering unprecedented performance across performance, efficiency, and reliability

metrics. This new standard establishes a comprehensive framework to address the escalating connectivity requirements of the digital transformation era, serving diverse applications from enterprise operations to consumer needs. Notably, Wi-Fi 7's enhanced capabilities can substantially reduce the burden on cellular networks by efficiently offloading data transmission and even voice calling to high-capacity Wi-Fi infrastructure, potentially offering superior throughput with lower network congestion compared to 4G, 5G, and emerging 6G networks.

The adoption of new amendments to communication standards, such as in Wi-Fi, invariably presents **challenges** for device manufacturers in meeting the stipulated requirements. The adoption of Wi-Fi 7 is no exception. The significant advancements, including wider channels, reduced latency, support for more complex coding schemes, and the utilization of new spectrum bands necessitate engineers to overcome increasingly complex hurdles in delivering capable products to end users. The challenges extend beyond just the need for increased Signal to Noise and Distortion ratio (SNDR) and bandwidth (BW) across all modules. Achieving the desired performance necessitates advancements in also power efficiency and footprint to comply with environmental regulations, green deal goals and user expectations for compact devices.

At the core of any over-the-air transceiver module are the converters, that translate signals between the digital domain, necessary for signal processing, and the analog domain, required for transmission and reception. A transceiver combines both a transmitter and a receiver within a single package. The transmitter, comprising from a signal chain, handles signal encoding and transmission to the medium, while the receiver, another signal chain, is responsible for receiving and decoding the signals from the medium. Despite the significant advancements in digital technology over the past 40 years, the world around us remains stubbornly analog and over-the-air transmission is still done in pure analog domain. The converter in the transmit (TX) chain, which transforms digital signals into analog form, is known as the Digital-to-Analog Converter (DAC). Conversely, the converter that transforms analog signals back into digital form sitting in the receive (RX) chain is called the Analog-to-Digital Converter (ADC).

This work focuses on the design of an ADC optimized for the stringent requirements of the Wi-Fi 7 standard. The design is undertaken within the constraints of real-world industry, where Design for Manufacturability (DFM) and yield are paramount. Consequently, uncompromising quality and reliability are fundamental design considerations. Given the imperative for high yield across varying process tolerances, design choices are dictated by manufacturability rather than solely by performance metrics. There is no chance of selecting just the best parts for sales, so the yield has to stay extraordinarily high regardless of the manufacturing process tolerances. The thesis focuses on achieving very high yields by designing converters capable of more than 3σ specifications compliance. Additionally, design choices are not exclusively driven by ADC performance but also influenced by the broader requirements of the RX chain.

Problem Statement

The Wi-Fi 7 standard requires that the converters are at least capable of handling 160MHz signal BWs, with the entire RX chain being able to decode 12-bit 4k-QAM modulated OFDMA signals. The specification is applicable for the entire RX chain and should not be taken as a stand alone for the ADC.

To attain IEEE 802.11be compliance, a complex interplay of technical, regulatory, and procedural factors must be considered. While the standard outlines comprehensive requirements, the specific quantification of EVM remains somewhat ambiguous, being indirectly correlated to Bit Error Rate (BER) through real-world channel modelling under varying Access Point (AP) and Station (STA) configurations. A general consensus within the industry posits that an EVM of -38dB is sufficient for compliance. Nevertheless, to mitigate certification risks and enhance competitive positioning, manufacturers often strive for superior EVM performance. Moreover, improved EVM correlates with reduced BER in challenging propagation environments, ultimately contributing to enhanced system robustness.

An ADC with superior SNDR minimizes its contribution to overall receiver chain degradation, thereby facilitating compliance with target specifications. However, balancing performance with power consumption is crucial. Given the contemporary trend of integrating four or eight antennas per band (2.4GHz, 5GHz, and 6GHz) within Wi-Fi transceivers, and the industry preference for passive cooling solutions, individual RX chain power consumption must remain significantly below half a watt, inclusive of digital processing, to achieve substantial market penetration.

There are several different types of ADCs available, with each type having many different topologies and countless design choices. The designers need to be both familiar with what is feasible but also what is practical and achievable with the time and budget constraints. With evolving technological capabilities, the boundaries are always shifted further and things not feasible in the past can suddenly become very appealing.

Time-to-market (TTM) is another critical factor in the semiconductor industry. To mitigate supply chain risks, many original equipment manufacturers (OEMs) necessitate dual-source strategies, often $\Sigma\Delta$ ADC involving porting designs between different fabrication facilities and process technologies. ADC architectures with a higher degree of digital integration generally exhibit greater porting flexibility, rendering them advantageous in such scenarios.

Based on the aforementioned problem statement, the following research questions have been formulated:

RQ1: Is it possible to develop a world class, power-efficient, high-performance ADC optimized specifically for Wi-Fi applications, with comprehensive reconfigurability to facilitate various calibration operations across the System on Chip (SoC).

RQ2: Can an ADC be engineered that demonstrates superior performance metrics compared to current market solutions in the Wi-Fi converter domain, specifically targeting excellence in the Schreier Figure of Merit (FOM) as benchmarked against state-of-the-art converters (detailed in Chapter 4, Figure 11).

RQ3: Can it be ensured that the manufactured ADC achieves exceptional production yield exceeding the industry 3σ threshold, with a targeted Process Capability Index (CPK) of 1.33. This stringent requirement is necessitated by the architecture's demand for eight parallel ADCs per Wi-Fi device to meet performance specifications.

Statement of Novelty and Author's Contribution

As a consequence of these multifaceted design constraints, Successive Approximation Register (SAR) ADCs have emerged as the predominant choice among different design houses for Wi-Fi applications. Their compact footprint, substantial digital content, and relatively low power consumption contribute to their widespread adoption. To establish a competitive edge within the industry, a novel approach is required to surpass industry benchmarks in ADC performance while simultaneously reducing footprint and power consumption. The objective is to develop an ADC that outperforms current state-of-the-art solutions in this field used by the competition.

While SAR ADCs have been the focus of extensive research and development over the past decade, with numerous advancements in digital algorithms presented annually at conferences, their efficiency still lags behind that of Sigma-Delta ($\Sigma\Delta$) converters. SAR ADCs do offer the additional option of interleaving to increase potential BW, a capability lacking in $\Sigma\Delta$ ADCs. However, within the target frequency band, $\Sigma\Delta$ ADCs generally exhibit superior efficiency and performance. Furthermore, Continuous-Time (CT) $\Sigma\Delta$ ADCs eliminate the need for a high-speed input buffer, a critical component in SAR and Discrete-Time (DT) $\Sigma\Delta$ ADC architectures. The often-overlooked power consumption and complexity of input buffers can distort performance comparisons, presenting a more favourable view of SAR structures than is accurate compared to $\Sigma\Delta$ ADCs.

This work presents a fully custom 6-bit, third-order CT $\Sigma\Delta$ ADC, capable of supporting all features needed for the Wi-Fi 7 platform, including support for the much-coveted 320MHz channels. Outperforming competing solutions in signal-to-noise ratio (SNR) and power consumption, the ADC also includes features that anticipate the demands of the forthcoming Wi-Fi 8 standard, with the potential compatibility with 480MHz channels. The feedback loop incorporates three custom-designed high-speed multi-path operational amplifiers (opamps). The 6-bit quantizer (QT) employs a one-time foreground calibration to enhance Differential Non-Linearity (DNL) accuracy, while the feedback DAC utilizes a background calibration scheme to precisely match all the DAC cells. Finally, a novel low-power Voltage-Controlled Oscillator (VCO) mode is integrated for even more efficient idle-mode operation and channel scanning functions.

This design represents an evolution of the commercially hugely successful CT $\Sigma\Delta$ ADC architecture, originally developed for Intel Corporation's Wi-Fi 6 access point product lineup. While the foundational Wi-Fi 6 implementation was developed, by the author, as a single-contributor effort, the Wi-Fi 7 iteration leveraged collaborative expertise, incorporating research contributions from the University of Madrid and coordination across a small specialized design team. Building upon the established Wi-Fi 6 ADC platform, the Wi-Fi 7 version incorporates several critical enhancements to meet next-generation wireless communication requirements. Primary improvements include a doubling of integrator bandwidth capacity and an upgrade from 5-bit to 6-bit quantization resolution. These architectural refinements enabled a systematic design progression that maintained both performance quality and development efficiency throughout the implementation cycle.

The work has been designed, manufactured, measured and tested as a part of the forthcoming MaxLinear Wi-Fi 7 MxL31 product line.

Thesis Outline

This thesis comprises an introduction, nine chapters, and a conclusion. The first five chapters establish foundational knowledge to familiarize readers with the essential aspects of the topic. Chapter six constitutes the core analysis of the thesis, while chapters seven through nine provide post-design analysis and synthesis of simulations and measurements.

Pre-Design – Background information needed for context

Chapter 1

Introduction to the IEEE 802.11 standard. Evolution from its humble early beginnings to the latest iterations of the standard.

Chapter 2

Explanation of the rudimentary terms and building blocks needed to construct a basic receiver chain. Introduction to wanted signals and blockers present in Wireless systems.

Chapter 3

Description of different ADC architectures, the plusses and minuses of different structures and the usage of different architectures.

Chapter 4

The usage and selection of different types of ADCs plus the state-of-the-art design of ADCs for applications sitting in similar space to Wi-Fi 7 converters.

Chapter 5

History of the Sigma-Delta Conversion and a quick overview of the basic $\Sigma\Delta$ operation.

Core Design – Main Body of the thesis

Chapter 6

The design of the complete CT $\Sigma\Delta$ ADC, including (but not limited to) the Filter, the QT, the DAC, the Biasing, the VCO, the PGA and the Supply.

Post-Design – Interpretation of the results

Chapter 7

Validation and top-level simulations of the designed CT $\Sigma\Delta$ ADC.

Chapter 8

Layout and Verification simulations of the designed CT $\Sigma\Delta$ ADC.

Chapter 9

Measurement results of the designed CT $\Sigma\Delta$ ADC.

Conclusions

Contains various points, including conclusions of the work, summary of the claims and suggestions for potential developments in the future.

1 IEEE 802.11 Standards (Wi-Fi)

This chapter covers the following topics:

- History of Wi-Fi in wireless communications.
- The evolution of the IEEE 802.11 standard.

The foundation for Wi-Fi can be traced back to the 1960s with the pioneering work on wireless packet networks. Driven initially by the University of Hawaii, a wireless packet network called ALOHAnet was first developed to connect the Hawaiian Islands in 1971. This innovative system served as the first wireless packet network specifically designed to connect the Hawaiian Islands. ALOHAnet utilized two unlicensed channels at 407.350 MHz and 413.475 MHz within the Ultra High Frequency (UHF) band at 100kHz BW, operating on separate frequencies for uplink and downlink communication [5]. This groundbreaking technology laid the groundwork for numerous future advancements, including industry standards like Global System for Mobile (GSM) for cellular networks, Ethernet for wired local area networks, and ultimately, Wi-Fi itself.

Technological advancements and the crucial allocation of unlicensed radio bands in the 1980s paved the way for large-scale wireless communication by the 1990s. NCR Systems Engineering, a pioneer in promoting wireless Local Area Networking (LAN) technology, introduced WaveLAN in 1991 as a precursor to the 802.11 standard. This innovative system used two unlicensed frequency bands: a single channel in the 900MHz band and eight channels in the 2.4GHz band. These bands, along with the 5GHz band, are collectively referred to as “Garbage Bands” due to their prior utilization for various applications, including microwave ovens, which employ radio waves for cooking. In collaboration with the Institute of Electrical and Electronics Engineers (IEEE), NCR standardized the WaveLAN protocol in 1997, resulting in the first iteration of the IEEE 802.11 standard [6]. This initial version offered a maximum throughput of 2Mbps but was quickly superseded by the more widely adopted IEEE 802.11b [7] standard, which boasted transfer rates of up to 11Mbps. The subsequent IEEE 802.11g [8] standard further enhanced speeds up to 54Mbps. The standardized IEEE 802.11 protocol utilized up to 14 overlapping 20MHz channels within the 2.4GHz band (typically limited to 13 channels in most regions). While the original standard and its 802.11b successor employed techniques like direct-sequence spread spectrum and frequency hopping spread spectrum for transmission, the 802.11g standard and all subsequent iterations have incorporated various forms of orthogonal frequency-division multiplexing (OFDM) modulations.

The limitations of 20MHz channels in the 2.4GHz band for maximum throughput soon became increasingly evident. To address this bottleneck, a need for higher bandwidth channels arose. Since the 2.4GHz band offered limited options for expansion, a new 5GHz unlicensed band was introduced for Wi-Fi communication. The IEEE 802.11a standard, while utilizing the same coding scheme and channel widths as 802.11g, did not offer immediate throughput advantages. However, it pioneered the use of the new unlicensed band, paving the way for future advancements. This was soon followed by the IEEE 802.11n amendment, which significantly increased channel width (up to 40MHz) and introduced the potential for much faster throughput (up to 600Mbps).

By 2014, the proliferation of standards and confusing nomenclature within the 802.11 family necessitated a more user-friendly approach. The industry adopted a new branding strategy, naming the IEEE 802.11ac standard, offering 80MHz channels and potential speeds of 7Gbps in the 5GHz band, as Wi-Fi 5. This trend continued with the IEEE 802.11ax standard, branded Wi-Fi 6, which boasts transfer rates up to 9.6Gbps and

introduces an even newer amendment with the use of the 6GHz unlicensed band. This additional band allows for further channel width expansion, overcoming the limitations of the 5GHz band, which presently is capped out at 160MHz. Past standards like 802.11b and 802.11g were also given backdated names to further simplify the Wi-Fi roadmap and nomenclature (Figure 1).

2024 was poised as the introduction of the latest amendment to the IEEE 802.11 Wi-Fi standard portfolio, designated as 802.11be and commercially branded as Wi-Fi 7. This iteration promises significant advancements, including support for channel widths up to 320MHz in the newly available 6GHz unlicensed band. Additionally, 802.11be leverages better spectrum utilization and 4k-QAM modulation schemes, facilitating lower latency and enabling theoretical data rates of up to 46Gbps. These enhancements hold particular promise for densely populated areas with a high concentration of nodes requiring low latency and concurrent communication, which is especially attractive to WLAN systems because it offers traffic offload for the explosively growing traffic of cellular networks for the possible adoption of IoT [9].

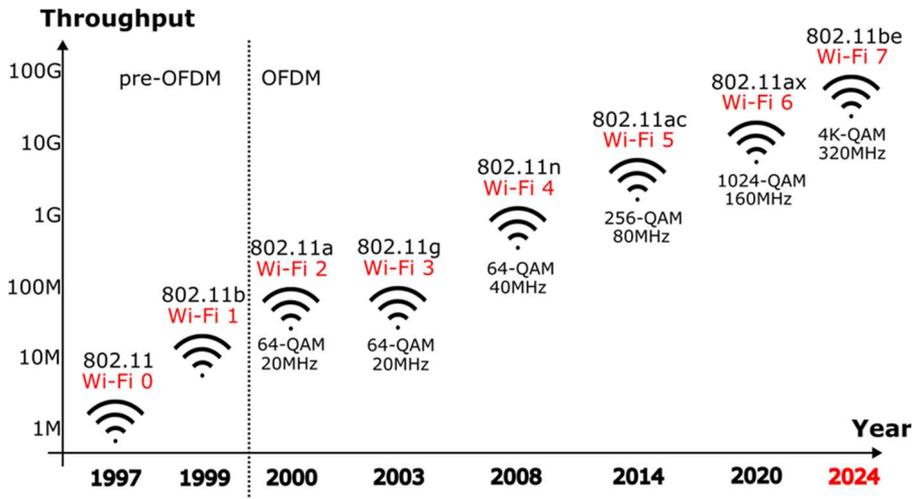


Figure 1. Evolution of Wi-Fi standards.

This work focuses exclusively on solutions designed specifically for the forthcoming Wi-Fi 6GHz band. This band encompasses carrier frequencies, or Local Oscillator (LO) frequencies, ranging from 5935MHz to 7115MHz and supports channel bandwidths between 20MHz and 320MHz. Solutions for other bands will be tailored to their distinct requirements, facilitating reduced complexity and optimized power consumption. Although Wi-Fi 7 certification does not mandate support for the 6 GHz spectrum, it is notable that only the 6 GHz spectrum can accommodate full 320MHz channels. The ADC examined in this thesis has been specifically engineered to address this particular use case.

It is important to note that the IEEE 802.11 standard also encompasses millimeter wave (mmWave) channels, collectively referred to as WiGig. However, these channels are typically excluded from discussions of Wi-Fi generations and will not be further addressed here. While WiGig protocols necessitate support for significantly wider channels compared to standard Wi-Fi, they also function with much simpler coding schemes. This shift in strategy prioritizes high speeds over maximizing performance in WiGig designs and generally uses different type of converters.

2 Wi-Fi Receiver Chain (RX chain)

This chapter covers the following topics:

- An introduction to the Wi-Fi transceiver.
- Potential implementation of the RX chain.

A commonly used transceiver architecture consists of two loosely intertwined half-channels called the Transmitter (TX) chain and the Receiver (RX) chain. The task of the TX chain is to process, convert, and transmit data to the chosen medium. Conversely, the RX chain acquires and decodes data received from the medium. While these chains operate independently for their core functions, they may share certain modules, such as biasing and supply circuits. Careful consideration must be given when selecting the most effective architecture to ensure that shared modules do not introduce interference or compromise performance.

This work focuses on an ADC situated within the RX chain. The ADC's primary function is to decode an analog-domain Wi-Fi signal received through the air as a complex Radio Frequency (RF) domain In-phase (I)/Quadrature (Q) signal to digital domain for easy signal processing. Traditionally, two separate ADCs would be required for independent I and Q conversion. To simplify the design and achieve superior matching between the I and Q paths, an I/Q ADC is employed. This module incorporates two parallel conversion channels while sharing certain sub-blocks, leading to improved power efficiency and enhanced gain accuracy. The designed ADC in this work is of an I/Q type, essentially integrating two ADCs within a single module. Like the considerations for shared modules between TX and RX chains, careful design is crucial to minimize I/Q interference within the I/Q ADC itself.

Wi-Fi signals propagating through the medium exhibit significant diversity. Due to the open nature of most channels, i.e. also other devices and wireless technologies can send and receive in the same frequency range, various scenarios can occur:

- 1) Sparsely populated environments (Suburban or rural areas).
- 2) Densely populated environments (High density areas, such as high-rises, concert halls or stadiums).

In sparsely populated environments, where the desired signal is the only detectable one and no interfering signals (blockers) are present, Out-of-Band (OoB) filtering becomes unnecessary. The sole task is detection and decoding of the transmitted signal. Consequently, the transceiver specifications can be simplified to two key parameters: Maximum Error Vector Magnitude (EVM) and sensitivity. EVM, derived from the Bit Error Rate (BER) mandated by the IEEE for Wi-Fi coding standards, serves as a primary specification. It is a commonly used performance indicator, that measures the average deviation between the transmitted (ideal) constellation points and the received ones (Figure 2). These ideal constellation points are uniformly distributed across the constellation plane, with the objective of hitting them to accurately decode the transmitted data. The smaller the EVM, the closer the data is to the ideal points and more sophisticated the modulation schemes can be employed. For instance, a 16-QAM scheme consists of 16 distinct points and carries a 4-bit word, whereas a 1024-QAM scheme contains 1024 points and carries a 10-bit word. The received points cluster around the target point (blue) and must consistently fall within a designated boundary (red) to ensure accurate decoding. Based on EVM measurements the maximum possible constellation and coding rate, i.e. Modulation Coding Scheme (MCS) for Wi-Fi, can be determined to

guarantee a defined BER. In a sense, EVM can be seen as approximation of the traditional SNR and further related to E_b/N_0 (Energy per bit over Additive white Gaussian noise (AWGN) Power Spectral Density (PSD)). Hence, EVM coming from the BER of the IEEE for Wi-Fi coding standards, serves as a primary specification. For the new Wi-Fi 7 standard, the transceiver must support a 4096-quadrature amplitude modulation (4k-QAM) scheme over a 320MHz channel.

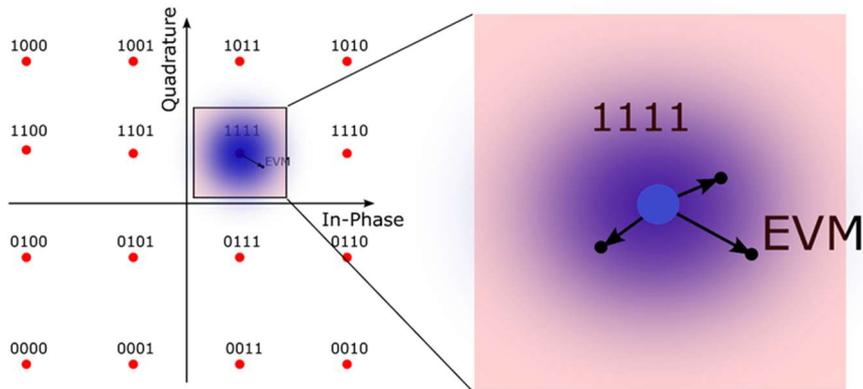


Figure 2. 16-QAM modulation scheme.

Sensitivity, on the other hand, quantifies the minimum received signal power level at which information can still be decoded using the standard's simplest MCS. Sensitivity is thus expressed in dBm, indicating the signal strength required for the receiver to achieve the EVM necessary for supporting a predefined BER using the simplest modulation scheme called Binary Phase-shift keying (BPSK).

In dense environments characterized by numerous active Wi-Fi channels, the potential for interference from blockers (unwanted signals) significantly increases. These blockers can manifest at any frequency, including neighboring channels or even the desired channel itself. To mitigate this interference, OoB filtering becomes a desirable and often necessary component of the RX chain. The primary objective of OoB filtering is to preserve the integrity of the desired signal while attenuating unwanted blockers as much as possible. This allows for the increased amplification of the desired signal, ultimately improving the resolution of the RX chain. While the OoB filtering capability, and hence performance of a receiver in presence of blockers, is not explicitly mandated by the IEEE specifications, Original Equipment Manufacturers (OEMs) may still prioritize its inclusion to gain a competitive edge. As a consequence, this becomes a very important performance criterion for the transceiver system and design.

Designing an optimal RX chain that offers the best performance in both dense and sparsely populated medium presents a significant challenge. The addition of more filtering inevitably leads to increased power consumption and potentially reduced maximum performance due to the introduction of noise from the filtering modules. Conversely, the absence of enough filtering limits the achievable gain due to potential signal overload caused by blockers, thereby compromising the RX chain's maximum potential performance in the presence of interference.

The Wi-Fi standard categorizes interfering signals (blockers) based on their spectral proximity to the desired signal. Unlike the 2GHz band, the 6GHz band eliminates the issue

of overlapping channels. However, improper access point configuration can lead to full or partial channel reuse, potentially introducing interference. The closest interfering signal to the desired signal, in either direction of the frequency spectrum, is classified as an Adjacent (Adj) blocker. Subsequent interfering signals are categorized as Alternate 1 (Alt1) and Alternate 2 (Alt2), etc. respectively. The likelihood of encountering blockers is equally distributed across these neighboring channels.

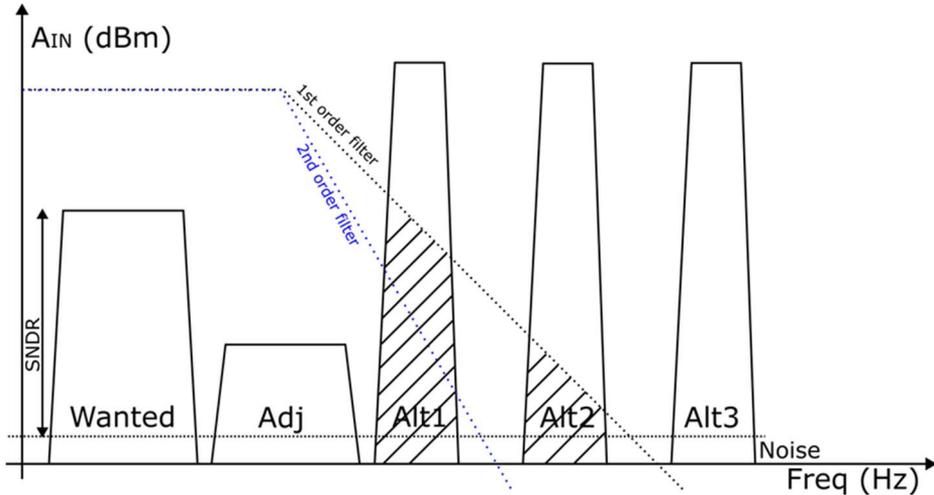


Figure 3. Effects of filtering on unwanted blockers with first order filtering (black) and second order filtering (blue).

Additionally, the blocker’s channel width may differ from the desired signal, leading to either a more concentrated or dispersed power distribution within the spectrum. Blockers can be seen as an unwanted addition to the wanted signal, eating up the available signal headroom and forcing the ADC to operate with a larger-than-necessary safety margin (backoff). While incorporating filtering within the RX chain can help suppress blockers, this strategy comes at the cost of introducing additional noise from the filtering process itself. The decision of incorporating filtering (and how much filtering) within the RX chain (Figure 3) represents a heavy system design trade-off where higher order filtering will lower the overall blockers but will also degrade the overall noise figure. The effectiveness of filtering is directly related to the spectral distance between the blocker and the desired channel. Consequently, Adjacent Channel blockers pose the most significant challenge. Fortunately, the IEEE standard acknowledges the presence of blockers in this specific scenario and relaxes BER and EVM performance specifications somewhat.

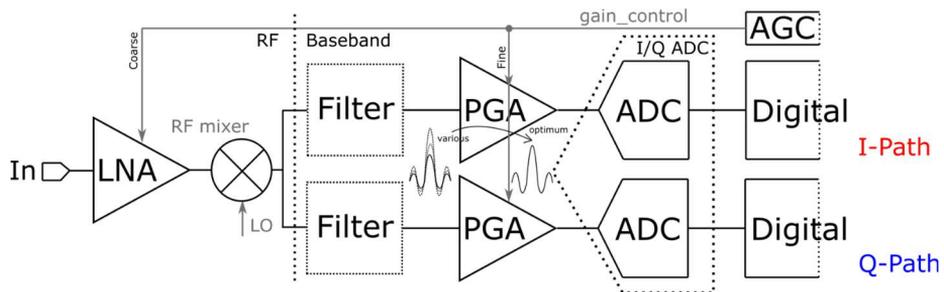


Figure 4. An example of an RX chain comprising of an LNA, Mixer, Filter, PGC (pictured PGA) and ADC.

A typical Wi-Fi RX chain consists of two distinct functional sections: the RF section and the baseband section (Figure 4). The RF section incorporates modules that operate within the RF domain, while the baseband section houses modules that function within the baseband (BB) domain. The RF section ends where the I/Q signal is separated into distinct I- and Q-paths and the BB section needs all modules doubled. The standard RF macro block typically includes at least a Low Noise Amplifier (LNA) and an RF Mixer. The LNA amplifies the weak signal received by the antenna, and the RF Mixer down-converts the amplified RF signal to lower BB frequencies, facilitating further processing. The BB section can encompass various modules, such as a filtering stage, a Programmable Gain Controller (PGC), and the ADC. The LNA typically offers coarse gain control with a limited number of steps, while the PGC provides fine gain control with a higher degree of accuracy. Both the LNA and PGC gain stages are typically controlled by an Automatic Gain Control (AGC) algorithm. The primary objective of RX chain design is to optimize the received signal for the ADC. This optimization entails adjusting the signal strength and eliminating interfering signals (blockers). By ensuring optimal signal conditions, the ADC can perform the desired analog-to-digital conversion process with the highest possible fidelity.

The practices in industry vary, but often modules such as Filters and PGCs are also shared between RX and TX chains to save physical die space in systems that use Time Division Multiplexing (TDM) such as Wi-Fi.

3 Analog to Digital Converters

This chapter covers the following topics:

- Background information about common ADCs used in different applications.
- The high-level description of different ADC architectures.

The past five decades have witnessed a remarkable transformation driven by the pervasiveness of digital data processing. However, the physical world around us remains inherently analog. ADCs play a critical role in bridging these two worlds, acting as essential intermediaries between the continuous realm of analog signals and the discrete nature of the digital domain. These ubiquitous devices are indispensable for processing signals from a wide array of sensors, microphones, cameras, and other analog sources. Selecting the optimal ADC for a specific application necessitates a careful evaluation of various factors, including conversion speed, resolution, power consumption, and cost. This section delves into the different classes of ADCs, exploring their key characteristics and their suitability for diverse applications.

The types of ADCs can be roughly divided into six distinguished subgroups [10]:

- 1) Flash ADCs
- 2) Pipeline ADCs
- 3) Sigma Delta ADCs
- 4) Successive Approximation Register (SAR) ADCs
- 5) Integrating ADCs
- 6) Counter ADCs

3.1 Flash ADCs

Among ADC architectures, the Flash ADC stands out for its simplicity, achieving the highest conversion speeds (up to several gigasamples per second). This remarkable speed stems from its parallel architecture, where numerous comparators operate simultaneously to convert all bits of the analog signal in a single clock cycle. This makes Flash ADCs ideal for applications demanding high-speed signal acquisition, such as oscilloscopes, high-bandwidth data acquisition systems, and advanced radar systems.

A Flash converter can also be cascaded in n -steps (typically no more than three), where conversion is executed in multiple stages. Each subsequent converter performs finer quantization. Figure 5 illustrates a typical two-step Flash converter: the analog input is initially sampled by a sample-and-hold circuit and the digital word obtained is again converted to the analog domain with an internal DAC. During the hold period, the first Flash ADC performs a coarse quantization on the held signal. The held signal is then subtracted from the output of the DAC, and the residue is amplified and passed down to the next stage for fine quantization to the converter's full resolution. Recent advancements have enabled latter stages to sometimes incorporate other ADC architectures, transforming the n -stage Flash converter into an n -stage hybrid-flash converter.

The unparalleled speed of Flash ADCs comes at a cost. Their fully parallel design necessitates significant power consumption. Additionally, calibrating the numerous comparators within the architecture can introduce complexity. Furthermore, Flash ADCs typically exhibit lower resolution compared to other ADC types. Despite these limitations, Flash ADCs remain the preferred choice for ultra-high-speed applications. Their unmatched conversion bandwidth surpasses the capabilities of other architectures, making them indispensable tools for specific high-performance domains.

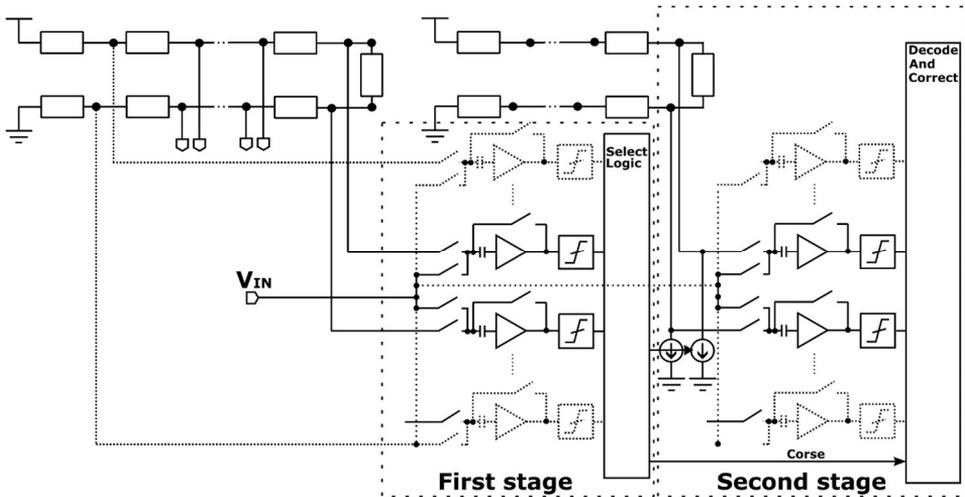


Figure 5. Basic principle of a Two-Step flash ADC.

3.2 Pipeline ADCs

Pipeline ADCs are segmented, with each stage performing a partial conversion as shown in Figure 6. This pipelined architecture enables high conversion speeds (up to Gigabit per second) while maintaining good resolution. The resolution can be always increased at a cost of a higher number of stages. The output codes from each stage are then computed in a digital block that provides the error correction. Pipeline ADCs are also suitable for high-speed data acquisition systems, different types of communication systems, and video processing applications. However, the advantages of pipeline ADCs come with inherent design challenges. Their segmented architecture can be complex, often requiring sophisticated digital calibration techniques to achieve optimal performance. Additionally, the multi-stage pipelined structure can result in higher power consumption compared to other conventional ADC types, as all stages must operate at very high speeds.

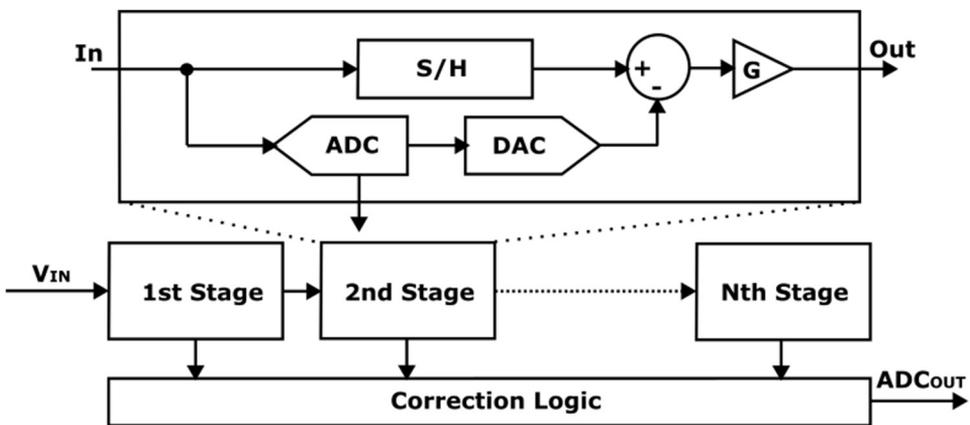


Figure 6. Basic principle of a pipeline ADC.

3.3 Sigma-Delta ADCs

Sigma-Delta ADCs prioritize high resolution over speed by employing oversampling and noise shaping techniques. They utilize a loop filter designed to create different transfer functions for the quantization noise input (Noise Transfer Function, NTF) and the signal input (Signal Transfer Function, STF). The STF typically exhibits an all-pass characteristic, while the NTF shows a high-pass characteristic, effectively implementing noise shaping (Figure 7). This process shifts most of the noise power to higher frequencies, allowing for increased resolution in the band of interest through oversampling [11]. The approach enables them to achieve exceptional resolution even with moderate conversion speeds, typically up to hundreds of MHz signal bandwidth. Their oversampling introduces latency, making them less suitable for real-time applications requiring very high-speed conversion.

There are two primary architectures for Sigma-Delta ADCs: discrete sampling and continuous time. Sampling architectures excel in applications requiring high resolution for low-bandwidth, high-resolution spectrums. Conversely, continuous-time architectures are better suited for scenarios where moderate resolution is desired at higher bandwidths. A significant advantage of Sigma-Delta ADCs is their relatively low power consumption compared to other ADC types. This characteristic makes them ideal for battery-powered devices like mobile phones, where extending battery life is crucial. Sigma-Delta ADCs are especially well-suited for audio processing, precision measurement instruments, and sensor interfaces where high resolution and low noise are critical, but also communication systems for mobile devices where high speed together with acceptable power consumption is necessary.

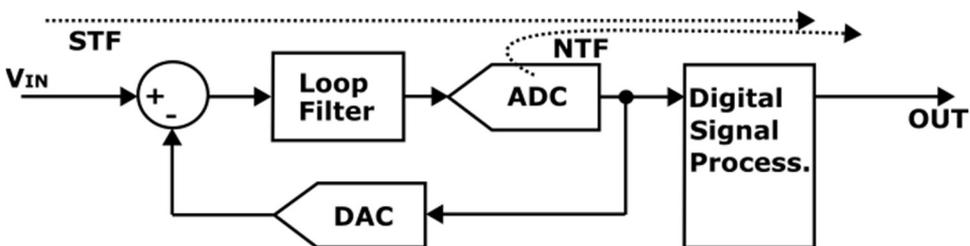


Figure 7. Basic principle of a $\Sigma\Delta$ ADC.

3.4 SAR ADCs

Successive Approximation Register (SAR) ADCs exhibit a commendable equilibrium between speed and resolution through the utilization of a binary search methodology for iterative approximation of the analog input. They boast commendable conversion speeds, reaching up to hundreds of MHz in signal bandwidth, coupled with remarkable power efficiency. The field of SAR ADC development has seen significant advancements in recent years. New algorithms and architectures are continuously emerging and being presented at leading conferences, reflecting the ongoing research efforts. This continuous innovation translates to a high degree of design flexibility, allowing engineers to tailor SAR ADCs to specific application needs by optimizing performance, resolution, and power consumption. This adaptability renders SAR ADCs indispensable in various domains, including data acquisition systems, battery-operated devices, and general-purpose analog signal processing applications. Nonetheless, despite their superior power efficiency, SAR ADCs

tend to exhibit slower pure maximum conversion speeds compared to Flash ADCs or pipeline ADCs and may offer lower resolution in comparison to Sigma-Delta counterparts.

A well-established technique to achieve higher conversion bandwidths in SAR ADCs (recently also in Pipeline ADCs) is time-interleaving. The approach utilizes multiple identical SAR ADCs that sample a common input signal with phase-shifted sampling periods. This effectively increases the overall system sampling rate while preserving the resolution of individual ADCs. Consequently, the measurement system can acquire input signals at a faster pace without compromising resolution. Theoretically, an infinite number of ADCs can be interleaved to achieve an infinite sampling rate. However, successful implementation hinges on precise synchronization between ADCs. The same clock can be used to generate the conversion signal for the ADCs, shifting only the phase of the clock. This phase shift must adhere to the following equation:

$$\text{ADC}_{\text{PHASE}} = 360^\circ * \frac{p-1}{n} \quad (1)$$

where n represents the total number of interleaved SAR ADCs and p signifies the order number of the individual ADC. For instance, a system with two interleaved ADCs necessitates a 180° phase shift, whereas three interleaved ADCs require a 120° phase shift relative to the clock. While interleaving offers a compelling avenue to expand the achievable bandwidth of SAR architectures, generating high-quality clocks with precise phase shifts and minimal phase noise presents a significant challenge. Additionally, each interleaved ADC must possess the minimum acquisition and conversion time for proper operation. It is crucial that the acquisition times of individual ADCs do not overlap when sampling the same input signal. These acquisition and conversion times ultimately limit the achievable maximum sampling rate for each individual ADC. Furthermore, each additional ADC contributes to increased power consumption and, more importantly, occupies a larger chip area.

With the advent of Extreme Ultraviolet Lithography (EUVL) and highspeed serial data interfaces, massive interleaving of SAR ADCs has recently become very popular in many segments of the industry, including next-generation DDR memory interfaces and custom data capture systems like high-bandwidth signal analyzers. Massive interleaving techniques can enable bandwidths reaching up to 100GHz by operating thousands of SAR ADCs in parallel, albeit with limited resolution. SAR ADCs, being relatively digital-intensive and easily portable from one node to another, depend largely on complex algorithms to achieve maximum performance improvements. This makes them well-suited for fast integration into high-performance computing modules, such as Accelerated Processing Units (APUs) and volatile memory systems, which utilize the most advanced semiconductor nodes available in the market.

3.5 Other ADCs

Alternative ADC architectures, such as Integrating or counting ADCs, employ methodologies like capacitor charging and discharging, or free-running binary counters to facilitate the conversion process. Distinguished by their exceptionally low power consumption and high resolutions, these ADC variants are frequently integrated into microcontroller systems, especially prevalent in battery-powered applications such as wearable electronics and low-power sensor interfaces. However, they are characterized by markedly slower conversion speeds relative to other discussed ADC types, rendering them suitable primarily for highly specialized applications within niche fields. Consequently, they will not be further explored within the scope of this work.

4 Wi-Fi 7 and State of the Art Converters

This chapter covers the following topics:

- The usage and selection of different type of ADCs.
- State of the art design of ADCs for Wi-Fi 7 applications.

4.1 RX Chain

The utilization of various non-interleaved ADC architectures is summarized in Figure 8. Flash and Pipeline ADCs offer the highest bandwidths at the expense of resolution and are employed in applications where complexity and power consumption are of lesser importance. At the opposite end of the spectrum, ADCs based on Counters and Integration architectures can achieve high resolutions with minimal power consumption and are primarily utilized for specialized applications such as audio or highly wearable devices. Occupying the middle ground are the Sigma-Delta and SAR ADCs, which cover the broadest range of frequency bandwidths. Consequently, these ADC types can be regarded as truly multi-purpose and successfully designed for numerous diverse applications. SAR ADCs tend to offer slightly better power consumption figures at the cost of maximum achievable performance. Sigma-Delta ADCs, however, additionally generally support a moderately higher performance.

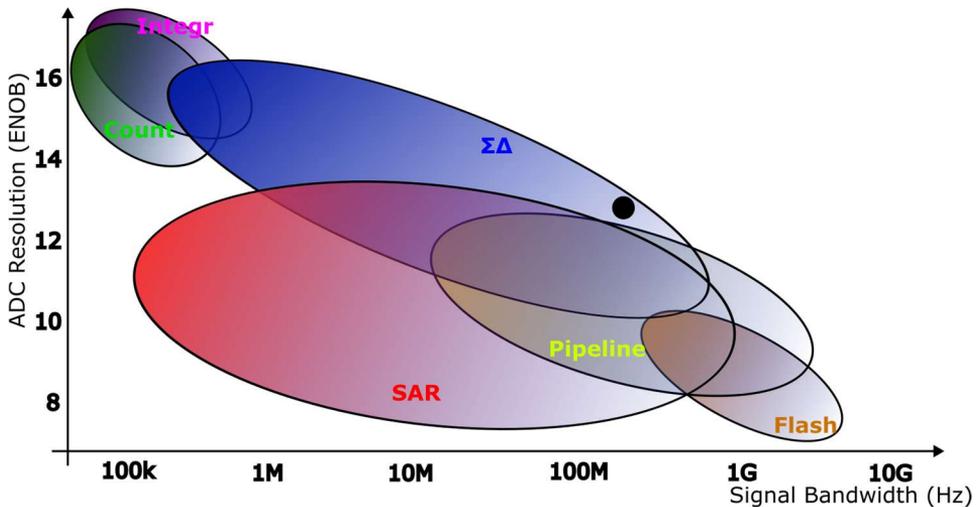


Figure 8. Different types of ADCs and their respectable potential non-interleaving performance over signal BW. Black dot represents the target performance needed for optimum Wi-Fi 7 operation.

When employing conventional RX chains, as illustrated in Figure 4, system architects are tasked with the allocation of noise figure (NF) budgets, among various submodules within the chain. The assignment of NF budgets is mostly contingent upon the expertise of the system design team, who must consider factors such as the potential for re-using existing building blocks, the anticipated complexity and power consumption of new designs. System architects may opt to prioritize the reuse of components from prior generations, necessitating enhanced performance from new designs to align with specifications. Alternatively, they may allocate NF budgets based on the challenges

associated with meeting targets and the anticipated power consumption required to fulfill the specified requirements.

Each module within the chain introduces its own share of noise and distortion, making it generally advantageous to incorporate greater gain upfront to enhance overall performance. Failure to do so can result in subsequent modules amplifying the noise and distortion introduced by preceding submodules. A critical bottleneck in any RF system lies in the phase noise of the local oscillator (LO), responsible for up/downsampling the RF channel. Hence, ensuring amplification before the introduction of LO noise proves highly advantageous, mitigating the risk of its amplification downstream. Consequently, prioritizing gain, particularly in components like the LNA rather than the PGC, is often preferred. The generation of LO phase noise is intricately linked to LO power consumption, thereby rendering RF systems predominantly constrained by the phase noise, which typically constitutes the primary contributor to the overall NF. The creation of exceptionally clean LO signals entails significant costs and is often impractical from a power consumption standpoint. Consequently, real-world RF systems are frequently designed around what can be feasibly achieved by the phase-locked loop (PLL) circuitry within reasonable power consumption parameters, with other modules tailored accordingly with much higher specs [12].

In recent years, there has been a surge of interest in exploring alternative approaches to traditional RF-sampled RX chains for analog-to-digital conversion. Techniques such as direct sampling or sub-sampling utilizing Two-step Flash ADCs or Pipeline architectures have garnered attention, presenting opportunities to remove or minimize the impact of LO generation on overall performance [13]. However, while these concepts hold promise, their practical implementation for Wi-Fi applications is still in its nascent stages. It is anticipated that they may find viability in future iterations of Wi-Fi standards, such as Wi-Fi 8 or subsequent Wi-Fi 9 protocols, once the technologies advance and topologies further mature.

4.2 Design Considerations for the Optimal ADC Type of Wi-Fi 7

The practical implementation of the Wi-Fi 7 RX chain needs to be engineered and tailored to meet the rigorous specifications of the latest Wi-Fi 7 standard, encompassing the new 802.11be protocol slated for release and widespread rollout in 2025. This ground-breaking standard ushers in a new era of wireless communication, boasting five distinct RF channel bandwidths spanning from 20MHz to a staggering 320MHz. This represents a significant leap forward from its predecessor, Wi-Fi 6, introduced in 2019. In addition, Wi-Fi 7 fully embraces the entire 6GHz band, a vast spectrum that can be partitioned into up to three non-overlapping 320MHz RF channels. This culmination of advancements translates into a staggering 46Gbps maximum throughput per client, a colossal more than 4X increase from Wi-Fi 6's 9.6Gbps limit.

The quality of the Wi-Fi signal and the performance of the communication chain are measured through a metric known as EVM. The EVM, measured across the entire chain, determines the appropriate coding scheme and the achievable maximum throughput. The signal quality in wireless systems exhibits significant variations, influenced by factors such as the distance between the access point and the client, intervening obstacles, quality of medium, interference, blockers from neighboring channels and the number of connected devices on the same channel. To derive the spec, a simulation test on 100 snapshots on channel model, with different environment conditions (by changing physical obstacles) was run, with the goal of reaching less than 20% Fail Rate (Figure 9).

Fail rate shows the number of configurations where 4K-QAM specification of less than 10% Packet Error rate was not achieved. To reach the project internally agreed PER, incorporate 4K-QAM in 80% of the different environments, and fully adhere to the Wi-Fi 7 standard, the RX channel should maintain an EVM of at least -42dB and the TX channel should maintain an EVM of at least -45dB, an improvement of -3dB of both from the -39dB/-42dB requirement used for the previous Wi-Fi 6 flagship project completed by the authors [14].

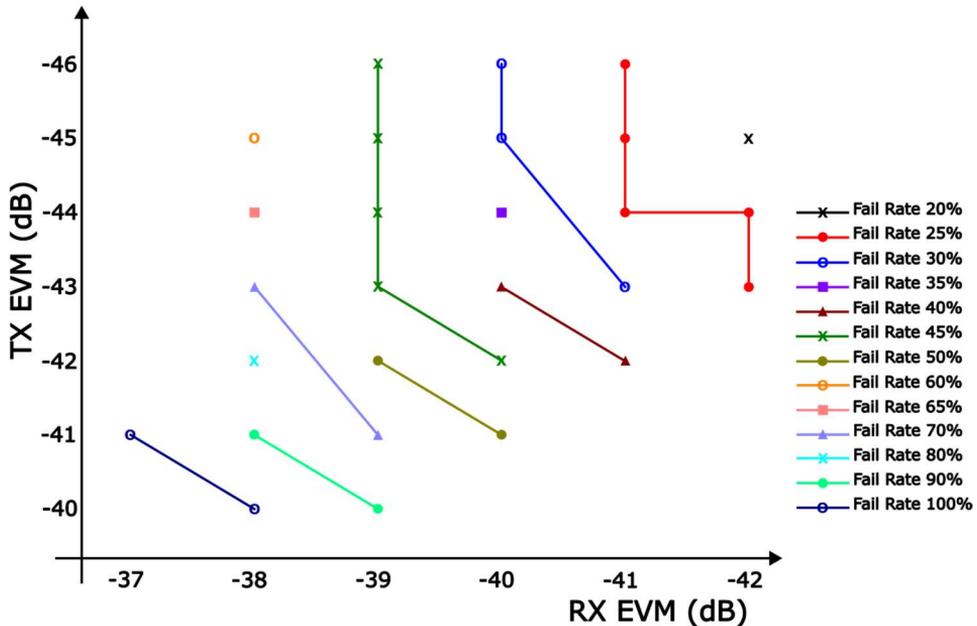


Figure 9. Fail rate simulation on 100 snapshots of channel model (100 different rooms and locations of the AP and STA).

Conversely, Wi-Fi 7 users also demand extended Wi-Fi range, prioritizing link stability over peak throughput. In latter use case scenarios (called sensitivity), simpler coding schemes are employed to prioritize connection reliability, even at the expense of maximum data transfer rates. In the end, EVM of a receiver chain is a sum of many imperfections coming from different modules like the non-linearities associated with the RF frontend and different PGCs used to relay the signal or phase noise coming from the PLLs and clocking. The ADC, sitting in the end of the RX chain, plays an important part both in sensitivity and high maximum throughput scenarios, and must offer ultimate performance in order to reach the required RX spec.

In light of these considerations, the entire system has to be emulated simultaneously both top down and ground up incorporating insights gleaned from previous generations of Wi-Fi transceivers. This approach enables a comprehensive understanding of the system’s capabilities and limitations, ensuring that the design aligns with both technological feasibility and power and area constraints. Allocating the optimal EVM budgets for both sensitivity and maximum throughput across different modules is a complex endeavor that demands extensive experience and rigorous testing.

As a rule of thumb, by far the biggest portion of the EVM budget in full performance modes is consumed up by the phase noise of the PLL and clocking, while noise from the

ADC and other linear amplifiers play a bigger role in sensitivity scenarios [12, 15]. That means that the ADC has a bigger impact in sensitivity and should be designed specially with the DR as the primary specification consideration. Standard OFDM signals, such as those employed in Wi-Fi, are inherently composed of multiple subcarriers, eliminating the likelihood of scenarios where the ADC faces a single-tone input. This somewhat simplifies the design process, as high single tone linearity, often used as a benchmark for comparing converters in literature, is rarely required in practical applications such as Wi-Fi or Mobile Communications.

OFDM signals utilized in all Wi-Fi standards starting from Wi-Fi 2 inherently manifest random phase variations across subcarriers, leading to unpredictable summations over time and varying peak-to-average ratios (aka crest factors) from symbol to symbol. This dynamic behavior induces fluctuations in the RMS of the signal across different symbols, necessitating the adjustment of the ADC to operate at varying backoffs from full scale (FS). A trade-off exists between employing a large or small backoff. A large backoff ensures that the ADC operates in a conservative region, minimizing the risk of overload and packet loss. However, this may come at the expense of performance optimization. Conversely, a small backoff may lead to optimal ADC performance but increase the likelihood of overload, particularly during scenarios where subcarrier phases align in an unfavorable manner. Wi-Fi standards accommodate a certain degree of packet loss, acknowledging the limitations of hardware and signal dynamics. Therefore, selecting the optimal backoff becomes an optimization problem, balancing performance and reliability. This challenge is more pronounced in low-bandwidth scenarios, where the number of subcarriers is limited. Conversely, higher bandwidth channels, with hundreds or thousands of subcarriers, offer a larger degree of phase diversity, reducing the likelihood of unfavorable summations and simplifying the backoff control. The precise backoff control is implemented through an Automatic Gain Control (AGC) algorithm, a complex topic that merits a dedicated thesis for thorough exploration. This discussion will not delve further into the intricacies of AGC algorithms and the fundamental principles governing backoff optimization in Wi-Fi OFDM receivers, focusing instead on practical aspects of the ADC design.

The average crest factor of a generic complex Wi-Fi signal varies from 7dB to 10dB, influenced by factors such as the specific modulation scheme employed, signal strength, and environmental conditions. An additional 3dB is required to transition from the complex RMS to the I- or Q-channel RMS value, bringing the average crest factor that a single ADC encounters close to 12dB. A widely accepted design practice dictates that the FS of analog circuitry should be positioned 3 dB below the supply voltage (VDD). This approach provides sufficient operating margin for analog stages, particularly the AB-stage in amplifiers, to function effectively and avoid excessive distortion. Given this consideration, the ADC has been designed to maintain a 3dB backoff from VDD, resulting in an FS of approximately 700mV peak. Accordingly, a generic operating point of -10dBm or 70mV RMS has been established at the ADC input, serving as the target for the AGC to set the signal chain amplification. This configuration leaves an average of 6dB for RMS variations, representing a favorable trade-off where the ADC operates near its optimal performance while still maintaining some backoff from full scale to prevent overloading and signal clipping, which could potentially compromise the analog signal, leading to packet loss.

4.3 Deriving the Specifications for the ADC

Instead of employing EVM, a complex signal that requires transformation into I and Q components for simulation, measurement, or evaluation of ADC performance, Missing Tone Power Ratio (MTPR) proves to be a more convenient and efficient metric. MTPR (Figure 10) is defined as the proportion of the signal amplitude to the amplitude of the unwanted signal in the unoccupied frequency bin that can be applied to a single-channel ADC. With a substantial number of subcarriers, MTPR and EVM yield nearly identical results and since an MTPR is simpler to generate in both CAD tools and measurement equipment, the former is often favored among mixed-signal designers. Since EVM is a complex signal, transformation into I/Q MTPR requires a 3dB addition.

After estimating that reasonable jitter caused by the phase noise achievable by the PLL, with sensible power consumption, is around 70 picoseconds, corresponding to approximately 50dB MTPR, the specifications for other modules were determined. To ensure 802.11be compatibility, the combined contribution of all other modules in the receiver should not degrade the MTPR by more than 5dB. Consequently, the ADC was assigned a design target of 57dB MTPR at an input power of -10dBm. At the other end of the spectrum, at sensitivity levels where phase noise from the PLL has a lesser impact on overall performance, the ADC's DR must be sufficiently high to prevent it from becoming the dominant factor. Therefore, the ADC was assigned a dynamic range (DR) specification of 77dB (equivalent to ~12.5 effective number of bits [ENOB]).

$$\text{ENOB} = \frac{\text{DR} - 1.76}{6.02} \quad (2)$$

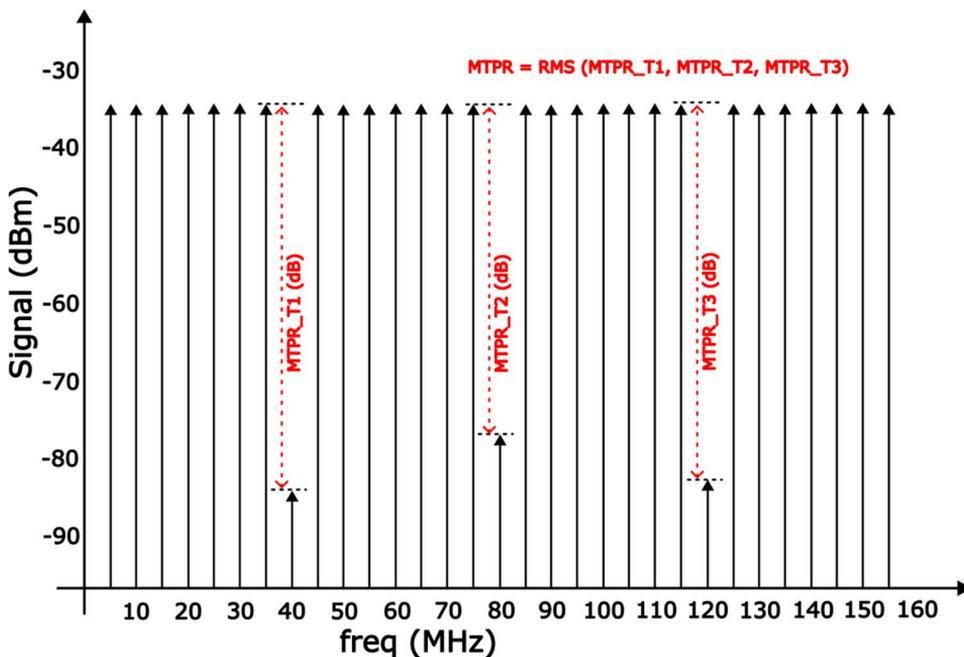


Figure 10. An example MTPR signal spectrum with a 5MHz subcarrier spacing.

Figure 8 illustrates several ADC topologies capable of achieving a DR of approximately 75dB (equivalent to ~12ENOBs) at a BW of 160MHz. These include pipeline, SAR and Sigma-Delta ADCs. Pipeline ADCs, while theoretically suitable, exhibit quite high-power

consumption and are typically reserved for applications where precise complex calibration is feasible. Since Wi-Fi can utilize configurations where 4 or even 8 antennas are used in parallel so that the access point could modulate data on all eight radio chains to a single client, which would result in some substantially high data rates, power consumption is of critical importance. Consequently, pipeline ADCs are generally not considered ideal for Wi-Fi applications. This leaves Sigma-Delta and SAR ADCs as the primary candidates for Wi-Fi 7 adoption. Industry practices vary, with most manufacturers favouring SAR-based architectures and only a few opting for Sigma-Delta based designs.

To maximize the noise budget available for other Wi-Fi 7 RX-chain modules while achieving superior performance, for this project a **Sigma-Delta ADC** ($\Sigma\Delta$ ADC) development plan was selected. While SAR architectures offer implementation simplicity, they cannot match the conversion efficiency and 13 ENOB performance target achieved by $\Sigma\Delta$ ADC architectures at the specified operational frequency. This architectural decision strategically positions the project to address all research objectives while establishing a significant competitive advantage in the market. Consequently, the ADC final specifications were derived, established, and agreed at a DR of 77dB and a MTPR of 57dB with an input power level of -10dBm. This translates to approximately 12dB below the total RX-chain noise budget. Under the assumption of uniform noise distortion across the operating band, the ADC's contribution to the overall MTPR degradation can be approximated by the following equation:

$$\Delta MTPR = 20 * \left(\log_{10} \left(\sqrt{10^{2 * \frac{MTPR_1}{20}} + 10^{2 * \frac{MTPR_2}{20}}} \right) \right) - MTPR_1 \quad (3)$$

where $MTPR_1$ is the total MTPR of the RX-chain and $MTPR_2$ is the MTPR of the ADC. If the ADC's SNR is 12dB higher than the total RX-chain SNR, the ADC's contribution to the total MTPR degradation is approximately 0.25dB. This implies that the ADC's noise and distortion should have by design a negligible impact on the overall performance of the total RX-chain, not so easily achievable with a SAR based structure.

A frequently utilized metric to evaluate the quality of ADCs in literature is the Schreier's Figure of Merit (FOM), depicted in Figure 11. This metric is measured in decibels per joule (dB/J), although the unit "joules" is often omitted for unspecified reasons. The Schreier FOM accounts for the power consumption of the converter, as well as the achievable BW and the ENOBs of the conversion, and is defined as follows:

$$FOM_{Schreier} = SNDR + 10 \log \left(\frac{BW}{P} \right) \quad (4)$$

where SNDR is the Signal-to-Noise and Distortion Ratio of the converter, P is the power consumption expressed in watts, and BW is the bandwidth of the converter (for non-oversampling converters, this is typically assumed to be $f_s/2$). The dotted line on the figure represents the envelope of theoretical limits given the current state of technology.

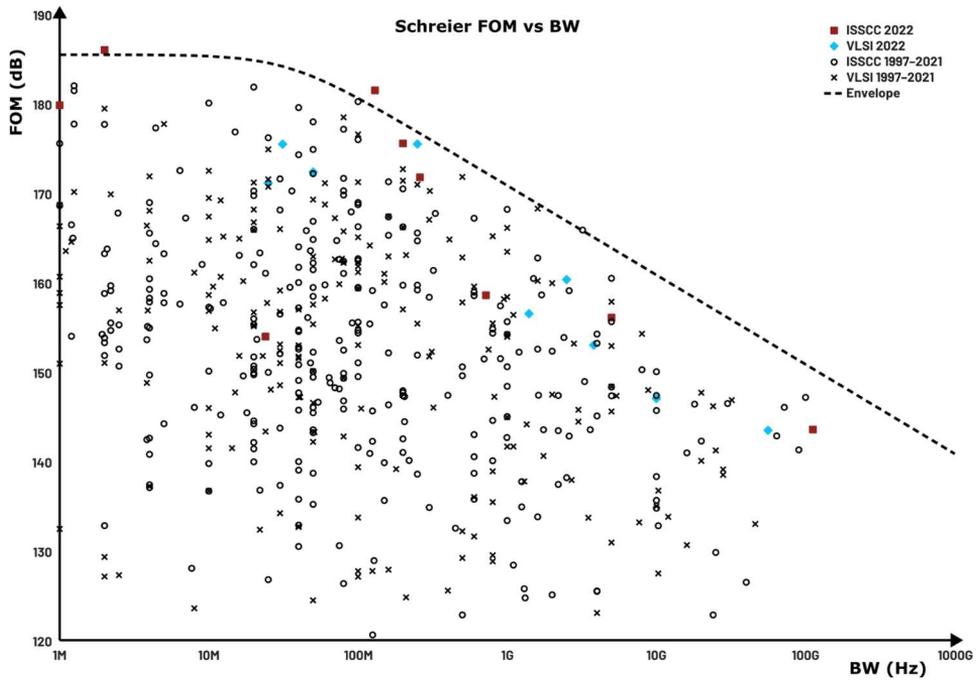


Figure 11. Claimed Schreier FOM vs BW in recent publications in ISSCC and VLSI (ADC Performance Survey – Compiled and collected by Willy Sansen and Boris Murmann).

5 Sigma-Delta Conversion

This chapter covers the following topics:

- History of Sigma-Delta Conversion.
- Quick overview of the Sigma-Delta operation.

Sigma-Delta ($\Sigma\Delta$) modulation is a widely used technique in digital signal processing and ADC conversion. It is particularly popular in applications where high-resolution conversion is required. The modulation scheme is known for its simplicity, robustness, and effectiveness in achieving high-resolution digital conversion.

The concept of Delta-Sigma ($\Delta\Sigma$) modulation can be traced back to a 1961 patent filed by Charles B. Brahm [16]. However, the formalization and naming of the technique is attributed to Hiroshi Inose of the University of Tokyo in 1962 [17]. His work provided a foundational framework and understanding for $\Delta\Sigma$ modulation, including the use of oversampling and noise shaping. Subsequent research delved deeper into the theoretical aspects of first and second-order $\Delta\Sigma$ modulators, further solidifying the theoretical underpinnings of the technology [18]. The term “Delta-Sigma” remained dominant for several years. However, in the 1970s, engineers at AT&T began using the alternative term “Sigma-Delta”. While both terms are widely used today, “Sigma-Delta” may hold a slight edge in technical accuracy due to the dominance of the sigma notation in signal processing and will be used throughout this work.

While $\Sigma\Delta$ converters have been discussed in open literature since the 1960s, their significant rise to prominence began in 1985 with the publication of an 18-bit $\Sigma\Delta$ ADC designed for audio applications [19]. This pioneering design, constructed entirely from discrete components, introduced multi-bit quantization – a novel concept at the time – and employed a third-order loop filter operating in continuous time. Despite its relatively high power consumption of 40 watts, the development marked a critical turning point in the widespread adoption of $\Sigma\Delta$ converters. Since then, substantial advancements have been made in integrating $\Sigma\Delta$ converters into CMOS technologies, significantly reducing power consumption.

At its core, $\Sigma\Delta$ modulation involves oversampling an analog signal at a frequency much higher than the Nyquist rate, typically from ten to a hundred times higher. This oversampled signal is then quantized using a quantizer (QT), resulting in a stream of high frequency modulated digital signals. The key principle of $\Sigma\Delta$ modulation lies in the feedback loop, where the difference between the actual and the desired output is fed back to adjust the quantizer’s input. This feedback loop pushes the QT-noise out of the band of interest, effectively trading signal bandwidth for increased resolution.

Delta modulation, a precursor to Sigma-Delta modulation, directly quantizes the difference between consecutive samples. While it’s simple, it suffers from high signal-to-noise ratio (SNR) due to slope overload distortion and granular noise. Sigma-Delta modulation, on the other hand, combines Delta modulation with an integrator in the feedback loop. This integration process, along with oversampling, allows Sigma-Delta modulators to achieve high resolution with low out-of-band noise.

Fundamentally, a typical $\Sigma\Delta$ ADC comprises from simple analog building blocks, such as comparators, voltage references, and one or more integrators and analog summing nodes, coupled with intricate digital computational circuitry. This digital circuitry incorporates a digital signal processor (DSP) that functions as a filter and a decimator, typically implementing a low-pass filter functions together with down-sampling. The realization of the $\Sigma\Delta$ ADC concept was facilitated in the 1980s by the advent of

large-scale integration and the widespread availability of signal processing capabilities, which enabled the implementation of the necessary digital processing components. The simplicity of the analog components within a $\Sigma\Delta$ ADC architecture is complemented by the complexity of the digital signal processing algorithms employed. This synergistic combination of analog and digital elements allows for the achievement of high portability in digital domain and high customizability in analog domain, rendering $\Sigma\Delta$ ADCs a compelling choice for various applications requiring high precision and robust performance.

Historically, most $\Sigma\Delta$ modulators have employed discrete-time (DT) loop filters implemented with switched-capacitor circuits. This approach facilitates precise mapping of system coefficients, leading to high conversion accuracy. However, for a given signal bandwidth, only two options exist to enhance modulator resolution: increasing the loop order or the oversampling ratio (OSR). While increasing the loop order offers benefits, it can introduce instability issues. Conversely, increasing the OSR is limited by the technology's transit frequency.

An alternative approach to achieve higher clock rates involves constructing the loop filter in the continuous-time (CT) domain, utilizing elements like transconductor-C or active-RC integrators. This methodology relaxes the requirements placed on the analog building blocks compared to their DT counterparts. Figure 12 illustrates the fundamental difference between DT and CT $\Sigma\Delta$ modulators. In the DT implementation, the sampling process occurs before the loop filter. Conversely, in the CT implementation, sampling takes place solely before the internal QT. The DT loop filter must operate on sampled data, enabling its construction using switched-current or switched-voltage techniques. This implies that only the analog value of signals at the clock's sampling instant is relevant, regardless of their intermediate behavior. In contrast, the CT implementation considers the entire history of internal signals, as it influences the subsequent sampling process by the QT.

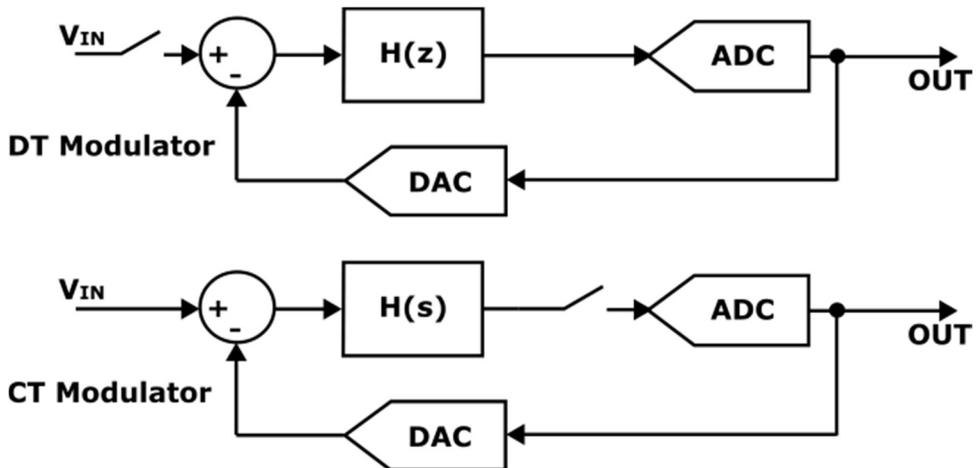


Figure 12. a) DT modulator b) CT modulator.

However, CT $\Sigma\Delta$ modulators offer three distinct advantages over DT $\Sigma\Delta$ modulators: the potential for higher clock frequencies, the absence of input switches, and the inherent anti-aliasing filter within the CT STF. These benefits are counterbalanced by several challenges, including increased design complexity, addition of excess loop delay (ELD), sensitivity to jitter, and component tolerance issues.

A typical CT design process begins with an equivalent DT model. Parameters such as the noise transfer function, the resolution of the internal QT, and the sampling frequency are initially determined in the DT domain. The design is then converted into an equivalent CT circuit, assuming precise knowledge of the transfer functions for each analog block. In reality, this assumption is often inaccurate, necessitating extensive simulations and iterative modeling to account for all analog circuit non-idealities and to achieve the desired performance. Given the requirement for signal bandwidths up to 500MHz (as will be further explained in the context of digital pre-distortion (DPD) modes), using a DT Σ ADC was not a feasible option for this work. Consequently, only CT Σ ADC implementations are discussed in the subsequent sections.

Due to their feedback-based systematic nature, Σ ADC can be prone to instabilities and oscillations if not meticulously designed. One particularly critical aspect unique to CT Σ ADC modulators, which is not encountered in DT Σ ADC designs, is the ELD. ELD refers to the nonzero delay between the QT clock edge and the time when a change in output bits is observed at the feedback point in the modulator [20]. In a CT modulator, ELD is defined as the timing delay between the sampling instant of the internal ADC and the corresponding updated output signal of the internal DAC. Since in a DT Σ ADC sampling happens at the input, ELD is not an issue, while in a CT Σ ADC it is one of the most critical aspects of design. The ELD usually requires an extra compensation system, called ELD compensation. The higher the ELD, the more compensation is required. If ELD is very large and excessive compensation is required, the entire system can nevertheless easily become unstable, especially when Dynamic-Element-Matching (DEM) techniques are used to improve linearity [21].

This covers the Pre-Design section of the thesis. The next chapter will focus on core design and is the main body of work for the thesis.

6 Design of the CTΣΔADC

This design builds upon the previous generation CTΣΔADC developed for the Wi-Fi 6 standard [14]. While the Wi-Fi 6 version was primarily designed by the author of the thesis, the Wi-Fi 7 iteration incorporates significant contributions from other collaborative efforts.

Specifically:

- Development of the toolbox meant for coefficient optimization in Wi-Fi 7 standard, enabling pushing the performance limits much higher – largely conducted by a separate team member, representing a body of work potentially suitable for independent thesis on its own.
- The quantizer concept and design needed to increase the resolution from 5-bits to 6-bits with reasonable power consumption, achieved through collaboration with a separate team member.
- The VCO mode conceptualized and developed by a research group at Charles III University of Madrid, where the author was primarily responsible for only integration and adaptation.

The concept and implementation of the remaining parts was originally designed and implemented by the author, with modifications and improvement added by the entire team to meet the stricter Wi-Fi 7 specifications.

All simulations and measurements were performed (unless otherwise noted) either with Cadence Virtuoso Studio Design Environment or Mathworks matlab or Simulink platforms.

6.1 Initial Draft of the CTΣΔADC Specifications

This chapter covers the following topics:

- The design requirements for the CTΣΔADC within a Wi-Fi receiver chain.
- The initial design considerations for the first draft of the CTΣΔADC.

The SNDR of an ΣΔADC can be decomposed into SNR and signal-to-distortion (SDR) components. SNR, akin to DR in sensitivity or single-bit operations, pertains to scenarios where the ΣΔADC processes signals of such minuscule amplitude that only a single cell of the feedback DAC is engaged in switching activity. SNR can be further partitioned into thermal and quantization noise components. Thermal noise stems from analog circuitry, while quantization noise arises from analog to digital conversion and the efficacy with which its error is mitigated through oversampling and noise shaping. Generally, intentional efforts are made to design quantization noise to be significantly lower than thermal noise. This objective is readily attainable when employing high OSR, facilitating efficient noise shaping from the desired band of interest to OoB. In instances where OSR is low, more aggressive noise-shaping coefficients are required, intensifying demands on amplifiers and consequently elevating their static power consumption. Alternatively, an increase in the number of quantization bits can be employed, necessitating additional comparators and DAC cells for digitization and feedback, thereby augmenting resolution at the cost of increased dynamic power consumption and complexity in design and calibration.

To ensure stability for mass production, the CTΣΔADC has been engineered with an ELD of no more than 0.5 in mind. While designs with ELDs even greater than 1 are theoretically possible and have been used in practice, they tend to exhibit very narrow stable tuning ranges and are prone to instability induced from process mismatch

variations and other imperfections. An ELD of 0.5 sets a solid baseline for DFM to achieve very high yields and reliability. In the context of Wi-Fi 7, which supports RF bandwidths up to 320MHz, necessitating a baseband section capable of handling bandwidths up to 160MHz, using an ELD of 0.5 imposes stringent limitations on clock rates due to signal propagation constraints during signal digitization and propagation. Despite advancements in process clocking speeds, the increase has not kept in pace with the shrinking process nodes. Even when employing the most advanced CMOS nodes, signal propagation limitations on chip severely restrict clock rates. Consequently, the ADC's maximum clock rate has been limited to 3.84GHz, corresponding to an OSR of 12 for 320MHz RF channels. This constrains the quantization (QT) noise budget to lie closer to thermal noise than necessarily desired in $\Sigma\Delta$ converters with the QT Noise floor being only a few dBs under the thermal noise. Given that the EVM requirement remains constant within the 802.11be standard across different RF bandwidths, oversampling is maintained when switching between bandwidths by adjusting the clock rate instead of the oversampling ratio. This strategic approach substantially lowers power consumption in the switching circuitry in lower BW modes, where power consumption is critical due to new European Union Code of Conduct rules.

The design targets and topology for the CT $\Sigma\Delta$ ADC were based on all these considerations and learnings from previous projects and are set to be:

- SNDR @ -10dBm: -57dB
- Dynamic Range: 77dB (~12.5ENOB)
- Power Consumption: < 45mA @ 0.95V per single 320MHz channel (double for I/Q)
- Area: < 0.1mm² per channel (double for I/Q)

Considering the specified design goals, a 3rd-order, 6-bit full feedback CT $\Sigma\Delta$ ADC operating at 3.84GHz, ELD of 0.5, and OSR of 12 was deemed the most suitable topology for the task of creating a Wi-Fi 7 W320-compatible converter that gives a competitive advantage while maintaining low power consumption. Smaller bandwidth modes (like W160 and W80) were created by reducing the clock rate, relaxing coefficients and reducing power, while maintaining the same OSR.

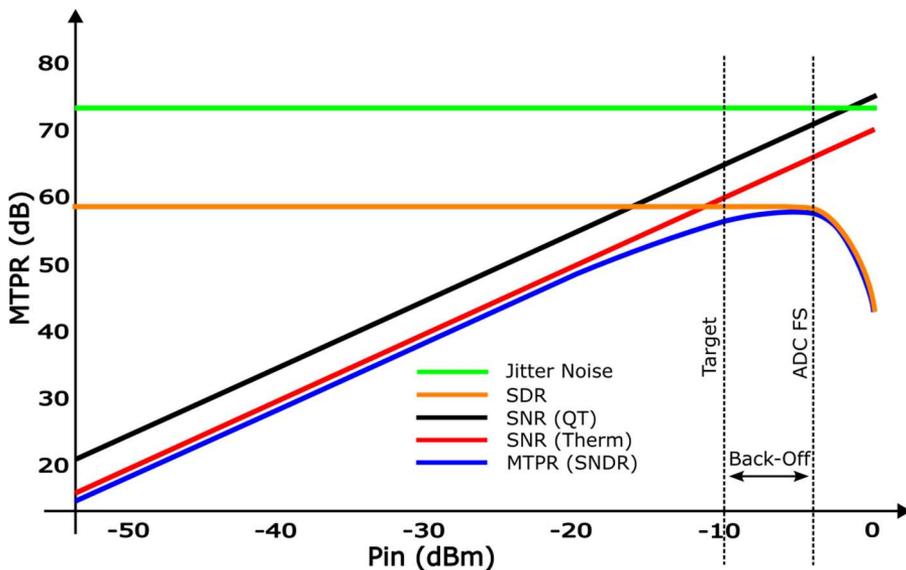


Figure 13. Designed matlab model of the $\Sigma\Delta$ ADC MTPR Performance vs input power.

To facilitate system simulation and validation, a simplified model (Figure 13) was developed that incorporated the four aforementioned non-idealities components: QT-Noise, thermal noise, distortion, and jitter, as a function of ADC input power. The matlab model, although not perfectly matching actual finally achieved performance, provides a good overview of the contributions from these components and served as a decent starting point to kick off the design activities. It also proved essential for system engineers to start evaluation on modulation schemes and potential throughput. The model, while simple and crude, gives a general estimation on how different non-idealities can effect the ADC performance with various input signal strengths. This was mostly used to kickoff system design based on best estimations and not relevant for further discussion.

It is important to note that while performance increases linearly with input power in low power scenarios (sensitivity), it levels off at peak input power, where RF PLL noise becomes the dominant factor. In such cases, even excellent SNDR performance from the ADC does not significantly enhance the RX chain because system performance is overwhelmingly dictated by RF PLL jitter. This should not be confused with the jitter noise illustrated in Figure 13, which originates from a completely different low-quality, low-power clock source. The impact of jitter performance on RF modules differs significantly from that on BB modules in terms of overall system performance. Rather than focusing on developing an ADC with high linearity for all cases including high-power input signals, the design prioritizes reducing power consumption and optimizing performance where the ADC operates closer to the overall noise floor, such as in sensitivity. This means that the designed CT $\Sigma\Delta$ ADC is not necessary conference publication material, as those works usually are measured with the input power concentrated in a single carrier, a scenario that unfortunately holds limited relevance in real-world applications.

6.2 Coefficient Development

This chapter covers the following topics:

- The design considerations of the filter coefficients within the $\Sigma\Delta$ ADC.
- Modes implemented in the $\Sigma\Delta$ ADC. Primary/DPD.
- Initial coefficient design for the $\Sigma\Delta$ ADC.
- Passives used to implement the coefficients.
- Resonator implementation and notch generation.

The design of optimally fitting coefficients for high-speed $\Sigma\Delta$ ADCs is a blend of art and science. Numerous factors influence the initial design, but in the end, the original design is almost always disregarded and tailored to match the capricious analog circuitry better. Nonetheless, coefficient design is an integral component of any CT $\Sigma\Delta$ ADC design and should be carefully considered from the very onset.

A variety of filter topologies exist, including Butterworth, Bessel, Elliptic, and different types of Chebyshev. They differ slightly in ripple behavior and the steepness of the roll-off, which is crucial in systems with very high OSRs. However, for systems with extremely high bandwidths but low OSRs, where analog linear circuitry is already stretched to its limits, perfectly fitting coefficients can never be fully realized in practice. Most common type – Butterworth filters, that are the simplest to implement, provide an additional benefit: they exhibit no ripple in the STF, thus preventing signal peaking in the $\Sigma\Delta$ ADC, what could

be caused by blockers from neighboring channels. Signal peaking is a genuine concern in wireless systems because it can lead to signal overloading in the presence of strong blockers. AGC algorithms favor predictability and aim to avoid signal peaking across various process corners. While other filter topologies can offer sharper roll-off and potentially more far-band OoB filtering, this does not outweigh the potential harm from near-band OoB peaking. The choice of filter topology, therefore, must carefully balance the desire for sharp filtering with the need to minimize signal peaking. Butterworth filters, with their smooth STF and inherent resistance to signal peaking, prove to be an effective initial choice for high-speed $\Sigma\Delta$ ADCs operating in scenarios with high bandwidths and low OSRs.

While the system's performance does not hinge on the $\Sigma\Delta$ ADC providing additional OoB filtering, it proves to be highly beneficial in providing extra headroom to safeguard the system overload from signal clipping. This is also the primary rationale behind opting for a filter topology with full-feedback, wherein each integrator has its own dedicated feedback DAC. The topology has superior attenuation compared to feed-forward filter structures, that in contrast provide wide-band STFs. This somewhat mitigates the reduced filtering inherent in Butterworth structures, compared to its competitors, thus offering the optimal combination of features.

Wi-Fi 7 standard encompasses RF channels with bandwidths of 20MHz, 40MHz, 80MHz, 160MHz and 320MHz. To facilitate a power-efficient converter with minimal overdesign, each of these modes should be accompanied by a dedicated ADC mode. As described earlier in the thesis, the $\Sigma\Delta$ ADC OSR is maintained constant across each designed mode, implying that quantization noise is also designed to remain similar across different modes. Each mode necessitates a set of distinct coefficients, which are implemented in circuitry comprising a specific combination of resistors and capacitors within the filter to construct integrators and current sources in the feedback DAC. More aggressive coefficients necessitate higher bandwidth integrators, which are realized using smaller RC values. This implies that lower bandwidth modes demand passive components that are larger in size. To minimize the overall layout footprint of the ADC within the SoC, only dedicated modes for 80MHz, 160MHz, and 320MHz RF bands have been implemented. Specific modes for 40MHz and 20MHz would demand capacitor sizes that would disproportionately expand the filter's size and were therefore omitted. For lower bandwidths, including modes like 40 MHz and 20 MHz, the $\Sigma\Delta$ ADC would operate either in the 80 MHz mode (minimum BW supported) or in the VCO mode, what will be discussed later in chapter 6.17.

The majority of the ADC's overall thermal noise budget is attributed to the passive resistors that comprise the integrators and the current sources that form the feedback DACs, particularly the resistor in the first integrator (R_1) and the current cells in the first DAC (DAC1) (Figure 37). This thermal noise component reduces in prominence when transitioning to lower bandwidth modes if noisy analog components like resistors, opamps and feedback DACs are kept the same. The integration bandwidth is reduced in tandem with the mode bandwidth, resulting in a 3dB reduction in the thermal noise component for each halving of the BW. It is however advantageous to maintain thermal noise in sync with quantization noise for optimal power consumption. For amplifiers, this is accomplished through BW programming, which also affects opamp noise performance. However, for other noisy passive components like resistors and current sources, such programming becomes challenging due to the RC-tuning employed within these components. The only feasible and efficient option is to either double or halve all

available resistors and double or halve all DAC cells, thereby modifying the bit-width of the ADC to sustain similar thermal noise performance across modes. This approach led to the implementation of four primary modes for the ADC: W320, W160, W160LP, and W80. These modes are created using only two sets of options for programming all resistors, capacitors and current sources:

- W320 is the ADC mode intended for the 320MHz RF bandwidth Wi-Fi mode, it uses resistors and capacitors that are both programmed to minimum nominal value available and the feedback DAC that is operating in 6-bit mode. The mode works with a 3.84GHz clock.
- W160 is the ADC mode intended for the 160MHz RF bandwidth Wi-Fi mode, it uses resistors that are programmed to minimum and capacitors programmed to the maximum available and the feedback DAC that is operating in 6-bit mode. The mode works with a 1.92GHz clock.
- W160LP is the ADC mode intended for the 160MHz RF bandwidth Wi-Fi mode, it uses resistors that are programmed to maximum and capacitors programmed to the minimum available and the feedback DAC that is operating in 5-bit mode. The mode works with a 1.92GHz clock.
- W80 is the ADC mode intended for the 80MHz RF bandwidth Wi-Fi mode, it uses resistors that are programmed to maximum and capacitors programmed to the maximum available and the feedback DAC that is operating in 5-bit mode. The mode works with a 0.96GHz clock.

mode clk	$\frac{\text{Therm noise}}{\text{QT noise}} \cdot \frac{DR}{DR}$	C	2C
R (6bit)	/	w320	w160
		3G84	1G92
2R (5bit)	/	w160LP	w80
		1G92	0G96

mode	Therm noise	QT noise	DR
w320	78dB	84dB	77dB
w160	81dB	84dB	79dB
w160LP	78dB	78dB	75dB
w80	81dB	78dB	77dB

Figure 14. Different modes of the $\Sigma\Delta$ ADC.

To achieve the target dynamic range of ~ 77 dB across all modes, while targeting that thermal noise ideally remains slightly more dominant than the QT noise, an initial noise budgeting exercise was undertaken. This approach involved optimizing the coefficient set for the most demanding mode (w320), and then utilizing those coefficients as a starting point for the remainder (Figure 14). Starting the profiling at the hardest achievable mode will optimize the coefficient set for that condition only, while the rest will be byproducts from the initial choices, meaning that the coefficient design is constrained to a limited number of degrees of freedom and is subject to significant limitations. The design target for thermal noise for the w320 mode was set at approximately 78dB, while QT noise was budgeted to be in the range of approximately 84dB. The w160 mode utilizes an identical set of coefficients to the w320, as both clock and integrator BWs are scaled similarly. The only distinction lies in the thermal noise, which is reduced by 3dB due to the halved integrated signal bandwidth. A reduction in ADC bit-width from 6-bit to 5-bit introduces an increase of 6dB in quantization noise (corresponding to loss of one ENOB). Consequently, the quantization noise budget for 5-bit modes is reduced to 78dB. This increased quantization noise is counteracted by a corresponding decrease in thermal

noise. For each halving of the signal bandwidth, total thermal noise reduces by 3dB under identical conditions. However, switching to 5-bit operating modes also incurs an additional 3dB of more thermal noise for the same bandwidth due to the doubled resistance. Essentially by employing such strategy, the number of coefficient sets was reduced from four to one, as all modes need to use the same set of coefficients. Nevertheless, while not the optimum solution in terms of flexibility, this strategy was implemented to facilitate analog design and layout, limit switching behavior, reduce parasitics, and suppress unwanted parasitic zeros in the transfer function.

The proposed third-order design also presents the opportunity to incorporate one resonator within the ADC loop-filter. This resonator modifies the ADC's NTF by shifting zeros away from the origin to distinct frequencies, thereby enabling the creation of deep notches or attenuations in specific frequencies of the noise spectrum. This approach effectively helps in shaping the quantization noise spectrum by concentrating noise energy in a more favorable manner, allowing for a more uniform noise distribution within the desired frequency band. Furthermore, it ensures a more efficient shaping of the noise, emphasizing certain frequency components more than others. The resonator is implemented by placing a resistor with a specific value between the output of the third integrator and the virtual ground node of the second integrator. This creates a feedback signal path that bypasses a portion of the two DACs and the third-order noise shaping property of the coefficients. This approach does not affect how much total noise is shaped out of the band (H_{inf}), but rather changes only its in-band shaping, that can be more beneficial in dedicated use cases.

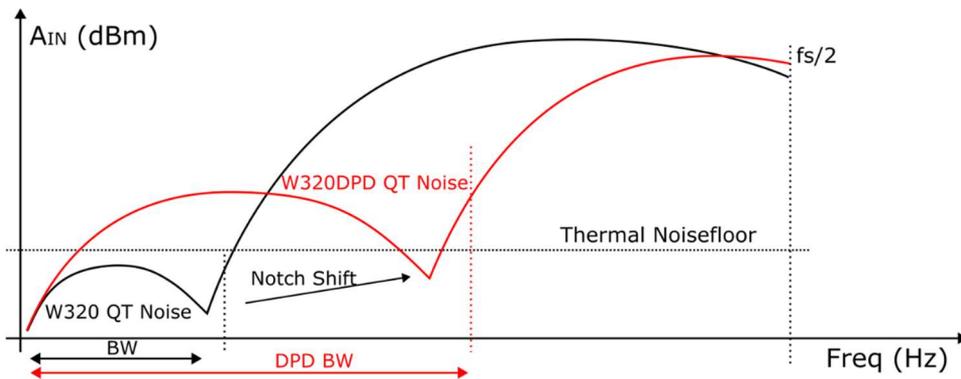


Figure 15. Shifting resonance frequency to create the DPD mode.

Resonators can be further leveraged by employing them with varying values, essentially trading noise shaping bandwidth for its effectiveness (Figure 15). Recognizing this potential, additional feature, designated as DPD modes, with the sole objective of extending bandwidth while compromising some performance, were introduced. DPD is a technique that enables digital linearization of power amplifiers in the transmit pathway by pre-distorting the transmitted signal. To implement DPD, the digital control algorithm must first be trained using a dedicated feedback receiver to minimize the third harmonic distortion (HD3) during transmission. This algorithm requires the ability to detect distortion during transmission and then apply appropriate correction. To accurately detect HD3, an ADC with three times the bandwidth of the transmitted signal is necessary. Rather than constructing a separate wideband ADC, which would unnecessarily complicate the system, this capability has been integrated into the CTΣΔADC by modifying the resonator

and shifting the notch such that the noise shaping bandwidth is three times greater than for the primary nodes. Each of the primary modes also includes a DPD mode in where extended notch shifting is applied, using the forementioned resonator approach, effectively expanding the $\Sigma\Delta$ ADC's functionality by adding four additional modes.

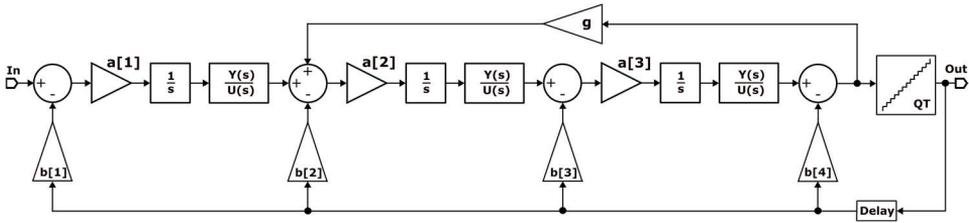


Figure 16. Simulink model of the implemented $\Sigma\Delta$ ADC.

Several freely available toolboxes and simulators are readily available for simplified coefficient design, with the works culminating in the Schreier toolbox and the Uni Ulm Sigma-Delta Synthesis Environment being particularly notable [22, 23]. The conventional approach to designing CT $\Sigma\Delta$ ADCs involves finding a DT $\Sigma\Delta$ ADC that is equivalent to the desired CT $\Sigma\Delta$ ADC [11, 24]. In this process, the differential equations describing the CT modulator are replaced by the difference equations of the DT modulator, greatly simplifying the time-domain simulation of the modulator and enabling the reuse of standard design software for DT systems. A DT delta-sigma modulator is considered equivalent to a CT modulator if the output sequences of both modulators are identical, assuming they operate at the same sampling frequency and receive the same input signal. The output sequence of the modulator is determined by the QT input sequence. Therefore, the output sequences of the two systems will be identical if their QT input sequences are equally identical.

For this design, an in-house coefficient design simulator using FFTs of pure analog signals was developed (Figure 16), that incorporated additional features to enhance its capabilities. The overarching design strategy involved maintaining a phase margin (PM) of over 30 degrees with ideal components while simultaneously minimizing the Hinf needed for Butterworth coefficients to achieve close to 84 dB of QT noise performance in the w320 mode (Figure 17). PM over 30 degrees has shown to be an effective trade-off between performance and stability, keeping some margin for non-idealities coming from analog circuitry while not overly constraining the complexity of design and ballooning the power consumption of the system.

The compromised coefficients attained to fit the set-out targets were:

- $vA = [0.26; 0.49; 1.32]$
- $vB = [1; 1; 1; 0.59]$
- $vG = [0.061]$
- $Hinf = [2]$

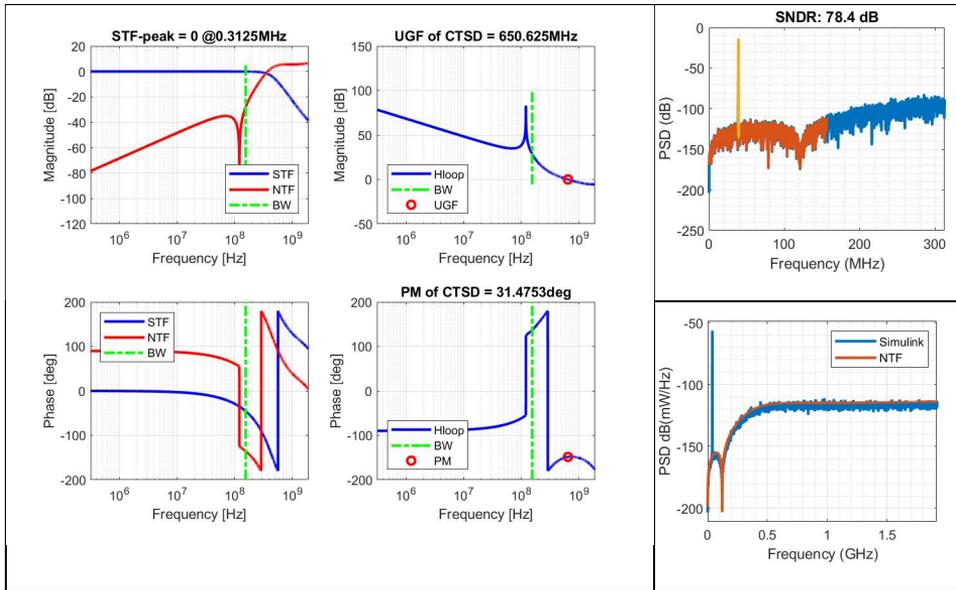


Figure 17. Transfer Functions (STF/NTF) phase and DR resulting from the designed coefficients.

The majority of the overall thermal noise budget in the ADC is contributed by the passive resistors that form the integrators and the current sources that form the feedback DACs, particularly the resistor in the first integrator (R_1) and the current sources in the first feedback DAC (DAC1). The resistor in the second integrator (R_2) and the current source in the second feedback DAC (DAC2) have the second highest noise contribution, while R_3 and DAC3 exhibit the least dominant impact. From a coefficient design perspective, this implies that R_1 should be the smallest and DAC1 current should be the highest. R_2 and DAC2 are less significant, while R_3 and DAC3 are the least relevant. To achieve the targeted thermal noise of 78dB, the initial coefficients were designed and partitioned according to the following scheme:

Coefficients for R_1 and DAC1: These passives were designed to minimize the thermal noise contribution from R_1 and DAC1, while ensuring adequate shaping of the quantization noise spectrum.

Coefficients for R_2 and DAC2: These passives were designed to focus on the shaping of the quantization noise spectrum, while also considering the thermal noise contribution from R_2 and DAC2.

Coefficients for R_3 and DAC3: These passives were designed to further refine the shaping of the quantization noise spectrum, with minimal emphasis on thermal noise.

Passives for the designed coefficients:

- **w320:** R = [250; 500; 750]; C = [4.00e-12; 1.06e-12; 2.63e-11]; DAC = [3e-3; 1.5e-3; 1e-3]; $R_{RESO} = 8.14e+3$; fbe = 0.58
- **w160:** R = [250; 500; 750]; C = [8.00e-12; 2.12e-12; 5.26e-11]; DAC = [3e-3; 1.5e-3; 1e-3]; $R_{RESO} = 8.14e+3$; fbe = 0.58
- **w160LP:** R = [500; 1000; 1500]; C = [4.00e-12; 1.06e-12; 2.63e-11]; DAC = [1.5e-3; 750e-4; 500e-4]; $R_{RESO} = 16.28e+3$; fbe = 0.58
- **w80:** R = [500; 1000; 1500]; C = [8.00e-12; 2.12e-12; 5.26e-11]; DAC = [1.5e-3; 750e-4; 500e-4]; $R_{RESO} = 16.28e+3$; fbe = 0.58

6.3 Top-Level Design of the ADC

This chapter covers the following topics:

- Top-level design and various sub-modules making up the entire $\Sigma\Delta$ ADC.

As a part of a wireless system employing an I/Q modulation, the proposed I/Q converter incorporates separate ADCs for the I-channel and the Q-channel. While most of the sub-modules within the ADC are dedicated to specific channels, some are shared for efficiency and performance (**Error! Reference source not found.**). This approach ensures that the I/Q converter can manage the complex signal processing requirements of wireless communication while maintaining high fidelity, with the lowest possible power consumption.

List of dedicated modules for one channel:

- 1) Loop-Filter
- 2) Quantizer
- 3) Overload-Detection
- 4) Excess-Loop-Delay Compensation
- 5) QT-Calibration assistance
- 6) Overload Masking
- 7) VCO

List of shared modules for I/Q:

- 8) Feedback DAC
- 9) Master Clocking
- 10) Calibration Clocking
- 11) Calibration Machine
- 12) Biasing
- 13) Active Capacitor
- 14) Digital Interface

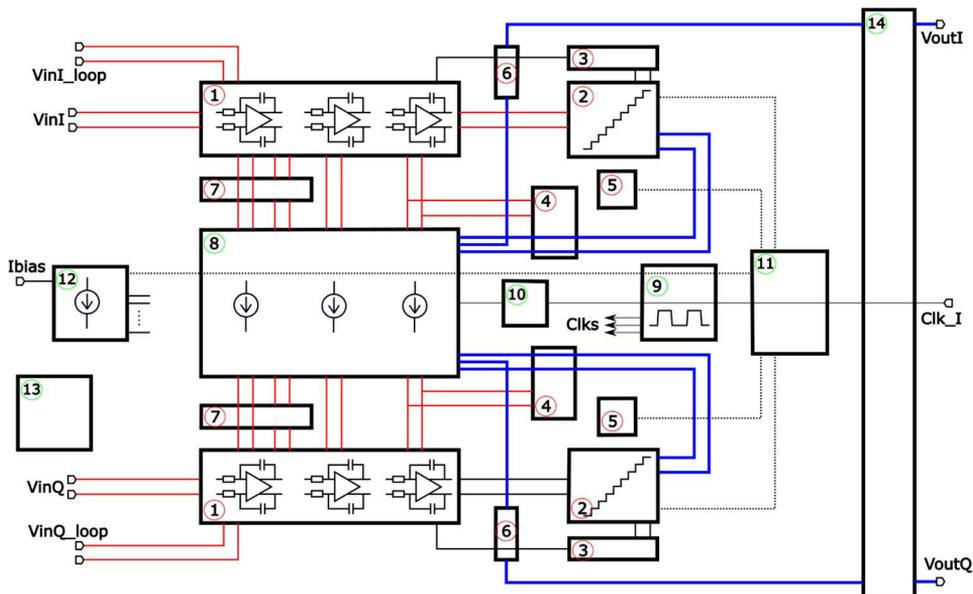


Figure 18. Top-level design of the CT $\Sigma\Delta$ ADC.

Everything not needed for a dedicated channel, such as clocking, biasing, and calibrating, have been integrated into shared sub-modules. This strategic consolidation serves to minimize the converter’s footprint while concurrently optimizing power efficiency. Nevertheless, it is imperative to carefully address considerations such as supply isolation and signal feedthrough during such fragmentation. Subsequent chapters will delve into a detailed examination of the design intricacies of each sub-module.

6.4 Design of the Loopback Filter

This chapter covers the following topics:

- Design considerations for the Loopback Filter.
- Topology of the Loopback Filter.
- The design of the integrators.
- The design and tuning of the passives.
- The design of the operational amplifiers.

6.4.1 Design Considerations for the Filter

The CT Σ Δ ADC incorporates a third-order full-feedback filter employing tunable capacitors for mode programming and tunable resistors for RC-tuning to compensate for process variation. The capacitor-based mode programming selects the desired bandwidth of the ADC, while the 5-bit resistor-based fine-tuning is used to refine the final RC product. The filter features two distinct input paths: the primary input, sourced from the PGC in the RX chain, and the secondary input, originating from the loopback path. This versatile loopback input serves multiple purposes, including an external test input, a test buffer for various SoC nodes, or a loopback input for performing calibrations necessary to optimize overall system performance.

Power consumption has emerged as a critical design consideration in many modern electronic systems. The CT Σ Δ ADC adopts a comprehensive approach to power optimization, particularly targeting the filter programming and opamp topologies. The strategy aims to minimize power consumption while ensuring optimal performance for each supported mode. The 3rd-order filter employs three active integrators, each meticulously crafted with a unique set of coefficients tailored to specific operating modes. This optimization ensures consistent SNDR per tone and DR across all selected modes, regardless of the bandwidth setting. As a result, low-bandwidth modes do not demand the same amplifier performance as high-bandwidth, greatly enhancing power efficiency. To achieve performance optimization, the opamps are designed with multiple power modes, each offering a distinct combination of open-loop bandwidth and noise performance. This flexibility enables the selection of the optimal opamp mode for a specific Wi-Fi mode, ensuring that power consumption aligns with the actual requirements. By incorporating mode-specific amplifier design and multiple power modes, the CT Σ Δ ADC successfully implements a multi-faceted approach to power optimization. This comprehensive strategy significantly reduces the overall power consumption, contributing to the overall energy efficiency of the device.

The CT Σ Δ ADC's ability to operate in both 5-bit and 6-bit configurations necessitates a scalable resistor bank that can accommodate both options. This design incorporates redundancy into the resistor bank, enabling it to be easily transferred between the two modes based on the selected configuration. The capacitors, on the other hand, are not tunable and are solely utilized for selecting the base Wi-Fi mode.

The resistors are therefore built with redundancy, so they can be either used in 5-bit mode or 6-bit mode, depending on what mode is selected. Since the entire RC-tuning is performed exclusively through resistors, the tunable range must encompass both resistor and capacitor process variations. This means that the designed resistor bank must cover a minimum range of:

$$x_{min,max} = \frac{1}{(1 \pm R_{var}) * (1 \pm C_{var})} \quad (5)$$

where R_{var} is the maximum variation of the resistor and C_{var} is the maximum variation of the capacitor in the process node.

Assuming a standard advanced node CMOS process with a general variation of 20% in both capacitors and resistors, the resistor bank must span from approximately 60% to 160% to fully cover the process variation. This ensures consistent performance across the full range of manufacturing conditions. With a 5-bit linear tuning range for the resistor bank, the overall accuracy of the tuning will become:

$$step = \frac{x_{max} - x_{min}}{2^n} \quad (6)$$

where x_{max} is the maximum of the tunable range of the, x_{min} is the minimum of the tunable range of the resistor bank and n is the tuning range number of bits. In order to achieve a reasonable accuracy of $\pm 2\%$, 5-bit RC tuning was used. This brings the step size is close to 3.125%, meaning the accuracy of the BW in the ADC can be potentially less than $\pm 1.6\%$ per step after final tuning.

6.4.2 The Design of Integrators

An integrator is a fundamental electronic component that performs the mathematical operation of integration on an input signal, yielding an output signal that is directly proportional to the integrated value. It functions as a low-pass filter, effectively attenuating high-frequency components while allowing low-frequency signals to pass through. The integrator's cutoff frequency is determined by system parameters, such as the values of the input resistor and the feedback capacitor.

To implement integrators in CMOS processes, two primary approaches can be employed: opamp based integrators and current source based integrators. While current source based integrators offer simpler implementation, they are accompanied by several drawbacks, including low gain, restricted output swing, and reduced linearity. Opamp based integrators, on the other hand, are more complex to construct and build, but they overcome these limitations. In the $\Sigma\Delta$ ADC presented in this thesis, all integrators have been implemented based on opamps.

An ideal integrator possesses an infinite DC gain and exhibits a 90° phase shift across the entire frequency spectrum, from the origin to the unity gain frequency (UGF) and beyond. However, real integrators, which exhibit finite DC gain and a finite open-loop UGF, deviate from this ideal behavior. While the finite DC gain (A_x) of the operational amplifier does not significantly impact the integrator's performance, if its more than 40dB, what is easily achievable, the finite Gain Bandwidth (GBW) of the opamp and the limited loop gain over frequencies alter the integrator's behavior, resulting in a steeper slope than the desired -20 dB/decade and introduce a phase shift that diverges from the ideal 90° (Figure 19). Consequently, the realized UGF shifts left from ideal UGF U_0 to U_x and the resulting phase shifts down from the ideal P_0 to P_x . The disparity between the ideal phase shift of 90° to P_x is denoted as Phase Loss (PL). To match the designed filter performance with the anticipated characteristics of the ideal filter, it is imperative that U_x closely matches U_0 , and P_x closely matches P_0 . Sometimes excessive drift is compensated by tweaking the coefficients post-design to shift U_x back right in the plot.

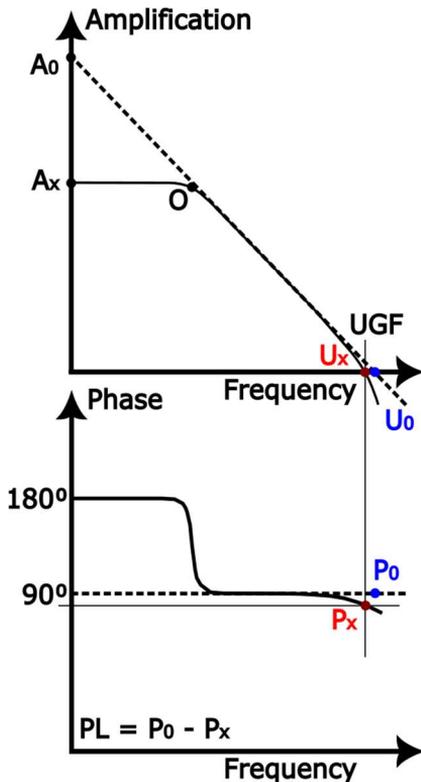


Figure 19. Integrator STF and equivalent PL.

This method does however entails moving the entire integrator slope, in a way that it does not overlap the ideal integrator characteristic anymore, introducing nuanced changes to the overall system behavior. An alternative method involves shifting the PX up by introducing a zero in the integrator transfer function via adding a series resistor. This approach, while somewhat imprecise due to limitations in accurate tuning, should only be employed as a supplementary measure and not as a primary design target. In this particular design, the technique is only selectively applied to the third integrator, where the total bandwidth (BW) and the resulting Phase Loss (PL) are most prominent.

The three integrators require different opamp BWs, noise performance and linearity specifications, dictated by the filter topology and resulting coefficient design. The integrator UGF is smallest for the first integrator, higher for the second integrator, and largest for the third integrator. A commonly employed design strategy involves utilizing a single opamp design for all three integrators. However, this approach fails to accommodate the unique needs and intricacies of each integrator. In recognition of

these complications and unique demands posed by each integrator, a deliberate design strategy has been pursued to develop three distinct opamps, each meticulously tailored to precisely fit with the specific requirements of its corresponding integrator.

The noise and linearity specifications for each integrator vary significantly. The first integrator sees no noise shaping, meaning that all thermal noise generated by the opamp is directly transmitted to the output of the filter. In contrast, the second integrator benefits from first-order noise shaping, and the third integrator profits from second-order noise shaping. This implies that the noise requirements for the second and third integrators are considerably relaxed compared to the first integrator. In a similar way, the linearity requirements for each integrator also differ. To achieve optimal performance, the first integrator must exhibit superior linearity, while the third integrator, primarily needed for compensation purposes, commands minimal linearity and instead should prioritize raw open-loop bandwidth. The second integrator falls somewhere in between, producing some non-linearities at the band edge but not as pronounced as those of the first integrator.

Another aspect to consider is the stability requirement for each opamp. A generally sound design approach involves creating each internal loop with faster settling characteristics (wider GBW) and better stability (higher PM) than the external loop in which it is embedded. This ensures that internal loops rapidly converge to their steady-state conditions, what is a strong prerequisite and essential for the external loops to successfully settle. In a third-order filter, the last integrator, in conjunction with DAC3, functions as the innermost loop, that needs the widest GBW and highest PM, while the

first integrator represents the outermost loop, that is more relaxed. This distinction becomes particularly important during overload conditions, when the $\Sigma\Delta$ -loop is temporarily disrupted, and the system must rapidly regain its proper operating point once restored. Having wrong order in setting can lead to prolonged times in overload recovery or in worst case systems that are prone to instability.

Given these considerations, the specifications for all three opamps differ substantially, rendering a single design suboptimal in terms of power optimization. Three distinct sets of specifications for each amplifier were therefore derived and implemented. The key parameters considered were:

- Noise: The first integrator requires the lowest noise performance, as it directly impacts the overall ADC noise floor.
- Linearity: The first integrator also demands best linearity to ensure accurate signal processing.
- Bandwidth: The third integrator commands high bandwidth for accurate loop stability.
- Stability: Internal loops prioritize higher BW and better loop stability for rapid convergence.
- Power Consumption: A consistent power budget was maintained for each opamp by carefully balancing trade-offs between various performance specifications considered above.
 - ✓ Specifications for opa1: Thermal Noise @2MHz: 3.5nV/√Hz, PM > 35°, PL < 1°
 - ✓ Specifications for opa2: Thermal Noise @10MHz: 5nV/√Hz, PM > 40°, PL < 2°
 - ✓ Specifications for opa3: Thermal Noise @10MHz: 8nV/√Hz, PM > 45°, PL < 5°

6.4.3 Design of Integrator1

Integrator1 (Figure 20a) comprises a resistor (R_1), a capacitor (C_1), and an amplifier (opa1). The product of C_1 and R_1 determines the integrator's integration frequency. C_1 selects the coarse frequency for mode programming, while R_1 is used to 5-bit fine-tune the accuracy to accommodate process variations.

All capacitors have been designed with the same design strategy, where the total capacitor has been divided into three parts, consisting of 75%, 25% and 100% of the capacitor needed to build the coefficients for the highest BW (W320) mode (Figure 20b). In that way that capacitor can be programmed to 75%, 100%, 175% and 200% of that value, enabling the usage of all planned modes. Programming the capacitor to 75% or 175% of the W320 value is not strictly necessary for the system's operation but is a good practice to maintain flexibility for testing purposes and to explore the potential of more aggressive coefficients in future development. In addition to the segmented capacitor, a resistor is placed across the capacitors to perform a reset function of the filter, that can be used in case of ADC overload, which will be discussed in a later section.

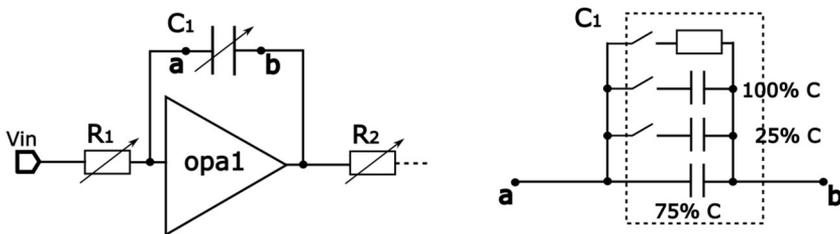


Figure 20. a) Integrator1 of the loop filter; b) C_1 of the loop filter used for mode programming.

Resistor R_1 has two inputs: a primary input from the RX path and a secondary “loopback” input for calibration and testing purposes (Figure 21). The primary input is connected to a 5-bit tunable resistor that can be configured to operate in either 250 Ω or 500 Ω mode. The loopback input, on the other hand, has a 2-bit tunable resistor that only supports 250 Ω mode. The loopback switch is limited to 2-bit control to reduce the number of switches connected to the virtual ground of the opa1, to keep the parasitics manageable. Each switch introduces unwanted capacitance to the virtual ground node, which degrades the bandwidth (BW) and phase margin (PM) of the opamp, resulting in reduced performance. As a result, the loopback resistor has been sacrificed to prioritize the full performance in the RX path. While less tuning has some impact on gain accuracy and large-signal linearity when using the loopback input, it is not a significant limitation for the necessary testing and calibration purposes.

The resistor matrix is constructed using 3.5k Ω unity resistors, each coupled with a proportionally sized switch to maintain a constant switch channel resistance. Typically, connecting switches to nodes with an input common-mode voltage requires the use of a transmission gate (T-gate), which incorporates both NMOS and PMOS devices. However, to further minimize parasitic capacitance to sensitive virtual ground nodes, only NMOS switches were employed in this design. Techniques like boosting the switch control signals with a charge pump have been successfully employed in other similar designs to restore the switch operating point to their linear region. However, this comes at the expense of increased complexity, as careful startup design and verification are necessary to ensure proper product life and reliability. This is because the gate voltage of the device would exceed the process node’s natural limits. Additionally, charge pump-boosted control signals introduce switching noise, which could potentially cause unwanted disturbances. As a result, it was decided to utilize non-boosted NMOS switches and size them such that the switch channel resistance is maintained at a constant 5% of the unity resistor. This choice did introduce some distortion to the ADC, but was deemed to be the better choice.

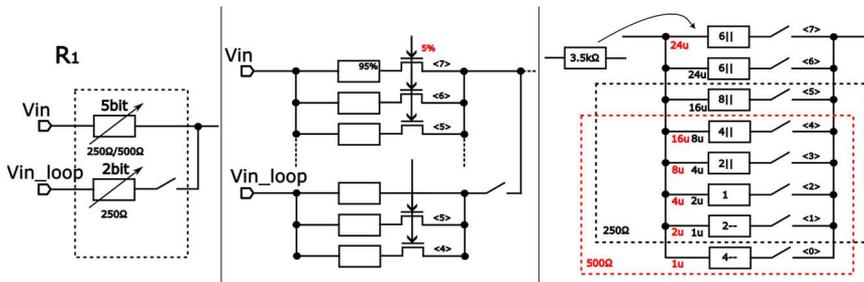


Figure 21. 5-bit, double input, tunable R_1 of the loop filter.

Resistor R_1 is constructed from 3.5k Ω unity resistors, enabling a total tuning range of 60% to 160% to accommodate the process extremes in manufacturing. The resistor is implemented using an internal 6-bit binary tuning scheme, which is externally controlled by a 5-bit word. When operating in 250 Ω mode switches <7> and <6> are closed and the 5-bit control is realized with branches <5>...<1>. In contrast, when operating in 500 Ω mode, switch <6> is opened and the 5-bit control is realized with switches <4>...<0>. This configuration enables the implementation of both the 250 Ω and 500 Ω modes using a single 5-bit control word. An intuitive way of looking at this is to view the resistor bank as 55 binary weighted parallel resistor branches, where a minimum of 24 branches (tuning word 00000) or a maximum of 55 branches (tuning word 11111) can be selected.

The nominal value is achieved with 38 branches (tuning word 01110), where branch resistance of one is equivalent to the resistance of 4 series unit resistors, or the LSB value of the overall resistor, equaling 14k Ω .

6.4.3.1 Design of opa1

A unified synthesis methodology that maintains a target transconductance/Drain-Source current ratio (gm/I_{ds}) versus the normalized current is applied to the design of all opamps [25]. This method aims to hit a target gm/I_{ds} ratio between 15 and 25, which provides a suitable compromise for definitely keeping the devices operating in weak inversion and gives a simple tool for optimizing the W/L ratios of individual devices. A perfectly sized MOS transistor exhibits a large gm without having a high C_{gs} or consuming excessive current. Therefore, gm/I_{ds} effectively represents the efficiency of a MOS transistor in converting current into equivalent transconductance. This parameter reaches a maximum in weak inversion and decreases as the operating point shifts toward strong inversion. While a lower gm/I_{ds} ratio is sometimes preferable for stages responsible for driving large capacitive loads, the forementioned gm/I_{ds} ratio serves as a sound design guideline.

The design of opa1 prioritizes low noise and high linearity, with the expense of extra BW and stability. This emphasis on linearity necessitates a more complex circuit topology. The four primary sources of nonlinearity in the integrator are the finite loop gain of the amplifying circuit, the inherent nonlinearity of the AB stage, the nonlinearity of the signal path switches used for mode and input programming, and the nonlinearity of the resistors. This is also why the number of capacitor tuning modes is kept to a minimum. Placing a switch in series with a capacitor introduces an unwanted zero in the transfer function of the integrator and adds nonlinearity to the filter. With proper attention to input resistor design, ensuring that the switches and resistors are physically sized correctly, the dominant nonlinearity will originate from the amplifier itself. This is a critical consideration in modern high-speed data communications, which require full-rail swing operation for maximum throughput. Therefore, advanced circuit topologies with optimized performance and power consumption are essential and have been employed in the design of the CT Σ ADC.

An n-stage opamp will have at least n high-impedance nodes, that will all introduce a pole. A common design strategy, regardless of the topology, has been to place a dominant pole in the low-frequency range, a second pole approximately three times the GBW, and push all other poles out to frequencies where they can be considered non-dominant or negligible. This approach yields a well-defined phase margin of around 60 degrees and a safe first-order roll-off before the GBW. An alternative design strategy involves employing more complex feed-forward (FF) structures to compensate in-band poles with zeros at the same location. This technique allows for the exact compensation of all poles at their source, while creating pole-zero doublets. However, these doublets sometimes do not match, leading to unpredictable roll-off characteristics and a less well-defined phase margin. However, both of these methods fail to fully leverage the potential performance advantages of multi-stage opamp topologies.

To achieve maximum performance, all in-band poles should be pushed just around the band-edge, maximizing loop gain within the signal bandwidth and leaving them uncompensated until just before the GBW. At this point, as many zeros as possible should be placed to regain the lost phase. The design of opamps with favorable performance-to-power ratios often necessitates a paradigm shift. On the one hand, good linearity requires loop gain throughout the entire signal bandwidth, while on the other hand, a lower GBW is desirable to conserve power by placing non-dominant poles at higher frequencies. The GBW significantly impacts the opamp's power consumption,

as non-dominant poles must be designed with margin and pushed to higher frequencies. These seemingly contradictory requirements pose a challenge in achieving efficient design. The true value of feed-forward structures lies in their ability to enable the designer to place two or more poles in close proximity, resulting in a higher-order (more than second-order) roll-off. This can be compensated for by one or more zeros from the feed-forward branches before the GBW [26]. This allows for the creation of systems where the loop gain remains relatively constant within the required signal bandwidth, experiences a rapid (more than second-order roll-off) roll-off between the required bandwidth, and levels out to a first-order (or similar) just before the GBW, keeping the GBW itself relatively low. The high number of stages required for the design of such opamps serves only to create extra high-impedance nodes in the signal path and not for amplification. The high gain often associated with multi-stages topology arises as a byproduct, without any additional effort, for free (Figure 22).

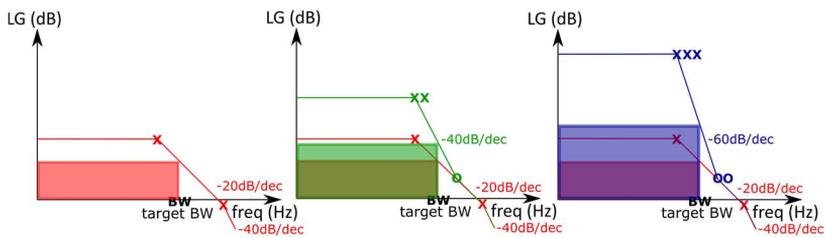


Figure 22. The difference of potential performance between systems with various roll-offs.

For opa1, the design targets were a $3.5\text{nV}/\sqrt{\text{Hz}}$ spot noise @ 2MHz, PL of less than 1° and a PM of more than 35° . The noise contribution of an opamp in an integrator configuration primarily originates from the input stage. It comprises both colored $1/f$ noise components and white noise components. In advanced CMOS processes, such as the ones used in this ADC design, the noise corner frequency, where the dominance of colored noise transitions to white noise, extends into the MHz range, indicating that both noise components play an equally crucial role in maintaining good thermal noise performance at 2MHz. Consequently, the input stage of the opamp must be both large in area, to mitigate colored noise, and consume a substantial share of the overall power consumption, to suppress white noise. The PL is a product of the open-loop gain of the opamp at the integrator's UGF, as seen from the figure above. As such, the ability to position several poles in close proximity, leading to a steep roll-off, can significantly enhance the open-loop gain over specific frequencies, ultimately increasing the performance of the integrator.

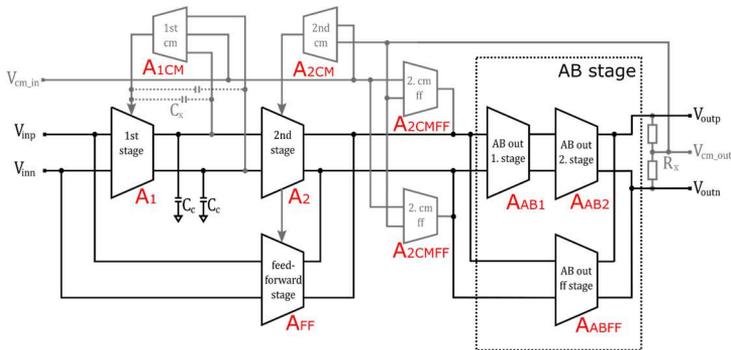


Figure 23. Topology of opa1.

The designed amplifier employs a fully differential four-stage double feed-forward compensated signal path, a two-stage Miller compensated common-mode loop, and a four-stage feed-forward compensated common-mode loop (Figure 23). The feed-forward compensation ensures a sufficiently high loop gain across the entire desired operating bandwidth, where poles from A_1 , A_2 and A_{ABOUT} offer a fast roll off and A_{FF} and A_{ABFF} introduce a zero before the UGF. At DC and low frequencies, the gain from A_1 , A_2 and A_{AB1} and A_{AB2} is dominant and provides the gain and linearity across the entire operating band of interest. The feed-forward stages A_{FF} and A_{ABFF} can be viewed as loads that reduce the total output impedance at the outputs of the second stage and at the AB stage, which in turn slightly diminishes the gain from that stage. The feed-forward stages become dominant once the gain of the cascaded stages drops below the gain of the single feed-forward stage itself. $A_{\text{FF}} > A_1 * A_2$ and $A_{\text{ABFF}} > A_{\text{AB1}} * A_{\text{AB2}}$. At frequencies approaching the GBW, the opamp essentially operates as a two-stage amplifier comprising of stages A_{FF} and A_{ABFF} .

The pole placement becomes extremely critical in this kind of complex systems. The poles introduced by A_1 , A_2 , A_{FF} , A_{AB1} , and A_{AB2} should be strategically positioned immediately after the desired signal bandwidth, as closely as possible, to maximize their beneficial effects and roll-off behavior. Stages A_2 and A_{FF} are intentionally designed to be similar and to have the same corner frequency. While this is not mandatory, it simplifies layout matching considerations and common-mode control scenarios, as they can now be controlled by a single stage $A_{2\text{CM}}$ (Figure 24). This means that transition from medium to high frequencies coming from the first and second stage being dominant to the feed-forward stage being dominant, will result in a change from second-order roll-off to a first order roll-off in the first part of the amplifier. In the AB stage however the feed-forward stage A_{ABFF} is designed to have a real non-dominant pole in a much higher frequency, so that it could provide a double zero very close to GBW. When crossing the 0dB gain point, the system can behave much like a single pole amplifier, with the single dominant pole coming from A_{FF} . This means the only current hungry stage is the A_{ABFF} , which in reality will consume close to half of the total current of the amplifier.

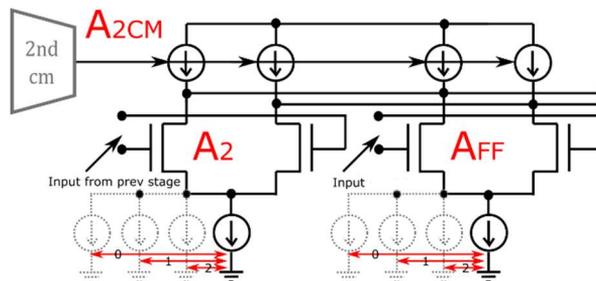


Figure 24. Second common-mode loop control.

The designed amplifier has two common-mode loops aimed at fixing the common-mode voltage of the opamp output. These loops are distinct in their composition, sharing certain stages with the signal path while also incorporating unique stages. The primary objective of the first loop is to translate the system-specified common-mode reference voltage ($V_{\text{CM_IN}}$) into a predefined common-mode voltage at the output of the initial stage. This loop is built as a two-stage Miller-compensated sub-amplifier comprising of stages $A_{1\text{CM}}$ and A_1 . The dominant pole within the loop is strategically placed at the first gain

stage (A_{1CM}), ensuring that the second pole stemming from the input stage (A_1), as previously discussed, integrated within, and devised for the signal path, remains at least three times further from the common-mode loop's GBW. This deliberate placement is intended to attain a highly stable loop with a phase margin exceeding 60 degrees. The GBW of the first common-mode loop is not critical for the overall system recovery in this instance. Its purpose is solely to generate and fix a common-mode voltage for the signal at the input of the second amplifying stage, A_2 . This loop does not contribute at BWs where the stability is determined, particularly when the opamp functions as a two-stage amplifier, deriving its initial amplification properties from A_{FF} rather than the recovery of the first common-mode loop. Consequently, the bandwidth of the first common-mode is kept rather low and compensated using Miller capacitors C_C . Given the substantial dimensions of the first stage of the signal path, driven by stringent noise requirements for opa1, the capacitive load seen by A_{1CM} is in general so large, potentially obviating the need for Miller compensation for the first CM-loop.

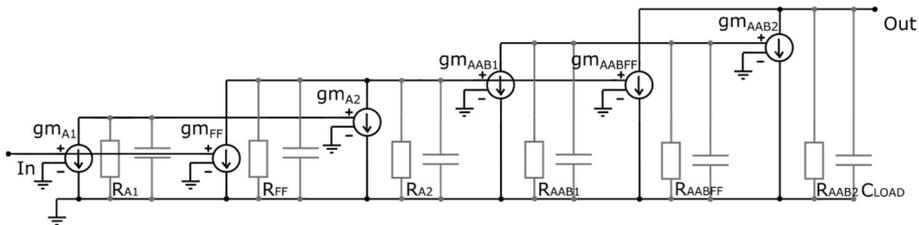


Figure 25. Small signal equivalent of the signal path in opa1.

The second CM loop bears a close resemblance to the designed differential path, employing a four-stage double feed-forward compensated loop comprising the four gain stages A_{2CM} , A_2 , A_{AB1} , and A_{AB2} , where the two feed-forward paths are now A_{2CMFF} and A_{ABFF} . The inputs for this loop are the system-provided common mode reference voltage (V_{CM_IN}) and the filtered common mode voltage of the outputs averaged across the R_{XS} (V_{CM_OUT}). Similar to the signal path, if A_{2CM} possesses a pole similar to A_1 and A_{2CMFF} a pole similar to A_{FF} , the system will exhibit a two-stage opamp-like behavior near the UGF, with a GBW closely resembling that of the signal path. In this configuration, A_{2CMFF} effectively acts as the dominant pole, while A_{ABFF} remains a non-dominant one. In practical scenarios, the phase margin of the common mode loop needs to be slightly higher than that of the overall amplifier to ensure robust stability and recovery. Therefore, the pole placement and bandwidth of A_{2CM} are intentionally not perfectly matched with those of A_1 . Instead, A_{2CM} 's design prioritizes a higher phase margin in the CM loop.

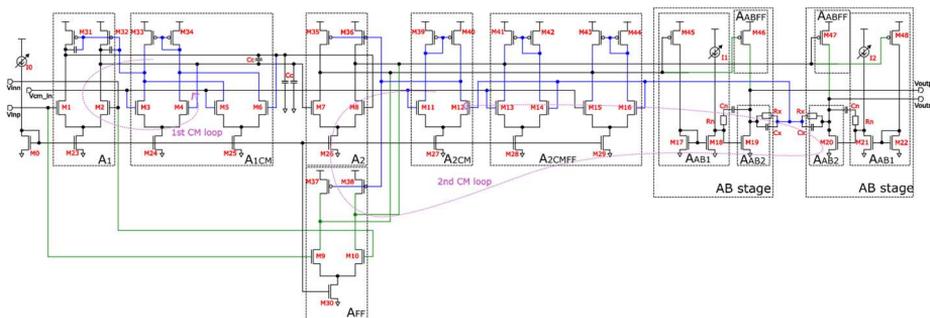


Figure 26. Schematic of opa1.

The fully-differential topology of the opamp (Figure 26), which is necessary for most high-performance systems to reject potential unwanted noise and disturbers needs also careful planning in layout. While all gain stages except for the AB-stage are composed of plain differential pairs with differential outputs, it gives the distinct advantage, that one does not need to worry about polarities of the signal as both inverted and non-inverted outputs are always available. AB-stage is built from a direct path in the P-MOS side, that also acts as a feedforward path and a dual amplifying stage N-MOS side, where the first stage is a standard fixed current source gain stage. All paths in the filter and in the opamps should form non-crossing loops with minimal length paths to be effective and high-performing (Figure 27a). The dimensions of the amplifier are drawn approximately 60um*60um in layout (Figure 27b).

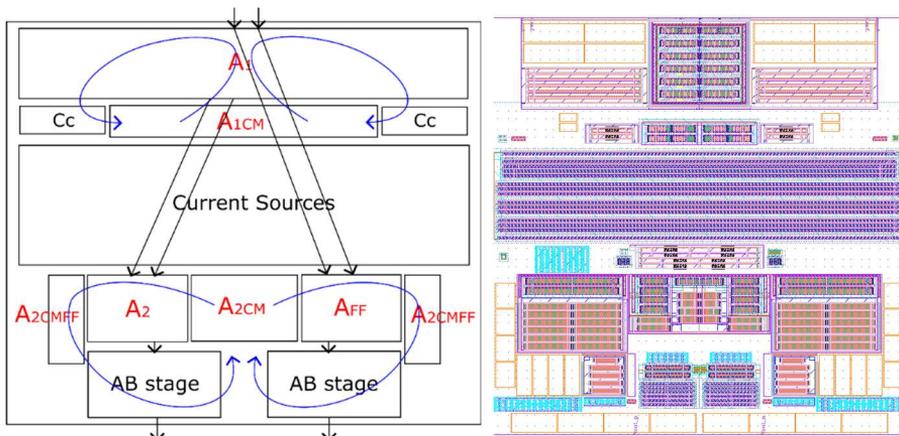


Figure 27. a) Floorplan of opa1; b) Layout of opa1.

6.4.3.2 Programming of opa1

Opa1 incorporates four built-in bandwidth modes, just like the other two amplifiers in integrator2 and integrator3, with two additional unique features tailored to its specific design. The BW programming of the amplifier is targeted to fulfill the PL specification of 1° for different primary Wi-Fi modes. The target opa1 current consumption with each programming step to reach PL of 1° is dedicated as 0:1:3, where:

- 0: used for W80 mode; Target $I_Q \sim 1.5\text{mA}$
- 1: used for W160, W80DPD and W160LP; Target $I_Q \sim 3.0\text{mA}$
- 2: used for W320 and W160DPD; Target $I_Q \sim 4.5\text{mA}$
- 3: used for W320DPD; Target $I_Q \sim 6.0\text{mA}$

Mode programming is typically achieved in all opamps by only changing the biasing current in sources I_0 , I_1 and I_2 . This approach changes the output impedance of all amplifying stages, simultaneously shifting the poles associated with each high-impedance node and the zeros linked to the FF-stage crossings. This technique maintains the amplifier's structure intact, avoiding the need to add or remove branches for individual device programming, thereby minimizing parasitic effects. The only notable exception to this method is the first stage programming in opa1. Due to the critical importance of noise performance, the tail current in this stage is kept constant over modes to preserve the much-needed noise floor. Instead, bandwidth programming is implemented by modifying the capacitive load C_c at the output of the first stage. The significance of programming the opamp without altering its topology is particularly evident in zero placement, which

is accomplished using the FF-branch A_{FF} . When A_{FF} perfectly matches A_2 , the location of the double poles arising from A_1 and A_2 compared to A_{FF} stays constant over distance in frequency domain. This ensures that the transfer function shape remains unchanged regardless of the biasing current for I_0 . This approach is further reinforced in the layout planning, where not only are all differential structures constructed in a fully differential manner, but also designed with a half-cell mindset, where stages requiring matching are delicately physically designed in the most efficient manner.

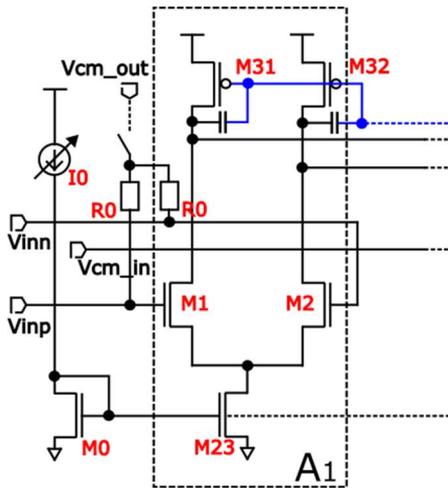


Figure 28. CM self-generation.

capabilities, as previous measurement approaches had limitations coming from noise, primarily stemming from the measurement equipment employed to generate the correct common-mode voltage for the ADC. While signal generation equipment typically demonstrates excellent noise performance when generating a differential signal without offset, adding a DC voltage includes the caveat of reduced performance, which is not ideal for accurate characterization. To fight that, often an external buffer is placed on the Printed Circuit Board (PCB), introducing more noise and degrading linearity during measurements. To circumvent these issues, an option was introduced to capacitively decouple the input on the PCB, eliminating the need to generate an external common-mode voltage. However, this option necessitates the ability to generate an internal common-mode voltage directly at the source. Therefore, the amplifier's self-generated V_{CM_OUT} signal was repurposed to bias the amplifier through a high-impedance resistor R_0 (Figure 28). This feature has proven to be remarkably effective in eliminating external noise originating from circuitry on the PCB, providing an opportunity to accurately assess the full performance of the standalone ADC.

6.4.3.3 Simulation Summary of opa1

The post-layout simulated results of the important parameters of opa1 are summarized in Table 1.

In addition to the bandwidth programming, there exists an optional current boosting mechanism that injects an additional 1 mA into the AB-stage across all modes, thereby enhancing its driving strength. This feature was primarily conceived as a debugging tool to ensure the amplifier's ability to drive the succeeding integrator in 6-bit mode, where the filter resistors are halved. However, this feature proved to be unnecessary in practice, as opa1 demonstrated the capability to handle all loads effectively without additional boosting needed.

The opamp's last programmable unique feature is the capacity to generate its own common-mode voltage. This addition was implemented to enhance the measurement

Table 1. Simulation summary of opa1.

W320					W160LP					
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used					Corners Used					
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	45.3	66.8	77	Phase Margin of the Amplifier	°	50.7	65.1	79.2	
GBW (signal path)	GHz	2.04	2.58	3.58	GBW of the Amplifier	GHz	1.22	1.51	2.01	
Gain (signal path)	dB	92.5	128	145	DC Gain of the Amplifier	dB	90	128	145	
PM (1. CM loop)	°	64.1	68.1	71.9	Phase Margin of the 1. Common Mode	PM (1. CM loop)	°	64.1	70.0	73.3
GBW (1. CM loop)	MHz	94.8	126	162	GBW of the 1. Common Mode	GBW (1. CM loop)	MHz	65.6	86.8	135
Gain (1. CM loop)	dB	88.2	94.4	97.5	DC Gain of the 1. Common Mode	Gain (1. CM loop)	dB	88.2	94.7	97.6
PM (2. CM loop)	°	52.9	56.9	75.2	Phase Margin of the 2. Common Mode	PM (2. CM loop)	°	53.4	51.5	65.5
GBW (2. CM loop)	MHz	552	797	919	GBW of the 2. Common Mode	GBW (2. CM loop)	GHz	381	566	662
Gain (2. CM loop)	dB	116	147	163	DC Gain of the 2. Common Mode	Gain (2. CM loop)	dB	117	149	164
Phase Loss	°	0.3	0.8	1.3	Phase Loss of the Integrator	Phase Loss	°	0.2	0.4	0.9
IQ	mA	3.5	4.5	6.3	Quiescent Current	IQ	mA	2.5	3.3	4.7
Noise @ 2MHz	nV/√Hz	3.1	3.8	4.5	Input referred Noise	Noise @ 2MHz	nV/√Hz	3.2	3.9	4.6
W160					W80					
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used					Corners Used					
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	49.2	61.5	76.5	Phase Margin of the Amplifier	PM (signal path)	°	45.2	56.7	68
GBW (signal path)	GHz	1.17	1.44	1.92	GBW of the Amplifier	GBW (signal path)	GHz	0.59	0.73	1.03
Gain (signal path)	dB	90.0	127.6	145.0	DC Gain of the Amplifier	Gain (signal path)	dB	84.6	126	144
PM (1. CM loop)	°	64.1	70.0	73.3	Phase Margin of the 1. Common Mode	PM (1. CM loop)	°	67.9	73	75.3
GBW (1. CM loop)	MHz	65.5	86.8	135	GBW of the 1. Common Mode	GBW (1. CM loop)	MHz	33.7	44.2	71.2
Gain (1. CM loop)	dB	88.2	94.7	97.6	DC Gain of the 1. Common Mode	Gain (1. CM loop)	dB	87.2	94.9	97.7
PM (2. CM loop)	°	49.3	52.5	66.6	Phase Margin of the 2. Common Mode	PM (2. CM loop)	°	42.1	44.8	52.5
GBW (2. CM loop)	GHz	365	541	637	GBW of the 2. Common Mode	GBW (2. CM loop)	GHz	209	312	373
Gain (2. CM loop)	dB	118	149	164	DC Gain of the 2. Common Mode	Gain (2. CM loop)	dB	117	149	164
Phase Loss	°	0.2	0.4	0.9	Phase Loss of the Integrator	Phase Loss	°	0.2	0.4	0.9
IQ	mA	2.5	3.3	4.7	Quiescent Current	IQ	mA	1.5	1.9	2.9
Noise @ 2MHz	nV/√Hz	3.1	3.8	4.5	Input referred Noise	Noise @ 2MHz	nV/√Hz	3.2	3.9	4.6

6.4.4 Design of Integrator2

Integrator2 (Figure 29a) comprises of a resistor (R_2), a capacitor (C_2), and an amplifier (opa2). The product of C_2 and R_2 determines the integrator's integration frequency. C_2 selects the coarse frequency for mode programming, while R_2 is used to 5-bit fine-tune the accuracy to accommodate process variations.

Just like in C_1 , capacitors have been designed with the design strategy, where the total capacitor has been divided into three parts, consisting of 75%, 25% and 100% of the capacitor needed to build the coefficients for the highest BW (W320) mode (Figure 29b). In that way that capacitor can be programmed to 75%, 100%, 175% and 200% of that value, enabling the usage of all planned modes.

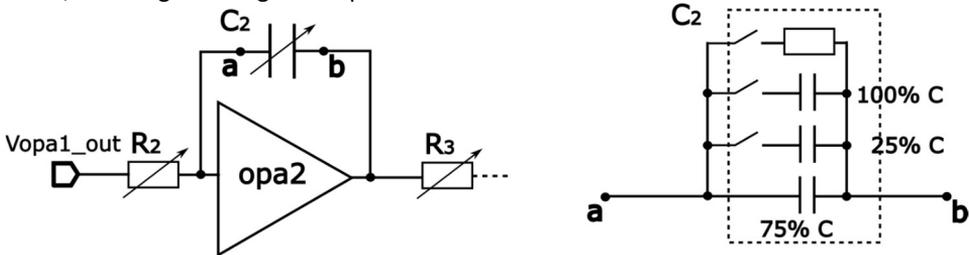


Figure 29. a) Integrator2 of the loop filter; b) C_2 of the loop filter used for mode programming.

Unlike R_1 , R_2 has only one input. It is connected to a 5-bit tunable resistor that can be configured to operate in either 500Ω or 1kΩ mode, depending if the ADC is operating in a 6-bit or 5-bit domain. In addition to R_1 , R_2 does have a built-in option to modulate the impedance of the resistor in frequency domain. This is achieved by adding a capacitance in series with a fixed resistance across the entire resistor bank. This test feature can be used to enhance the maximum BW of the ADC by shaping noise away in a more favorable manner.

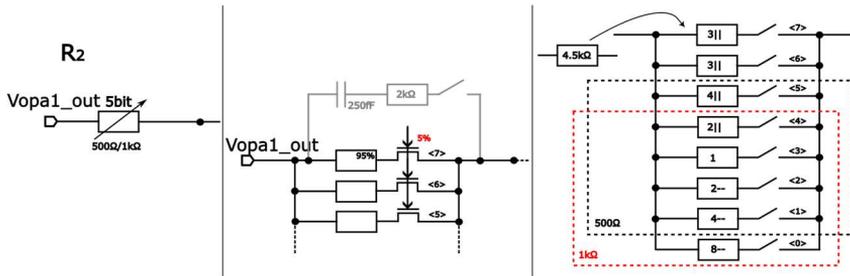


Figure 30. 5-bit, tunable R_2 of the loop filter.

Resistor R_2 is constructed from $4.5k\Omega$ unity resistors, enabling a total tuning range of 60% to 160% to accommodate the process extremes in manufacturing (Figure 30). The resistor is implemented using an internal 6-bit binary tuning scheme, which is externally controlled by a 5-bit word. When operating in 500Ω mode switches <7> and <6> are closed and the 5-bit control is realized with switches <5>...<1>. In contrast, when operating in $1k\Omega$ mode, switch <6> is opened and the 5-bit control is realized with switches <4>...<0>. This configuration enables the implementation of both the 500Ω and $1k\Omega$ modes using a single 5-bit control word.

6.4.4.1 Design of opa2

The design of opa2 tries to strike a balance between linearity, bandwidth, stability and power consumption, aiming for an optimal compromise among all aspects. Similar to opa1, opa2 employs a fully differential four-stage double feed-forward compensated signal path, one two-stage Miller compensated common-mode loop, and one four-stage feed-forward compensated common-mode loop. However, opa2 distinguishes itself with the extra ability to transform itself into a three-stage structure by disabling the input stage A_1 , the first common mode stage A_{1CM} , and the second stage A_2 (Figure 31). In such a case the Feed-Forward stage A_{FF} is used as the input and the opamp turns into a fully differential three-stage feed-forward compensated structure, with a four-stage feed-forward compensated common-mode loop. The feasibility arises from the relaxed noise specifications, eliminating the need for an overly large input stage. By switching off the second stage, the first common stage which provides it with V_{CM} also loses all meaning, so it too can be turned off. The input stage of the second common mode loop A_{2CM} in needs in this case only to control the FF-stage A_{FF} . The deactivation of these stages is accomplished by shutting down the tail current sources associated with the mentioned stages.

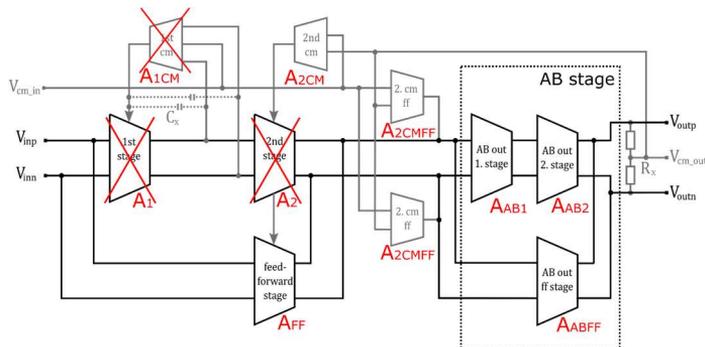


Figure 31. Topology of opa2.

A significant advantage of this topology lies in the amplifier's phase margin (PM) being relatively unaffected by the choice between four-stage and three-stage configurations. The "two poles with a zero" transfer function inherent in the four-stage configuration transitions smoothly to a single-pole transfer function characteristic of the three-stage configuration, particularly in the vicinity of the GBW. Even the extra parasitics remain the same. The primary distinction emerges from the elimination of the second stage, A_2 , which no longer influences the output impedance, resulting in a marginal PM shift when switching between the two topologies (Figure 32). In actual applications, the amplifier was exclusively operated in a three-stage configuration. Both simulation and measurement results indicated a negligible enhancement in overall performance, failing to outweigh the additional power consumption associated with the four-stage configuration.

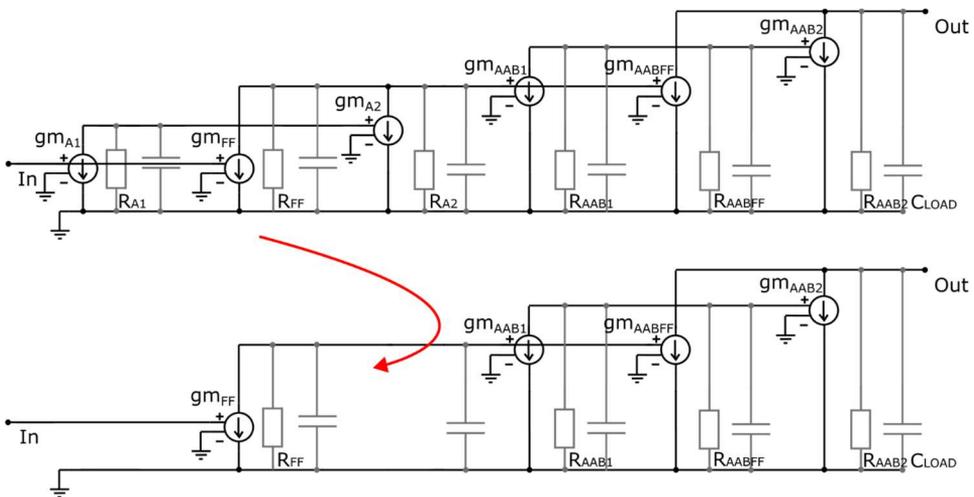


Figure 32. Different configuration small signal equivalent of the signal path in opa2.

Layout of opa2 is done in a similar manner as the layout for opa1, with the main difference being the size of the input stage, which is now considerably smaller. Opa2 measures approximately 50um*60um in layout.

6.4.4.2 Programming of opa2

Opa2 also incorporates four built-in bandwidth modes, just like opa1, with the switching to four-stage topology adding roughly 25% of additional current in the power budget of the opamp. The BW programming of the amplifier is targeted to fulfill the PL specification of 2° for different primary Wi-Fi modes. The target opa2 current consumption with each programming step to reach PL of 2° is dedicated as 0:1:3, where:

- 0: used for W80 mode; Target $I_Q \sim 1.5\text{mA}$
- 1: used for W160, W80DPD and W160LP; Target $I_Q \sim 3.0\text{mA}$
- 2: used for W320 and W160DPD; Target $I_Q \sim 4.5\text{mA}$
- 3: used for W320DPD; Target $I_Q \sim 6.0\text{mA}$

In contrast to opa1, opa2's mode programming is entirely accomplished by adjusting the biasing currents in sources I_0 , I_1 , and I_2 , encompassing the first stage as well. This approach introduces minor variations in the amplifier's noise contribution across different modes. However, these changes are not significant due to the first-order noise shaping employed by the amplifier, ensuring minimal impact to the overall noise budget.

6.4.4.3 Simulation Results of opa2

The post-layout simulated results of the most important parameters of opa2 configured in 4stage topology are summarized in Table 2 and the most important parameters of opa2 configured in 3stage topology are summarized in Table 3.

Table 2. Simulation summary of opa2, configured in 4stage configuration.

W320				W160LP						
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	45.1	59.7	70.5	Phase Margin of the Amplifier	PM (signal path)	°	41.0	55.4	70.0
GBW (signal path)	GHz	3.32	3.91	5.11	GBW of the Amplifier	GBW (signal path)	GHz	1.91	2.31	2.98
Gain (signal path)	dB	97.2	127	142	DC Gain of the Amplifier	Gain (signal path)	dB	96.1	128	143
PM (1. CM loop)	°	74.1	77.9	80.2	Phase Margin of the 1. Common Mode	PM (1. CM loop)	°	75.0	78.7	80.9
GBW (1. CM loop)	MHz	55.6	74.4	117	GBW of the 1. Common Mode	GBW (1. CM loop)	MHz	38.7	52.2	82.9
Gain (1. CM loop)	dB	90.5	96.9	100	DC Gain of the 1. Common Mode	Gain (1. CM loop)	dB	91.7	97.8	101
PM (2. CM loop)	°	43.8	49.2	66.4	Phase Margin of the 2. Common Mode	PM (2. CM loop)	°	39.9	49.8	58.3
GBW (2. CM loop)	MHz	0.61	0.93	1.06	GBW of the 2. Common Mode	GBW (2. CM loop)	GHz	0.43	0.56	0.77
Gain (2. CM loop)	dB	118	146	160	DC Gain of the 2. Common Mode	Gain (2. CM loop)	dB	118	147	161
Phase Loss	°	0.5	1.1	2.0	Phase Loss of the Integrator	Phase Loss	°	0.3	0.6	1.1
IQ	mA	4.1	5.3	7.2	Quiescent Current	IQ	mA	2.6	3.4	4.7
Noise @ 10MHz	nV/√Hz	2.8	3.4	4.0	Input referred Noise	Noise @ 10MHz	nV/√Hz	3.3	4.1	4.8
W160				W80						
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	40.3	54.8	69.5	Phase Margin of the Amplifier	PM (signal path)	°	42.1	52.5	63.1
GBW (signal path)	GHz	1.84	2.24	2.90	GBW of the Amplifier	GBW (signal path)	GHz	0.77	0.96	1.35
Gain (signal path)	dB	96.1	128	143	DC Gain of the Amplifier	Gain (signal path)	dB	91.5	127	143
PM (1. CM loop)	°	75.1	78.8	80.9	Phase Margin of the 1. Common Mode	PM (1. CM loop)	°	75.8	79.2	81.1
GBW (1. CM loop)	MHz	38.6	52.1	82.9	GBW of the 1. Common Mode	GBW (1. CM loop)	MHz	19.6	26.9	43.6
Gain (1. CM loop)	dB	91.7	97.8	101	DC Gain of the 1. Common Mode	Gain (1. CM loop)	dB	92.1	98.8	102
PM (2. CM loop)	°	39.5	48.8	56.7	Phase Margin of the 2. Common Mode	PM (2. CM loop)	°	42.2	50.4	55.9
GBW (2. CM loop)	GHz	0.42	0.54	0.75	GBW of the 2. Common Mode	GBW (2. CM loop)	GHz	0.3	0.39	0.56
Gain (2. CM loop)	dB	118	147	161	DC Gain of the 2. Common Mode	Gain (2. CM loop)	dB	117	148	162
Phase Loss	°	0.3	0.5	1.0	Phase Loss of the Integrator	Phase Loss	°	0.4	0.8	1.5
IQ	mA	2.6	3.4	4.7	Quiescent Current	IQ	mA	1.5	1.9	2.9
Noise @ 10MHz	nV/√Hz	3.3	4.1	4.8	Input referred Noise	Noise @ 10MHz	nV/√Hz	5.3	6.6	7.8

Table 3. Simulation summary of opa2, configured in 3stage configuration.

W320				W160LP						
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	63.8	74.1	83.4	Phase Margin of the Amplifier	PM (signal path)	°	61.1	72.9	84.8
GBW (signal path)	GHz	3.53	4.10	5.52	GBW of the Amplifier	GBW (signal path)	GHz	1.95	2.35	3.00
Gain (signal path)	dB	75.5	94.8	106	DC Gain of the Amplifier	Gain (signal path)	dB	75.1	95.7	107
PM (2. CM loop)	°	56.5	61.3	81.2	Phase Margin of the 2. Common Mode	PM (2. CM loop)	°	51.3	63.4	73.5
GBW (2. CM loop)	GHz	0.79	1.24	1.43	GBW of the 2. Common Mode	GBW (2. CM loop)	GHz	0.56	0.74	1.04
Gain (2. CM loop)	dB	118	143	160	DC Gain of the 2. Common Mode	Gain (2. CM loop)	dB	118	146	161
Phase Loss	°	2.2	3.3	4.4	Phase Loss of the Integrator	Phase Loss	°	1.6	2.3	3.1
IQ	mA	3.2	4.1	5.6	Quiescent Current	IQ	mA	2.0	2.6	3.8
Noise @ 10MHz	nV/√Hz	3.2	3.8	4.4	Input referred Noise	Noise @ 10MHz	nV/√Hz	3.7	4.3	5.0
W160				W80						
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	60.6	72.3	84.3	Phase Margin of the Amplifier	PM (signal path)	°	58.5	67.9	77.4
GBW (signal path)	GHz	1.87	2.26	2.90	GBW of the Amplifier	GBW (signal path)	GHz	0.78	0.96	1.36
Gain (signal path)	dB	75.1	95.7	107.0	DC Gain of the Amplifier	Gain (signal path)	dB	74.3	95.3	106
PM (2. CM loop)	°	51.4	55.1	72.7	Phase Margin of the 2. Common Mode	PM (2. CM loop)	°	42.7	55.6	64.9
GBW (2. CM loop)	GHz	0.54	0.87	1.01	GBW of the 2. Common Mode	GBW (2. CM loop)	GHz	0.38	0.5	0.71
Gain (2. CM loop)	dB	118	146	161	DC Gain of the 2. Common Mode	Gain (2. CM loop)	dB	115	146	160
Phase Loss	°	1.5	2.2	3.0	Phase Loss of the Integrator	Phase Loss	°	1.7	2.6	3.4
IQ	mA	2.0	2.6	3.8	Quiescent Current	IQ	mA	1.3	1.6	2.6
Noise @ 10MHz	nV/√Hz	3.7	4.3	5.0	Input referred Noise	Noise @ 10MHz	nV/√Hz	5.7	6.9	7.9

6.4.5 Design of Integrator3

Integrator3 comprises a resistor (R_3), a capacitor (C_3), and an amplifier (opa3), with the product of C_3 and R_3 determining the integrator's integration frequency. C_3 selects the coarse frequency for mode programming, while R_3 is used to 5-bit fine-tune the accuracy to accommodate process variations.

Because the third integrator needs to support a very large integration frequency, its proper implementation is considered the most critical aspect of the entire $\Sigma\Delta$ ADC design. To achieve the desired integration bandwidths, the R_3 and C_3 product must be very small. However, if the resistor is chosen too small, it becomes difficult for opa2 to drive the resistive loads. Conversely, if the capacitor is too small, parasitic effects become dominant and can significantly impact signal propagation and system integrity (Figure 33a). To address this challenge, special care is dedicated to minimizing the distances between components connected to opa3's virtual ground, thereby minimizing parasitics and ensuring signal integrity. **This design consideration serves as the foundation for proper layout planning in the entire $\Sigma\Delta$ ADC, emphasizing its importance.**

Similar to C_1 and C_2 , the design strategy for C_3 employs a straightforward approach, dividing the total capacitor into three distinct segments: 75%, 25%, and 100% of the value required for the highest bandwidth W320 mode. This enables programming the capacitor to 75%, 100%, 175% and 200% of its original value, accommodating the desired operating modes. In addition, a small 100 Ω resistor has been placed in series with the 75% part of the capacitor, which is always connected. This resistor introduces a desired zero in the transfer function, intended to compensate for phase loss caused by the opamp (Figure 33b). The placement of the zero is inherently inaccurate due to the absence of fine-tuning and the influence of process variations. However, it serves as a helpful and useful aid rather than a definitive calibration mechanism.

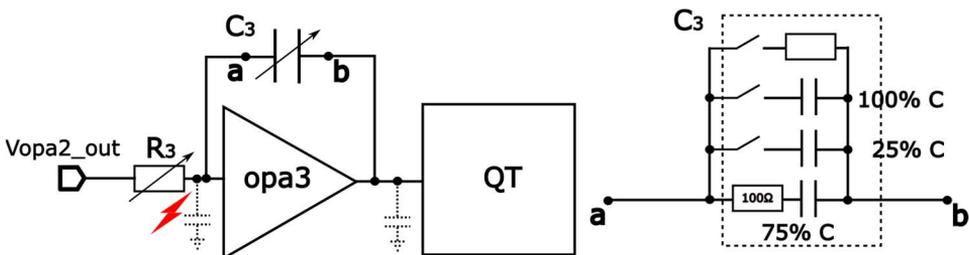


Figure 33. a) Integrator3 of the loop filter; b) C_3 of the loop filter used for mode programming.

Like R_2 , R_3 also has only one input. It is connected to a 5-bit tunable resistor that can be configured to operate in either 750 Ω or 1.5k Ω mode, depending if the ADC is operating in a 6-bit or 5-bit domain. Like R_2 , R_3 also does have a built-in option to modulate the impedance of the resistor in frequency domain. The same RC combination can be enabled as a test feature to enhance the maximum BW of the ADC by shaping noise away in a more favorable manner.

Resistor R_2 is constructed from 3.5k Ω unity resistors, enabling a total tuning range of 60% to 160% to accommodate the process extremes in manufacturing. The resistor is implemented using an internal 6-bit binary tuning scheme, which is externally controlled by a 5-bit word (Figure 34). When operating in 750 Ω mode switches <7> and <6> are closed, while switch <0> is opened and the 5-bit control is realized with switches <5>...<1>. In contrast, when operating in 1.5k Ω mode, switch <7> is closed, switches <6>

and <5> are opened and the 5-bit control is realized with switches <4>...<0>. This configuration enables the implementation of both the 750Ω and 1.5kΩ modes using a single 5-bit control word.

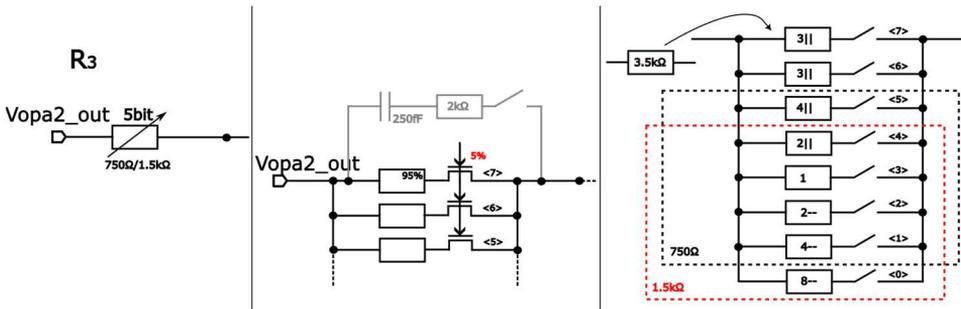


Figure 34. 5-bit, tunable R3 of the loop filter.

6.4.5.1 Design of opa3

The design of opa3 prioritizes extraordinary BW and high stability over linearity and noise. It forms the innermost loop of the $\Sigma\Delta$ ADC feedback system, requiring the highest signal bandwidth among all three amplifiers, along with the highest common-mode bandwidth. To achieve this while minimizing power consumption, the opamp comprises of only two stages in the signal path, with a three-stage feed-forward compensated common-mode loop. Unlike the two-stage AB stage utilized in opa1 and opa2, opa3 employs a pseudo-AB stage that gradually activates its push-pull functionality over frequency. In the DC and low-frequency regions, the stage behaves as a conventional A-class amplifier, with the NMOS side being biased through M18/M21 and acting solely as a current source (Figure 35). As the frequency increases, the NMOS side is progressively activated, based on the RC product of RW and CW, enabling the full push-pull functionality. This structure exhibits limited swing at lower frequencies, rendering it unsuitable for applications requiring rail-to-rail operation. However, it possesses a unique feature of increasing gain over frequency, occurring when the NMOS side of the AB stage becomes active. This property aligns perfectly with applications with heavy noise shaping, such as the third integration within an SD filter, where performance is particularly crucial in the vicinity of the band edge.

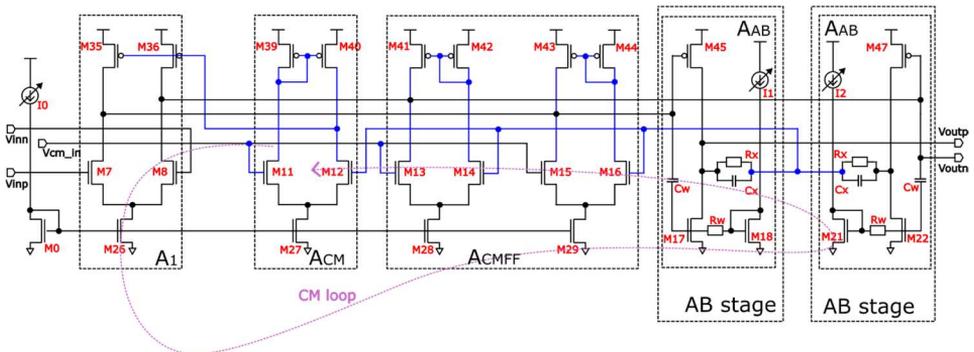


Figure 35. Schematic of opa3.

6.4.5.2 Programming of opa3

Opa3 not only boasts the simplest topology but also exhibits the most modest feature set. It incorporates the same two-bit bandwidth programming as found in the other two opamps. The BW programming of the amplifier is targeted to fulfill the PL specification of 5° for different primary Wi-Fi modes, without the help of the series resistor in the feedback path of the integrator. The target opa3 current consumption with each programming step to reach PL of 5° is dedicated as 0:1:3, where:

- 0: used for W80 and W80DPD mode; Target IQ ~2mA
- 1: used for W160, W160DPD and W160LP; Target IQ ~4mA
- 3: used for W320, W320DPD; Target IQ ~8mA

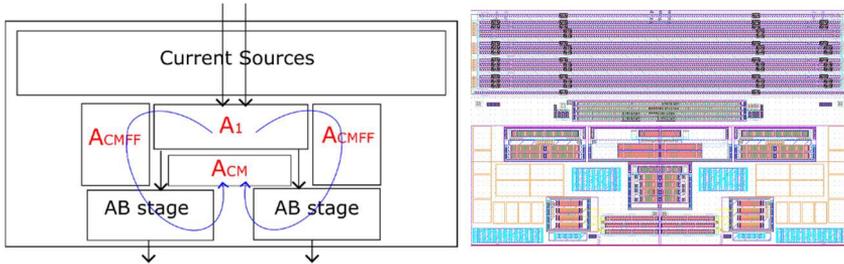


Figure 36. a) Floorplan of opa3; b) Layout of opa3.

The layout of opa3 is simplified due to the lack of additional amplifying stages. Since the amplifier exhibits only extremely high BWs, component placement and parasitic reduction is of utmost importance. The current source bank has for that reason been placed on one edge of the floorplan (Figure 36a), so that the internal loops would only form very short spirals with minimum extra capacitances. Opa3 measures approximately 40um*60um in layout (Figure 36b).

6.4.5.3 Simulation Summary of opa3

The post-layout simulated results of the important parameters of opa3 are summarized in Table 4.

Table 4. Simulation summary of opa3.

W320				W160LP						
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used				Nominal, Fast, Slow, SlowMOS_FastrC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	Corners Used				Nominal, Fast, Slow, SlowMOS_FastrC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	44.6	47.2	51.4	Phase Margin of the Amplifier	PM (signal path)	°	52.3	55.9	61.5
GBW (signal path)	GHz	10.6	11.7	12.9	GBW of the Amplifier	GBW (signal path)	GHz	6.53	7.02	8.03
Gain (signal path)	dB	52.5	60.7	65.5	DC Gain of the Amplifier	Gain (signal path)	dB	53.3	62.2	66.5
PM (CM loop)	°	78.1	85	88.4	Phase Margin of the Common Mode	PM (CM loop)	°	72.3	79.5	86
GBW (CM loop)	GHz	3.97	5.28	7.43	GBW of the Common Mode	GBW (CM loop)	GHz	2.05	2.33	3.04
Gain (CM loop)	dB	91.7	112	120	DC Gain of the Common Mode	Gain (CM loop)	dB	84.8	112	120
Phase Loss	°	4.3	5.0	5.3	Phase Loss of the Integrator	Phase Loss	°	3.5	4.3	4.7
IQ	mA	6.8	8.8	11.5	Quiescent Current	IQ	mA	3.5	4.5	6
Noise @ 10MHz	nV/√Hz	3.4	3.8	4.5	Input referred Noise	Noise @ 10MHz	nV/√Hz	3.6	4.1	4.9
W160				W80						
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment	
Corners Used				Nominal, Fast, Slow, SlowMOS_FastrC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	Corners Used				Nominal, Fast, Slow, SlowMOS_FastrC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	
		Min	Typ	Max			Min	Typ	Max	
PM (signal path)	°	48.9	52.9	56.5	Phase Margin of the Amplifier	PM (signal path)	°	57.7	62.6	69.3
GBW (signal path)	GHz	6.40	6.90	7.90	GBW of the Amplifier	GBW (signal path)	GHz	3.99	4.37	5.14
Gain (signal path)	dB	53.3	62.2	66.5	DC Gain of the Amplifier	Gain (signal path)	dB	53.6	62.7	66.8
PM (CM loop)	°	73.3	80.2	86.6	Phase Margin of the Common Mode	PM (CM loop)	°	58.3	61.5	67.6
GBW (CM loop)	GHz	1.97	2.24	2.89	GBW of the Common Mode	GBW (CM loop)	GHz	1.11	1.28	1.59
Gain (CM loop)	dB	84.8	112	119	DC Gain of the Common Mode	Gain (CM loop)	dB	78.3	110	120
Phase Loss	°	3.3	4.0	4.5	Phase Loss of the Integrator	Phase Loss	°	2.6	3.3	3.8
IQ	mA	3.5	4.5	6	Quiescent Current	IQ	mA	2.2	2.7	3.9
Noise @ 10MHz	nV/√Hz	3.6	4.1	4.9	Input referred Noise	Noise @ 10MHz	nV/√Hz	4	4.7	5.5

6.4.6 Design of the Resonator

The proposed third-order design also presents the opportunity to incorporate one resonator within the ADC loop-filter. This resonator modifies the ADC's NTF by shifting zeros away from the origin to distinct frequencies, thereby enabling the creation of deep notches or attenuations in specific frequencies of the noise spectrum. For most effective implementation, the resonator has been realized by placing a feedback resistor between the output of opa3 and the virtual ground of opa2.

Creation of notches, trades away the third order noise shaping of the ADC with a first order noise shaping at lower frequencies. The implementation needs to create gain from the resonator feedback configuration in conjunction with R_2 . R_{RESO} , that is used to build the resonator, that shift filter zeroes away from origin in a constant manner over process corners, need to be RC tuned similarly as the integrator resistors, so that the ratio between R_2 and R_{RESO} stays constant. Meaning that the value of R_{RESO} usually needs to be in order of magnitude higher than R_2 . Building large n-bit tunable resistors in a similar way as described for the integrators can be challenging, because of the large footprint they will exhibit. Each extra bit would require a resistor with twice the nominal value of the previous one. For that reason, an R2R resistor network has been utilized in the creation of the R_{RESO} (Figure 37). An R2R network allows us to configure a small value n-bit tunable resistor to a larger one, by redistributing the units in the 5-bit resistor as a resistive divider before the reference resistor R_x . Since R_2 is already connected to the virtual ground of the integrator2, an exact copy of R_2 , together with a larger resistor R_x has been duplicated and connected from the virtual ground of integrator2 to the output of integrator3. By doing this, a 5-bit RC tunable R_{RESO} can be easily built from a 5-bit RC tunable R_2 by just adding a much larger untuned resistor R_x .

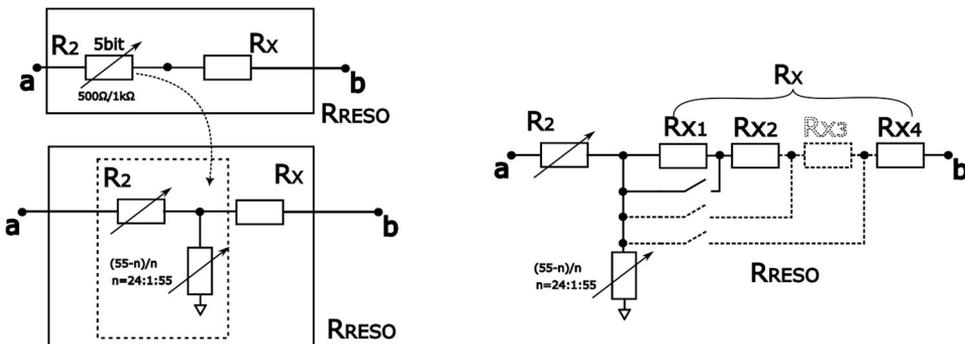


Figure 37. 5-bit, tunable R2R resistor used for creating the resonator.

If R_x is much greater than R_2 , the total resistance of the resonator becomes:

$$R_{RESO} = \frac{R_x}{\left(\frac{n}{55}\right)} \quad (7)$$

where R_x is the physical value of the placed resonator resistance and n is the number of parallel branches selected by the tuning word, where $n = \text{tuning_word} + 24$. R_x value can be made easily programmable afterwards, to fit a wide range of notch placements. This allows for an optimum value resonator selection for each primary Wi-Fi mode and the implementation of DPD modes for PA linearization discussed earlier. Switching from primary mode to DPD mode is only done with gain programming for the resonator, with

rest of the ADC configuration staying intact. Each of the four primary Wi-Fi modes will therefore have an accompanying DPD mode, bringing the total supported modes to eight.

Just like R_2 and R_3 , R_{RESO} also does have a built-in option to modulate the impedance of the resistor in frequency domain. This is achieved by adding a capacitance in series with a fixed resistance across the entire R_{RESO} . This test feature can be used to enhance the maximum BW of the ADC by shaping noise away in an even more favorable manner.

6.4.7 Summary of the Designed Filter

The designed fully differential 3rd order filter is built of from three unique opamps, three sets of programmable capacitors for mode selection, three sets of 5-bit tunable resistors for RC-tuning and includes a 5-bit tunable resonator made from R_{RESO} for switching between primary Wi-Fi modes and DPD-modes (Figure 38).

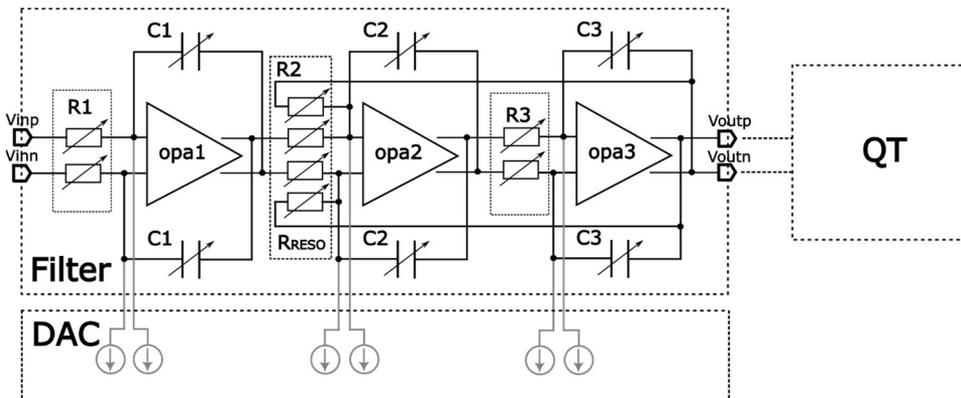


Figure 38. Designed Filter in the $\Sigma\Delta$ ADC.

Summary of the filter features:

Support for Wi-Fi modes W320 (**6-bit**), W160 (**6-bit**), W160LP (5-bit) and W80 (5-bit).

- R_1 : **250 Ω** /500 Ω , R_2 : **500 Ω** /1k Ω , R_3 : **750 Ω** /1.5k Ω .
- R_{RESO} : Supports optimized noise shaping for Standard/DPD modes.
- opa1: 4 BW modes, AB boost mode, Input CM self-generation.
- opa2: 4 BW modes, switchable topology.
- opa3: 4 BW modes.
- Capacitors C_1 , C_2 and C_3 can be tuned to 75%/100%/175%/200% of nominal value.
- Resistors R_2 , R_3 and R_{RESO} have a built-in option to modulate their resistance over frequency for more optimum noise shaping.

6.5 Design of the Quantizer

This chapter covers the following topics:

- Function and topology of the QT.
- ELD, Feedback path of the $\Sigma\Delta$ ADC and signal propagation challenges.
- Optimization of the QT.

6.5.1 Function of the Quantizer

The quantizer is the core of any analog-to-digital conversion. It periodically converts the analog domain signal to the digital domain, which can be then:

- a) fed back to the loop via a DAC.
- b) read out in digital domain for signal processing.

As a first order approximation, the QT can be conceptualized as an ultra-fast flash ADC, with a crucial emphasis on rapid decision-making, and an acceptably small static signal Differential Non-Linearity (DNL). Rapid decision-making is of paramount importance, even if it occasionally leads to minor inaccuracies. Small quantization errors can be noise-shaped and corrected over subsequent conversions, that at most will cause some OoB NTF peaking. Conversely, slow conversions could induce metastability, diminishing the effectiveness of noise shaping while increasing the quantization noise floor, significantly reducing the overall SNR of the $\Sigma\Delta$ ADC.

Flash ADCs can broadly be categorized into two types:

- Zero-Crossing (ZC) ADC: Employ a comparator positioned at the origin, ensuring continuous toggling within the SD loop even in the absence of an input signal, provided there is no offset. Usually the preferred solution, due to clearly defined gain.
- Non-Zero-Crossing (NZC) ADC: Exhibit a dead zone within the quantizer, where the gain of the SD loop becomes undefined. If an input signal is absent, can lead to breaking the loop for an undefined time causing problems in signal processing.

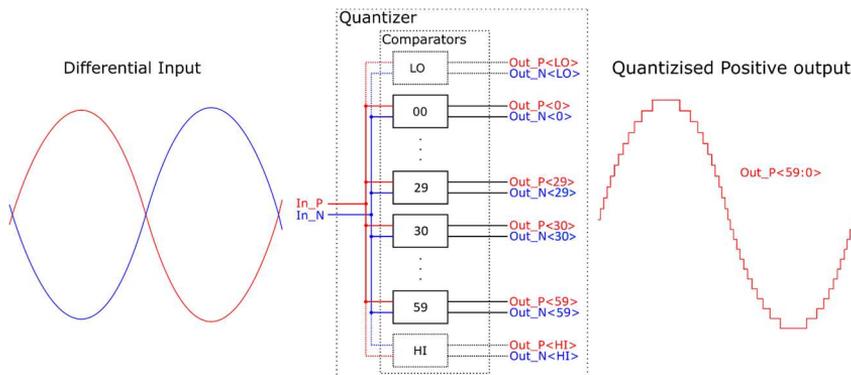


Figure 39. Function of the Quantizer.

The designed quantizer consists of 60 comparators and two additional HI/LO comparators, quantizing the signal into 60 digital outputs, resulting in a total of 61 distinct output states (Figure 39). This even number of comparators classifies the QT as NZC, with all comparators positioned symmetrically equidistant from the origin to facilitate pairing. The rationale behind this grouping will be elaborated upon in the subsequent chapter dedicated to the Feedback DAC. The thermal noise of the system and the aggressiveness

of the coefficients has to guarantee no extended time in the dead zone. Fortunately, with a 6-bit QT, the quantization step is sufficiently narrow to minimize the likelihood of this issue. The HI/LO comparators are dedicated to signal overload detection and are not directly integrated into the system loop. Instead, they trigger a reset signal for the loopback filter and generate a digital warning signal for the system, as will be discussed in the dedicated overload detection chapter.

6.5.2 Comparator

The core of the QT is the comparator – a critical component responsible for making rapid decisions regarding the input signal voltage. Comparators compare the input signal voltage to a predefined reference and determine whether the input signal voltage exceeds or falls below the reference at the sampling edge of the clock.

The comparator itself is composed of three essential elements (Figure 40):

- Pre-amplifier (Pre-amp): The pre-amplifier serves three primary functions. It amplifies the voltage difference between the input and reference signals, performs a level-shifting operation to ensure optimal common-mode voltage for the latch input and introduces a programmable offset for calibration purposes. This offset is introduced through a local 3-bit memory cell embedded within each pre-amp.
- Latch: The latch serves as the decision-making unit, determining the binary output of the comparator based on the input signal comparison. It stores the decision until the next clock cycle, ensuring the integrity of the quantization process. The latch’s design prioritizes speed over all other considerations, enabling ultra-fast decision-making.
- Signal Buffers: Signal buffers are used to convey the output of the QT to the feedback DAC and the ELD DAC. They are designed for optimal signal propagation between the QT and the DAC, that have considerable physical distance in between.

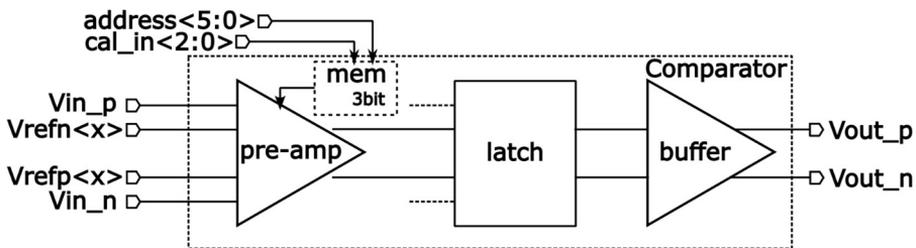


Figure 40. Comparator slice in the QT.

The comparator’s physical implementation adopts a slice-based approach, whereby distinct cells are arranged adjacently to form vertical columns. Each slice is designed to be similar in size and parasitic characteristics to maintain a manageable total column width. To achieve this, the width of each slice is restricted to a maximum of 4um. In a 6-bit QT, this translates to over 200um per channel, potentially affecting the I/Q converter’s geometry. Consequently, each sub-module within a slice must adhere to the specified 4um footprint and refrain from excessive complexity. This slice-based design methodology ensures that the comparator operates at the highest possible speed, enabling the ADC to achieve its target performance specifications. Moreover, the compactness of the design minimizes parasitic effects and contributes to the overall efficiency of the ADC.

6.5.3 Reference Generation and Pre-Amplifier

The reference for each comparator is derived from a resistor ladder composed of 79 unity resistors (Figure 41). This arrangement produces a potential of 78 unique voltage references, from which 60 inner ones are used for the signal path and two are used for overload detection. This sets the full-scale voltage at 60/79 of VDD or roughly 0.72V peak (VDD = 0.95V). Each pair of voltage references serves as one differential input to the pre-amp. The pre-amp is not a sampling circuit, but an open loop amplifier with a passive load (Figure 42). It consists of a differential quadruple input made from devices M₁, M₂, M₃ and M₄ and passive load composing of tunable resistors R_i. The second pair of inputs are connected directly to the loop-filter output, implying that the third integrator sees a capacitive load consisting of all the pre-amplifiers connected to it. This forms a differential amplification circuit that amplifies the difference between the differential loop-filter output to the differential reference generated by the resistor ladder. As a single-stage amplifier, this structure offers limited amplification but is power-efficient, straightforward to implement, does not cause much delay and is compact enough to be physically placed within the specified layout width constraints (4 μ m).

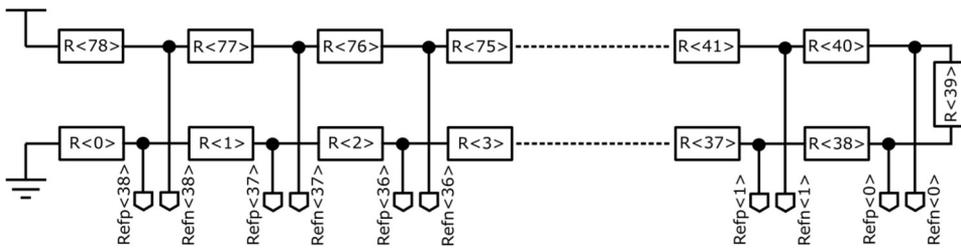


Figure 41. Resistor ladder of the QT.

The passive load for the differential quadruple input is realized with two programmable resistors (R_i), that can be separately adjusted to approximately 70%, 80%, 90% or 100% of the nominal value. The nominal value of the resistors is chosen to level-shift the input signal to a voltage best suited to the succeeding latch module, ensuring the fastest possible decision-making. In addition to the programmable resistors, each pre-amplifier integrates a local 3-bit memory cell that independently controls the corresponding tuning values. This feature introduces the capability to intentionally introduce an offset at the pre-amplifier output, known as QT trimming. The deliberate introduction of a reverse offset serves to counterbalance for inherent offsets arising from device mismatch. Access to the memory cell is facilitated by selecting a unique address corresponding to the specific pre-amplifier and then setting the desired calibration code. Each pre-amplifier has its own dedicated address, accessible through a specialized calibration machine. The calibration functionality is discussed in detail in Section 6.13 dedicated to the calibration machine.

6.5.4 Latch and Buffers

The latch serves as a sampling circuit, periodically capturing a snapshot of the input states with the rising edge of the clock input. Its inputs are directly connected to the outputs of the pre-amplifiers, ensuring that the latch receives the amplified and level-shifted input signals, latching the difference and holding their comparison until the next rising edge of the clock (Figure 42). The latch's behavior resembles that of an open loop amplifier with infinite amplification. To retain the captured state until the next clock cycle, the latch

employs a cross-coupled NAND pair that acts as an SR-trigger. This arrangement forms a memory cell that will not be overwritten until the next sampling event. The cross-coupled NAND pair ensures that the stored state remains stable, preventing any transients or signal degradation.

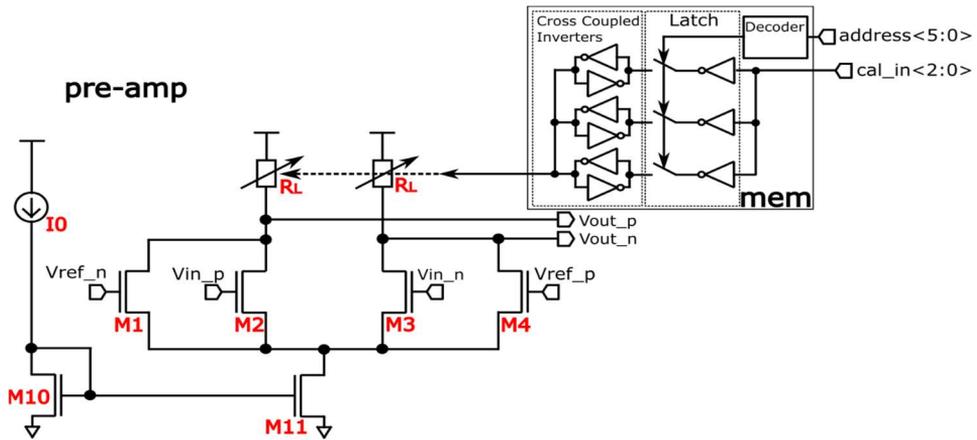


Figure 42. Pre-Amplifier of the QT.

The latch has two phases of operation:

- Reset – Sampling switch ST is opened and pull-up switches S_N and S_P are Closed. This will pull all internal nodes, including \bar{S} and \bar{R} to VDD, making the SR-trigger acting as a memory cell, that keeps the previous Q at the output of the SR-trigger.
- Latching – All pull-up switches are opened and the sampling switch ST is closed. As the latch's tail node is pulled to ground, the input quadruple begins to actively pull both nodes A and B towards ground. The side that exhibits stronger inversion, characterized by a lower channel resistance, will react faster and succeed. Since the latch has a cross-coupled feedback through an NMOS guards M_3 and M_4 , only one side can succeed at a time and the other side is forced back to VDD.

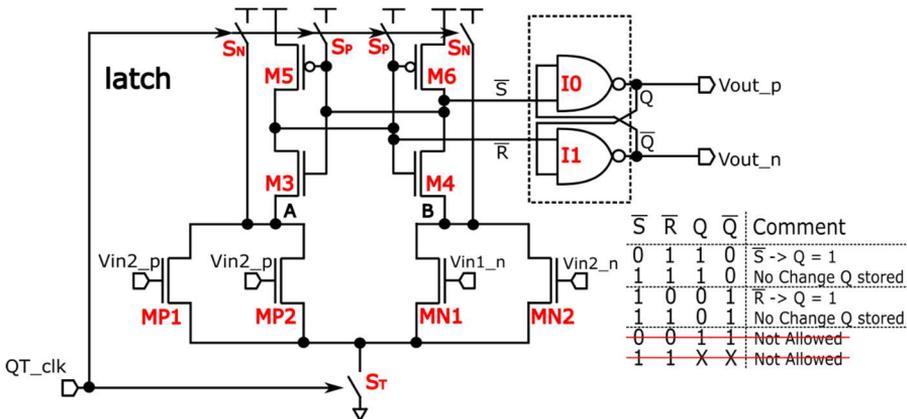


Figure 43. Latch of the QT.

The QT latch is a crucial submodule, where the ultimate decision regarding the input voltage is rendered. In instances where the input voltages to the latch are closely spaced, indicating that the analog input was nearly equal to the reference voltage established by the resistor divider, decision latency can be significantly extended. This arises from the pre-amplifier's inability to adequately amplify the differential inputs of the latch, preventing it from making a definitive decision. While substantial efforts are directed towards mitigating these scenarios, they cannot be entirely eliminated and do occasionally occur.

The final sub-block comprises of a series of buffers that transmit the resulting quantization values to the Feedback DAC. As this path extends beyond the quantizer and traverses over the ELD, inverters are strategically positioned throughout the path to optimize signal speed and ensure sharp signal edges. While seemingly insignificant, the physical placement of these inverters plays a crucial role in minimizing the delay of the small signal path.

6.5.6 Signal Propagation Challenges

The comparator's decision-making process must operate with exceptional speed to ensure that the output data is promptly re-latched with the opposite edge of the ADC clock, effectively confining the entire operation to half a clock cycle (Figure 44). This phenomenon is known as Excess Loop Delay. ELD is defined as the delay incurred between the quantizer latching and the loop filter feedback input. It can also approximately be interpreted as the duration the signal spends in the digital domain relative to the entire clock cycle. In this $\Sigma\Delta$ ADC design, the ELD is strictly constrained to less than 0.5. This stringent requirement aims to maintain a simplified and precise clocking scheme, utilizing the opposite edges of the same clock for signal processing, ensuring maximum stability over process variations. While higher ELD values are theoretically possible, implementing a robust tuning mechanism over process variations has proven challenging, thus restricting ELD to a reasonable value has led to improved manufacturing yields without any extra calibration needed.

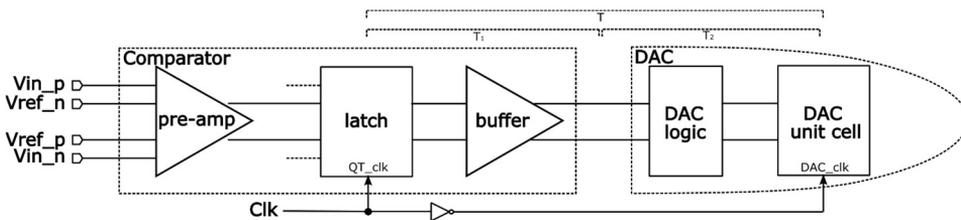


Figure 44. Feed-back path of the $\Sigma\Delta$ ADC.

The portion of the loop, where the signal is transversing in digital domain is known as the feedback-path (FB-path). Part of the comparator is in the analog domain and is not included in the FB-path. The FB-path commences with the latching process in the quantizer and concludes with the latching of the DAC unit cell, which will be discussed in the subsequent chapter. The FB-path comprises the quantizer latch, signal buffers for propagation between the quantizer and DAC, and the DAC logic extending up to the DAC unit cell, where the signal is ultimately re-latched. The quantizer latch marks the starting point of the FB-path, and the latching in the DAC signifies its termination. The ELD of 0.5 imposes a stringent constraint on the propagation time, ensuring that the signal

completes its journey within half a clock cycle (T). The total time which the system has allocated for the propagation therefore is:

$$T = \frac{1}{D * f_s} \quad (8)$$

where f_s is the clock rate used and D is the duty cycle of the clock.

In case of the highest mode supported (w320), where the system is working with 3.84GHz clock, the ELD budget is approximately 130ps. Of this budget, approximately 60% is allocated to the QT decision-making process (T_1), while the remaining 40% is dedicated to signal propagation from the QT output to the DAC latching (T_2). This allocation strategy is only designed for modes operating at the highest clock rates. The reason for this is that the latching and buffering characteristics remain consistent across different modes, and signal propagation becomes a critical factor only in modes where the time budget is most constrained. In modes operating at lower clock rates, the decision-making time (T_1) remains constant, but the total available time for signal propagation (T) increases.

If signal propagation cannot be guaranteed within the allocated time budget, the onset of metastability will gradually increase, leading to a quickly elevating noise floor. That means that in roughly 80ps the comparator must make a decision, in order not to violate the strict meta-stability requirements. If the input signal was nearly equal to the reference voltage established by the resistor divider, intrinsic decision latency in regular latches can be significantly extended. To address this challenge, the latch incorporates an internal cross-coupled inverter preceding the final SR-trigger. This inverter mechanism forces a decision to be rendered by the time the input crosses the half VDD threshold, even if it is not entirely accurate.

If an erroneous decision is made and the error in the loop filter exceeds expectations for the entire cycle, the loop filter's integration speed will increase, eventually noise-shaping out the error without the loss of SNR. The only consequence will be some extra out-of-band peaking. However, if the decision-making process takes longer, and the final value, including signal propagation exceeds the allocated time budget T , meta-stability occurs, where the digital output code is not present in the loop for the entire length of the cycle. When this happens, the SNR will degrade depending on how long the time limit was violated. Therefore, it is preferable for the $\Sigma\Delta$ ADC to prioritize quick decision-making, even if it sometimes leads to inaccuracies, rather than accurate decision-making that compromises the decision-making speed.

In rare instances, over many conversions, the 80ps time budget may still be surpassed, potentially violating the metastability constraints. This scenario typically manifests in a single comparator at a time, depending on the level of the input signal and its proximity to the reference voltage levels. Violating the time budget does not automatically mean meta-stability as the DAC logic is built with redundancy and will in most cases use a faster path. In rare unfortunate cases however, when the particular slice where slow decision making occurs, happens to be in calibration and the signal needs to be re-routed the allocated time budget might not be enough. This cannot be totally avoided and the possibility of meta-stability happening is a certainty over many conversions. The DAC latch is therefore built transparent to minimize the total effect of it and will be discussed in the following chapter.

6.5.7 Optimization of the Quantizer

Minimizing area and power consumption was a paramount consideration in the ADC design, driven by the need to adhere to emerging industry standards and secure a competitive edge in the highly competitive market. Any advantage over the competition could prove to be a decisive factor, making it crucial to optimize the ADC's resource utilization. Different techniques, such as switching the $\Sigma\Delta$ ADC between 6-bit and 5-bit modes and reducing opamp BWs, have been implemented across the module to accomplish this.

One of the areas where substantial power savings can be made, is the 6-bit quantizer. A pure 6-bit QT consists of 60 rows of comparators, so making any reduction in this count translate into a substantial gain. The total current consumption of the QT can be translated into static and switching currents. While static current remains constant, dynamic current is dependent on the clock frequency. Transitioning to lower clock rates only reduces switching current. The resistor ladder and pre-amplifiers contribute to static current consumption, while latches and buffers account for the dynamic component. Reducing the number of latches has a proportional current benefit among all modes. Reducing the number of pre-amps is over-proportionally beneficial in modes with lower clock rates, as it reduces a bigger proportion of over all current consumption.

Quantizer resolution is particularly important for the mid-levels of the QT, when the ADC is operating around input levels that are close to differential zero. This happens around 90% of the time, since even when the ADC is not operating in sensitivity, OFDM signals are typically characterized by relatively low RMS power levels and moderately high crest factors. Therefore, power optimization efforts should prioritize maintaining full performance in the mid-levels of the QT, while allocating less resources to the outer-levels.

One approach to reducing power consumption is to lower the resolution of the QT near the outer edges by expanding the step size of the references (Figure 45). This allows for the removal of entire slices of QT comparators, significantly reducing both static and dynamic power consumption. When partitioned effectively, the SNR is minimally impacted due to the ADC's backoff mechanism. These lower-resolution comparators at the edges are only employed in rare instances when subcarrier phases align in an unfavorable manner. For the majority of the time, the ADC would operate at full performance without any power penalty.

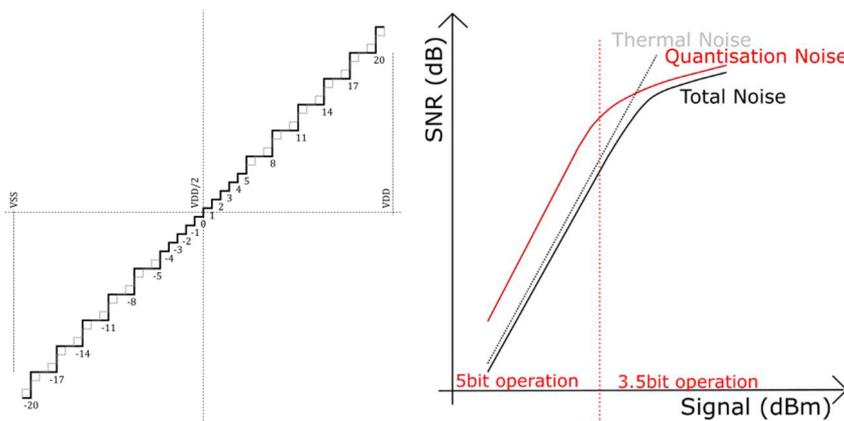


Figure 45. a) Non-uniform QT; b) Performance estimation of a non-uniform QT.

Another avenue to minimize power consumption lies in intelligently leveraging existing information for multiple purposes within the QT. One such strategy involves interleaving a single reference to generate input stimuli for multiple latches. The latch possesses a segregated input that is already partitioned into two positive inputs and two negative inputs. Instead of using a single differential input as stimulus for the latch, two sets of inputs that are then compared to each, can be used. This functionality mirrors that of the differential quadruple in the pre-amplifier, except instead of performing subtraction, it serves as a summer. This approach eliminates the need for every other reference and pre-amplifier, effectively halving static current consumption.

6.5.8 Quantizer Calibration

An ideal comparator transitions its output state when the input signal crosses a predefined threshold value. If the comparator's output changes at a voltage other than the intended threshold, it is said to exhibit an offset. For an ideal 6-bit QT, the output is divided into 60 uniform steps, with each step having a constant and linearly spaced width. However, real-world comparators may exhibit non-idealities due to mismatches within their internal circuitry, resulting in nonlinearly spaced output steps. Any deviation from the ideal step width is called differential non-linearity (DNL). DNL quantifies the difference between the analog values corresponding to adjacent input digital values. It is an important specification for measuring errors in ADC. Ideally, any two adjacent digital codes correspond to output analog voltages that are exactly one Least Significant Bit (LSB) apart. While moderate DNL errors can be mitigated by the sigma-delta modulation process, minimizing DNL is essential to prevent an increase in QT noise, mismatches between the STF and NTF and, most importantly, I/Q imbalance. To correct the DNL deviation, calibration schemes have been developed and implemented. For an ideal ADC, the DNL is zero LSB. In a practical ADCs, DNL errors come from its architecture and cannot be avoided.

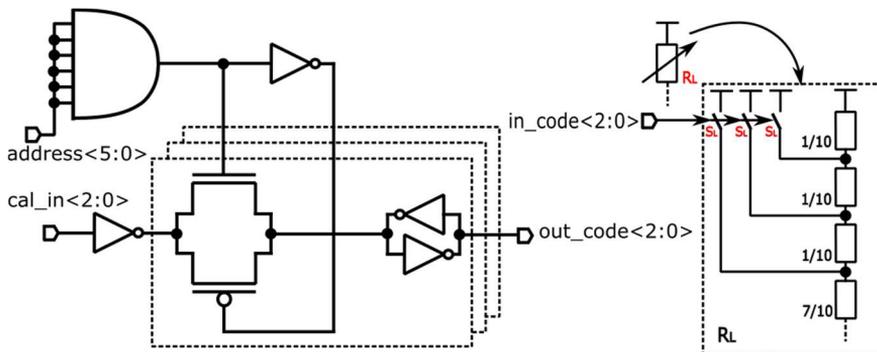


Figure 46. a) Memory cell of the offset generation; b) Resistor tuning of the QT offset generation.

To calibrate the DNL errors each pre-amp has a built-in 3-bit trim feature (Figure 46a). By changing the value of the passive load R_L in the pre-amp, a deliberate offset can be introduced. The load resistor R_L can be programmed to roughly 70%,80%,90% or 100% of the nominal value, resulting in an approximately 7mV step, depending somewhat on the gain of the preamp, which is process dependent (Figure 46b). A 7mV step will determine the total tunable range of $\pm 21\text{mV}$ per comparator, meaning the generically occurring offset should not exceed $\pm 24.5\text{mV}$ to reach an error of better than plus/minus half step.

6.5.9 Switching between 5-bit/6-bit Modes

Switching the QT between 5-bit and 6-bit operating is the same as switching between 30 outputs and 60 outputs (Figure 47a). Therefore in 5-bit operation every second latch unit (all the units that have no dedicated pre-amps and use interpolated inputs) can be turned off (Figure 47b). This will automatically reduce half of the dynamic current from the QT.

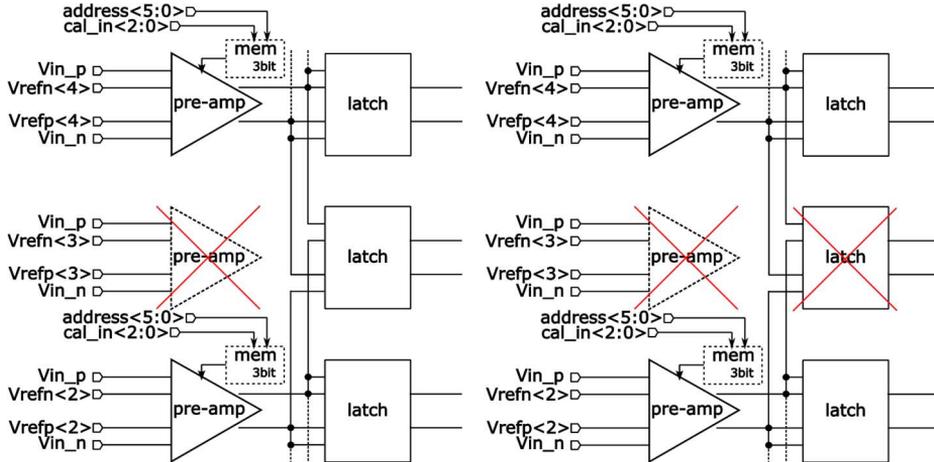


Figure 47. a) 6-bit topology of the QT; b) 5-bit topology of the QT.

Having a NZC QT, also necessitates the configuration change in the resistor ladder when moving from the 6-bit mode to the 5-bit mode, as switching off every second comparator would otherwise cause an offset. This is done by changing the resistor ladder configuration by removing one unity resistor from the origin (Figure 48). A side-effect of switching the QT operation to 5-bit and removing the unity resistor is that it will slightly also change the full-scale range of the ADC, that will now become 60/78 of VDD instead of 60/79 (6-bit mode).

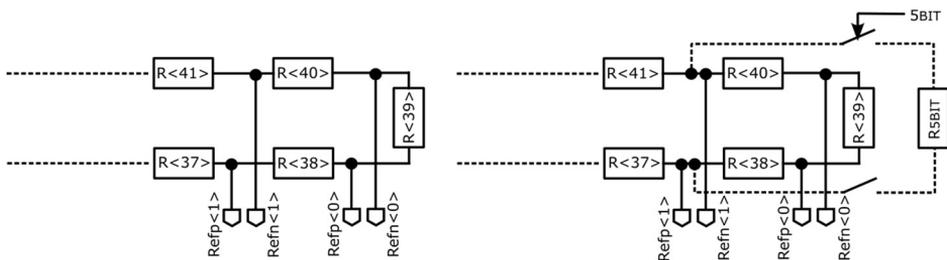


Figure 48. a) Resistor Divider in 6-bit configuration; b) Resistor Divider in 5-bit configuration.

6.5.10 Simulation Results

The comparator design underwent careful optimization to achieve both QT delay and offset reduction. The objective was to maintain a delay below 80ps while minimizing the aggregate offset arising from the pre-amplifier and latch. The full comparator slice raw offset deviates marginally from the interpolated slice raw offset, so both are presented in the summary table for comprehensive assessment. Comparator offset after calibration is discussed in the chapter dealing with the Calibration Machine. The post-layout simulated results of the important parameters are summarized in Table 5 and detailed in Figure 49.

Table 5. Simulation summary of the important parameters of the QT latch.

Comparator	Parameter	Unit	Value				Comment
Corners Used			Min	Mean	Max	StdDev	Method used: Latin Hypercube (1000 runs)
Delay		ps	57.8	65.0	84.4	3.82	Delay of the comparator
Offset 1 (Total)		mV	-20.2	0.2	21.2	5.18	Offset of the full comparator
Offset 1 (Pre-Amp)		mV	-11.9	0.2	16.4	3.85	Offset of the full comparator (Pre-Amp mismatch only)
Offset 1 (Latch)		mV	-13.2	0.2	11.3	3.30	Offset of the full comparator (Latch mismatch only)
Offset 2 (Total)		mV	-18.5	-0.1	14.4	5.81	Offset of the interpreted path comparator
Offset 2 (Pre-Amp)		mV	-12.3	-0.1	10.7	3.89	Offset of the interpreted path comparator (Pre-Amp mismatch only)
Offset 2 (Latch)		mV	-15.3	0.0	13.6	4.59	Offset of the interpreted path comparator (Latch mismatch only)

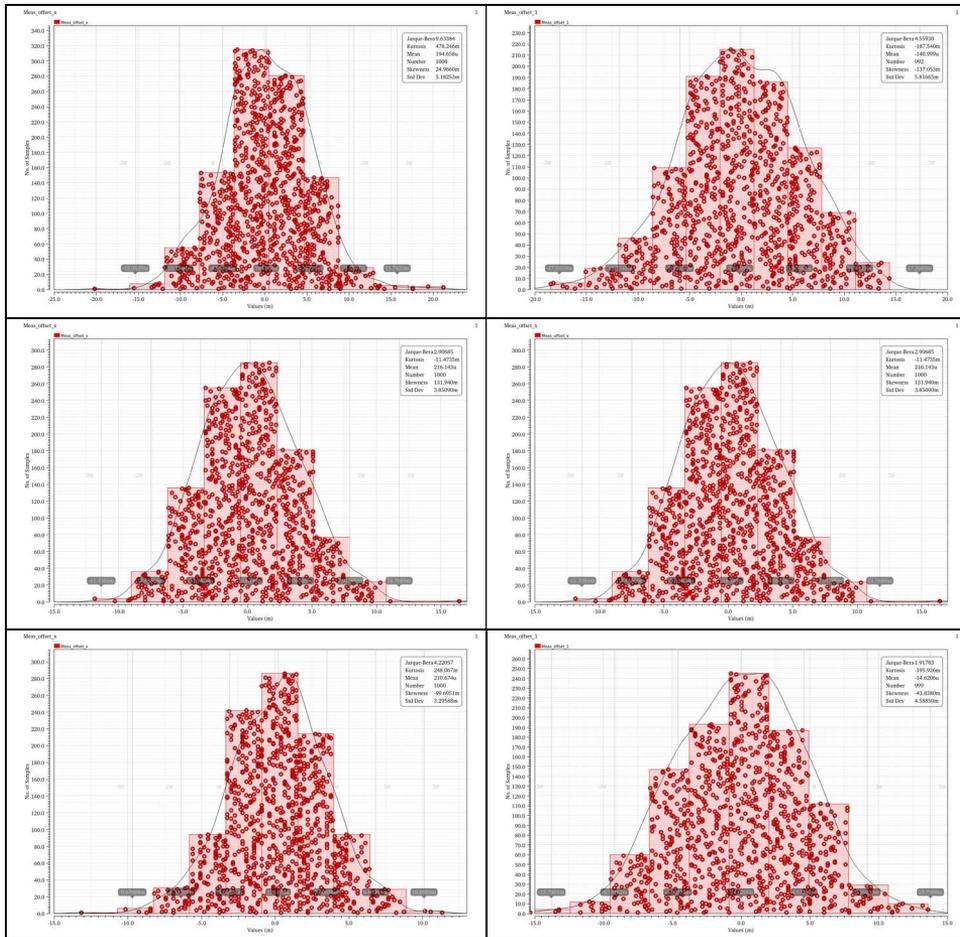


Figure 49. Uncalibrated offset distributions of the Comparator modules simulated in Cadence design environment.

6.6 Design of the Feedback DAC

This chapter covers the following topics:

- Design considerations for the Feedback DAC.
- Topology of the Feedback DAC.
- Function of DAC Calibration.

6.6.1 Design Considerations for the Feedback DAC

The feedback DAC forms the third and final element of the sigma-delta core, completing the feedback loop that iteratively refines the output by continuously comparing it to the original input. It serves as a crucial component within this loop, translating the quantized output of the converter into an analog signal that can be effectively compared with the input. Feedback DACs can be designed in various ways, including resistor-strings, capacitive arrays or current-steering topologies. This design leverages two of these topologies: a current-steering DAC for the feedback DAC module and a capacitive array-based DAC for ELD compensation. The latter will be elaborated upon in chapter 6.10.

The feedback DAC is designed to operate in both 5-bit and 6-bit configurations, based on which mode the ADC is currently operating in. It is built up from 60 current steering cells when operating in 6-bit mode. In 5-bit mode, every second cell is disabled, operating with only 30 cells. This halves the total current output while maintaining compatibility with the filter resistors. The topology enables the DAC to generate $60 + 1$ distinct outputs in 6-bit mode and $30 + 1$ outputs in 5-bit mode, making the feedback DAC in actual life closer to a 5.9-bit DAC in 6-bit mode or a 4.9-bit DAC in 5-bit mode, respectively. This design decision was driven by a specific consideration. To ensure symmetry in the layout and allow for efficient cell grouping, the number of current-steering cells must be an even number. An odd number would necessitate either halving the current in the middle cell or leaving a cell ungrouped, both of which would significantly complicate the layout strategy. The smallest even number that meets the requirement of not exceeding the 5-bit bus width is 30. If 32 cells were used, the number of outputs would increase to 33, requiring a bus width of 6-bit, in the 5-bit mode. Similarly, in the 6-bit mode, design strategy to keep things simple was implemented by just doubling of all cells. For that reason, 62 was not used, although that would be the smallest even number in a 6-bit configuration, where the number of outputs does not exceed the maximum bus width of 6-bit.

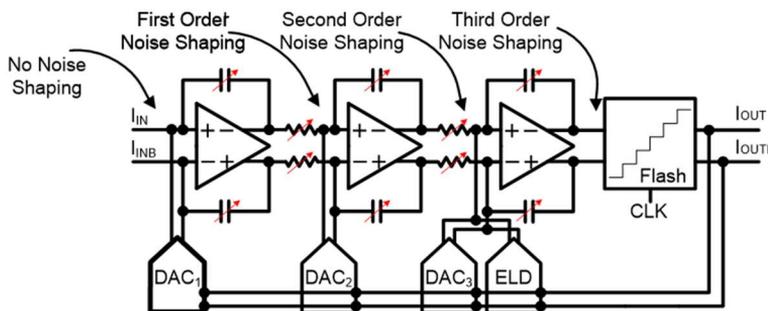


Figure 50. Noise shaping properties of the system.

The current-steering feedback DAC generates an analog signal by converting a digitized feedback code into a proportional current, which is subsequently fed back to the summing node. This process is iterative and occurs with the clocking rate of the system.

The accuracy of the feedback current plays a pivotal role in the overall linearity of the $\Sigma\Delta$ ADC. Clocking current cells and rapidly turning them on or off at high speeds necessitates fast response and efficient settling time. These two aspects are somewhat incompatible. High accuracy in current cells demands excellent matching, which in turn necessitates a large area and a high output impedance that is not conducive to rapid settling. DAC1 nonlinearities are directly subtracted from the input signal. If an error is made, this is directly seen at the output of the ADC (Figure 50). DAC2 imperfections are first order noise shaped. It is possible that in some cases, DAC2 accuracy does not need to be excellent, depending on ADC resolution required and the aggressiveness of the coefficients used. DAC3 and ELD DAC imperfections are second order noise shaped and do not need to be accurate for good SNDR but could play a pivotal role in problems associated with OoB peaking. Since blocker rejection and I/Q balance is important for the system the ADC is designed for, all DACs, except the ELD DAC were calibrated with an operation called linearization.

6.6.2 NRZ and RTZ Coding Schemes

Non-equal rise time and fall time and settling effects of the DAC cells causes harmonic distortion, even if a single-bit architecture is chosen for the implementation. This phenomenon, known in the scientific literature as inter-symbol interference (ISI), [27] results from the differing energies of identical feedback codes when preceded by different sequences. This signal-dependent memory effect introduces distortion. Two types of coding are available for DAC control: Return to Zero (RTZ) and Non-Return to Zero (NTZ). NTZ is significantly affected by the finite and often unequal rise and fall times of the current pulse during the transition phases, as well as clock feed-through and charge injection from the switches that direct the current of the unit cell to the DAC's positive or negative output. Clock feed-through and charge injection generate switching spikes at each transition, superimposing on the output pulse and causing data-dependent error signals, thus leading to distortion. Although RTZ coding mitigates these issues, it comes at the cost of reduced dynamic range and increased logic complexity. Despite the challenges associated with NTZ, such as the potentially daunting finite rise and fall times, NTZ was selected for the implementation of this ADC. This choice was driven by its simpler implementation, which requires less space, reduces parasitics, and offers a higher potential dynamic range, outweighing the negative aspects.

NTZ coding (Figure 51a) is characterized by:

- The current remains constant during periods of high or low state and only changes during transitions.
- The signal remains at the high or low state until it needs to change.
- The transitions between high and low states occur only when the signal changes.
- This results in a continuous current profile that only changes at the clock edges where the state transitions.

RTZ coding (Figure 51b) is characterized by:

- The signal returns to zero between each bit, regardless of the previous state.
- Each bit period includes a return to zero state, resulting in a pulse-like current profile for each bit.
- This means that even if consecutive bits are the same, the signal will still return to zero in between, removing the memory effects.

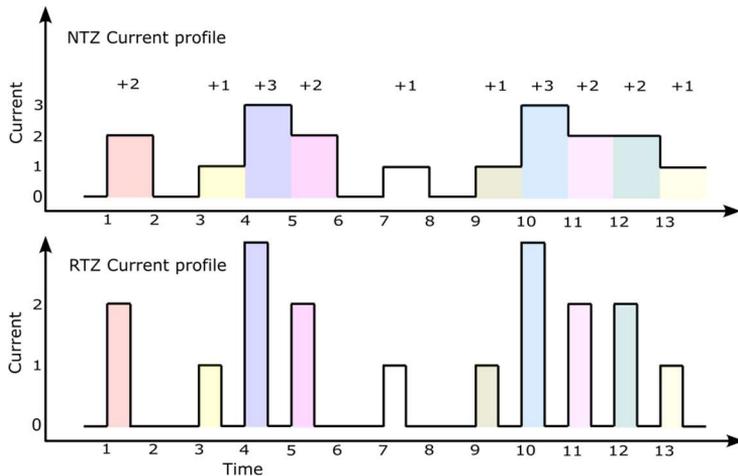


Figure 51: a) Current profile of the NTZ DAC; b) Current profile of the RTZ DAC.

6.6.3 Linearization of the Feedback DAC

To address the challenge of building accurate fast switching devices, DAC cells are actively linearized during operation. This is typically achieved through DEM or calibration. DEM, more popular in low BW SD converters like the ones employed in audio, where the OSR is exceptionally high, utilizes redundancy by adjusting or rearranging the connections of the elements requiring linearization during the operation of the DAC. DEM algorithms often incorporate randomization or scrambling techniques to vary the sequence in which the elements are used. By introducing randomness, the systematic errors induced by specific element mismatches can be noise shaped out over time, diminishing the impact of overall DAC imperfections. It is extensively employed in applications such as audio digital-to-analog converters, where maintaining high fidelity is crucial for achieving accurate sound reproduction. The challenge with DEM is that it needs extra circuitry, which in return introduces extra latency to the critical signal path. Since the design target of the current design was to maintain an ELD below 0.5, incorporating DEM was deemed impractical due to the additional delay it would introduce.

A more suitable linearization approach for high-speed $\Sigma\Delta$ ADCs is DAC calibration. This technique involves comparing DAC cells to a reference cell and adjusting them accordingly. DAC calibration schemes can be categorized into two main types: static and dynamic. Static calibration is typically performed only once before operation or with occasional recalibration when operating conditions change. The calibration data is stored in memory and used to adjust the cells. It often requires heavy redundancy through programming, which complicates the cell design and introduces additional parasitics to the summing node or the virtual ground of the integrator, an undesirable effect that was strived to mitigate in this design. Additionally, $\Sigma\Delta$ ADCs with static calibration have proven to be challenging to simulate, requiring complex algorithms and extended simulation times to account for imperfections in the feedback DAC. In contrast, dynamic calibration operates continuously in the background, even when the ADC is in active use. It dynamically adjusts each cell individually while the converter is performing its normal $\Sigma\Delta$ functionality. This is typically achieved by removing the cell in calibration from the feedback circuitry, temporarily replacing it with redundancy during the operation, performing the necessary tuning, and then placing it back in the feedback path. Dynamic calibration is a switching algorithm, and if not carefully designed, can introduce unwanted

tones into the signal. However, it is relatively straightforward to simulate and does not introduce parasitics to sensitive nodes. For these reasons, dynamic calibration was the preferred choice for this particular design.

6.6.4 Design of the Feedback DAC

To achieve consistent calibration of the DAC cells and maintain identical gain and transfer functions in both the I and Q loops, the feedback DAC was consolidated into a single subblock rather than being depicted as separate I and Q modules. This integration ensures that all DAC cells are calibrated to a single reference voltage, thereby eliminating any potential discrepancies between the I and Q channels.

The feedback DAC (Figure 52) consists of:

- 1) 1 Reference Generation
- 2) 1 DAC1 Reference Cell
- 3) 1 DAC2 Reference Cell
- 4) 1 DAC3 Reference Cell
- 5) 60 DAC1 unit Cells (30 for I, 30 for Q)
- 6) 60 DAC2 unit Cells (30 for I, 30 for Q)
- 7) 60 DAC3 unit Cells (30 for I, 30 for Q)
- 8) 2 DAC1 Redundant Cells (1 for I, 1 for Q)
- 9) 2 DAC2 Redundant Cells (1 for I, 1 for Q)
- 10) 2 DAC3 Redundant Cells (1 for I, 1 for Q)
- 11) 2 DAC1 Dummy Cells (1 for I, 1 for Q)
- 12) 2 DAC1 Dummy Cells (1 for I, 1 for Q)
- 13) 2 DAC1 Dummy Cells (1 for I, 1 for Q)
- 14) 1 Calibration logic
- 15) 1 DAC Clock generation
- 16) 2 Thermal to Binary (1 for I, 1 for Q)

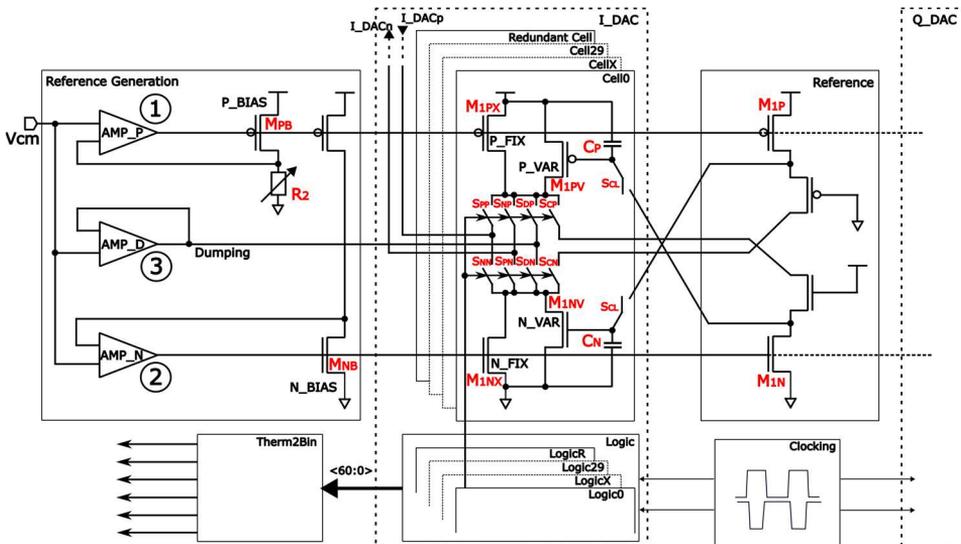


Figure 52. Topology of the Feedback DAC.

6.6.5 Reference Generation

The Feedback DAC incorporates a dedicated reference generation subblock specifically designed to provide a pristine voltage reference to the current sources that constitute the module. This is achieved through the implementation of three amplifiers, each tailored to its unique purpose. AMP_P is responsible for generating the voltage reference for PMOS current sources. It is a single-ended three-stage operational amplifier employing a Miller compensation scheme (Figure 53). The first stage features an N-MOS differential pair formed by devices M₁, M₂, M₃, M₄, and M₁₁. The second stage constitutes a standard common-source stage formed by devices M₅ and M₁₂. The third stage also adopts a common-source configuration and utilizes devices M_{PB} and R₂. R₂ serves as a replica of the R₂ resistor from the loop filter, which in this instance is employed exclusively in 5-bit mode to maintain consistency in reference generation regardless of the ADC mode. Miller compensation is achieved through resistors R_M and C_M, which are connected via a source follower stage consisting of devices M₆ and M₁₃. The fundamental principle behind this amplifier is to maintain a constant voltage V_{cm} across the RC-tuned resistor R₂. This establishes a precise current through M_{PB}, which is then mirrored to all PMOS current sources within the common DAC cells, ensuring proportional tracking of resistor tuning in the loopback filter. However, the presence of 183 PMOS sources (60 DAC1 cells, 60 DAC2 cells, 60 DAC3 cells, and 3 reference cells) connected to the node can lead to a substantial capacitive load, imposing limitations on the maximum GBW that the amplifier can operate at. It is somewhat mitigated by the fact that the DAC matrix is spread out over considerable distance introducing parasitic resistance over reference lines, providing shielding for the opamp.

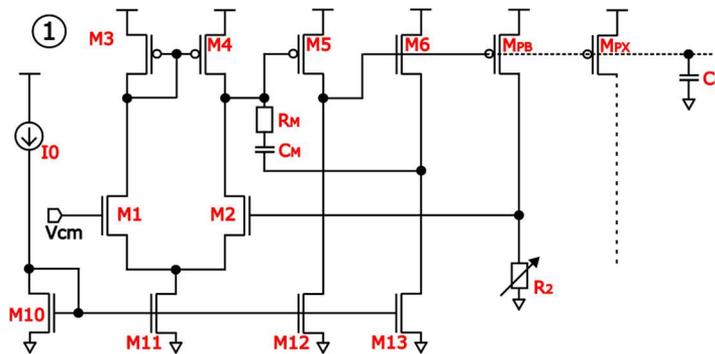


Figure 53. Topology of PMOS reference generator (AMP_P).

To save some current and leverage the bias already provided by AMP_P, the N-side reference is created from its mirror M_{PX}. It is a simple 2 stage Miller compensated amplifier intended to provide reference for all 183 NMOS current sources (Figure 54a). The same that was said about AMP_P applies also to AMP_N. GBW is limited due to a large capacitive load, that is somewhat mitigated with parasitic resistances due to long distribution lines in the feedback DAC.

The final amplifier (AMP_D) performs the task of creating a low impedance dumping node where all current sources can be connected when not actively used for closing the loop (Figure 54b). This allows the disconnection of DAC cells from the feedback loop, effectively eliminating the noise generated by devices not actively participating in the loop. AMP_D is a two-stage Miller-compensated amplifier, featuring a differential pair comprising devices M₃₁, M₃₂, M₃₃, M₃₄, and M₁₅, and a common-source second stage

formed by current source M_{36} and the amplifying device M_{35} . The current source establishes a one-sided limit on the current allowed to be dumped into the node.

When a cell dumps current into the dumping node, both NMOS and PMOS currents are dumped simultaneously. Under ideal conditions, these currents perfectly match, and no imbalance is created so the impedance of the dumping node remains constant. However, process variations and device mismatches inevitably exist, so one side biasing tends to dump more than the other. If the NMOS currents are, on average, larger, it does not significantly impact performance, as the active amplifying device M_{35} can regulate for excess discharge. However, if process variations alter the circuitry such that the average PMOS currents are larger, the impedance of the node can no longer be regulated solely by M_{35} and could become imbalanced. That happens when the excess dumping from PMOS side is larger than the current supplied by current source M_{36} . This negatively impacts performance because the virtual grounds of the cells in dumping are no longer aligned with the virtual ground of the opamp, showing up as noise in the spectrum. Therefore, the absolute value of current through M_{36} is crucial and should closely reflect the number of cells utilized, as well as the maximum allowable process variations and device mismatches. By carefully designing the value of M_{36} , the impedance of the dumping node can be maintained within acceptable limits, ensuring optimal performance and minimizing noise generation.

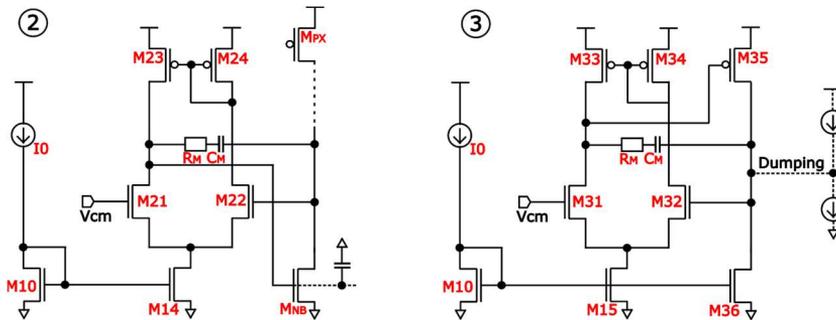


Figure 54. a) Topology of NMOS reference generator (AMP_N). b) Topology of the low impedance node (Dumping) generator.

6.6.6 DAC Calibration

The ADC employs a dynamic calibration mechanism that operates continuously in the background. One by one, each cell is periodically extracted from its normal operation and recalibrated to align with the reference cell. During the calibration time, a redundant cell, with equal properties is substituted in its place. Each DAC current source incorporates two distinct components: Fixed and Variable (Figure 55). The fixed part of each current source is connected directly to the reference voltage generated by the reference generation, while the variable part can be controlled with calibration and is connected to charge storing capacitors (C_P , C_N). During the calibration process, switches S_{CP} , S_{CN} and S_{CL} (Figure 52) are closed, enabling the comparison of the total current through the N-side ($M_{1NX} + M_{1NV}$) to the PMOS reference M_{1P} and the current through the P-side ($P_{1NX} + P_{1NV}$) to the NMOS reference M_{1N} in DAC1 and similarly in DAC2 and DAC3. If there are any discrepancies between the cell and reference currents, the voltage resulting from the charge stored in the capacitors is adjusted until the total current of fixed and variable parts matches again the reference. The charge stored on the capacitor gradually drifts over time and must be designed to be large enough to maintain stable charge between

calibration cycles. However, it must also be small enough to allow for rapid adjustments during calibration. This balancing act necessitates careful optimization, typically performed as the final stage of the ADC design process.

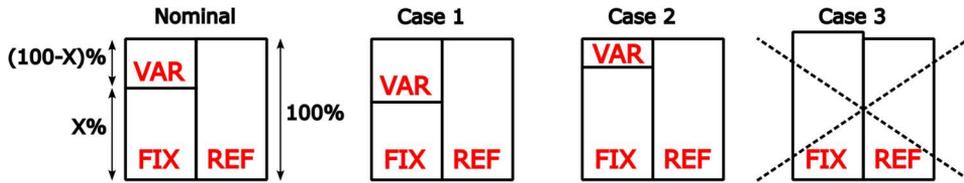


Figure 55. The idea behind DAC cell calibration.

The sizing of the fixed and variable portions of the cell is of extreme importance. Due to the circuit's inherent design, currents cannot be subtracted, but only added. The objective is to calibrate the total cell to 100% of the reference current (Nominal/Case1/Case2). Therefore, the current in the fixed portion X% should never surpass the reference current. If it does, the cell would become incapable of calibration back to 100% (Case 3). However, if the fixed portion is unreasonably small and the variable portion dominates the total current (Case 1), any defects, such as the aforementioned charge drift, would exert an excessively significant influence on the linearization process. For optimal design, the fixed part should be sized in a manner that, even under extreme mismatch conditions, it never exceeds the reference but gets to as close as possible (Case 2). This approach leverages the maximum stability of the fixed part while minimizing the variation of the variable part, leading to optimum overall performance.

The designed $\Sigma\Delta$ ADC, based on the implementations of the coefficients, uses a different feedback current for each of the three DACs. DAC1 current is triple of DAC3 current and DAC2 current is double of DAC3 current. The ratios of currents match the ratios of the resistors used in the loopback filter, giving a feedback coefficient of one for each of the loopback paths. Each of the three DACs is biased by the same reference generation, described earlier, so the physical sizes of DAC cells also match the currents they need to provide, giving a ratio of 3:2:1. This, in return, also determines the ratio of sizes between fixed and variable parts of the devices in unit cells. The fixed part can potentially be larger in DAC1, smaller in DAC2 and has to be the smallest in DAC3, to fit to the requirement of not exceeding 100% of reference current. The cells were sized to fit a simulation of 1000 Monte Carlo runs in Cadence design environment with sampling done with Latin Hypercube algorithm, where the only condition was that none of the runs should exceed 3 sigma target yield of the current in the reference cell. All PMOS unit cells were compared to NMOS reference cells and all NMOS unit cells were compared to PMOS reference cells.

In FitFET processes, the Fins, which make up the effective width of a MOS device are discrete by nature and can only take fixed values. Any sizing between fixed and reference device sizes can only be done with fractions where the denominator is the value of the fins used in the reference cell. In this design the reference cell consisted of 12 fins and the fixed part therefore was sized to be:

- $10/12 = 83\%$ for DAC1 P_FIX
- $9/12 = 75\%$ for DAC2 P_FIX
- $9/12 = 75\%$ for DAC3 P_FIX
- $10/12 = 83\%$ for DAC1 N_FIX
- $9/12 = 75\%$ for DAC2 N_FIX
- $9/12 = 75\%$ for DAC3 N_FIX

Sizing of the variable part does not have to be 100% minus the fixed part, since the variable part is controlled by its own analog domain reference potential stored on the capacitor. To have some matching, the length of the device is however kept similar to the

fixed part. The optimum value of the size of the variable part is in fact mostly determined by the simplicity of the calibration loop, that has very limited gain to be able to operate with high speed and efficiency. Assuming similar impedances amongst all current sources in the unit cell and in the reference cell, ideal comparison would yield voltages around half VDD stored at the capacitor. Designing the variable part in a way, where such potential would also equate the missing part of the current would lead to optimum sizing. In real life smaller W/L ratio seemed to fit best, so all variable parts were sized in a similar manner to:

- $1/12 = 17\%$ for DAC1 P_VAR
- $1/12 = 25\%$ for DAC2 P_VAR
- $1/12 = 25\%$ for DAC3 P_VAR
- $1/12 = 17\%$ for DAC1 N_VAR
- $1/12 = 25\%$ for DAC2 N_VAR
- $1/12 = 25\%$ for DAC3 N_VAR

6.6.7 Unit Cell Operations

The unit cell can operate in five dynamic modes i.e.:

1. Calibration
2. Zero/Dumping
3. Plus
4. Minus
5. Disable

Calibration is an operational mode where the unit cell is removed from the feedback loop. All inputs to the cell are discarded and re-routed to a redundant cell, which is used during the process of calibration. The redundant cell is an exact copy of the unit cell and will be periodically calibrated in a similar fashion. There is a separate redundant cell for both I and Q paths. This is mostly practical for layout symmetry reasons as in theory the ADC could get by with only one. During the calibration process, switches S_{PP} , S_{NP} , S_{DP} , S_{PN} , S_{NN} and S_{DN} will be force opened, switches S_{CP} , S_{CN} and S_{CL} will be closed, enabling the comparison and calibration routine, where the voltage over C_P and C_N will be adjusted. After calibration is done, switches S_{CP} , S_{CN} and S_{CL} will be opened, and normal operation can resume until the next calibration cycle. One by one, in order, each cell will perform the same procedure.

Zero is an operational mode where total current added in the summing nodes is net zero. This happens either when the unit cell neither adds or subtracts total current to loop summing node or adds and subtracts the same amount. In this implementation, when working in operational mode zero, the unit cell will cut the feedback by opening the switches S_{PP} , S_{NP} , S_{PN} and S_{NN} , keeping all calibration switches open and closing switches S_{DP} and S_{DN} , effectively dumping the current from the sources to the dumping node. This simple trick removes the thermal noise coming from the current sources that function in operating mode zero. This mode is also referred to as dumping.

Plus is an operational mode where total current added in the summing nodes is one unit. In this implementation, when working in operational mode plus, the unit cell will open the switches S_{PN} , S_{NN} , S_{DP} and S_{DN} , keeping all calibration switches open and closing switches S_{PP} and S_{NP} , routing the current from the P-Side to the negative input of the opamp and the current from the N-Side to the positive input of the opamp.

Minus is an operational mode where total current subtracted in the summing nodes is one unit. In this implementation, when working in operational mode plus, the unit cell will open the switches S_{PP} , S_{NP} , S_{DP} and S_{DN} , keeping all calibration switches open and closing switches S_{PN} and S_{NN} , routing the current from the N-Side to the negative input of the opamp and the current from the P-Side to the positive input of the opamp.

Disable is an operational mode where all switches are opened. References stay intact, but all current paths are broken. Mode “Disable” is similar to mode “Zero”, except the current is not dumped to the dumping node.

6.6.8 Logic for DAC

The control logic for the DAC is a full custom digital design, carefully built and verified in the analog domain, with the target to switch between the five operational modes. Each unit cell has a control logic called LogicX, tasked with selecting one of the four operational modes in which the unit cell functions. LogicX decodes the received input, checks if the unit cells should be in calibration instead and synchronizes the outputs by simultaneously latching the chosen operational mode for each cell (Figure 56). When a unit cell undergoes calibration, the input needs to be re-routed to the redundant cell. This is done with the help of a parallel path composed of LogicZ and LogicR. LogicZ acts as a slave to LogicX, substituting it as a transmitter (if the LogicX is disabled) and rerouting the input via a parallel path. This transmission is received by a receiving cell called LogicR, which relays the re-routed signal to the redundant cell control logic LogicXR. LogicXR is similar in functionality with LogicX, except it has been optimized for signal propagation speed. LogicD serves as a logic cell that dedicates a single unit cell to operate solely in dumping mode, effectively emulating a dummy cell.

The information indicating which unit cell is currently undergoing calibration is generated using two non-overlapping low-frequency clocks, Clk_Cal1 and Clk_Cal2. Those clocks are alternated between LogicX cells, ensuring that each subsequent cell utilizes a different calibration clock. This handover delay serves as a warning signal to the next LogicX cell that a calibration cycle is about to commence. During this handover window, the previous unit cell prepares to transition back to normal operations while the next unit cell prepares to initiate the calibration procedure. This handover process is continuously repeated between the unit cells, seamlessly transitioning between the I_Logic and Q_Logic paths, effectively creating an endless infinite-shaped signal propagation path. The dummy cell LogicD was needed to keep the number of slices per channel even. This is needed for correct calibration handoff, since every cell with Clk_Cal1 has to follow a cell with Clk_Cal2 and vice-versa.

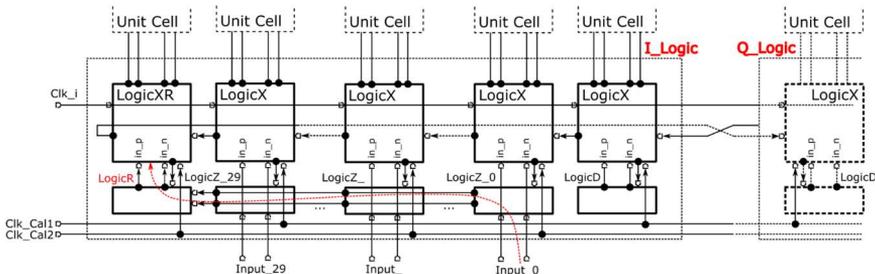


Figure 56. DAC calibration logic.

The primary function of the LogicX cell is to generate a clear and synchronized control signal for the unit cell, ensuring that the unit cell operates in a single mode at a time and that only one set of switches is active at any given instance. This is accomplished through the use of four sets of latches followed by a strong-arm memory cell (Figure 57). With every rising edge of the DAC clock, the latch switch is closed, and the buffer overrides the previously stored configuration in the strong-arm memory cell with the current logic state. If all buffers possess identical driving strengths, all switches have comparable on-resistances, and all traces have equivalent lengths, synchronization of latching among different LogicX cells is ensured.

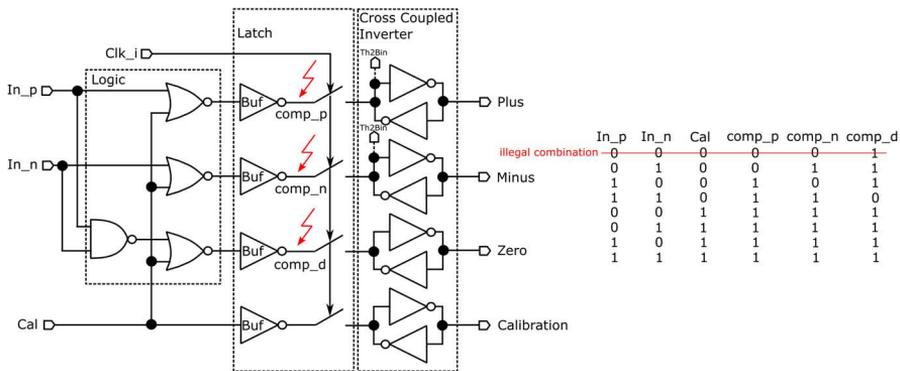


Figure 57. Topology of the DAC latch.

The latch clock is generated from an inverted quantizer clock, a simple and robust approach that reduces circuit complexity and avoids extra jitter, laying the foundation for the 0.5 ELD specification established from the outset. This is because the next stable delay easily achievable would be 1. The data must reach the nodes just before the latch (comp_p, comp_n, comp_d) prior to the positive edge of the DAC clock, avoiding metastability. Metastability is a transient condition that can occur in digital circuits when a signal is transitioning between two stable states. During this time, the signal can exist in a “forbidden” or “gray” state, where it does not adhere to the logic levels of either high or low. The switch remains transparent for the duration of half clock cycle, and if propagation delay is excessive, resulting in undefined input, that alters the amount of charge integrated from cycle to cycle. Metastability is the one thing $\Sigma\Delta$ ADC designers try to avoid at all costs, as it severely compromises ADC performance, causing noise floor to rise substantially. It is better for the quantizer to make a wrong, but fast conversion, as the error will be shaped out in following cycles, than to suffer from metastability. Therefore, the quantizer must be designed for fast latching rather than precise conversion. Timing simulations that assess signal integrity between critical nodes will be discussed in subsequent chapters.

6.6.9 6-bit/5-bit Modes

The DAC has the built in possibility to switch between 6-bit and 5-bit modes. This is done automatically with overall $\Sigma\Delta$ ADC mode programming with w320 and w160 modes requiring the DAC to work in 6-bit configuration and w160LP and w80 requiring the DAC work in 5-bit configuration. Every second unit cell in the DAC array is switched to operational mode “Disable”. LogicX cells, that are connected to cal_clk1 will not relay any information anymore. This will effectively remove those DAC cells from the loop without connecting them to the dumping node, as it would be in the 6-bit mode. It is done to save the current used by those cells, as all potential current paths have now been removed. While these unused unit cells continue to undergo calibration, they remain disconnected from the feedback loop.

6.6.10 Thermal to Binary

Thermal to Binary (Th2Bin) is a module that converts the coding that is done in thermometric mode inside the $\Sigma\Delta$ ADC loop to binary mode, for easier transmission, manipulation and processing (Figure 58). It consists of two 5-bit Th2Bin modules and a 6-bit adder combinatory logic. If operating in 5-bit mode, the inputs of one 5-bit module will receive all zeros, turning the 6-bit Th2Bin simply into a 5-bit Th2Bin.

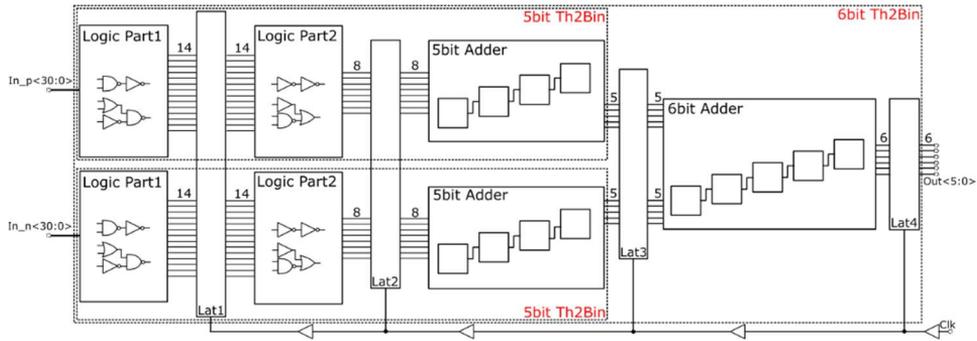


Figure 58. Topology of the Thermal to Binary.

The construction of a 6-bit Th2Bin logic poses a significant challenge, because the block needs to be fast enough to work with the highest clock rate applied by the system. A straightforward 6-bit Th2Bin implementation is too complex to achieve the necessary conversion within a single clock cycle. As a result, latches must be strategically placed within the data chain to temporarily store intermediate values. These latches cost a considerable amount of power, especially when working at highest clock rates. It becomes an optimization task to build the simplest topology, where the number of latches required is the smallest for safe signal propagation. In this design, each 5-bit Th2Bin needs 27 latches and the 6-bit adder requires another additional 6 latches. Consequently, 60 latches are utilized during 6-bit operation, while only 32 latches are employed in 5-bit mode. The clock signal runs in reverse to the data to ensure glitch-free operation.

6.6.11 Simulation Summary

The three amplifiers utilized in the DAC generation have been simulated over the standard 19 corner set:

Table 6. Simulation summary of Amplifiers used for Bias generation.

AMP_P				AMP_N					
Parameter	Unit	Value		Comment	Parameter	Unit	Value		Comment
Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)	Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)
PM (signal path)	°	Min	Typ	Max	PM (signal path)	°	Min	Typ	Max
		33.0	46.0	66.6			38.9	48.5	55.9
		Phase Margin of the Amplifier					Phase Margin of the Amplifier		
GBW (signal path)	MHz	226	311	375	GBW (signal path)	MHz	153	195	251
		GBW of the Amplifier					GBW of the Amplifier		
Gain (signal path)	dB	85.7	97.2	102	Gain (signal path)	dB	65.5	78.3	81.2
		DC Gain of the Amplifier					DC Gain of the Amplifier		
IQ	mA	1.2		1.2	IQ	mA	0.6		0.6
		Quiescent Current					Quiescent Current		
Noise @ 10MHz	nV/√Hz	1.6	1.9	2.4	Noise @ 10MHz	nV/√Hz	2.8	3.3	4.1
		Input referred Noise					Input referred Noise		
AMP_D									
Parameter	Unit	Value		Comment					
Corners Used				Nominal, Fast, Slow, SlowMOS_FastRC, FastMOS_SlowRC, FS, SF (-40°, 25°, 125°)					
PM (signal path)	°	Min	Typ	Max					
		82	95.5	105					
		Phase Margin of the Amplifier							
GBW (signal path)	MHz	207	339	595					
		GBW of the Amplifier							
Gain (signal path)	dB	67.9	74.5	76.6					
		DC Gain of the Amplifier							
IQ	mA	0.3		0.3					
		Quiescent Current							
Noise @ 10MHz	nV/√Hz	9.2	11.2	14.2					
		Input referred Noise							

The DAC cell mismatch was simulated over 1000 Monte Carlo runs in Cadence design environment with sampling done with the Latin Hypercube algorithm (Figure 59). All PMOS unit cells were compared to NMOS reference cells and all NMOS unit cells were compared to PMOS reference cells. DAC1 variable parts in the unit cell were sized to be

on average 17% smaller than reference cell and DAC2, DAC3 variable parts in the unit cell were sized to be on average 25% smaller than the reference cell.

Table 7. Simulation summary of DAC cell mismatches.

DAC				
Parameter	Unit	Value		Comment
1000MC Runs		Min	Mean	Max
Method used: Latin Hypercube				
DAC1 PMOS	%	7.33	16.4	25.8
DAC1 NMOS	%	0.02	11	19.9
DAC2 PMOS	%	8.7	23.6	38.3
DAC2 NMOS	%	5.07	17.5	29.9
DAC3 PMOS	%	5.69	23.6	38.5
DAC3 NMOS	%	3.86	17.4	33.1

In simulations, the mean difference shifted closer together than expected by sizing differences. Especially for the NMOS, where difference between expectation and reality was statistically significant.

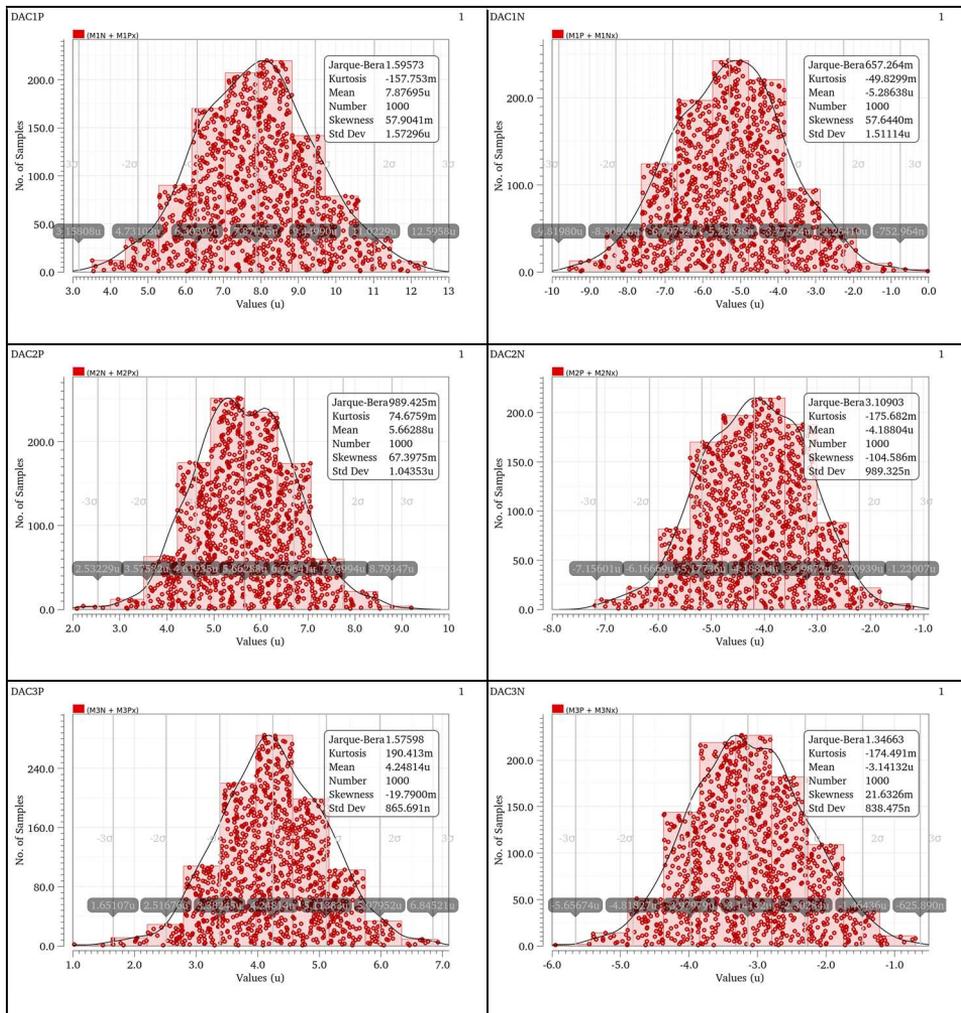


Figure 59. Difference between the Fixed part and Reference currents over different DAC cells.

6.7 Design of the Biasing

This chapter covers the following topics:

- Biasing Design.
- Constant/PTAT current.
- Biasing programming and the enabling of various test features.

The biasing of the ADC is implemented using two reference currents that are supplied from the central part of the SoC. The ADC receives one constant current and one proportional-to-absolute-temperature (PTAT) current, which are then mixed and distributed throughout the biasing circuitry. This approach provides a reliable and stable biasing source for the ADC, ensuring accurate and consistent performance across varying operating conditions (Figure 60). This is crucial for maintaining the ADC's accuracy and performance. By using two reference currents and carefully controlling the distribution of them, the system can be optimized to find the best power-to-performance sweet spot.

Modules needing of local biasing are (9 in total):

- 2 Opa1s in the loop filter (1 for I, 1 for Q); Address 00000
- 2 Opa2s in the loop filter (1 for I, 1 for Q); Address 00001
- 2 Opa3s in the loop filter (1 for I, 1 for Q); Address 00010
- 1 Reference Generation in the DAC; Address 00011
- 2 pre-amp slice banks in the QT (1 for I, 1 for Q); Address 00100
- 1 Active Capacitor; Address 00101

The Reference Generation, the QT pre-amp slice banks and the Active Capacitor of the ADC require only constant currents for proper operation, and the biasing circuitry simply replicates the correct absolute values. However, all of the six opamps can significantly benefit from being biased with a PTAT current, particularly in hot corners, where the gm of the devices can be compromised. PTAT biasing currents can considerably mitigate the impact of temperature variations on opamp performance. For that reason, an option has been built-in to select the best fitting current profile for opamps in the loop filter.

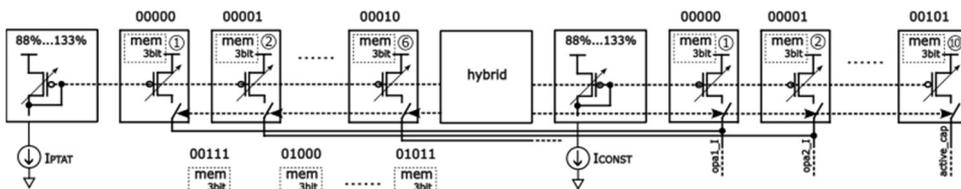


Figure 60. Topology of the $\Sigma\Delta$ ADC local biasing.

The PTAT current side comprises six current sources, while the constant current side utilizes ten current sources. The selection of which current profile is employed is controlled by a switch within the biasing cell, which is managed by the Hybrid module. The Hybrid module determines the mode in which the biasing circuitry operates. The operational amplifiers in the loop filter can operate in:

- Constant mode
- PTAT mode
- Hybrid mode

Constant mode provides a fixed current to the opamps, regardless of temperature. PTAT mode utilizes a current that varies proportionally with temperature, ensuring

relatively consistent performance across temperature ranges. Hybrid mode dynamically selects the current profile that offers the highest performance at a given temperature. This is useful, as current consumption is rarely characterized in inclement conditions and lowering total consumption does not bring any extra benefits. This is achieved by comparing the constant and PTAT currents and selecting the higher one. The Hybrid module incorporates a temperature Schmitt trigger to prevent the biasing mode from switching repeatedly over small temperature variations (Figure 61a). The Schmitt trigger ensures that mode transitions only occur when the temperature difference exceeds a predefined threshold, typically 20 degrees Celsius (Figure 61b).

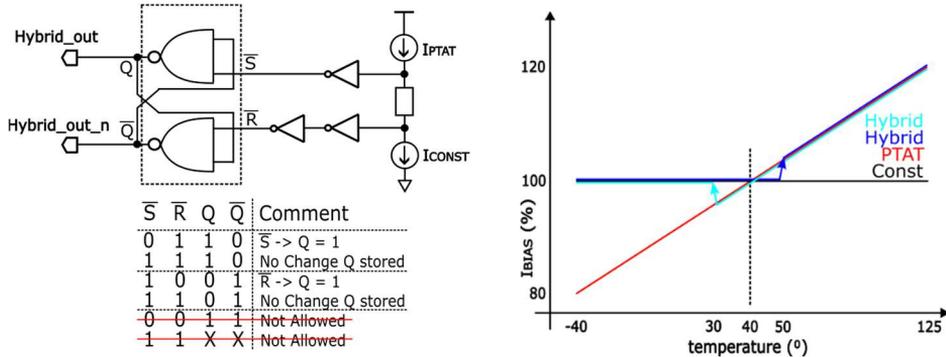


Figure 61. a) Triggering of Const/PTAT change; b) Hybrid Current profile.

Each current source incorporates a built-in 3-bit memory cell, enabling independent control of its absolute value. The 3-bit memory cell employed is identical to those used in the quantization (QT) modules. This local programming capability serves as a valuable debugging feature for pinpointing performance discrepancies or identifying weak spots within the biasing circuitry. Access to each memory cell is exclusively granted through the Calibration Machine, responsible for trimming the QT modules. Each cell possesses a dedicated address that must be selected to gain control of its programming. Once access is obtained, the cell can be reprogrammed using a sequence of low-level bit swaps. It is crucial to note that the content of the memory cell is volatile and is lost upon power loss. Moreover, an internal power-on reset (POR) mechanism is implemented to restore all memory cells to their initial default values upon system startup. Reprogramming the current cell is a time-consuming process that requires manual coding of each step followed by latching the changes using the Calibration Machine. As such, this feature should be employed solely for debugging purposes and should not be incorporated into the ADC's default settings. A detailed explanation of the Calibration Machine's operation will be provided in subsequent chapters.

The ADC's biasing circuitry supports four programmable modes for both constant and PTAT current profiles, ranging from 88% to 133% (88%, 100%, 116%, 133%) of the default value. This global bias programming enables flexible customization of the ADC's biasing behaviour to suit specific mode requirements. To maintain the integrity of the hybrid biasing profile, the constant and PTAT current mirrors are programmed simultaneously. Moreover, individual current cells within the biasing circuitry can also be programmed to a range of values from 66% to 150% of their nominal value using the local 3-bit memory. That means that together with global current programming, each cell can theoretically range between 55% and 200%. The programmable biasing capability extends to all opamps within the ADC, regardless of whether they utilize constant or PTAT current sources.

Some addresses are shared, meaning opamps within the I and Q filters and preamp banks in the QT are programmed as pairs to ensure symmetrical biasing. Opa1 biasing in the I filter is programmed together with opa1 in the Q filter, opa2 biasing in the I filter is programmed together with opa2 in the Q filter and opa3 biasing in the I filter is programmed together with opa3 in the Q filter. The pre-amp biasing in the I QT are also programmed together with the pre-amp biasing in the Q QT. Distinct programming options are available for DAC reference biasing and Active Capacitor biasing, as these components lack the corresponding I/Q counterparts. This ensures independent control of these critical components. In summary, the ADC’s comprehensive biasing programming capabilities provide a versatile tool for debugging the ADC’s performance to test future application requirements and optimize its operation under diverse operating conditions.

All-in-all, only six addresses from a possible 32 were used up, so in order not to waste potential flexibility, several extra 3-bit memory cells including unique addresses, were placed in the Biasing module without a corresponding current cell. The outputs of these memory cells were buffered across the chip to control different built-in test features, like controlling the 75% capacitor option in the integrator capacitors or enabling the resistor modulations. These features were categorized as “experimental” and were not intended for public release or customer access. To maintain their confidential nature, they were accessed solely through backdoor overrides provided by the Calibration Machine. In essence, this approach reserved a portion of the biasing circuitry for future exploration and experimentation, allowing for the testing of innovative concepts without compromising the stability and integrity of the core ADC functionality. This strategy ensured that these experimental features remained under controlled access, preventing their inadvertent activation or exposure to unauthorized parties.

6.8 Design of the Master Clocking

This chapter covers the following topics:

- Clock generation for different modules.
- Delay and duty cycle programming.

The Clocking module generates four of the five clocks needed across the entire SDADC (Figure 62). The clocks generated are:

- QT_clk: Clock for the QT latch
- DAC_clk: Clock for the DAC latch
- sync_clk: Clock used for synchronization
- early clock: Clock used for therm2bin

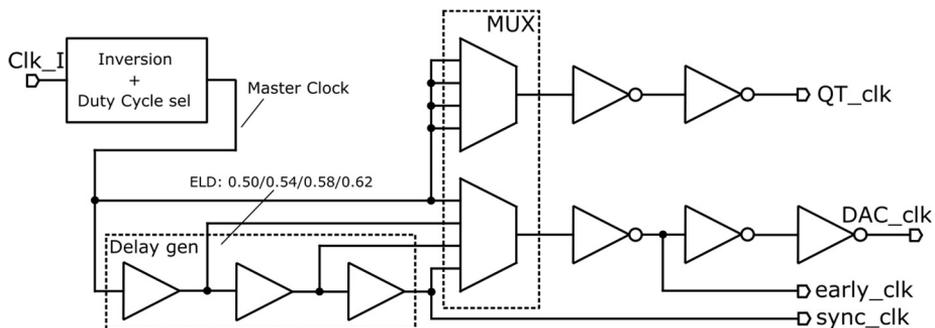


Figure 62. Topology of the clock generation.

The primary function of the Clocking module is to produce a pair of complementary clocks (QT_clk and DAC_clk) essential for the ADC's operation. The phase difference between these clocks can be subtly adjusted through a delay in the DAC_clk path, resulting in an ELD variation from 50% to approximately 62% as perceived by the ADC. This ELD difference is an approximation, as it is susceptible to process corner variations due to its reliance on the total inverter delay time. The QT_clk path, serving as a reference, remains non-programmable, generating a constant-delay comprising of the MUX delay in conjunction with the buffers. The remaining two clocks, sync_clk and early_clk, are employed internally within the DAC for data synchronization and alignment. This is necessitated by the MUX delay's significant process dependence, requiring resynchronization to maintain signal propagation integrity.

The Clocking module has also two extra features:

- The master clock can be inverted. This is useful in case of issues with clock spurs, as the sampling instance can be moved to a potentially cleaner phase of the clock period.
- The master duty cycle can be also programmed to 25% or 12.5%. Needed for the VCO mode. In practice, changing the master clock duty cycle to 25% changes the QT_clk duty cycle to 25% and DAC_clk duty cycle to 75%. But since the entire ADC works with positive clock edges, the effective duty cycle equals out to 25%.

6.9 Design of the Calibration Clocking

This chapter covers the following topics:

- Calibration clock generation.

The last of the five clocks needed for the ADC operation is the low frequency clock needed for the DAC calibration. This is the only clock that is generated separately from the Master Clocking. The Calibration clock can be generated as an:

- LFSR clock: DAC calibration clock is generated with a 11-bit LFSR logic. Calibration time is randomized from cell to cell.
- Constant clock: DAC cells are calibrated with a constant clock. Each cell is in calibration for the same time.
- One Loop clock: Every DAC cell is calibrated only once during startup.
- No Cal: DAC cells are not calibrated.

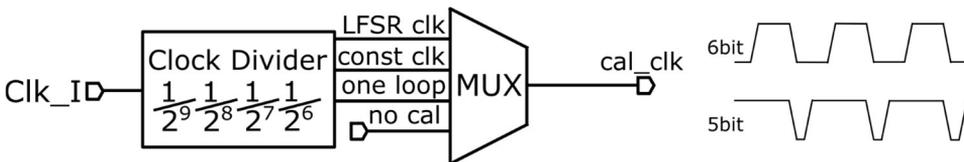


Figure 63. Concept of calibration clock generation.

The Calibration clock cal_clk is derived from the ADC input clock (Figure 63). It can be programmed to four different division ratios to work in different clock speeds, ranging from division by 64 up to division by 512 for both LFSR or constant clock. The LFSR clock speed is average and approximate. The optimum division is unique to each SDADC mode, due to the difference in ADC clock frequency.

At system initialization, the cal_clk frequency is temporarily set to the maximum available, to run through the first DAC calibration cycle as fast as possible. This establishes

an initial, albeit not precise, DAC cell current, which serves as a workaround to ensure proper system startup and a working state closer to the desired operating conditions. After the initial loop is completed, `cal_clk` switches back to the programmed division speed and continues the calibration with the optimal clock speed, best fitted for each mode. The full potential performance of the SDADC is only realized after the first calibration cycle has been conducted using the optimal clock speed.

The calibration clock serves as a critical input for the DAC calibration handover process. Two distinct sets of non-overlapping clocks are generated within the DAC. While operating in 5-bit mode, every alternate DAC cell is disconnected from the loop and rendered inactive. However, due to the calibration topology and the handovers that must be established in fixed order, the calibration of unused cells cannot be omitted. To optimize calibration efficiency and avoid unnecessary time expenditure, the calibration clock duty cycle within the Calibration Clocking module is dynamically adjusted. The unused cells that do not participate in the loop are only calibrated for a very short time, allowing to allocate more time to the cells that are actually used. Consequently, half of the current cells undergo comprehensive calibration, while the remaining half receive a more abbreviated calibration process.

6.10 Design of the ELD DAC

This chapter covers the following topics:

- Concept of the ELD.
- Compensation of the ELD.

ELD is defined as the latency incurred between the quantizer sampling and the data arriving at the loop filter. It can also approximately be interpreted as the duration the signal spends in the digital domain relative to the entire clock cycle. In a perfect $\Sigma\Delta$ modulator, the signal would travel directly from the quantizer, through the DAC, up to the loop filter with no delay. However, in real-world systems, there are always some delays in sampling, processing and propagation. These delays are introduced by the physical characteristics of the circuit, such as sampling, the propagation delays of logic gates and parasitics of the wires.

To counteract the delay inherent in the sigma-delta module, a fourth DAC is incorporated into the coefficient design to introduce a high frequency zero, that recovers the phase lost, via supplying an additional feedback. This additional DAC necessitates the implementation of an additional summing node, which in turn calls for the deployment of an extra summing amplifier together with the fourth feedback DAC. The utilization of a complete amplifier solely for the purpose of establishing a summing node is inefficient and unnecessary, as it entails significant current penalty (Figure 64a). A more efficient alternative is to leverage an existing summing node to introduce the missing feedback to compensate for the excess loop delay (ELD). The nearest available summing node is the virtual ground of the third integrator (Figure 64b). However, since ELD compensation does not require integrating functionality, the quantized signal must first be digitally differentiated to equate back to the original summing operation.

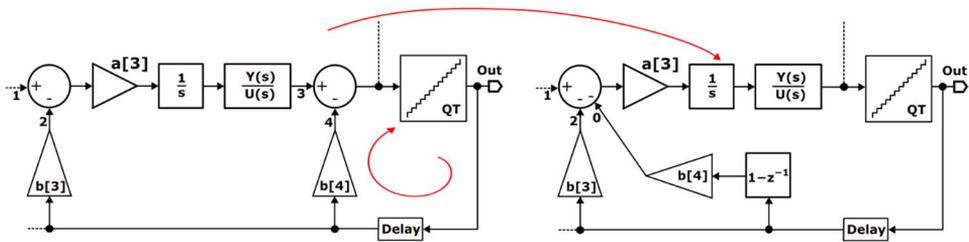


Figure 64. a) ELD compensation with an additional summing node; b) ELD compensation using an exciting summing node.

An even more efficient option is to bypass the integrator functionality altogether for the ELD compensation. This can be achieved by employing opa3 as a capacitive amplifier, effectively bypassing the need for a differentiator (Figure 65). By converting the ELD DAC from a current steering DAC to a capacitive DAC, opa3 ceases its role as an integrator. Inputs 1 and 2 still serve as inputs for the integrator, while input 0 would act as an input for the capacitive amplifier. The most compelling aspect of this implementation lies in the preservation of the ELD coefficients, which remains relatively similar to the design generated using the toolbox. The summing operation remains intact, with coefficients translating into a capacitor ratio between ELD DAC input C_{ELD} and the C_3 in the loopback filter. The only deviation arises from the ratio between the full-scale voltage of the ADC and the supply voltage, as the inverters driving the capacitors operate from supply different than full-scale. Additionally, switching between 5-bit modes and 6-bit modes does not necessitate any modifications to the ELD programming, as half of the cells remain unused due to their control via shutdown comparators.

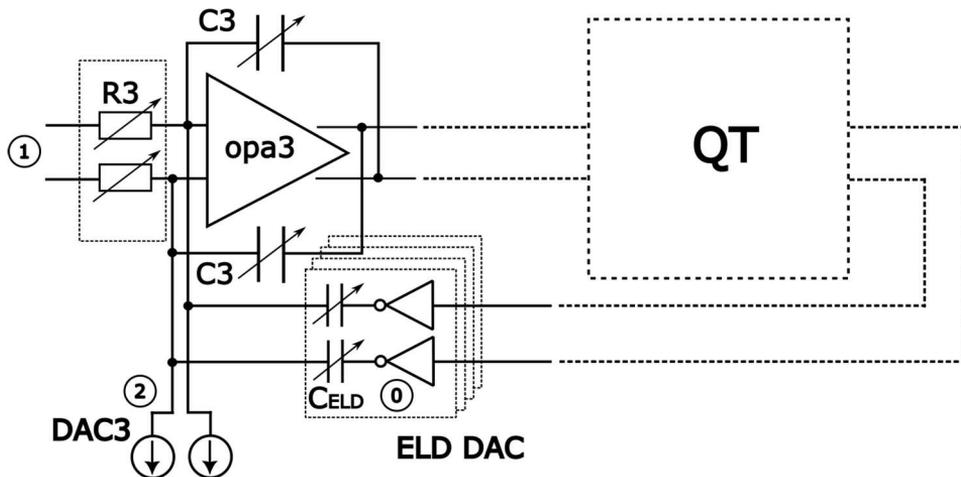


Figure 65. Implemented ELD compensation.

The designed ELD DAC incorporates a 4-bit programming word to select the optimum capacitor value for precise compensation. It employs a binary weighted capacitor array, supplemented by a constant term. This configuration minimizes the quantization step, enabling accurate ELD compensation that significantly improves the stability and performance of the ADC.

6.11 Design of the Overload Recovery

This chapter covers the following topics:

- Definition of overload in the SDADC.
- The concept of assistance in overload recovery.

Overload Recovery serves as a redundant safety mechanism to restore loop stability after the ADC experiences overload and the SD-loop is temporarily disrupted. Once overload recedes, the system must rapidly regain its proper operating point to continue operation. In instances where aggressive coefficients hinder the ADC's ability to self-correct, leading to oscillations, the Overload Recovery module intervenes by triggering a forced reset of the ADC to restore equilibrium. The QT module incorporates two dedicated comparators specifically tasked with overload detection. Upon detecting overload, these comparators can trigger a filter reset signal, initiating the ADC's recovery process. The Overload Recovery module governs the reset function, controlling the duration for which overload must persist before triggering the reset and the duration of the reset pulse itself. This ensures that the reset is initiated only when absolutely necessary and that the system transitions back to normal operation with limited loss in time and SNR.

Options for overload duration trigger selection (Figure 66):

- d0: No Reset (default).
- d1: Trigger reset with the next rising edge of the QT clock.
- d4: Trigger reset with the fourth rising edge of the QT clock. *
- d8: Trigger reset with the eighth rising edge of the QT clock. *

* If the overload duration is shorter than the programmed delay, it will be discarded.

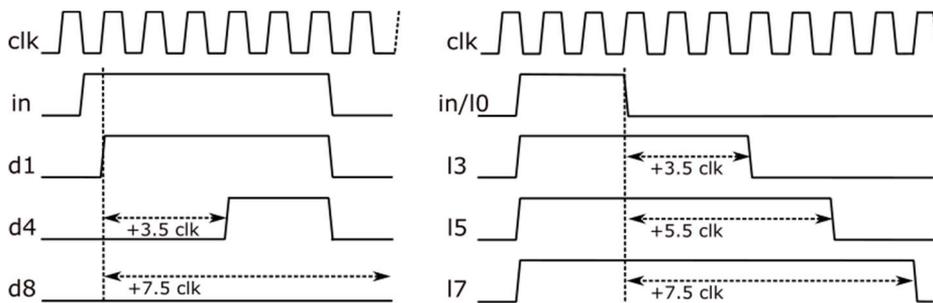


Figure 66. Programmability options of the overload recovery assistance.

Options for the total reset pulse length:

- I0: No extension
- I3: The total reset time will be extended by 3.5 clock cycles after original reset ends.
- I5: The total reset time will be extended by 5.5 clock cycles after original reset ends.
- I7: The total reset time will be extended by 7.5 clock cycles after original reset ends.

While clipping leads to SNR degradation in the ADC, it is considered a less detrimental outcome compared to initiating a forced reset of the loop filter. Force resetting the loop filter can have a severe impact on the overall SNR, and the longer the reset pulse,

the greater the degradation. Therefore, whenever possible, forced resets should be avoided entirely or, if unavoidable, limited to short pulses. Fortunately, the designed ADC does not need the reset function in any of its operating modes, and the module is solely retained as a backup redundancy measure.

6.12 Design of the Overload Masking

This chapter covers the following topics:

- Use case and function for overload masking.

The Overload Masking sub-module complements the Overload Recovery module, addressing the drawbacks associated with forced resets of the loop filter. Wi-Fi channel training protocols commence with the assumption of low initial signal strength, initiating the gain settings to their maximum values. If the RMS of the signal exceeds a pre-established threshold, the gain is dynamically adjusted downwards. This iterative process continues until the desired target input RMS of -10dBm is attained during the preamble. This process is known as the AGC algorithm. It is executed during the gain training phase, preceding the establishment of the Wi-Fi channel. However, in instances where the initial signal strength significantly exceeds the assumed level, the SDADC output experiences overload. If the Overload Recovery function is enabled and the filter is being reset during the duration, the AGC algorithm malfunctions. This is because resetting the filter simultaneously also nullifies its outputs, causing the RMS of the output to register as zero during the entire reset duration. Consequently, the AGC misinterprets the RMS and fails to provide accurate estimations. This can lead to a vicious positive feedback loop, as the AGC may attempt to increase the gain again in the subsequent step.

A dedicated masking mechanism has been implemented to temporarily transform the ADC outputs to full-scale upon overload detection. This intervention is solely necessary when the force reset function is enabled. While the SNR of the signal is further degraded during this masking period, the AGC functionality is preserved. To minimize layout complexity and conserve space, the masking functionality is integrated within the signal path buffers that transport the output from the core of the DAC Th2Bin modules, that are otherwise located in the geographical center of the $\Sigma\Delta$ ADC, to the edge.

6.13 Design of the Calibration Machine

6.13.1 Calibration Machine Functionality

This chapter covers the following topics:

- Working principle of the Calibration Machine.
- Offset generation for the Pre-Amp.
- Impact of calibration on the comparators.

The Calibration Machine is a full-custom analog state machine designed to identify the optimal trim value for each 3-bit memory unit within the pre-amplifier of the QT. It operates on the fundamental principle that if a comparator exhibits no offset and the input signals are at the same potential, the probability of generating a high (HI) or low (LO) output should be precisely 50%. This translates to an equal distribution of positive and negative outputs over multiple conversions. Each memory cell is assigned a unique address, and during the initial startup sequence of the ADC, the inputs of each comparator are sequentially shorted to their corresponding reference voltage, establishing an ideal

equilibrium. Subsequently, various input stimuli are applied to the 3-bit memory cell within the pre-amplifier, intentionally introducing an offset between the p and n branches (**Error! Reference source not found.**). During each offset combination, the comparator executes 512 comparisons, and the Calibration Machine records the number of times the output returns a high (HI) output state.

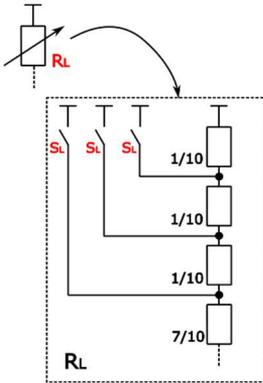


Figure 67. QT offset generation.

The trim value is swept from minimum to maximum, where:

- 0: maximum offset to the negative side (7 units vs 10 units)
- 1: medium offset to the negative side (8 units vs 10 units)
- 2: minimum offset to the negative side (9 units vs 10 units)
- 3: no offset (10 units vs 10 units)
- 4: no offset (10 units vs 10 units)
- 5: minimum offset to the positive side (10 units vs 9 units)
- 6: medium offset to the positive side (10 units vs 8 units)
- 7: maximum offset to the positive side (10 units vs 7 units)

(In brackets is the configuration of series resistors that makes up the passive load for the pre-amp.)

The optimal trim is the value where the number of HIs is closest to 50% of all conversions, or 256. Upon identifying the best trim value, it is permanently stored within the corresponding 3-bit memory unit before proceeding to the next address. This process is repeated for each of the 32 pre-amplifiers within both QTs. Once the calibration sequence is complete, the Calibration Machine returns to a dormant state. The overall duration of a single calibration cycle equals to:

- 512 clock cycles for conversions plus 512 clock cycles for various calculations
- x8 for each trim code
- x32 for each pre-amp
- x2 for both QTs

meaning a full calibration cycle takes 2^{19} clock periods, or 136us utilizing the ADC's highest-rate clock. If necessary, the calibration sequence can be reinitiated to refine the trim values, but for most applications, a single calibration cycle suffices throughout the ADC's power-up cycle.

The Calibration Machine incorporates an additional manual override functionality, allowing direct access and modification of individual trim values. Each address can be accessed and overwritten through backdoor override. It is done with a series of low-level bit swaps, where the address and trim code is written to a reserved register and latched in with a manual override. This function is accessible for both the I QT and Q QT, as well as for the Biasing circuitry, each available with maximum 32 distinct address locations.

Another feature embedded in the full-custom design of the Calibration machine is the debug capability. When using debug, after each calibration step, the state machine will write the counter value to a read only public access register and proceed to a break point in the algorithm. That register can read back, fetching the counter value and the current step. After that, the algorithm can be restarted, by using the unbreak function done with low-level bit-swaps, until stopping again in the next calibration break point. After the last trim code has been programmed to the 3-bit memory and all conversions are done, the last read back will fetch the counter value and the best code chosen by the algorithm. This process is repeated until all addresses are once gone through and both quantizers are calibrated, after which the Calibration Machine will again go dormant.

6.13.2 Calibration Impact on the QT

The QT employs two distinct comparator slice types: full and interpolated. The full slice incorporates all modules within the signal chain, while the interpolated slice utilizes two adjacent pre-amps from full slices to infer the reference voltage level. When left uncalibrated, both full and interpolated comparators contribute comparable offset levels, with simulated 1σ values ranging between 5mV and 6mV. This is deemed acceptable, as the step between two ideal neighboring comparators is 24mV, minimizing the likelihood of quantization errors called bubbles even in the uncalibrated state.

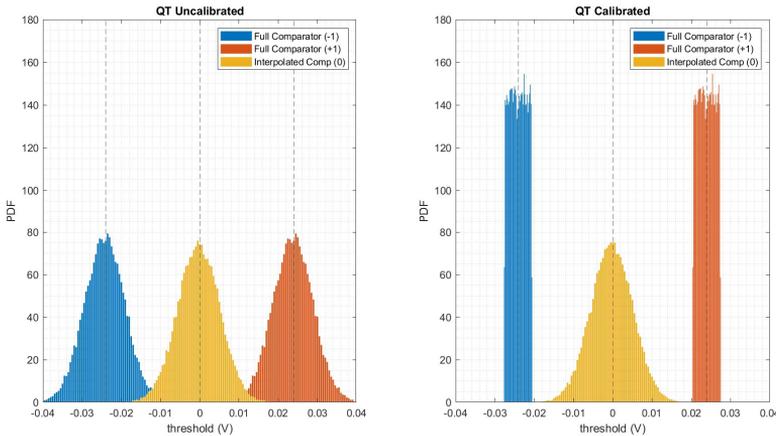


Figure 68. The probability density function of offsets a) pre-calibration; b) post-calibration.

Upon completion of calibration, the picture alters slightly. All full comparators undergo calibration via a manual offset cell embedded within the pre-amplifier. Ideally, with a flawless calibration, all full comparators should attain $\pm 3.5\text{mV}$ accuracy. Since the pre-amplifier and latch offsets are adjusted together, they could theoretically exhibit substantial offsets in opposing directions but effectively cancel each other out. After perfect calibration, the offset at the input of the latch equals the offset of the latch itself. This means, post calibration the interpolated comparator will no-longer see the offset of the adjacent pre-amps, but rather the adjacent latches and equals the average offset of the neighboring latches plus the offset of the latch in the interpolated path. Since the latch and pre-amps have roughly the same offset, the probability density function (PDF) does not change pre- and post-calibration for the interpolated QT slice (Figure 68).

6.14 Design of the Digital Interface

The Digital Interface is a crucial yet straightforward module that serves as the juncture in the signal chain where the supply domain transitions from analog to digital. To prevent the formation of destructive current loops within the substrate or PCB, a concerted effort has been made to confine all supplies, inputs, and control signals within a single domain, eliminating the flow of current across domain boundaries. This means, that together with the supply domain change also the clocking signal changes and the clock used in the digital interfaces stems from a different branch as the analog clock. The phase difference between these two clocks can vary substantially from chip to chip.

The Digital Interface translates the data in the analog domain inside the ADC to digital domain that will be transferred to digital processing. Its design needs to guarantee that data integrity stays intact, no matter how much the two clocks differ in phase. It is also the last place on the chip where forced clock inversion can be utilized as an emergency measure to restore signal integrity in worst case scenario, as rest of the signal path is purely synthesized.

6.15 Design of the Power Supply

This chapter covers the following topics:

- The requirements and design considerations for the Power Supply of the converters.
- The design and performance of the designed LDO.

6.15.1 The Need for On-Chip Power Supplies

High quality supply voltage is important for the realization of top-tier performance A/D and D/A converters. Modern electronics devices require a variety of different voltage levels to cater to the specific requirements of their various functional blocks. One SoC chip can easily have in the excess of hundreds of different analog/mixed-signal modules, each requiring its own regulated power supply. Almost as important is the area and the overall cost. External components introduce additional elements to the Bill of Materials (BOM), increasing complexity and expenses, thus warranting their minimization whenever feasible.

Efficiency of a power supply is a metric that shows how much power the supply is able to deliver as a proportion to the amount it receives:

$$\eta = \frac{W_{OUT}}{W_{IN}} \quad (9)$$

In a power supply with a single input power domain, efficiency is defined as the ratio between the input power and the output power. For a power supply with multiple input power domains, efficiency is calculated as the ratio between the total input power and the output power:

$$\eta = \frac{V_{OUT} * I_{OUT}}{\sum_{k=1}^n (V_{IN_k} * I_{IN_k})} \quad (10)$$

DC-DC Converters are power supply modules easily capable of delivering regulated voltages exceeding a few amperes with efficiency surpassing 90%. While ubiquitous in many chip designs and applications, their utilization in RF systems presents three primary drawbacks. Firstly, their internal clock-driven operation introduces a distinct triangle-shaped clock ripple onto the regulated voltage. This ripple necessitates robust power supply rejection (PSR) capabilities from the downstream analog blocks, which can be challenging to deal with in the pursuit of ultra-high performance. Secondly, the large switching currents generated by DC-DC converters can inadvertently interfere with adjacent analog IPs sharing the same silicon. To mitigate this, additional isolation circuitry is required, which adds to area and complexity overhead. Thirdly, every high-quality DC-DC converter necessitates at least two off-chip passive components. In systems with multiple channels and power domains, this additional BOM cost often proves prohibitive for OEMs. DC-DC converters that require no off-chip components suffer in performance and efficiency and are not suitable for high performance supply. The large amount of current a single DC-DC converter can offer is also not very useful, because very few

single RF/Mixed-Signal modules really require it. Consequently, isolating individual IPs with high-performance LDOs to supply high-quality regulated voltages proves more advantageous in high-performance RF SoCs.

An alternative power supply option is an Low Dropout Regulator (LDO) with an external ceramic capacitor for filtering. This approach can be effective in scenarios where modules within the chip share a common voltage level. However, if different IPs demand distinct voltage rails, each LDO would necessitate an independent blocking capacitor, leading to an exponential increase in external components plus package pins and consequently, again a substantial rise in BOM costs. Another potential issue lies in the package itself – the voltage generated by the LDO must traverse the package pins and bonding wires to reach the external capacitor positioned on PCB. These package structures introduce parasitic capacitance and inductance, creating resonances in the 100MHz to 1GHz range. This renders the external capacitor's effectiveness questionable at frequencies above the potential resonance, necessitating the placement of an additional on-chip blocking capacitors to lower impedance around clock frequencies. LDOs with external capacitors typically also exhibit a dominant pole arising from the pass device, resulting in a bandwidth that is limited to the lower frequency range, falling short of the entire bandwidth required by the A/D and D/A converters.

SoCs with numerous RF blocks that demand exceptional performance typically employ separate voltage regulators to ensure each block receives the optimal regulated voltage and isolation from neighbouring modules sharing the chip. Capless LDOs emerge as a compelling solution for such systems, eliminating the need for external components and minimizing on-chip disturbances. When properly designed and employed effectively, capless LDOs can achieve remarkably high efficiencies (exceeding 85%), rivalling those of DC-DC converters, while simultaneously eliminating the detrimental effects of disturbances and correlated noise inherent in DC-DC converters.

Efficiency of a power supply is a metric that shows how much power the supply is able to deliver as a proportion to the amount it receives:

High performance state-of-the-art advanced A/D and D/A converters demand a very clean and constant power supply, to achieve peak performance. An ideal power supply node can be envisioned as possessing an infinitely low output impedance. The power supply adds unwanted noise to the converters. The noise can be roughly divided into two main types: correlated and uncorrelated. Correlated noise, arising from periodic events such as clocking, signal periodicity and insufficient BW or external disturbances, is the more significant source of concern. Uncorrelated noise, characterized by the converter's intrinsic thermal noise, is generally of lesser importance.

When designing an LDO, it is therefore preferable to prioritize correlated noise performance over uncorrelated noise. Discrete tones generated by clocking or other disturbances tend to manifest in the frequency spectrum of the converters, significantly impacting the subcarrier performance at those frequencies. Excellent regulation and power supply rejection (PSR) should be preferred over general thermal noise characteristics of the supply, as they directly address correlated noise issues. To prevent performance degradation stemming from the power supply and maintain a sufficiently low output impedance within the converter's operating range, the power supply's BW should be equal to or greater than the converter's BW, while simultaneously providing exceptional rejection from external supplies. This ensures that the power supply operates seamlessly and consistently with the converters, exploiting the realization of their full potential.

The capless LDO used in powering the A/D and D/A started out as a project tailored especially for the needs of the converters only. Initially designed for the converters themselves, its outstanding performance and configurability paved the way for its adaptation to meet the diverse requirements of various modules across the SoC. From powering LO propagation modules demanding exceptional thermal noise performance to RF Phase Locked Loops prioritizing PSR and spur cancellation, the LDO's extreme adaptability extended its usage to over 50 instances within the SoC, ultimately ending up as a general purpose LDO useful for all needs.

6.15.2 Options of the LDO

The final LDO, was designed for exceptional adaptability to cater to the diverse requirements of various integrated circuits such as ADCs, DACs, voltage references, LNAs, PLLs, LO generators, and buffers [28]. The maximum output current is 300mA (600mA in some custom design variants), and the LDO leverages two power rails: 1V8 and 1V2. To achieve high efficiency, bulk of the current will be drawn from the 1V2 rail. The nominal output voltage of the LDO is 0.95V, aligning with the targeted VDD for A/D and D/A converters. Despite being a capless LDO, the output capacitance contributed by the supplied IPs themselves and the surrounding blocking capacitors can reach several hundreds of picofarads. The LDO's robust design effectively addresses this impedance challenge, ensuring stability over different surrounding environments.

Operating conditions of the LDO:

- Input of 1V2: 1.10V-1.65V (default 1.2V)
- Input Reference: 0.60V-0.90V (default 0.85V)
- Output: 0.87V-1.03V (Nominal Range); 0.62V-1.09V (Extended Range*)
- Input of 1V8: 1.70V-1.90V (Nominal Range**); 1.60V-1.90V (Extended Range**)
- Output Current: 0-300mA (600mA in some versions)
- Output Capacitance: 1fF-1nF

* Extended range can be used with changing the input reference voltage (default 0.85V)

**Output current maximum 300mA: when input of 1V8 > 1.70V

**Output current maximum 100mA: when input of 1V8 > 1.60V

Specification of the LDO:

- 4 Bandwidth modes: 80M / 120M / 160M / 200M #
- 4 Output load modes: <50pF / 50pF-200pF / 200pF-600pF / >600pF
- 16 Output voltage selections (Output_Selection<3:0>):
 - $Output = Input\ Reference * (82 + Output_Selection<3:0>) / 80$
- Minimum output current: 0
- Maximum output current: 300mA
- LDO Bypass mode: Pass device can be fully closed, so that the output will be shorted through the pass device to the input.
- PSRR: >50dB @ <2MHz; >20dB @ < 100MHz; >0dB at all frequencies @ 40mA output current (SDADC target current for w80) from the 1V2 input.
- PSRR: >30dB @ <100MHz; >0dB at all frequencies from the 1V8 input.
- Noise: <40nV/√Hz ** @ >100kHz; <10nV/√Hz @ >2MHz; <5nV/√Hz @ >10MHz;

Assumed GBW with output current of 40mA (The approximate current consumption of the converters in w80 mode)

Can be improved by adding more capacitance to the filtering input (V_{FILT}).

Configuration options for the LDO:

- output_selection<3:0>: Output voltage selection (0: minimum voltage; 15: maximum voltage)
- BW_selection<1:0>: Bandwidth selection (0: maximum bandwidth; 3: minimum bandwidth)
- load_selection<1:0>: Output load selection (0: minimum load; 3: maximum load)
- LDO_bypass: Bypass of the LDO (0: LDO not bypassed; 1: LDO bypassed)
- enable_load: Enable a 2.5kOhm load at the output (0: load disabled; 1: load enabled)
- fast_recovery: Improve the over- and undershoot recovery performance of the LDO. Will approximately half the over- and undershoot settling time.* (0: Disabled 1: Enabled)
- noise_enhancement: Special noise enhanced mode for the LDO. Will enable a faster current feedback. Maximum current should be limited to 50mA and max output cap to 300pF. Should be used together with enable_load. (0: Disabled 1: Enabled)
- de-regulation_capacitor: Option to lower potentially harmful spur emissions. (0: Disabled 1: Enabled)
- active_capacitor<1:0>: a module used to detect load current and compensate for over- and undershoot.

* When using the fast recovery mode, the maximum output cap allowed is reduced with the following formula: $\text{new_load} = (\text{old_load} - 50\text{p})/2 + 50\text{p}$

6.15.3 LDO Topology

Commonly used capless LDO's are designed to work under strict load conditions, that are known and will not change in a very wide range. This simplifies design considerably, because unused load conditions can just be ignored. Designing an LDO where load conditions can vary considerably can, if not done properly, restrict the LDO's effectiveness and stability. Either the LDO must be significantly oversized, potentially leading to excessive power consumption, or it may not be suitable for all load scenarios. Capless LDOs are usually designed from two separate voltage rails. One rail for the pass device, where bulk of the current is drawn from and one for the error amplifier, that regulates the pass device (Figure 69). This configuration enhances efficiency, as only a minimal current is required from the higher voltage 1.8V rail (VDD1V8) to generate the necessary control signal, allowing the dropout voltage across the pass device to remain minimal.

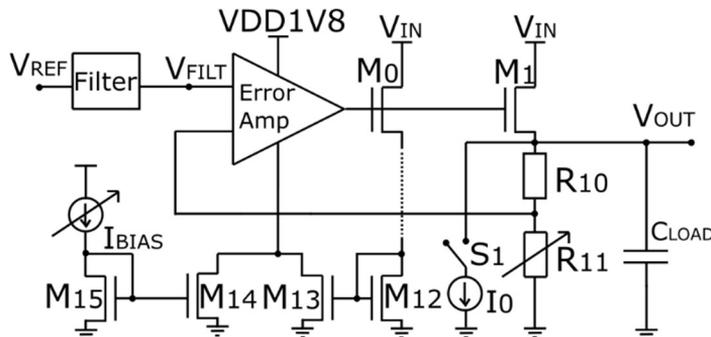


Figure 69. Topology of a proposed LDO with the resistive divider after the error amplifier.

$$V_{REQ} \cong 0.13 \sqrt{R_{EQ}} \quad (11)$$

where R_{EQ} is $R_{10} || R_{11}$, that will floor the noise at $13\text{nV}/\sqrt{\text{Hz}}$ with a $10\text{k}\Omega$ R_{EQ} and $4\text{nV}/\sqrt{\text{Hz}}$ with a $1\text{k}\Omega$ R_{EQ} . While $1\text{k}\Omega$ resistance floors the noise at $4\text{nV}/\sqrt{\text{Hz}}$, the current wasted from a 0.9V regulated output is already considerable, making the classical approach either inefficient with low loads or somewhat noisy.

Moving the resistive divider before the error amplifier and filtering the noise with a RC filter will remove the noise of the divider at higher frequencies. It will not do anything with the noise of the pre-amp and resistive ladder at low frequencies. If the RC Filter is larger more noise in the lower frequencies will be filtered. Therefore noise can be improved by adding extra capacitance outside the module in layout to the node V_{FILT} , that will move the corner frequency lower, filtering more noise in lower frequencies.

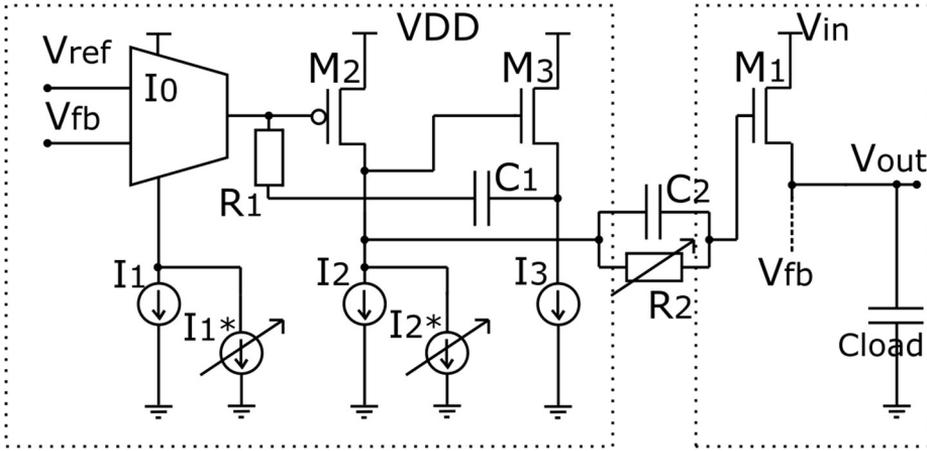


Figure 71. Topology of the error amplifier together with the pass device.

6.15.4 Error Amplifier Topology

The challenge of the design is to keep the phase margin constant over all load conditions. For that, internal compensation schemes have been developed to provide the LDO with some basic intelligence (Figure 71). The Error Amplifier itself consists of two stages plus a follower stage for Miller compensation. The first amplifying stage is a common differential pair, that consist of an NMOS input pair and PMOS mirrors for active load. The second stage is a little more than a single stage PMOS amplifier M_2 with NMOS current source. The circuit uses a Miller compensation scheme, that is deployed through a buffer composed of M_3 , that acts like an NMOS follower and a current source I_3 . The buffer serves two purposes. First, it level shifts the voltage seen by capacitor C_1 , making it possible to use much smaller non-linear MOS capacitor in the place of a linear capacitor. Secondly, it decouples the Miller compensation from the signal path. The Miller compensation creates a pole that comes from the output impedance of the first stage and the Miller capacitor. Assuming that the Miller capacitor is much larger than the gate capacitance of M_2 , a constant gain of one from the NMOS follower composed of M_3 and I_3 and that the current sources I_2 and I_2^* have a much larger output impedance than M_2 , the location of the pole is at:

$$P_0 = \frac{1}{2\pi * R_{OUT_{I_0}} * C_1 * R_{OUT_{M_2}} * gm_{M_2}} \quad (12)$$

The resistor R_1 adds a Miller zero to the transfer function. The location of the Miller zero is at:

$$Z_0 = \frac{1}{2\pi * R_1 * C_1} \quad (13)$$

The second stage combined with resistor R_2 and capacitor C_2 will add another pole and another zero. Assuming that the output resistance of I_2 and I_2^* are much higher than the output impedance of M_2 , the location of the second pole becomes:

$$P_1 = \frac{1}{2\pi * (R_{OUT_{M_2}} + R_2) * C_{g_{M_1}}} \quad (14)$$

The location of the second zero is at:

$$Z_1 = \frac{1}{2\pi * R_2 * C_2} \quad (15)$$

The last pole comes from the pass device. Since the pass device is an NMOS follower, the output impedance of it will be relatively low. However, the output capacitance can be – depending on the topology of the A/D, D/A or any other IP – quite high, spanning up to couple of hundreds of picofarads. The pole coming from the pass device is given by:

$$P_2 = \frac{1}{2\pi * R_{OUT_{M_1}} * C_{load}} \quad (16)$$

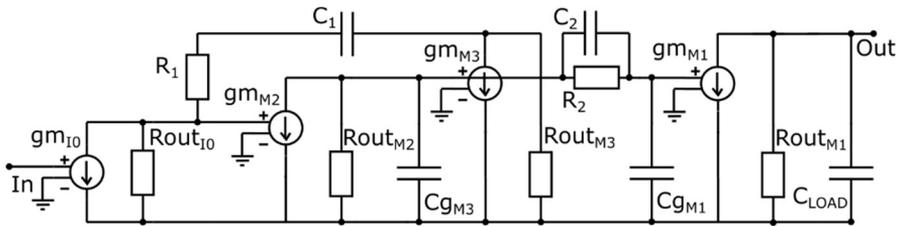


Figure 72. Small signal equivalent of the LDO.

A thorough examination of the design equations reveals that the two zeros are independent of load conditions and solely determined by the values of passive components. In contrast, the poles exhibit a dependency on the output impedances of various stages. By employing meticulously engineered current feedback, the system can dynamically adjust the output impedances by injecting varying amounts of current through I_1^* and I_2^* . This approach effectively lowers the output impedances of I_0 and M_2 , enabling further pole shifting into higher frequency ranges. Consequently, the feedback loop's bandwidth expands, enhancing the regulation speed. The LDO incorporates four distinct bandwidth modes, each employing specific current injection levels governed by the `BW_sel<1:0>` option. The design is specifically tailored to ensure that the feedback loop's bandwidth surpasses or equals that of the converters in Wifi320 mode, considering the approximate estimated current consumption of the A/D and D/A converters. The `BW_selection=2` mode serves as the default selection for all simulations and measurement results.

6.15.5 Stability of the LDO

Maintaining robust stability and phase margin (PM) across varying load conditions is crucial for optimal LDO performance. The pole at the LDO's output is intrinsically linked to the output impedance of the pass device, which is highly dependent on the current

flowing through it. As an NMOS follower, the pass device exhibits relatively low output impedance. However, the output capacitance can range from a few hundred picofarads up to a nanofarad, depending on the blocking capacitor configuration and the topology of the supplied module. This output capacitance can significantly impact stability, particularly under low current load conditions. Larger output capacitors cause the pole introduced by the pass device to migrate to lower frequencies, leading to a reduction in PM. To address this issue, the LDO's design enables internal programming to compensate for load capacitor values (Figure 73). When the output capacitor is substantial, a small current source I_0 is activated at the output to limit the minimum current that the pass device can supply, limiting the output impedance of the pass device to a fixed maximum. This current source operates only when the pass device's current is low and automatically deactivates when the current increases. The LDO offers four distinct capacitive ranges that are controlled by the `load_selection<1:0>` signal. The PM across different simulation corners over maximum capacitor values in all `load_selection` modes, with `BW_selection<1:0>=2`, and `fast_recovery=0` is illustrated in the following figures:

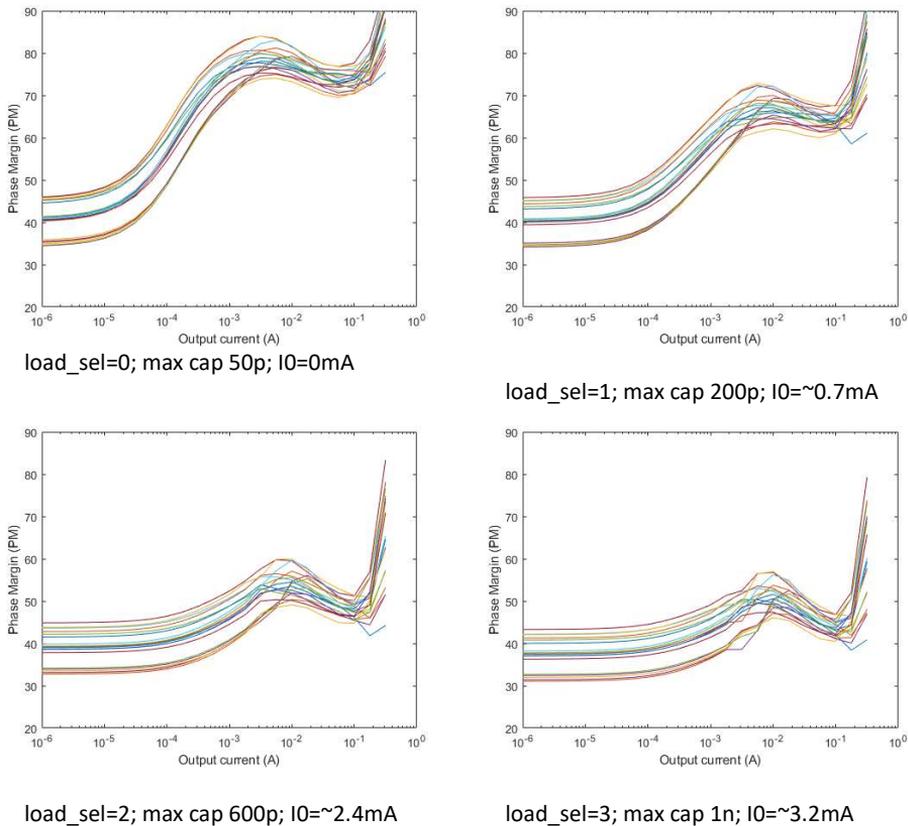


Figure 73. Simulated PM of the LDO with different capacitive loads over various outputs.

The LDO also has a built-in fast recovery mode, where resistor R2 has been halved. This roughly halves the time required for the LDO to recover when it goes out of operating conditions in case of heavy load jumps but also has an impact on stability. Therefore the maximum output cap used over different mode must be reduced.

6.15.6 Efficiency of the LDO

The LDO utilizes a portion of the current supplied to the load for its own internal operation. Part of this current is drawn from the 1V8 domain, while the majority originates from the 1V2 domain. In low-load scenarios, the LDO can activate an internal current source ($\sim 0.7\text{mA}/2.4\text{mA}/3.2\text{mA}$) from the 1V2 domain with the `load_selection<1:0>` (Figure 74). This internal current source is gradually deactivated when the output current surpasses a pre-defined point of $\sim 3\text{mA}$. At an output current of approximately 20mA and `BW_sel=2`, the self-consumption of the LDO is around $\sim 600\mu\text{A}$ and $\sim 200\mu\text{A}$ from the 1V8 and 1V2 domains respectively, translating to an efficiency of 82% with an input voltage of 1V1. As the output current increases to approximately 300mA , the LDO's self-consumption rises to $\sim 4\text{mA}$ and $\sim 300\mu\text{A}$ from the 1V8 and 1V2 domains respectively, translating to an efficiency of 85% with an input voltage of 1V1. Different `BW_sel` settings influence the 1V8 current consumption, as illustrated in Figure 74:

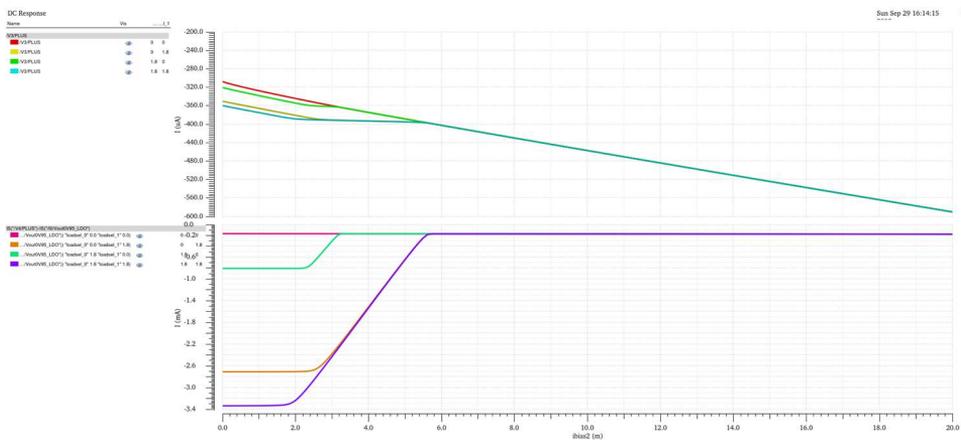


Figure 74. Simulated current consumption of the LDO with different load cap settings over various output currents.

6.15.7 Noise of the LDO

As elaborated in the preceding subsection, the LDO employs a structure where the reference voltage is amplified by a preamplifier to the desired output level and then buffered by an error amplifier featuring 100% voltage feedback. This arrangement enables us to relocate the resistor divider, a major source of noise, upstream of the error amplifier, allowing for subsequent filtering. The error amplifier's noise dominates medium-low ($>100\text{kHz}$) to high frequencies, while the preamplifier and resistor divider contribute significantly to low ($<100\text{kHz}$) frequencies (Figure 75). For enhanced noise performance in the low-frequency range, additional blocking capacitors can be added to the V_{FILT} node.

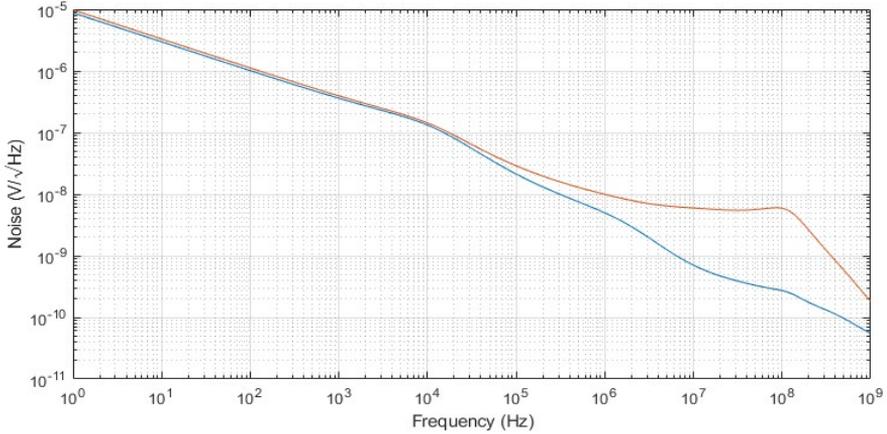


Figure 75. Simulated noise profile of the filter resistor (blue) and the full LDO (orange) with 20mA load.

High-frequency noise can be mitigated by selecting higher bandwidth modes in the error amplifier. * This enhances the speed of current feedback and injects more current back to the LDO's current feedback path (Figure 76). It also causes the LDO's bandwidth to increase more rapidly, effectively suppressing high-frequency noise. In an exceptional case for demanding low-noise applications (<2nV/√Hz), a special mode has been implemented to further optimize noise performance (programmable via the noise_enhancement option – a feature that would not have been possible with the resistive divided left downstream from the error amplifier). This mode restricts the maximum allowable output capacitor to 300pf and the maximum output current to 50mA.

*Some noise performance is lost in medium frequencies with higher BW modes.

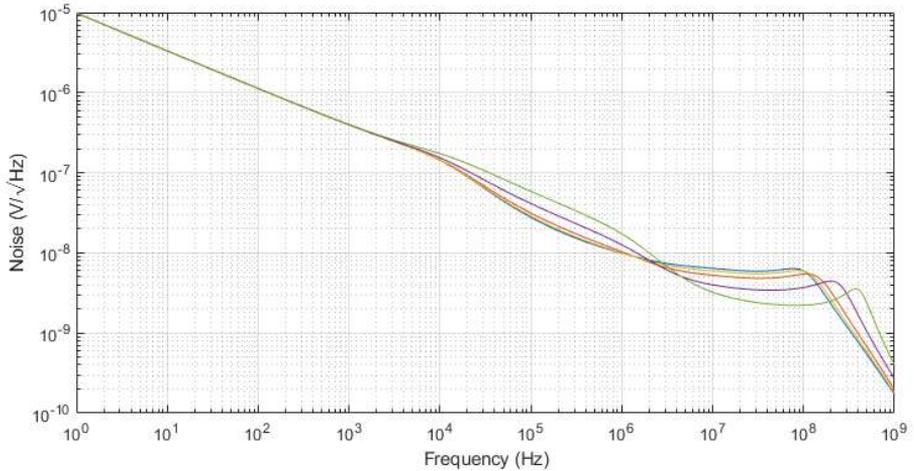


Figure 76. Simulated Noise profile of the LDO over different LDO BW settings. (Green represents the noise_enhancement mode).

6.15.8 PSRR of the LDO

Since the LDO uses two different power domains, PSRR is important for both of them (Figure 77, Figure 78). The main interest is the performance around 1–2MHz, where most commercially available DC-DC converters work. With 20mA output current the BW_sel has an impact on PSRR from 1V8/1V2 domains as shown in the following figures:

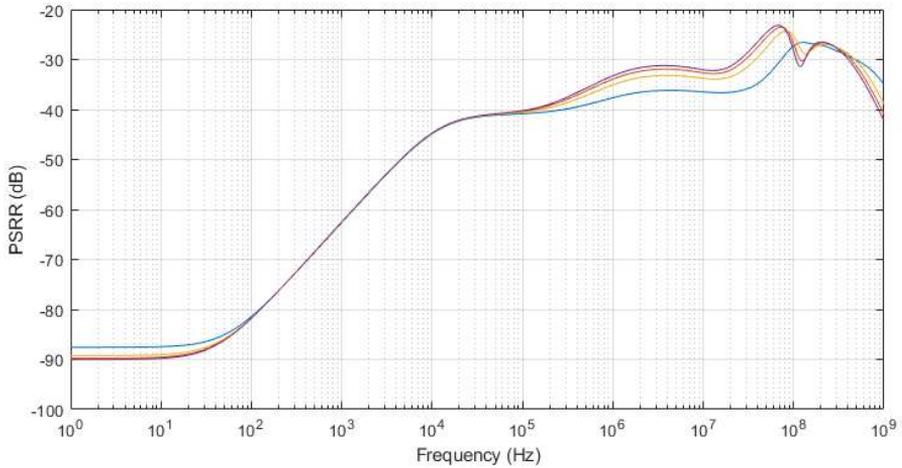


Figure 77. Simulated PSR ratio of the LDO from the 1V8 power rail over different BW settings with a 20mA output current.

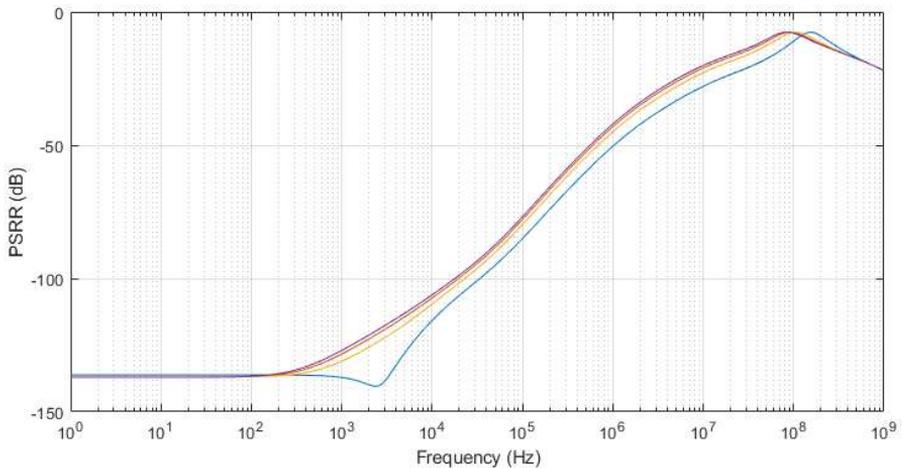


Figure 78. Simulated PSR ratio of the LDO from the 1V2 power rail over different BW settings with a 20mA output current.

6.15.9 Extra Features of the LDO

Fast Recovery

The LDO has a built in fast_recovery mode, where resistor R_2 in Figure 71 has been halved. This roughly halves the time required for the LDO to recover when it goes out of operating conditions in case of heavy load jumps, but also has an impact on stability. Therefore the maximum output cap used over different mode must be reduced. The transient response of load jumps with the fast_recovery bit enabled and disabled has been plotted in Figure 79:

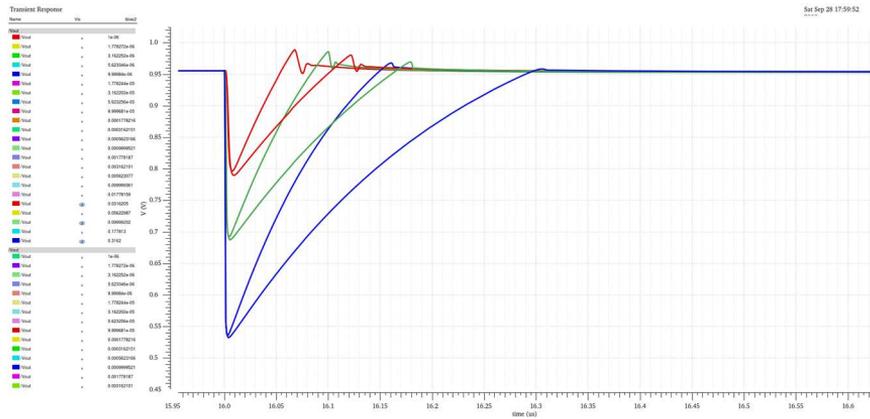


Figure 79. Impact of fast_recovery in the LDO regulation.

Enable Load

The LDO also incorporates a built-in manual 3mA load that can be activated at the output (Figure 80). This feature simply introduces a fixed 3mA load current at the output that remains active even when the LDO's automatic load detection is disabled. The option can be particularly beneficial with the noise_enhancement enabled, where the feedback current tends to turn off the automatic load excessively, leading to stability issues in extremely low-load scenarios. If such low-load conditions are not encountered, there's no need to activate the 3mA option. However, if they do arise, enabling this extra 3mA can effectively address stability concerns.

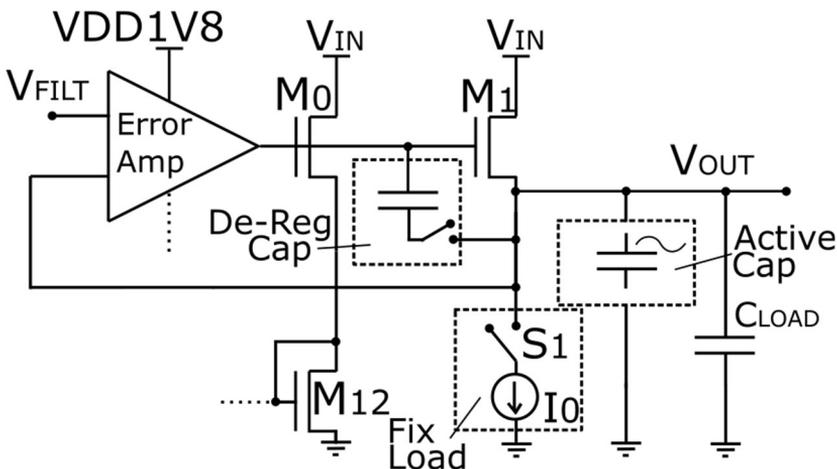


Figure 80. Extra options of the LDO. a) Fixed load b) De-Reg Cap c) Active Capacitor.

De-Regulation Capacitor

A de-regulation capacitor has been strategically positioned between the Error amp output and the LDO output (Figure 80). The capacitor can be used to break down a carrier frequency, mitigating spur emissions generated by the modules powered by the LDO, operating at specific clock frequencies and causing disturbances across the chip. This approach can be particularly effective at the aggressor side, for modules operating between the frequency range of 50MHz to 500MHz, such as timing-to-digital converters (TDCs) and certain oscillators.

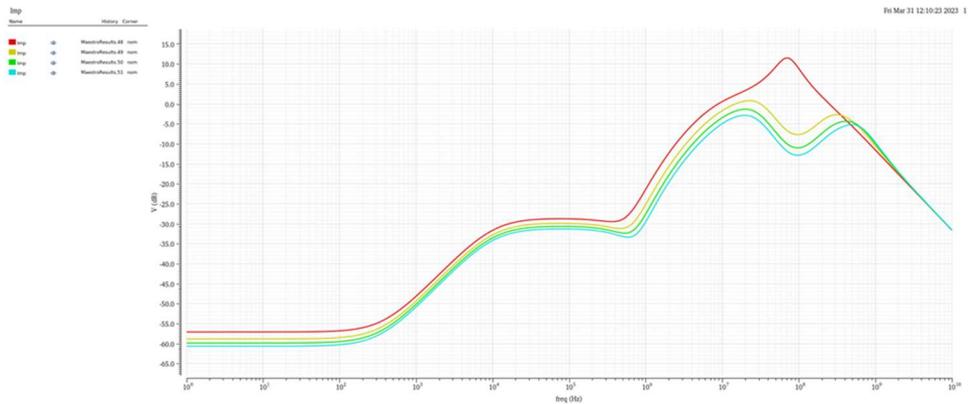


Figure 81. Impact of the active capacitor on the output impedance of the LDO.

Active Capacitor

To enhance the output impedance of the LDO at frequencies beyond the bandwidth of the Error amp, before the blocking capacitor becomes effective, an Active Capacitor module (aka Shunt) has been incorporated (Figure 80) [29]. This approach can lower the impedance of the power supply greatly in certain frequencies, improving the decoupling features and enhancing the spur resistance properties to protect the powered module (Figure 81). While the deregulation capacitor works on the aggressor side, the Active Capacitor has a dual benefit, effectively combating spurs both at the aggressor and victim sides. It can be programmed in four modes:

- 0: Active Capacitor off
- 1: Minimum active capacitor used (current penalty 7mA)
- 2: Medium active capacitor used (current penalty 14mA)
- 3: Maximum active capacitor used (current penalty 21mA)

6.16 Design of the Active Capacitor

In a similar fashion to the LDO, active capacitors have been strategically positioned and connected to various power rails throughout the chip, including integration within the $\Sigma\Delta$ ADC. These active capacitors serve as valuable tools for debugging and mitigating spur issues. Their operation is based on a sensing mechanism that detects and counteracts changes in supply voltages. M_0 is biased through R_1 to operate as a load, and upon detecting a voltage variation through the high-pass filter formed by M_1 , R_0 , and C_0 , a countervailing signal is relayed to M_0 via C_1 to neutralize current fluctuations that induce voltage changes, similar to a real capacitor (Figure 82). In order to be truly effective, the M_0 current needs to be in a similar range to the current ripple introduced by the module

the LDO is tasked with supplying. The Active Capacitor is designed to be effective in frequency ranges where both the regulator and the blocking capacitor are ineffective, particularly in the 100MHz range, where numerous oscillator crystals operate, creating a potential source of spurs all over the SoC. The Active Capacitor has been designed to operate in four distinct modes, dependent on the switching activity associated with the supply node, encompassing a current range from 0 to 21mA. It has demonstrated its versatility as a debugging tool for identifying disturbers across the SOC.

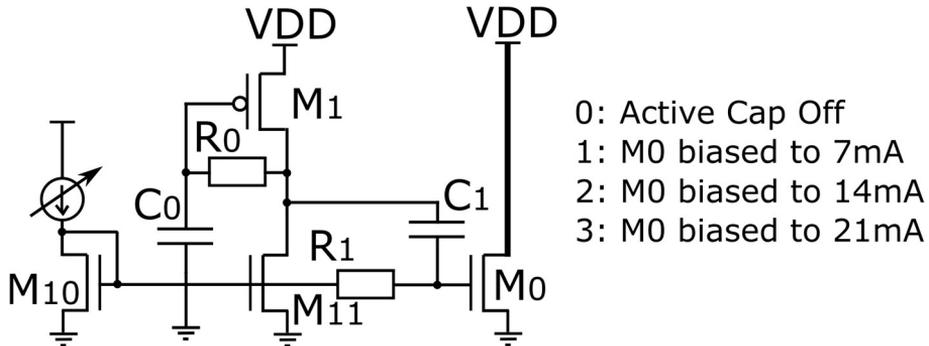


Figure 82. Topology of the Active Capacitor.

6.17 Design of the VCO

This chapter covers the following topics:

- The idea behind the VCO based sigma delta converters.
- Integration of the VCO-modes in the previously designed SDADC.

6.17.1 Background of VCO-based Time Encoding ADCs

A new and exciting recent development in sigma-delta modulation in the last decade has been the adoption of VCO-based conversion. A VCO-based delta-sigma converter leverages a voltage-controlled oscillators to encode the analog input signal into the frequency domain using $\Sigma\Delta$ modulation. It is highly digital in nature, easy to implement and port, plus has the added benefit of introducing a first order noise shaping function through the integrating properties of the VCO already in the sampler.

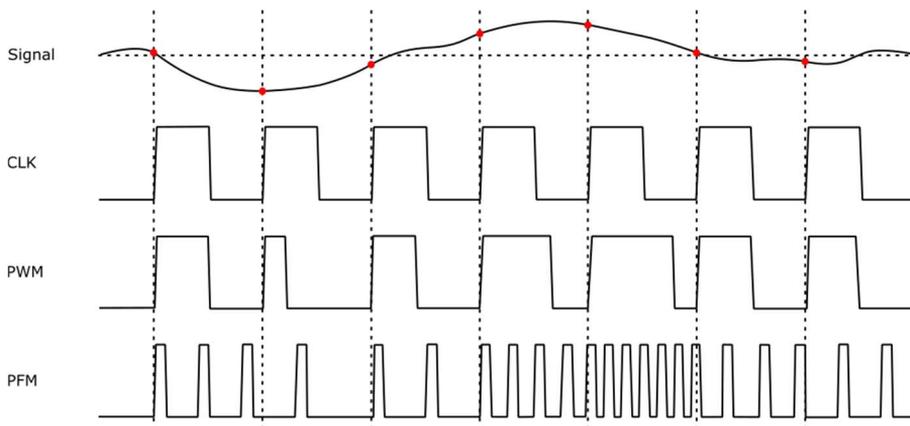


Figure 83. PWM and PFM modulation techniques.

At the heart of time encoding lies the concept of representing an analog signal using a modulated square wave, where the signal's information is encoded not in its instantaneous amplitude but rather in the transitions between its high and low states. This approach offers several advantages, including ease of handling with digital circuitry and resilience to noise and distortion [30]. The representation of analog signal information is achieved through pulse modulation techniques (Figure 83). Pulse-width modulation (PWM) stands as a prominent example and finds widespread application in various domains, particularly in power electronics. Time-encoding analog-to-digital converters can generate PWM-like outputs, but their implementation typically requires complex and highly linear circuitry, making them less attractive. A more ingenious approach is to utilize VCO-based ADCs, which can be efficiently implemented using a simple and scalable ring oscillator. This approach also shifts the pulse modulation type from PWM to pulse-frequency modulation (PFM), that is totally asynchronous and enables first-order noise shaping properties when directly sampled, unlike other pulse modulation techniques (such as PWM). The hardware simplicity, digital nature, and noise-shaping properties of PFM have propelled VCO-based ADCs lately to great prominence.

6.17.2 Functionality of VCO-based ADCs

A slowly varying input signal modulates the frequency of the VCO, encoding the analog signal's information into the VCO's oscillation frequency. The number of oscillator pulses generated during a fixed sampling period can then be accurately counted using a digital counter. The digital counter's value directly corresponds to the VCO's frequency and, consequently, provides a precise representation of the analog input signal. It is evident that a higher free running oscillation frequency relative to the sampling frequency translates to improved ADC resolution. This defines an important design relation in VCO-based ADCs: the resolution is proportional to the oscillator frequency relative to the sampling frequency (which corresponds to the number of quantization bits in a traditional amplitude-based ADC).

Conventionally, this implies that the sampler undersamples the rapidly oscillating VCO, which may seem counterintuitive, given that it suggests lower resolution. However, the quantization error arising from discretizing the oscillator's frequency into an integer value (the count) exhibits first-order noise shaping. This implies that the error introduced after and before sampling at the mid-point of a VCO cycle always adds to one cycle, and the quantization error can be expressed as the first-order difference of a finite-power sequence. This first-order noise shaping, coupled with oversampling and digital decimation following the counter, significantly enhances the signal-to-quantization-noise ratio (SQNR) compared to merely increasing the oscillator frequency.

The following approximate equation gives the maximum SQNR of a VCO-based ADC with analog signal bandwidth BW , sampling frequency f_s , and oscillator rest frequency is f_0 :

$$\text{SQNR} = 6\text{LOG}_2 \frac{2f_0}{f_s} - 5.17 + 9\text{LOG}_2 \frac{f_s}{2BW} \quad (17)$$

This equation demonstrates that the SQNR increases by 6dB whenever the VCO frequency f_0 is doubled relative to the sampling frequency f_s , corresponding to an increase in quantization resolution by one bit. If the oversampling ratio is doubled simultaneously with the quantization resolution, meaning the sampling frequency f_s is doubled alongside the oscillator frequency f_0 , the SQNR improvement is a remarkable

9dB. Additionally, if the sampling frequency f_s is doubled while maintaining a fixed oscillator frequency f_0 , the SQNR is improved by 3dB. The selection of oscillator and sampling frequencies allows for the first-order optimization of the SQNR.

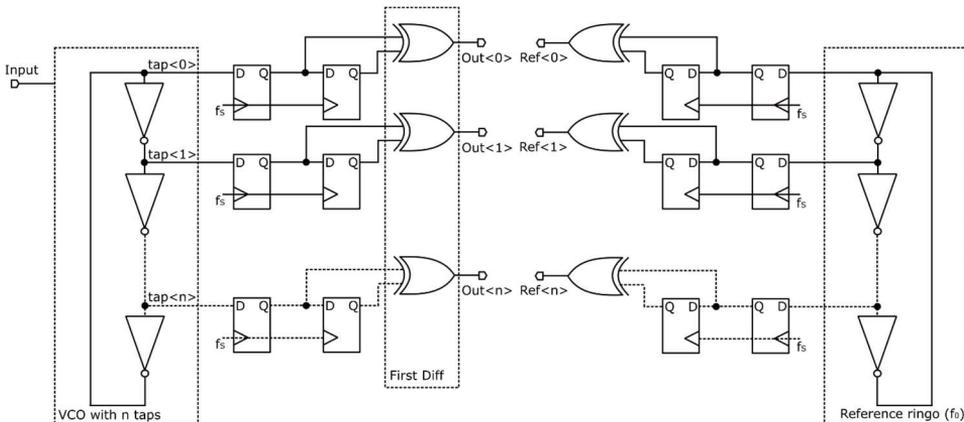


Figure 84. Topology of a VCO based Sampler offering first order noise shaping.

Most practical VCO-based ADCs are implemented with a ring oscillator, which is a circuit that produces many similar square wave signals slightly shifted in time, known as taps (Figure 84). By comparing the taps of the VCO to an equivalent reference circuit operating at a resting frequency, a straightforward mechanism for determining the count increment is achieved, eliminating the need for frequent counter resets during each sampling clock cycle. This exemplifies the inherent simplicity of VCO-based ADCs: they produce a high-resolution, first-order noise-shaped sequence using just a few digital logic gates and no operational amplifiers. Early VCO-based ADC converters were considered alternative implementations of $\Sigma\Delta$ modulators, and the explanation for their operation aligns with the principles of $\Sigma\Delta$ modulation.

6.17.3 VCO Design

The $\Sigma\Delta$ ADC design presented a compelling opportunity to integrate a low-power, low-bandwidth VCO mode into the B-step of the ADC by effectively leveraging a significant portion of the existing $\Sigma\Delta$ ADC circuitry. By introducing a VCO after the first integrator and feeding back the signal through DAC1, a second-order $\Sigma\Delta$ loop could be implemented. The VCO configuration allows for the shutdown of most of the $\Sigma\Delta$ modules, effectively reducing the active components to opa1, DAC1, and the clocking generation circuitry. This approach significantly curtails the converter's power consumption, as the opa1 can be programmed to its lowest power mode, and the overall power consumption of the VCO's ring oscillator is negligible compared to the remaining components. Consequently, the VCO mode emerges as a very attractive solution from a power-efficiency standpoint.

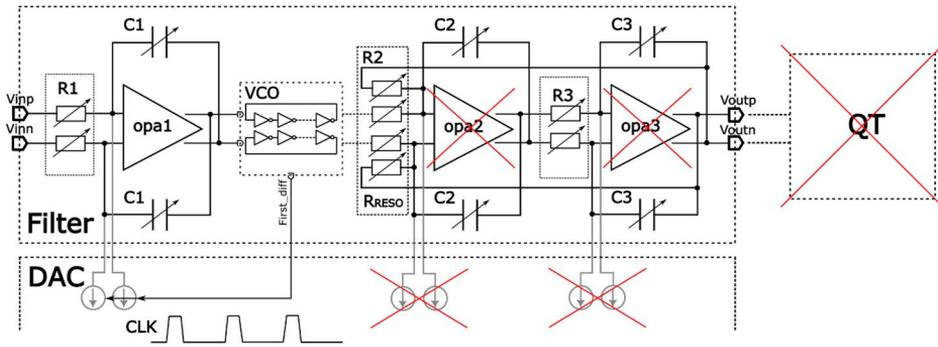


Figure 85. Integration of the VCO mode in the previously designed SDADC.

As an initial step, a 10MHz w20 mode was implemented, capable of operating effectively in both the lowest bandwidth Wi-Fi channels and in idle mode used for scanning the environment. The successful integration of this mode could pave a way for further development of progressively higher bandwidth VCO modes. The implementation was constrained to utilize existing modules, resulting in a VCO design with limited degrees of freedom and notable limitations. One primary restriction stems from the coefficient set embedded in the loop filter. To minimize the number of switches, the decision was made to retain the validated filter structure while introducing new functionalities. This constraint implies that the integrator bandwidth attainable by the first opamp is bounded by the passives already incorporated into the filter, necessitating the design of the remaining VCO properties around this constraint.

Another obstacle for the successful implementation of the VCO-mode is the lack of good tools building up coefficients surrounding the topic. Proper VCO design had to involve lots of trials of the modelling level to find the correct properties. In the end, a second-order, 4-bit VCO-based ADC mode working with 480MHz was implemented.

The VCO properties attained to best fit the set-out targets were:

- f_s – 480MHz (sampling frequency)
- f_0 – 110MHz (oscillator frequency)
- K_{VCO} – 710MHz/V (VCO gain)
- taps – 60 (6-bit ringo)
- ELD – 0.25 (duty cycle of clock 25%)

6.17.4 Integration of the VCO-mode

The inclusion of the VCO-mode in the $\Sigma\Delta$ ADC requires some minor modifications within other modules in the ADC (Figure 86):

- Master Clocking – Needs to implement an option for smaller duty cycles than 50% to achieve proper ELD compensation. 3 new modes were implemented: 25%/12.5%/6.25%.
- Loop-Filter – The implementation necessitates the inclusion of an attenuator to restrict the K_{VCO} range, enabling linearization of the ADC's operation. This addition introduces an extra load on the opa1. Fortunately, the opa1 possesses the capability to augment the AB stage current specifically for such scenarios.
- Feedback DAC – The feedback DAC requires two sets of inputs. One set coming from the QT and the other coming from the VCO. Inputs coming from the VCO do not have an illegal combination (unlike the QT), meaning that cells in the

feedback DAC used in the VCO-mode cannot be grouped in a similar way as it is done in w80/w160/w320. Effectively meaning the DAC loses one potential bit and can only be operated in only either 4-bit or 5-bit modes, allowing unit cell operate in six dynamic modes i.e.:

- 1) Calibration
- 2) Zero/Dumping
- 3) Plus
- 4) Minus
- 5) Disable
- 6) **Zero/Zero**

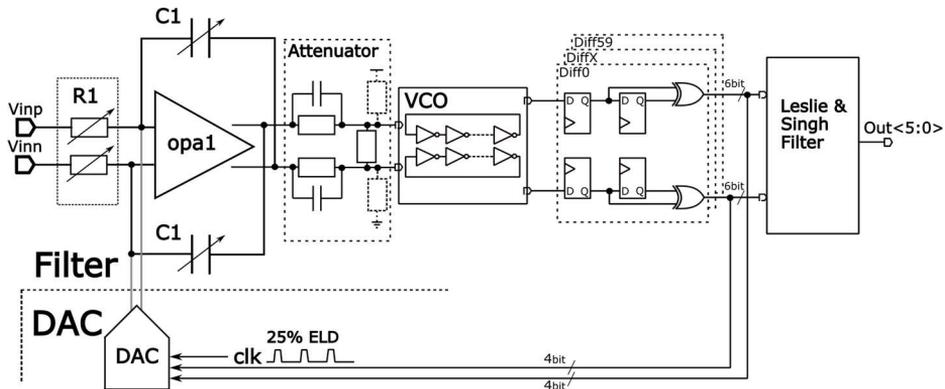


Figure 86. Topology of the designed VCO-based Sigma-Delta converter.

Zero/Zero is an operational mode where total current added in the summing nodes is net zero. This happens either when the unit cell neither adds or subtracts total current to loop summing node or adds and subtracts the same amount. In this implementation, when working in operational mode zero/zero, the unit cell will apply feedback by closing the switches S_{PP} , S_{NP} , S_{PN} and S_{NN} , and keeping all other switches open, effectively adding and subtracting the same amount of current from each summing node. This will not remove the thermal noise coming from the current sources as it does like functioning in operating mode Dumping. In practical terms, this means that the VCO-mode will have two types of operating mode Zeros, both equally probable, one type adding noise, the other type not.

6.17.5 Leslie & Singh Filter

A well-established method to reduce the OSR of a $\Sigma\Delta$ ADC is to enhance the bit-width of the QT. Leslie & Singh proposed a novel approach to achieve this goal in conjunction with a single-bit DAC, aptly named the Leslie & Singh filter [31]. This approach offers significant advantages in DAC linearization, as a single-bit DAC is inherently linear by design, eliminating the need for multibit DAC linearization techniques. Theoretically, this method is not limited to single-bit DACs and can be applied to any DAC with a smaller bit-width compared to the QT. In the proposed implementation, a 4-bit DAC was utilized in conjunction with a 60tap ring oscillator. To compensate for the inherent limitations of a 4-bit DAC, a Leslie & Singh filter was employed to convert 4-bit data within the feedback loop to an SNR equivalent to a 6-bit sampler. This approach effectively mitigates the performance limitations imposed by the DAC structure, potentially regaining an impressive 12dB of performance.

6.17.6 Usage of the VCO-mode

The K_{VCO} , a parameter that broadly reflects the gain of the ring oscillator, exhibits significant process dependence and necessitates calibration to the intended value prior to operation. This calibration process resembles the QT calibration, performed once before the ADC is deployed. However, instead of fine-tuning each pre-amplifier, only one transconductor requires calibration. This is accomplished through a coarse transconductance selection followed by a 5-bit binary weighted finetuning step [32]. In essence, one of three coarse-value transconductors is chosen, and the 5-bit fine-tuning is swept. The sweep, where the desired K_{VCO} lands closest to the middle code in the finetuning step will be stored and used. The coarse tuning is subsequently left untouched, and for any temperature variations, only the fine tuning is re-swept to identify the optimal calibration value. This adjustment can be executed swiftly, even during handover between transmit and receive modes. With a minor addition of logic, the Calibration Machine described in Chapter 6.13 can be repackaged for this purpose. Both calibrations can even be performed concurrently, as the VCO calibration only necessitates comparing the outputs of the first difference and does not require read/write access to any registers. Once the correct frequency is detected, tuning values can be stored in an internal memory and fine-tuned as needed. Owing to the inherently digital nature of the VCO, calibration and usage can be implemented in a highly straightforward manner.

The integration of the VCO into the $\Sigma\Delta$ ADC opens up a plethora of unexplored possibilities. While its primary application lies in ultra-low-power modes like w20 and idle, the VCO-mode holds promise for enhancing the AGC algorithm as well. Currently, a separate parallel signal path with a dedicated ADC is employed to assess the input power of the ADC. This is due to the $\Sigma\Delta$ ADC's third-order STF, which introduces discrepancies between the output and input signal power in the presence of strong blockers, implying that the $\Sigma\Delta$ ADC's STF is excessively steep (usually a good thing) for accurate power estimation in the presence of blockers. However, the VCO mode signal path, characterized by a single analog integrator, exhibits only first-order STF. This facilitates enhanced power estimation within the ADC itself. Operating the VCO in conjunction with the primary modes could therefore potentially simplify the power estimation and improve the AGC algorithm. Additionally, as a future enhancement it is proposed to feed back the oscillator output to the ADC input via the loopback path. This would enable the accurate measurement of the $\Sigma\Delta$ ADC's STF, as PCB parasitics approximation has always been challenging, often requiring estimating channel equalization and finally resulting in poor STF accuracy.

6.18 Design of the Programmable Gain Controller

This chapter covers the following topics:

- Design of the PGC.
- Different modes supported by the PGC.

The Programmable Gain Controller (PGC) serves as a crucial module positioned upstream of the ADC within the signal chain. It is built as a transimpedance amplifier, converting the output of a current domain mixer to voltage domain. The primary functions of the PGC include providing the fine-tuning step for ensuring accurate gain adjustment as dictated by the AGC algorithm, providing first order OoB filtering and driving the load associated with the ADC. The PGC offers eight distinct gain steps, together with five different BW modes.

The gain steps are defined as:

- 0dB
- 2dB
- 4dB
- 6dB
- 8dB
- 10dB
- 11dB
- 12dB

The BW modes are defined as:

- w40 (-3dB corner frequency 25MHz)
- w80 (-3dB corner frequency 50MHz)
- w160 (-3dB corner frequency 100MHz)
- w320 (-3dB corner frequency 200MHz)
- DPD (-3dB corner frequency 400MHz)

meaning a total support for 40 different configurations.

The gain programming is done with a 5-bit binary weighted resistor array, that is controlled with an internal 3bit decoder, meaning only eight combinations of the resistor array can be chosen externally (Figure 87a). The AGC algorithm operates with gains between 0dB and 10dB, while 11dB and 12dB gain steps are only used for DPD training and further debug.

The corner frequency is programmed employing a 9-bit binary-weighted capacitor array, controlled by a 6-bit tuning word in conjunction with the BW mode selection (Figure 87b). This approach enables the utilization of a single shared tuning word for each of the BW modes. Mode switching merely involves a binary shift, allowing for the re-employment of the same tuning word.

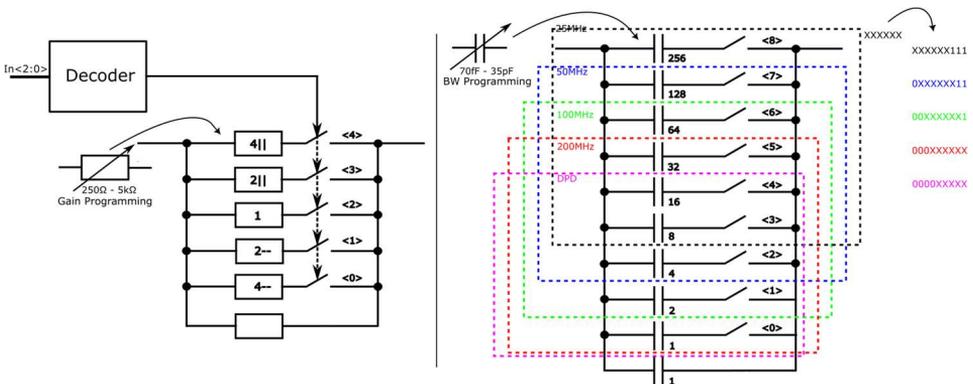


Figure 87. a) Gain programming of the PGC; b) BW programming of the PGC.

The designed amplifier employs a fully differential three-stage feed-forward compensated signal path, a two-stage Miller compensated common-mode loop a three-stage feed-forward compensated common-mode loop and a two-stage Miller compensated biasing loop inside the AB-stage (Figure 88). It is designed to have extremely low noise and exceptionally high linearity, and for this reason the AB stage utilized in the opamps of the $\Sigma\Delta$ ADC filter have been replaced with a more linear version

of the AB stage. The push-pull stage used in the PGC opamp offers less gain than the two-stage topology utilized in opa1 and opa2 of ADC filter but has equal and balanced operating gains from both the n-side and the p-side, resulting in HD3 performance well exceeding 100dB [33].

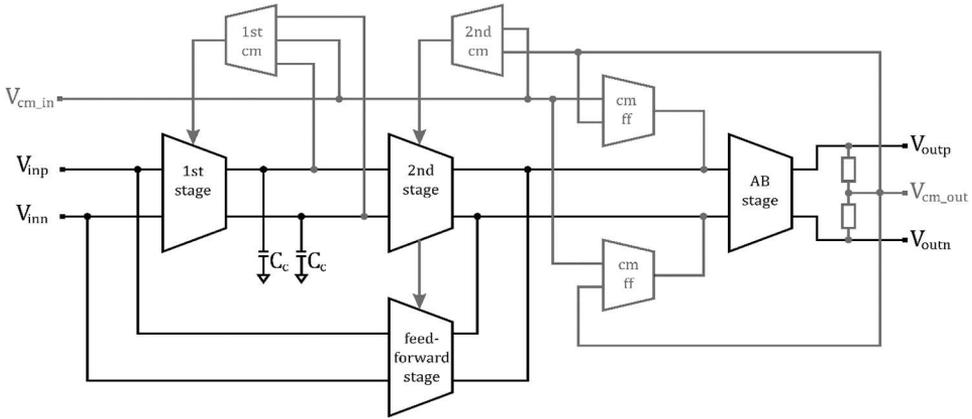


Figure 88. Topology of the PGC.

The feed-forward structure is employed to ensure a sufficiently high loop gain within the desired operating bandwidth. The exceptionally high DC gain of the multistage configuration is not specifically intended nor designed for, but rather emerges as a serendipitous byproduct. The true value of the feed-forward structure lies in its ability to enable the placement of multiple poles in close proximity, resulting in a higher-order (beyond second-order) roll-off. This enhanced roll-off is then compensated for by the addition of one or more zeros arising from the feed-forward branch before the unity gain frequency. If not implemented correctly, this approach can lead to potential stability issues, as the phase can easily shift beyond 180 degrees before the GBW. Moreover, the GBW's location varies with the gain configuration of the PGC. For instance, in buffer mode, the amplifier exhibits more than two clearly observable uncompensated poles between 30dB and 50dB (Figure 89b). Nevertheless, since the gain of the amplifier is inherently controlled by the system and cannot be freely programmed by the user as with a generic standalone amplifier, stability can be readily achieved through the utilization of distinct predefined biasing programming currents, which effectively shift the singularities for different gains.

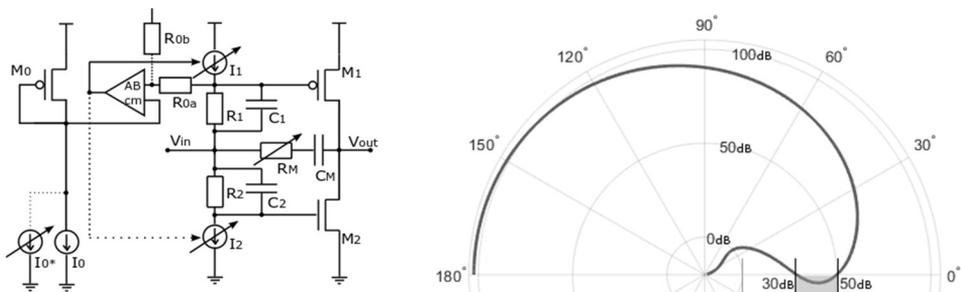


Figure 89.a) AB-stage of the designed PGC; b) Nyquist plot of the PGC.

A more balanced AB-stage employed in the PGC ensures extremely high linearity even with single tone Total Harmonic Distortion (THD) measured well above 100dB (Figure 89a). It uses a replica circuit to set a desired biasing for the AB-stage through a two-stage amplifying circuit (AB cm). The replica circuit composing of M0 and I0 sets a desired biasing point, that is then mirrored through the amplifier to M1. The DC current of the AB stage can easily then be set to the desired value with a designed ratio between M0 and M1.

The PGC is designed to function as both an on-chip gain stage and an off-chip measurement buffer. To mitigate the impact of parasitic capacitances introduced by external traces, pads, bumps, and other off-chip components, a 50 Ω resistor is placed at the PGC's output. This would change slightly the gain of the $\Sigma\Delta$ ADC if left uncompensated, therefore the input resistance of the $\Sigma\Delta$ ADC module has been partitioned between the PGC and the ADC and 50 Ω has been removed from the input resistor of the ADC.

7 Optimization and Validation of the Design

This chapter covers the following topics:

- Optimization of the coefficients during design phase.

7.1 Optimization of the Coefficients

Upon the first completion of circuit design, it became plainly evident that the w320 mode exhibited excessive out-of-band (OoB) peaking, stemming from parasitics associated with the virtual ground of the third integrator. This peaking would have rendered the AGC algorithm inoperative for the w320 mode, necessitating adjustments to the coefficients to counteract it. Excessive peaking can be remedied in two primary ways: either by introducing additional attenuation within the inherent coefficient design and utilizing peaking to restore it to the original state or by augmenting the feedback within the most internal loops. For this design, a combination of these approaches was implemented for the w320 mode, with corresponding coefficient adjustments. Modifying the coefficients for w320 mode also affects the coefficients for w160LP mode, leading to the creation of two sets of coefficients: one set for w320 and w160LP modes and a separate set for w160 and w80 modes. The first set of coefficients exhibits a slightly less aggressive characteristics compared to the latter.

The modified set of coefficients used for w320/w160LP becomes:

- $vA = [0.31; 0.49; 1.32]$
- $vB = [1; 1; 1.13; 0.59]$
- $vG = [0.061]$

Due to design limitations, the vB coefficients cannot be altered back to original, so the other set for w160/w80 is also changed and becomes:

- $vA = [0.26; 0.49; 1.32]$
- $vB = [1; 1; 1.13; 0.59]$
- $vG = [0.061]$

7.2 Validation of the $\Sigma\Delta$ ADC

To validate the designed $\Sigma\Delta$ ADC and ensure all operating modes function as intended, close to the modelled range, while simultaneously verifying the anticipated thermal noise characteristics, comprehensive top-level simulations were conducted with the thermal noise contribution both enabled and disabled. However, this approach only permits a rough estimation of all noise components, as the quantization noise can be simulated with or without thermal noise, but not vice versa, necessitating estimation. The aggressiveness of the first set of coefficients was reduced due to the modifications implemented, while the second set remained essentially unchanged. As a result, the first set is expected to exhibit a slight performance degradation compared to the initial given, while the second set should yield comparable results.

The complete top-level I/Q ADC has two inputs. For validation purposes two different stimuli were used. I channel has a single sine input placed at roughly third of the BW, while the Q channel has a wide-band multitone signal, composing of 26 sub-tones spaced equally apart with a distance of third of the BW. A Wide-Band Multi Tone input gives a good approximation on the STF of the ADC, while keeping the feedback-DAC operating in multi-bit mode to get an approximate reading on the potential SNR, when operating close to target input power. A single sinewave gives a worst-case estimation on the SDR, while

also giving an approximate reading on the potential SNR, when operating close to target input power. Each simulation also incorporates a start-up verification, including VDD ramp-up, proving that no maintained oscillations occur. Both stimuli have an input target power of -10dBm to mimic true operating conditions. This is done to include true thermal contributions while not including the potential linearity impairments.

The validation results are compared to the expected performance coming from the coefficient design and summarized in table 8:

Table 8. Performance comparison of designed ideal coefficients compared to schematic validation.

Top-Level Validation					
Parameter	Unit	Value			Comment
		QT	Therm	Total	
w320_coefs	dB	84	78	77	DR of w320 mode
w320_validation	dB	83	77	76	DR of w320 mode
w160_coefs	dB	84	81	79	DR of w160 mode
w160_validation	dB	84	83	81	DR of w160 mode
w160LP_coefs	dB	78	78	75	DR of w160LP mode
w160LP_validation	dB	79	77	75	DR of w160LP mode
w80_coefs	dB	78	81	77	DR of w80 mode
w80_validation	dB	81	80	77	DR of w80 mode

The actual simulation results are presented in Figure 87 (w320), Figure 88 (w160), Figure 89 (w160LP) and Figure 90 (w80). The upper plots display the Fast Fourier Transform (FFT) of the transient simulation without thermal noise, along with the corresponding DR calculated over a sliding window. The bottom plots show the FFT of the simulation incorporating thermal noise and the associated DR over a sliding window. The FFTs plotted are always calculated from the last 8k samples available from the simulated transient results. The x-axis of the sliding window represents the starting sample number utilized for the calculation.

It is evident that for both 6-bit modes, the initial rapid DAC calibration performed was not sufficiently accurate, necessitating a full calibration cycle to be executed at the clock rate of the operational system to attain maximum performance. This translates to approximately 14k (3 μ s) clock cycles from the enabling of the $\Sigma\Delta$ ADC to achieve full performance in the w320 mode and 7k (3 μ s) clock cycles from the enabling of the $\Sigma\Delta$ ADC to achieve full performance in the w160 mode. Conversely, for both 5-bit modes, the initial rapid calibration of the DAC appears to be sufficiently accurate to reach full potential from the outset. However, the 3 μ s duration would exceed the 2 μ s handover limits defined by the Wi-Fi standard between TX/RX modes, thus necessitating adjustments to the calibration clocking speeds prior to the verification process.

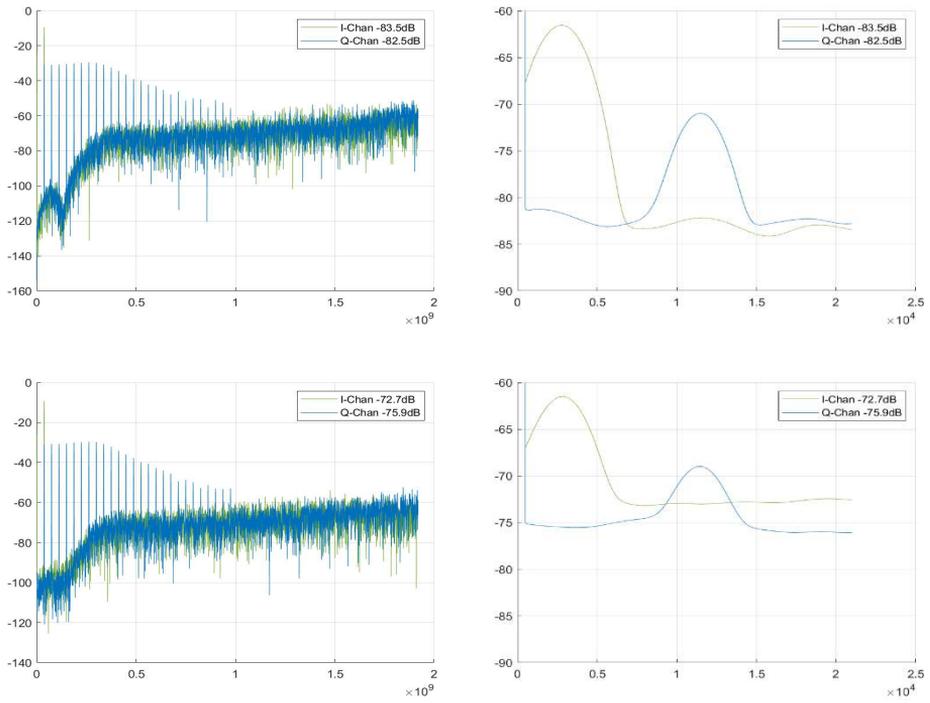


Figure 90. Validation of w320 mode. wo/w thermal noise. a) FFT; b) DR of a sliding window.

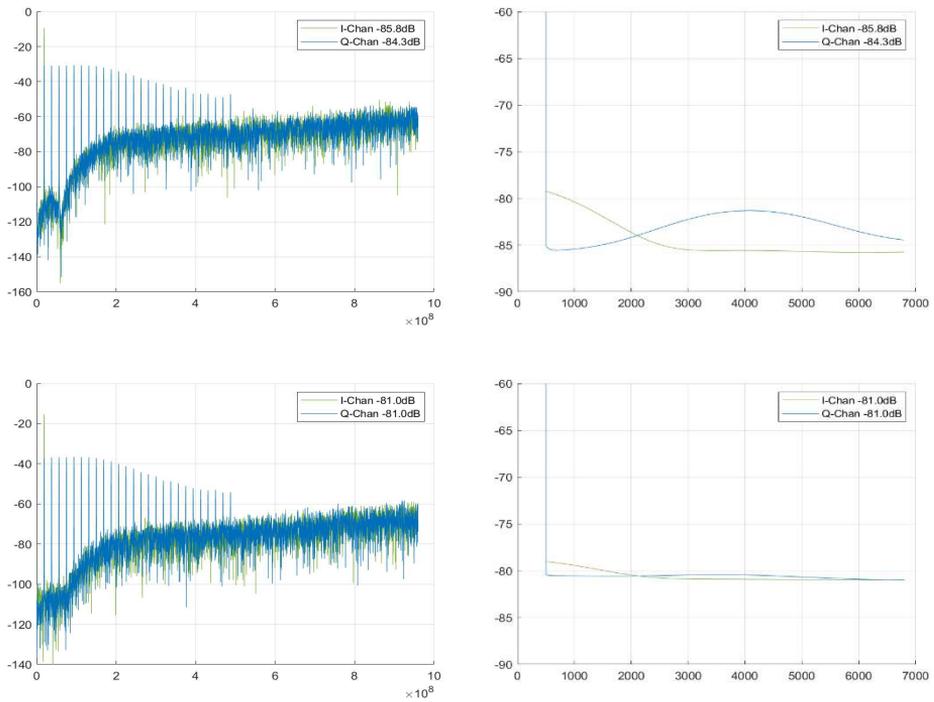


Figure 91. Validation of w160 mode. wo/w thermal noise. a) FFT; b) DR of a sliding window.

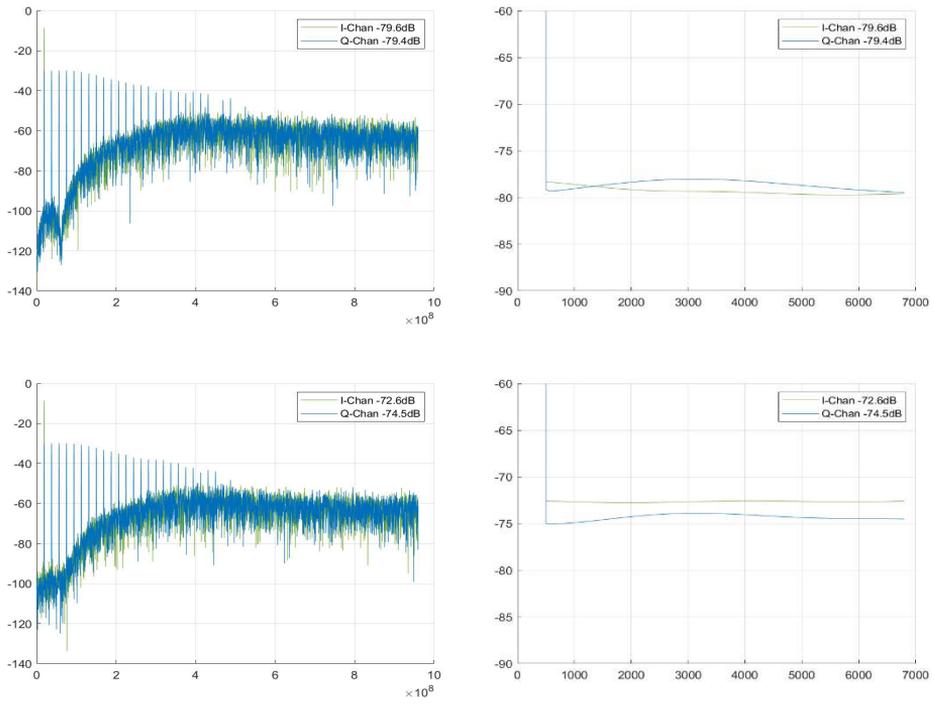


Figure 92. Validation of w160LP mode. wo/w thermal noise. a) FFT; b) DR of a sliding window.

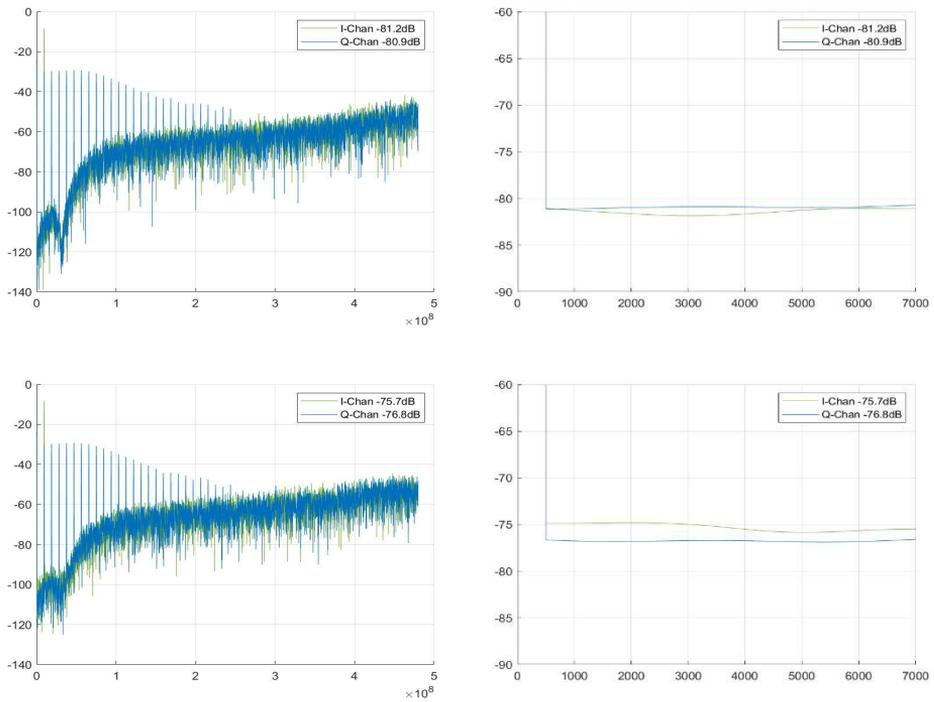


Figure 93. Validation of w80 mode. wo/w thermal noise. a) FFT; b) DR of a sliding window.

8 Layout and Verification

This chapter covers the following topics:

- Layout considerations and post-extraction verifications of the $\Sigma\Delta$ ADC.

8.1 Layout of the $\Sigma\Delta$ ADC

Similar to the opamps, the $\Sigma\Delta$ ADC has been carefully designed to minimize the lengths of critical signal paths, primarily in the digital domain, which plays a significant role in determining the ELD. To achieve maximum efficiency, a slice-based methodology has been employed (Figure 94). This approach ensures that all signal propagation paths remain straight and unobstructed, minimizing parasitic elements and leading to an optimal floor plan. Additionally, particular attention has been paid to the feedback loop between DAC3 and opa3 in the filter (Figure 95). The Filter and Feedback DAC have been carefully aligned only to minimize parasitics at this crucial node.

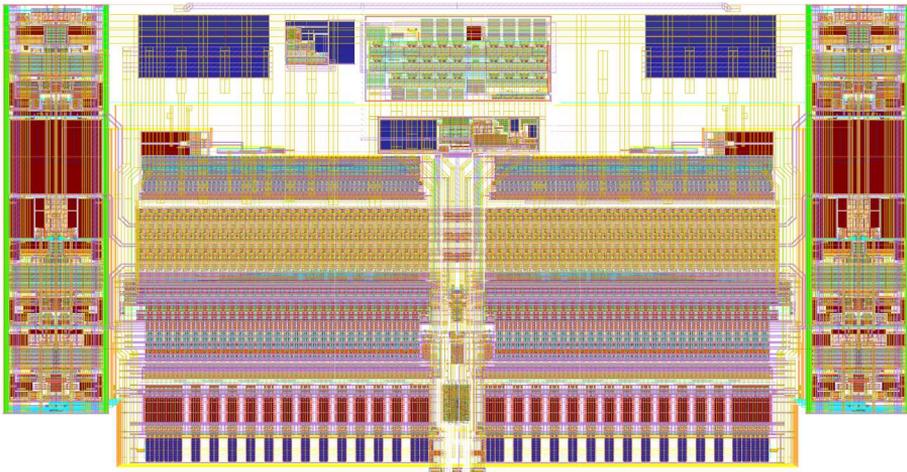


Figure 94. Layout of the designed $\Sigma\Delta$ ADC.

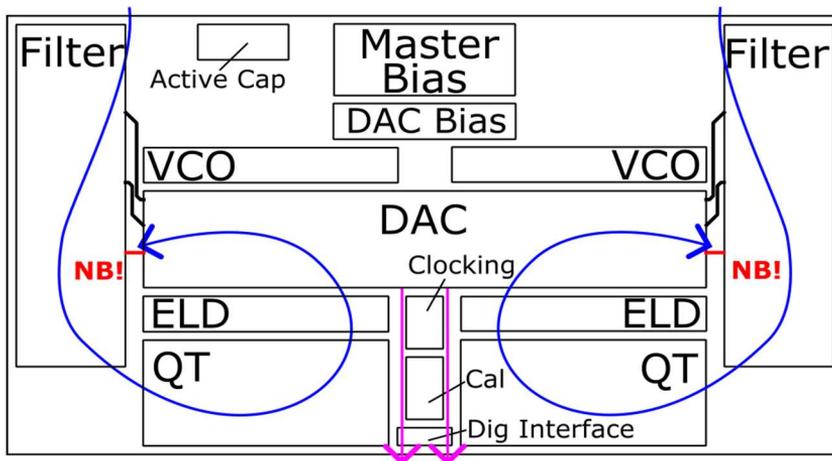


Figure 95. Floorplan of the designed $\Sigma\Delta$ ADC.

8.2 Verification Methodology

To verify the full $\Sigma\Delta$ ADC performance before taping, a series of top-level extracted simulations were performed. Since the ADC is an extraordinarily complex analog system with hundreds of thousands of components, simulation time is a major concern. Not only is the physical time considerable, but also the necessary computational resources required enormous plus corresponding license usage immense. To minimize overhead, key steps were taken to limit the simulation complexity to lowest set required. All simulations were performed before adding filling in layout and all extractions were reduced with a dedicated standalone netlist reducer to 20GHz. This does add some erroneous to the results, but also reduces the physical time needed for simulation to mere weeks, rather than months, which was still considered acceptable.

To further reduce the set of specific simulations needed, process corners used for sign-off were lowered from the usual 19 used in sub-IP verification to trivial five and modes reduced to the highest and lowest BW supported by the ADC (w320/w80). In addition, the w320 mode was simulated in DPD configuration and all nominal modes were also simulated with thermal noise included. Corners used, were Nominal, Fast_Cold, Fast_Hot, Slow_Cold and Slow_Hot, where Cold = -40° and Hot = 125° . This brought the number of simulations down to a manageable 21.

List of Top-level simulations done:

- 1) w320 nominal
- 2) w320 Fast -40°
- 3) w320 Fast 125°
- 4) w320 Slow -40°
- 5) w320 Slow 125°
- 6) w80 nominal
- 7) w80 Fast -40°
- 8) w80 Fast 125°
- 9) w80 Slow -40°
- 10) w80 Slow 125°
- 11) w160 nominal
- 12) w160LP nominal
- 13) w320 nominal (w transient noise)
- 14) w160 nominal (w transient noise)
- 15) w160LP nominal (w transient noise)
- 16) w80 nominal (w transient noise)
- 17) w320DPD nominal
- 18) w320DPD Fast -40°
- 19) w320DPD Fast 125°
- 20) w320DPD Slow -40°
- 21) w320DPD Slow 125°

8.3 Verification of the $\Sigma\Delta$ ADC

Since the ADC has two channels (I/Q), a similar strategy as in performance validation with two sets of different input signals were used for performance verification. To achieve high coverage with minimum resources, the most difficult or meaningful scenarios need to be considered. For that reason, a single sinewave input @ approximately 1/3 of the channel BW was used for the I-channel and a Wide-Band Multi Tone signal with the tones placed

equally apart with the distance of approximately 1/3 of the channel BW was used for the Q-channel. A Wide-Band Multi Tone input gives a good approximation on the STF of the ADC, while keeping the feedback-DAC operating in multi-bit mode to get an approximate reading on the potential SNR, when operating close to target input power. A single sinewave gives a worst-case estimation on the SDR, while also giving an approximate reading on the potential SNR, when operating close to target input power. Each simulation also incorporates a start-up verification, including VDD ramp-up, proving that no maintained oscillations occur.

The actual simulated performance is compared to both estimated performance coming from the coefficient design and from the performance assumed from schematic validation in the following table:

Table 9. Performance comparison of designed ideal coefficients compared to extracted simulations.

Top-Level Extraction					
Parameter	Unit	Value			Comment
		QT	Therm	Total	
w320_coefs	dB	84	78	77	DR of w320 mode
w320_validation	dB	83	77	76	DR of w320 mode
w320_extraction	dB	80	77	76	DR of w320 mode
w160_coefs	dB	84	81	79	DR of w160 mode
w160_validation	dB	84	83	81	DR of w160 mode
w160_extraction	dB	83	81	78	DR of w160 mode
w160LP_coefs	dB	78	78	75	DR of w160LP mode
w160LP_validation	dB	79	77	75	DR of w160 mode
w160LP_extraction	dB	78	78	75	DR of w160LP mode
w80_coefs	dB	78	81	77	DR of w80 mode
w80_validation	dB	81	80	77	DR of w160 mode
w80_extraction	dB	79	79	76	DR of w80 mode

In general, performance degradation from the actual perfect world coefficient design to the full-blown extraction is only a very manageable 1dB, which could be considered a very good alignment.

The actual simulation results are plotted in Figure 96 (w320), Figure 97 (w160), Figure 98 (w80) and Figure 99 (w320DPD). The upper row plots show the FFT of the full extraction transient simulation wo thermal noise and the corresponding DR over a sliding window and the bottom row plots show the FFT of the full extraction with thermal noise and the corresponding DR over a sliding window. The FFTs plotted are always calculated from the last 8k samples available from the simulated transient results. The x-axis of the sliding window represents the starting sample number utilized for the calculation.

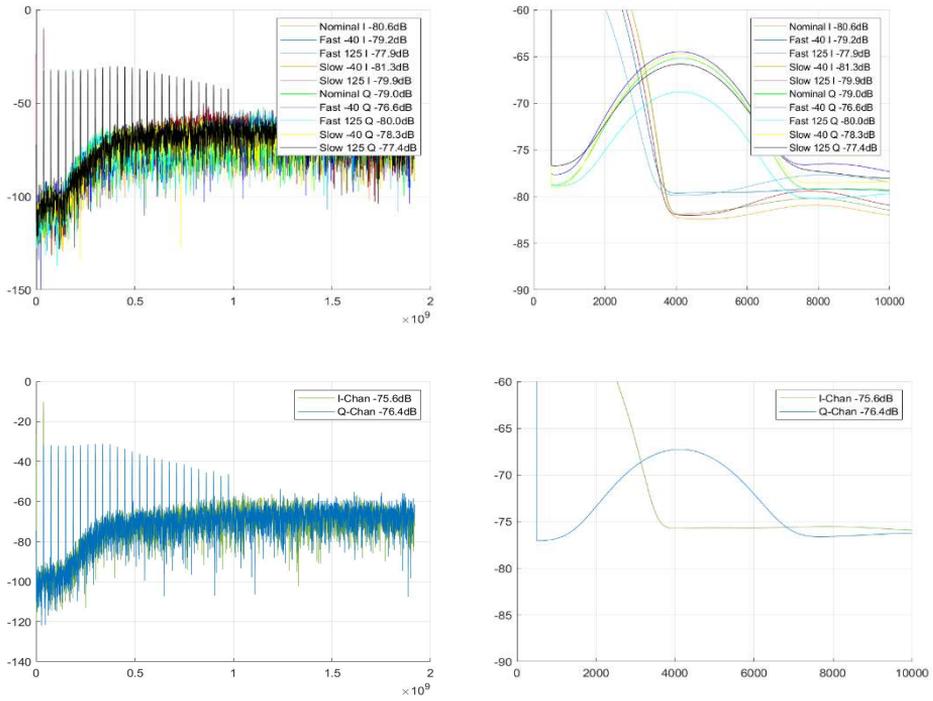


Figure 96. PEX simulations of w320 mode. wo/w thermal noise. a) FFT; b) DR of a sliding window.

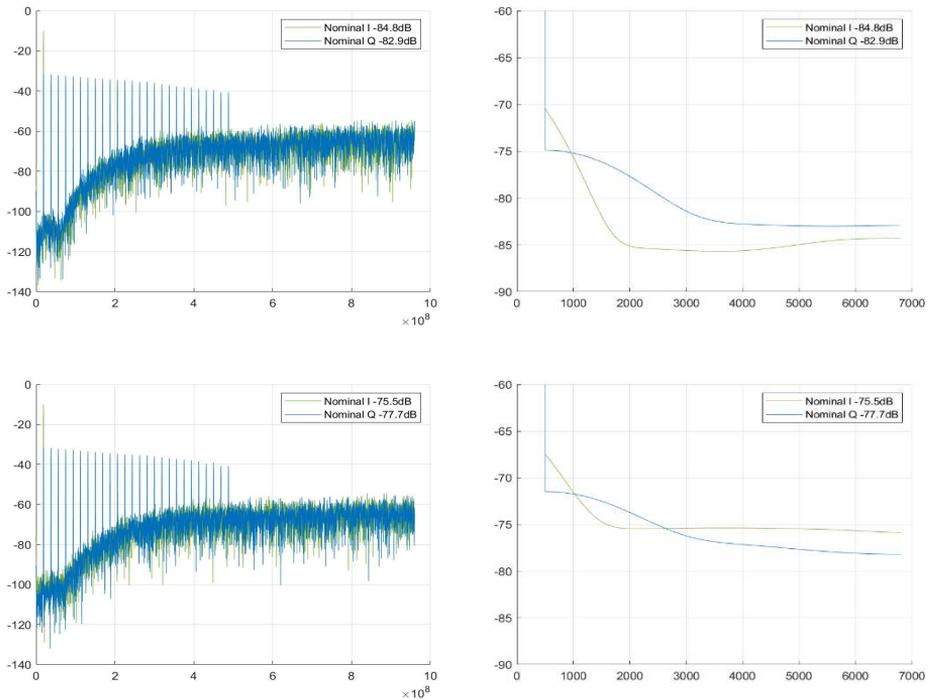


Figure 97. PEX simulations of w160 mode. wo/w thermal noise. a) FFT; b) DR of a sliding window.

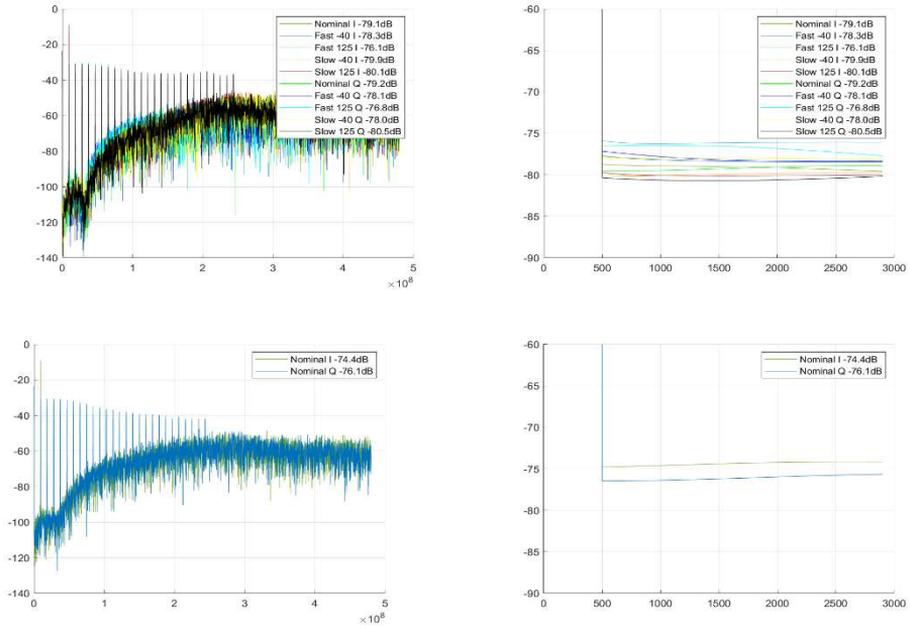


Figure 98. Post-Layout simulations of w80 mode. wo/w thermal noise. a) FFT; b) DR of a sliding window.

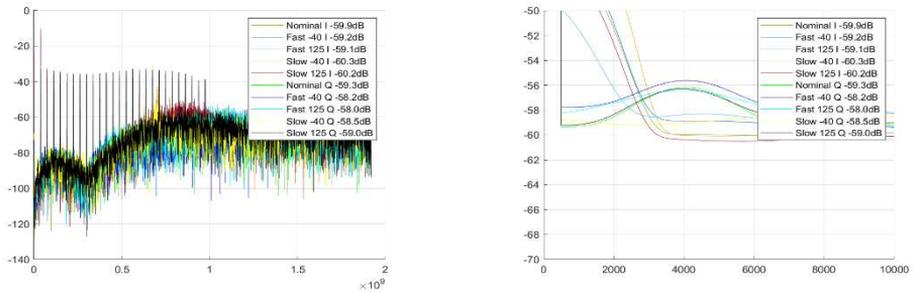


Figure 99. Post-Layout simulations of w320DPD mode. wo thermal noise. a) FFT; b) DR of a sliding window.

9 Testing and Measurements of the ADC

This chapter covers the following topics:

- Testing of the $\Sigma\Delta$ ADC in standalone configuration and within the RX half chain.

To comprehensively characterize the designed ADC on silicon, two distinct measurement sets were performed:

- Standalone SDADC Measurements: These measurements accessed the ADC directly through the loopback input, as detailed in sub-chapter 6.4.3.
- RX Half-Channel Measurements: This set evaluated the combined performance of the $\Sigma\Delta$ ADC together with the RX half-chain, including the LNA, Mixer, and PGC.

9.1. Measurements of the Standalone $\Sigma\Delta$ ADC Module

While the standalone $\Sigma\Delta$ ADC is physically accessible only through a multi-purpose input/output test-port, its characterization is subject to limitations due to shared connections with other on-chip modules, including a test buffer derived from the PGC described in chapter 6.18. This means, the capacitive load seen from the PCB and the testing equipment is not negligible and the impedance of the test-port becomes a factor and must be partly equalized with proper channel equalization. The equalization process is inherently estimative, involving subtle adjustments to the generated MTPR signal to achieve a flat STF at the $\Sigma\Delta$ ADC output. This approach relies on the assumption of an inherently flat in-band STF of the $\Sigma\Delta$ ADC, attributing any deviations to the test-port itself. These factors collectively constrain the test-port's effectiveness, potentially introducing linearity limitations, restricting maximum input power, and ultimately reducing achievable performance compared to ideal testing conditions. Consequently, it is reasonable to infer that the $\Sigma\Delta$ ADC's actual performance likely exceeds measured values by a few dBs.

To evaluate the $\Sigma\Delta$ ADC's performance under realistic conditions, MTPR patterns with 12dB CF were applied as inputs to both channels. This generates an SNDR result that incorporates:

- **SDR**: defined in the pattern as the RMS average MTPR value between a missing tone and the preceding present tone.
- **SNR**: representing the noise of all bins present due to oversampling.

The combined SNDR value closely correlates with complex domain EVM value, with a typical complex to I/Q domain conversion difference of 3dB. For each measurement, the repeating MTPR pattern is generated with five different seeds, and the captured directly at the ADC output to a RAM capable of holding 393216 points, five different times (Figure 100). This process is repeated for 20 different fixed missing tone locations. All the captured SNDR results (100 together for each input power) are afterwards RMS averaged. Finally, the input power of the signal is swept across a range of values. This process was repeated for 5 different chips, with the results averaged.

Equipment used for measurements:

- Input signal generation: N8241A AWG (Keysight, USA)
- Synchronization clock generation: HP 8644B Signal Generator (Keysight, USA)

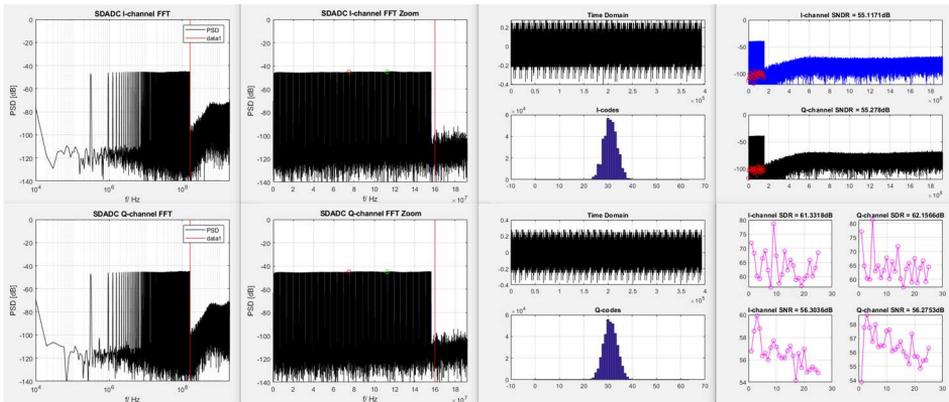


Figure 100. An example of a single I/Q SDADC w320 stand-alone measurement with $\sim -12\text{dBm}$ @ ADC input power having 393216 6-bit output captures.

The performance vs input power for different primary modes has been plotted in the following graph (Figure 101). Higher BW modes have limitations in measurements, because of the issues described previously. Due to equalization and input matching, the AWG is not able to generate input stimuli with sufficient power, limiting the w320 mode measurements to maximum -10dBm and w160 mode measurements up to -7dBm . The 5-bit modes can be measured up to full scale, with the test-port linearity limitation still showing impact. Estimated results plotted as a dotted continuation are based on the interpolation of simulations and RX chain measurements.

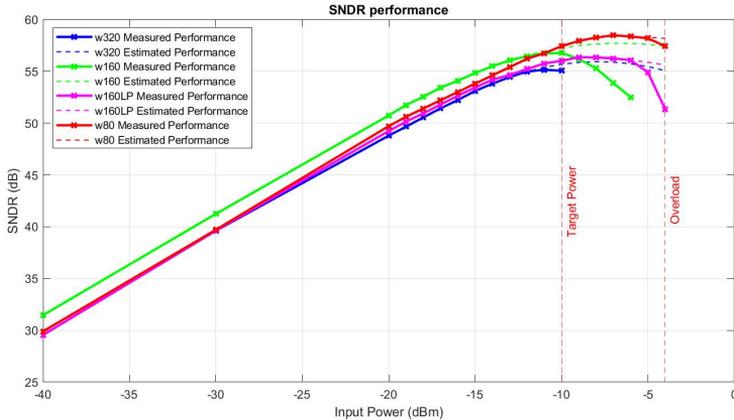


Figure 101. Input power vs performance sweep of the primary $\Sigma\Delta\text{ADC}$ modes.

9.2. Measurements of the RX Chain

Measuring the RX chain utilizes a complex MTPR input stimulus. This stimulus, comprising both in-phase and quadrature components, is applied to the RF port of the RX half-channel. The RX half-channel encompasses all modules within the signal path, including the LNA, Mixer, PGC together with the ADC.

Equipment used for measurements:

- Input signal generation: Rohde & Schwarz SMW200A Vector Signal Generator.
- Synchronization clock generation: HP 8644B Signal Generator.

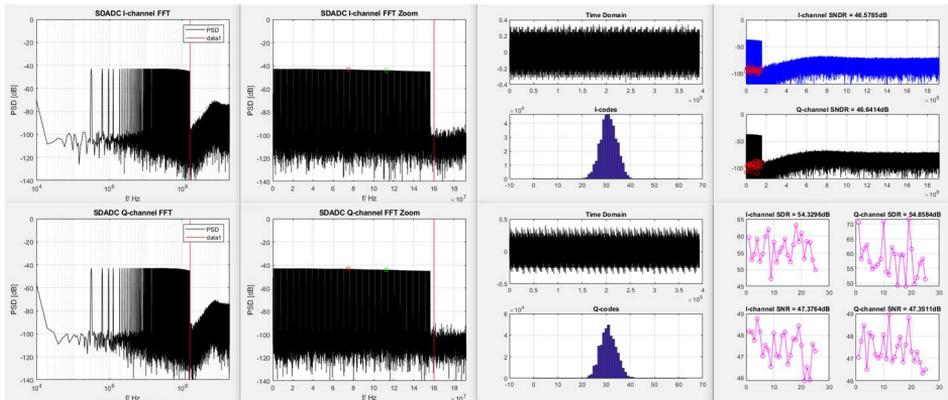


Figure 102. An example of a single I/Q SDADC w320 RX half-channel measurement with $\sim -12\text{dBm}$ @ ADC input power having 393216 6-bit output captures.

Characterization of the RX half-chain involved four distinct LNA and PGC gain configurations. Optimal design practice dictates early gain application in the chain to minimize overall NF. The AGC algorithm typically prioritizes LNA gain before PGC, leveraging the LNA's upstream position to amplify signals before significant noise accumulates. As a rule of thumb, PGC gain is used only sensitivity scenarios where extended range is required and simpler coding schemes can be employed.

Ideally, the ΣADC should not be the primary contributor to overall performance. When the PGC is at minimum gain, the ΣADC best case NF lies roughly 7-8 dB below the chain's total (Figure 103). This increases to 14 dB below at maximum PGC gain across the input power sweep. Notably, at the target power of -10dBm at the ADC input, where other distortions like PLL phase noise become dominant, the ΣADC 's NF falls even further below the total.

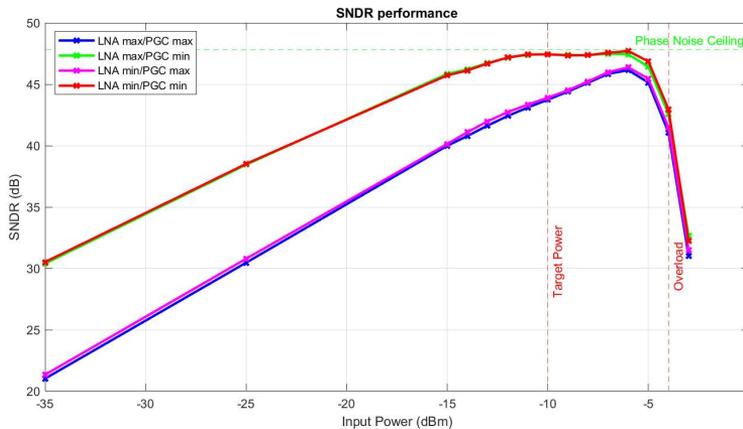


Figure 103. Input power vs performance sweep of the RX chain in w320 mode with different PGC/LNA gain configurations.

It is important to note that the ADC's intrinsic performance remains independent of LNA/PGC gain settings. These settings solely influence the ADC's relative contribution to the overall chain NF. The closest it lies to the overall performance is when other modules upstream perform the best. For that reason, other modules have been configured to work

under optimal conditions for measurements. For example, during the w320 measurements, the selected Wi-Fi channel was configured to the lowest LO frequency allowed (Channel 31 – 6.105GHz), minimizing expected phase noise contributions. Configuring the Wi-Fi channel to Channel 191, will lower the maximum ceiling, thus reducing the overall contribution of the ADC (Figure 104).

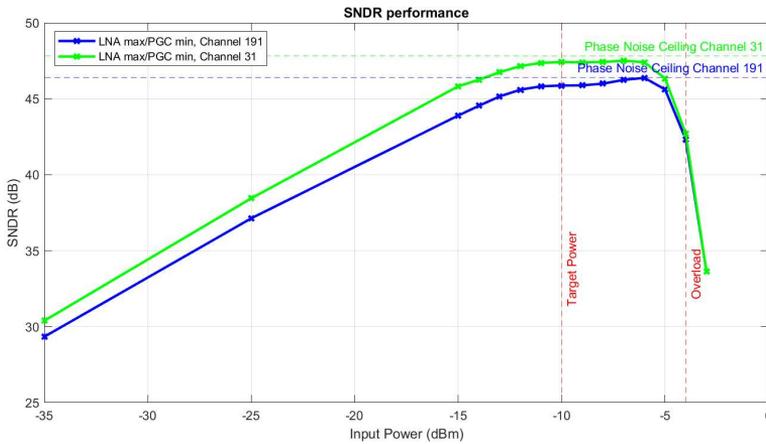


Figure 104. Input power vs performance sweep of the RX half-chain in w320 mode with different Wi-Fi channels used.

In conclusion of the testing, all features demonstrated successful operation and were fully functional. The performance of the ADC was on par or sometimes even exceeded expectations, enabling true Wi-Fi 7 experience. With its unparalleled performance, this Wi-Fi 7 chip is poised to revolutionize wireless connectivity, empowering an array of emerging technologies, including low-latency extended reality (XR), cloud-based gaming, and 8K video streaming.

9.3. Further Nuggets from Measurements

The designed $\Sigma\Delta$ ADC has many more modes that can be enabled. Some of them, used for DPD calibration were also briefly discussed in chapters dealing with coefficient design. While verifying the functionality of those modes is simple enough, accurate measurements with OFDM signals is more difficult, as signal generator BW and the forementioned test-port becomes a concern. Nevertheless, different modes, built in the ADC were verified.

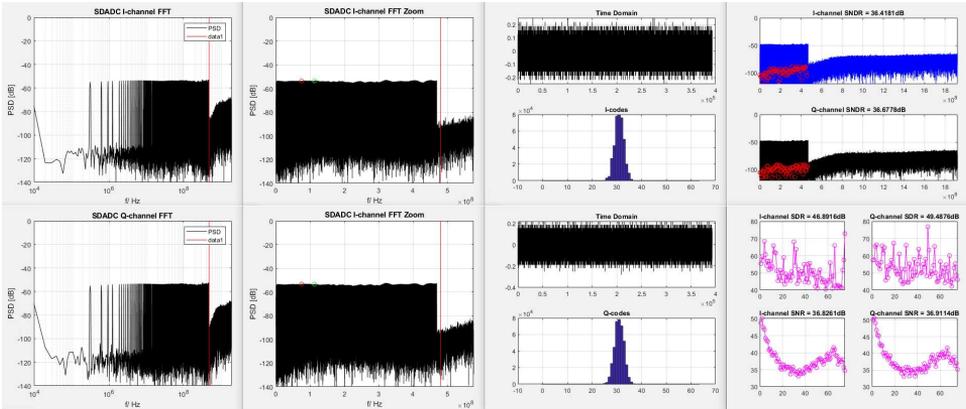


Figure 105. An example of a single I/Q SDADC w320DPD stand-alone measurement with ~ 18 dBm @ ADC input power having 393216 6-bit output captures.

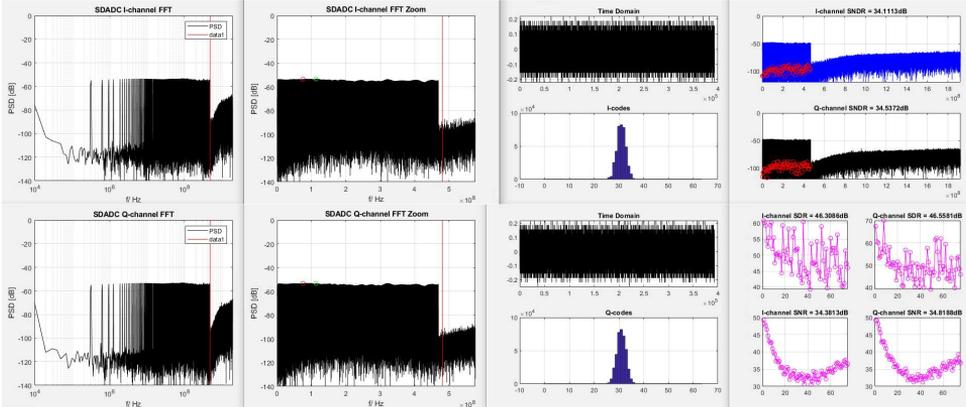


Figure 106. An example of a single I/Q SDADC configured with 550MHz BW stand-alone measurement with ~ 18 dBm @ ADC input power having 393216 6-bit output captures.

10 Conclusion

10.1 Summary

The PhD thesis provides a comprehensive summary of the design process for a Wi-Fi 7-compatible Analog-to-Digital Converter. Prior to delving into the design specifics, the first five chapters establish a basic foundation by defining key concepts, presenting a brief history of the Wi-Fi standard and explaining the rationale behind critical design decisions. The introduction also offers an overview of the RX half chain fundamentals, comparing different types of converters, their potential applications, and their respective strengths and weaknesses.

Chapter 6 was dedicated to the detailed design overview of the CT $\Sigma\Delta$ ADC. The initial sections of this chapter focus on system-level considerations, including the specification draft and the design of coefficients to meet the established requirements. Sub-chapter 6.3 introduces the 14 sub-modules that comprise the ADC, which are then thoroughly examined in sub-chapters 6.4 to 6.16. The analysis begins with the three core components of any $\Sigma\Delta$ ADC: the Loop Filter, Quantizer, and Digital-to-Analog Converter (DAC), followed by an exploration of the auxiliary sub-modules.

Additional sub-chapters are devoted to other critical elements of the RX half chain, such as the power supply and Programmable Gain Amplifier. These components significantly influence the overall performance of the ADC, and their inclusion in dedicated sub-chapters underscores their importance within the design process.

Chapter 7 did focus on the optimization and validation of the design. As is often the case with assembling complex systems like this high-speed CT $\Sigma\Delta$ ADC, it becomes apparent that initial design choices may not fully align with real-world performance, necessitating adjustments. For $\Sigma\Delta$ ADCs, this typically involves fine-tuning the coefficients derived from the initial design process. This chapter details the specific modifications made, quantifies the extent of these adjustments, and compares the outcomes of ideal system simulations using the original coefficients with top-level schematic simulations incorporating the refined coefficients.

Chapter 8 did extend the validation process discussed in the previous chapter by incorporating layout considerations and post-layout verification. Additionally, it addresses the increasing resource challenges faced by designers when developing complex systems that require comprehensive verification prior to fabrication. The chapter concludes by presenting the final verification results, alongside initial system estimates and simulated validation outcomes.

Chapter 9 did present the laboratory measurements conducted on the final manufactured and packaged die, dividing the analysis into two sections: one focused on the standalone CT $\Sigma\Delta$ ADC and the other on the complete RX half chain. The results clearly demonstrate that the ADC's performance aligns with the original design objectives, with the ADC contributing minimally to the overall NF, as predicted by Equation 3. While certain standalone measurements were limited by equipment constraints and auxiliary measurement path limitations, all critical tests necessary to validate the quality and maturity of the silicon were successfully completed.

The final ADC achieved a dynamic range of approximately 77dB, corresponding to an effective number of bits of around 12.5, an SNDR of -55dB @ -12dBm (measurements with -10dBm turned out impossible to execute due to measurement equipment limitations) and consumed approximately 80mA of current for I/Q from the OV9 rail in

W320 mode. Current consumption was approximately 80 mA of current for the I/Q channels from the OV9 rail in W320 mode. In W160 mode, the current consumption was measured at approximately 50mA, and in W80 mode, it was approximately 35mA.

Projecting the final results onto the Schreier Figure of Merit (FOM) graph (Figure 107) will illustrate the efficiency of the ADC in comparison to state-of-the-art designs. While this graph is a widely recognized tool and provides valuable insights, it can be somewhat misleading. The graph only reflects converter efficiency, without conveying information about maximum performance, making it difficult to discern whether the converter achieves high performance or low power consumption for a given bandwidth (BW).

Most converters presented in the graph that operate in the 100MHz range are based on non-interleaved SAR architectures, lacking an input buffer, resulting in impressive power consumption numbers. Although their FOM appears impressive, they do not achieve the dynamic range (DR) required for our application. Generally, SAR ADCs offer greater efficiency for lower performance requirements, but they struggle to deliver when a higher number of effective number of bits (ENOBs) is needed.

The designed converter demonstrates peak efficiency around 100MHz, which is expected given that the design was specifically tailored for optimal performance in the W320 mode, requiring a BW of 160MHz. The coefficients were custom-tuned exclusively for this mode, with the rest coming as a byproduct from power scaling. The graph further illustrates that efficiency declines sharply outside this BW range. On one hand, pushing $\Sigma\Delta$ ADC architectures to handle higher frequencies is inherently challenging due to clock speed limitations that restrict the OSR. On the other hand, optimizing power consumption for a single mode inevitably compromises efficiency in other modes, where additional programming introduces unnecessary parasitics. Designing an ADC solely for one mode would undoubtedly have resulted in a higher FOM.

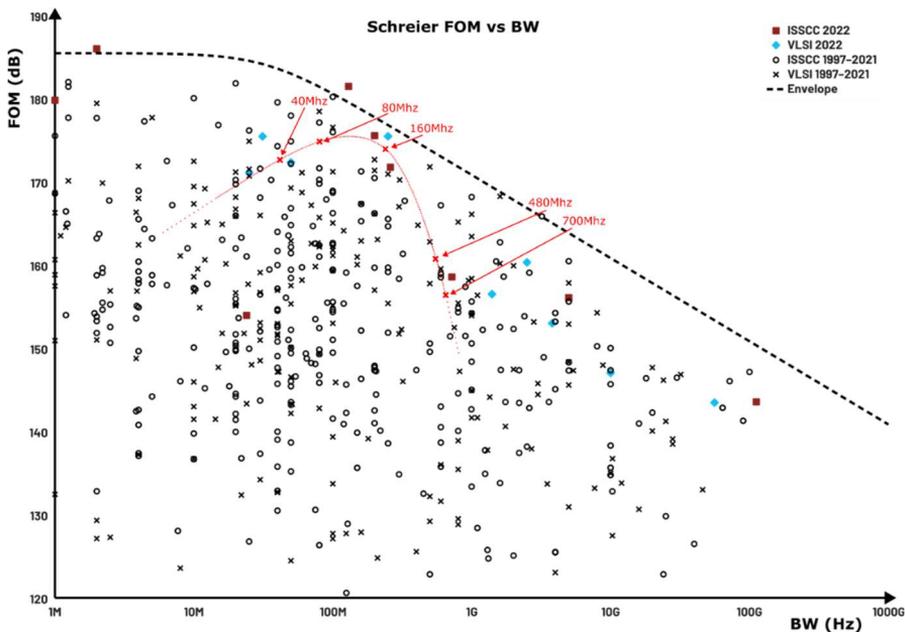


Figure 107. Final results projected onto the Schreier FOM graph.

10.2 Future Research and Wi-Fi 8

Although this work is grounded in the Wi-Fi 7 standard, the continuous evolution of technologies and standards is an ongoing process. Given the substantial costs associated with mask and silicon manufacturing for advanced technologies, the ADC team has strategically incorporated features that, while not essential for the current standard, will be valuable for future generations. While details about the Wi-Fi 8 standard remain scarce, there is speculation that it will not double the channel bandwidth like in previous amendments, instead supporting channels only up to 480MHz. This approach is based on the rationale that accommodating multiple 640MHz channels within the 6 GHz band would be impossible. However, by capping the maximum channel width at 480MHz, the possibility of fitting two separate channels becomes again feasible.

With this in mind, the ADC also supports a 1.5x bandwidth mode linked to each of the four primary modes, expanding the officially supported configurations to 12 (Figure 108). These include four primary modes, four DPD modes, and four 1.5x modes, which may prove to be potentially useful for Wi-Fi 8 exploration. The actual implementation mirrors the approach used for DPD modes, where the resonator notch is shifted, though to a lesser extent, to a more optimal frequency. While this method further reduces the OSR, bringing the quantization noise closer to thermal noise and thus increasing its influence on overall noise contribution, it offers a unique opportunity for initial practical tests and observations at the silicon level – a potentially valuable attribute for future advancements.

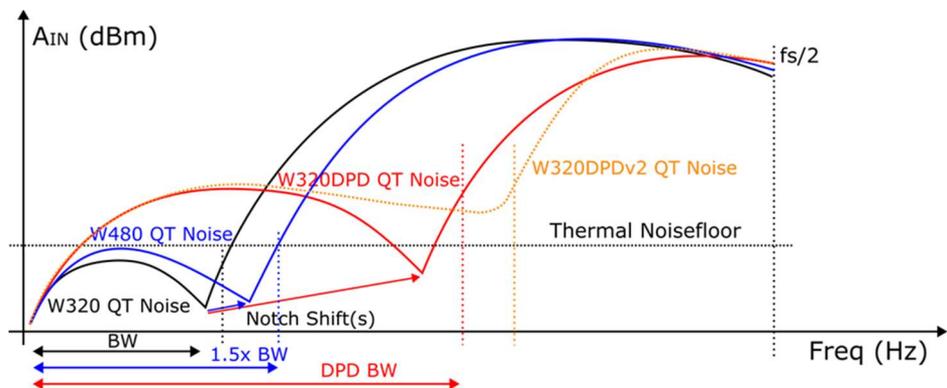


Figure 108. Additional modes available for the ADC.

Although the performance in this mode does not quite match that of the primary mode, it should be at least equivalent to the best SAR-based ADCs. Additionally, the rest of the RX chain is already compatible with higher bandwidth channels, as the RF modules do not include filtering, and the PGC has a special DPD mode that disables filtering for even higher bandwidths. While the Wi-Fi 8 standard and the potential implementation of 480 MHz channels and modulation schemes within it are still unconfirmed and speculative, this feature may ultimately prove to be less critical. However, it could also become a valuable addition, potentially providing backward compatibility once the new standard is released.

10.3 Experimental Features

As detailed in sub-chapter 6.7, the biasing circuitry contains six concealed 3-bit memory cells, commonly referred to as “hidden bits”. These experimental features are not intended for public or customer use, and access to them is highly restricted. To preserve their confidentiality, they can only be accessed through specific backdoor overrides provided by the Calibration Machine. Of the six cells, only four are currently in use, with the remaining two reserved as backups. Some of these features were also briefly discussed in the chapter addressing the $\Sigma\Delta$ ADC design. In the last sub-chapter, I would like to use this section to briefly go over the practicality and usage of them.

The four memory cells used could do the following:

- Cut the coefficients of the $\Sigma\Delta$ ADC to 75%. (*by reducing the integrator capacitors to 75% of the original value*)
- Modulate R_2 , R_3 and R_{RESO} over frequency. (*by adding a parallel capacitive path*)
- Change DAC biasing noise settings.
- Introduce clocking delay to the redundant cell in the DAC calibration.

While the DAC biasing settings option is of limited utility in this context, the other three options have proven to be quite valuable. In the W480 mode, shifting the notch 1.5x further as described in sub-chapter 10.2 introduces more in-band noise, as illustrated in Figure 108. To mitigate this, more shaping must be applied, effectively using more aggressive coefficients to move in-band noise OoB and reducing the quantization noise further. The 75% coefficients option is a feature specifically designed together with the use with the 1.5x BW modes. In this way 1.5x BW modes look much more similar to the primary modes from the performance point of view.

Modulating R_2 , R_3 or R_{RESO} all serve a similar purpose to extend the BW of the converter further, especially in the DPD mode. Utilizing standard resistance as a resonator typically results in a pronounced deep notch at a specific frequency. This often leads to a scenario where part of the QT noise within the frequency spectrum lies above the thermal noise floor while another part falls below it (Figure 108: Red). Consequently, the in-band noise is unevenly distributed, being heavily coloured rather than uniformly shaped. This is suboptimal because the signal sub-carriers are dispersed across the entire bandwidth, with the SNR constrained by the region where the QT noise is most prominent. By modulating the resistance over frequency, specifically by reducing impedance at higher frequencies, the QT noise can be shaped more favourably, remaining relatively flat across the in-band frequency spectrum except at DC and very low frequencies (Figure 108: Orange). This approach can lead to a potentially higher signal bandwidth, which was previously constrained by less optimal QT noise shaping. Modulating R_2 and R_3 produces a similar effect. When these techniques are combined, the bandwidth of the designed $\Sigma\Delta$ ADC in DPD mode can be extended from 480MHz to up to 700MHz.

The final hidden feature involves introducing a specific clocking delay to a particular cell during DAC calibration. The critical challenge in the $\Sigma\Delta$ ADC design, around which the entire architecture has been developed, is maintaining the ELD within the converter at a reasonably low level, approximately 0.5, to ensure high production yields. This constraint limits clocking speeds to a modest 3.84GHz, ensuring that signal propagation from the QT latch to the slowest DAC cell, known as the redundant cell (Figure 109: LogicXR), remains below 130ps. Exceeding this threshold introduces metastability into the system, manifesting as increased in-band noise. This issue only arises when the original current

cell is under calibration and the signal needs to be rerouted. Signals not currently being rerouted do not experience propagation problems – only the cell that is rerouted to the redundant cell is affected. This concept allows for the opportunity to delay the clocking of the redundant cell by a similar amount as the signal delay incurred during rerouting. In a 6-bit DAC, this would mean that 29 current cells are latched with the original clock, while one is latched with a delayed clock. This approach could mitigate the propagation speed challenges inherent in the DAC topology, potentially paving the way for further increases in clock speeds.

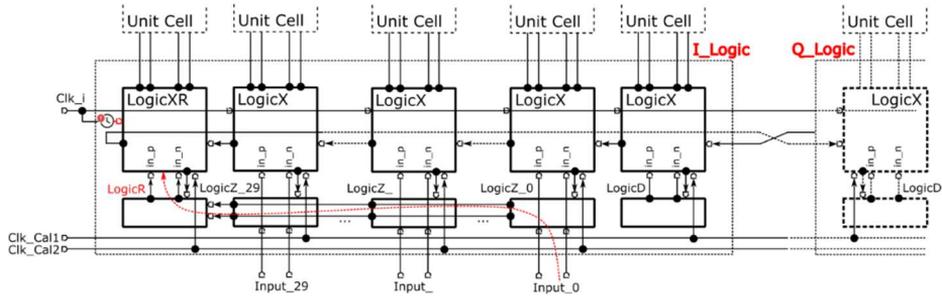


Figure 109. DAC Logic.

All experimental test features were integrated during the design phase but were neither simulated nor verified. This was primarily due to the limited simulation resources available, as accessing the hidden bits via the Calibration Machine and conducting extensive simulations would have required considerable time. Despite this, all features demonstrated successful operation and were fully functional during testing and verification.

By utilizing the R_2 , R_3 and R_{RESO} modulation schemes in conjunction with the 75% coefficient option, the $\Sigma\Delta$ ADC was able to achieve a maximum bandwidth of 700MHz. While the performance and efficiency at this bandwidth no longer match those attained by other ADC topologies, the mode shows great promise for niche applications such as the DPD.

During measurements, clocking speeds of 5.12GHz could be easily achieved using clock delay techniques. The $\Sigma\Delta$ ADC maintained its noise floor, indicating that metastability was avoided. However, a signal leakage phenomenon was observed, which degraded performance with high input signals. Despite this, the results highlight the potential of this approach for future exploration in pushing clocking speeds even further.

10.4 Final Words

This thesis concludes with a comprehensive overview of the design process for constructing a Wi-Fi 7 compatible Continuous-Time $\Sigma\Delta$ ADC (CT $\Sigma\Delta$ ADC). The work includes a concise historical context, an explanation of the design considerations, and a detailed account of the step-by-step design decisions. This is followed by the results obtained from validation, verification, and measurements. Additionally, the thesis presents some forward-looking experimental results. The designed ADC represents a significant advancement, potentially serving as a gateway for many in further exploration and innovation in the realm of high-speed Wi-Fi communications enabled by Wi-Fi 7.

References

- [1] IEEE Approved Draft Standard for Information technology--Telecommunications and information exchange between systems Local and metropolitan area networks--Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment: Enhancements for Extremely High Throughput (EHT), IEEE P802.11be/D7.0, pp.1-1089, Oct. 2024
- [2] D. Cailian, X. Fang, X. Han, X. Wang, L. Yan, R. He, Y. Long, and Y. Guo. "IEEE 802.11be Wi-Fi 7: New challenges and opportunities." IEEE Communications Surveys & Tutorials 22, no. 4 (2020): 2136-2166
- [3] S. Chauhan, A. Sharma, S. Pandey, K. N. Rao, P. Kumar "IEEE 802.11 be: A review on Wi-Fi 7 use cases." 2021 9th International Conference on Reliability, Infocom Technologies and Optimization (ICRITO), pp. 1-7. IEEE, Sep. 2021
- [4] G. Naik, D. Ogbe, J. M. Park "Can Wi-Fi 7 Support Real-Time Applications? On the Impact of Multi Link Aggregation on Latency." IEEE International Conference on Communications (ICC), Jul. 2021
- [5] N. Abramson, "Development of the ALOHNET", IEEE Transactions on Information Theory, Mar. 1985
- [6] Y. Asai, "Advanced Progress in IEEE 802.11 WLAN Standardization", IEEE 014 Asia-Pacific Microwave Conference, Nov. 2014
- [7] IEEE Standard for Information Technology - Telecommunications and information exchange between systems - Local and Metropolitan networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Higher Speed Physical Layer (PHY) Extension in the 2.4 GHz band, IEEE Std 802.11b-1999, pp.1-96, 20 Jan. 2000
- [8] ISO/IEC/IEEE 8802-11:2005/AMD4 [IEEE Std 802.11g-2003] Information technology--Local and metropolitan area networks-- Part 11: Wireless LAN Medium Access Control (Mac) and Physical Layer (PHY) Specifications--Amendment 4: Further Higher Data Rate Extension in the 2.4 GHz Band, ISO/IEC 8802-11:2005/Amd.4:2006(E) IEEE Std 802.11g-2003 (Amendment to IEEE Std 802.11-1999), pp.1-83, 15 Aug. 2006
- [9] Y. Inoue et. al., "The better spectrum utilization for the future WLAN standardization", doc.: IEEE 802.11/11-1464r2, Nov. 2011
- [10] M. J. M. Pelgrom, "Analog-to-Digital Conversion", Springer, New York, 2012
- [11] R. Schreier, G. C. Temis, "Understanding Delta-Sigma Data Converters", John Wiley & Sons, New Jersey 2005
- [12] B. Razavi, "Jitter-Power Trade-Offs in PLLs", IEEE Transactions on circuits and systems, Regular Papers, Vol. 68, No. 4, Apr. 2021
- [13] N. Yoshino, K. Norishima, M. Motoyoshi, S. Kameda, N. Suematsu, "A 28 GHz-Band Direct RF Undersampling S/H CMOS IC with 40 dB SNR" IEEE International Symposium on Radio-Frequency Integration Technology (RFIT), Aug. 2018
- [14] Maxlinear Inc, WAV600 Series-2 Wi-Fi 6 Chipset: 6th Generation Wi-Fi Chipset for Wi-Fi 6 and 6E Based on 802.11ax, Jan. 2022, accessed 05 Jan 2025, URL: <https://www.maxlinear.com/document?id=23302>

- [15] B. Razavi, "The End Is Near: The Problem of PLL Power Consumption", Youtube, uploaded by IEEE Solid-State Circuits Society, Sep. 2021
- [16] C. B. Brahm, "Feedback integrated system", U.S. Patent 3,192,371, Sept. 1961
- [17] H. Inose, Y. Yasuda, and J. Murakami, "A telemetering system by code modulation— Δ - Σ modulation", IEEE Trans. Space Electron. Telemetry, vol. 3, pp. 204–209, Sept. 1962
- [18] H. Inose and Y. Yasuda, "A Unity Bit Coding Method by Negative Feedback", IEEE Proceedings, Vol. 51 pp. 1524-1535, Nov. 1963
- [19] R. Adams, "Design and implementation of an audio 18-bit analog-to-digital converter using oversampling techniques", Journal of The Audio Engineering Society, pp 153-166, Mar. 1986
- [20] J. A. Cherry, W. M. Snelgrove, "Excess Loop Delay in Continuous-Time Delta-Sigma Modulators", IEEE Trans. Circ. Syst. II, pp 376-389, Apr. 1999
- [21] H. T. Jensen, I. Galton, "A reduced-complexity mismatch-shaping DAC for delta-sigma data converters", IEEE International Symposium on Circuits and Systems, volume1, pp 504-507, May 1998
- [22] S. R. Norsworthy, R. Schreier, G. C. Temes, "Delta-Sigma Data Converters. Theory, design and simulation", IEEE Press, Oct. 1996
- [23] Institute of Microelectronics, Uni Ulm Sigma-Delta Synthesis Environment, Universität Ulm, accessed 05 Jan 2025, URL: <https://www.sigma-delta.de/>
- [24] J. A. Cherry, W. M. Snelgrove, "Continuous-Time Delta-Sigma Modulators for High Speed A/D Conversion", Kluwer Academic Publishers, Boston 2000
- [25] F. Silveira, D. Flandre, P.G.A Jaspers, "A gm/ID based methodology for the design of CMOS analog circuits and its application to the synthesis of a Silicon-on-Insulator micropower OTA", IEEE Journal of Solid State Circuits. Vol. 31, pp 1314-1319, Sept. 1996
- [26] V. Kampus, M Trojer, R. Teschner, "Unleashing the full power of feed-forward opamps: a 200MHz, fully differential, conditionally stable, 36dB gain PGA, using a four-stage multi-path 2.5V amplifier with double feed-forward compensation", IEEE Nordic Circuits and Systems Conference, Oct. 2018
- [27] R. Adams, "Design and implementation of an audio 18-bit analog-to-digital converter using oversampling techniques." Journal of The Audio Engineering Society, pp 153-166, Mar. 1986
- [28] V. Kampus, T. Rang, "A smart capless voltage regulator for very high bandwidth A/D and D/A converters in a standard 28nm CMOS process", 15th Biennial Baltic Electronics Conference, Oct. 2016
- [29] M. Fulde, A. Belitzer, Z. Boos, M. Bruennert, J. Fritzin, H. Geltinger, M. Groinig, D. Gruber, S. Gruenberger, T. Hartig, V. Kampus, B. Kapfelsberger, F. Kuttner, S. Leuschner, T. Maletz, A. Menkhoff, J. Moreira, A. Paussa, D. Ponton, H. Pretl, D. Sira, U. Steinacker, N. Stevanovic, "A Digital Multimode Polar Transmitter Supporting 40MHz LTE Carrier Aggregation in 28nm CMOS", IEEE International Solid-State Circuits Conference, Feb. 2017
- [30] G. Gielen, L. Hernandez, P. Rombouts "Time-encoding analog-to-digital converters: bridging the analog gap to advanced digital CMOS – Part 1: Basic principles", IEEE Solid-State Circuits Magazine 12 p.47-55, Nov. 2020

- [31] T. Leslie, B. Singh, "An improved sigma-delta modulator architecture", IEEE International Symposium on Circuits and Systems, May 1990
- [32] L. M. Alvero-Gonzalez, V. Medina, V. Kampus, S. Paton, L. Hernandez, E. Gutierrez, "Ring-Oscillator with Multiple Transconductors for Linear Analog-to-Digital Conversion", MDPI Electronics Journal, Jun. 2021
- [33] V. Kampus, T. Rang, D. Knaller, C. Fleischhacker, M. Korak, J. Kiss "A fully differential, 200MHz, programmable gain, level-shifting, hybrid amplifier/power combiner/test buffer, using pre-distortion for enhanced linearity", 14th Conference on Ph.D. Research in Microelectronics and Electronics, Aug. 2018

Acknowledgements

I would like to express my sincere gratitude to Dr. Antonio Di Giandomenico for his invaluable contributions to the design, evaluation, and measurement process. His technical expertise was essential to the project's great success. Beyond his professional guidance, his mentorship and unwavering support have been invaluable. Not only was he an amazing mentor and colleague, but also a real friend, always available to lend an ear or offer encouragement. His optimistic spirit and collaborative nature fostered a truly positive and productive work environment. I am extremely grateful for his friendship and pollyannish way of thinking, which have been instrumental in overcoming personal challenges and will definitely set a rock-solid standard to aim for in all my future endeavors.

Even more, I am incredibly grateful for Professor Toomas Rang for sticking with me and believing in me and in my ability to complete this journey, even if sometimes I myself did not anymore. I sincerely believe that he is the most important reason for the completion of this work. Despite facing prolonged periods of inactivity and uncertainty on my part, he consistently supported and encouraged me and was always willing to proceed as everything was on track when I suddenly again showed up on his doorstep. His patience and willingness to accommodate my fluctuating progress is deeply appreciated. Professor Rang's invaluable guidance, expertise, and extensive network have been instrumental in bringing this project to completion. His wisdom, which surpasses his own recognition, has been a guiding light to me throughout this process.

I would also like to extend my sincere gratitude to my amazing colleagues from work, Maria Grazia Presicce and Dr. Stefan Trampitch, whose substantial contributions were equally vital to the success of this project and Dr. Alar Kuusik who helped me crossing the final hurdles in order to complete this challenging task under EU structural funds supported TemTA138 project. Without their expertise and dedication, the design of the ADC and ensuing thesis would definitely not have reached the same level of maturity and quality.

Chapeau! to everyone who guided and helped me on this path.

Abstract

Design of a Highly Efficient Analog-to-Digital Converter for Wi-Fi 7 Devices

This thesis explores the selection and design process of an ADC specifically tailored for an application within the Wireless Communications industry, with the objective of gaining a competitive advantage over the field. The research addresses the specification considerations for a Wi-Fi 7 receiver, offering a comprehensive system-level perspective and extending to feasibility studies of various topologies and their integration within the signal chain. The feasibility analysis underscores the advantages and disadvantages of selecting a $\Sigma\Delta$ -based architecture over the more commonly employed SAR ADCs, which are prevalent in the industry.

The design process is systematically broken down to the ADC sub-module level, with each distinct function elaborated upon in dedicated sub-chapters. These sections provide in-depth explanations and schematics, illustrating the design intricacies. For larger and more complex sub-modules, key simulation results are presented, focusing on the most critical outcomes. Additionally, the design section includes modules that significantly influence the ADC's performance, such as the power supply and the preceding PGA.

Following the chapter dealing with design, validation, verification, and measurement results are presented. Although some modifications between design, validation and verification steps were necessary due to discrepancies between the initial schematic design and the back annotated tape-out stream sent for fabrication, all design targets were ultimately achieved. The final product successfully met all established goals, with the final product achieving the required high-performance standards.

The ADC is designed with a forward-looking approach, equipped to meet the following next-generation challenges. It incorporates features that ensure compatibility with anticipated advancements in the forthcoming Wi-Fi 8 standard, necessitating maybe only minimal design modifications. Measurement results demonstrate that the ADC's efficiency in its segment is on par with the most advanced architectures presented at leading microelectronics conferences. Being part of one of only less than ten Wi-Fi 7 certified chips on the market, this ADC will in the future surely be in the homes of some people reading this work.

Continuous work is already in progress to integrate more and more wide-band modes under the umbrella of low power VCO operation to achieve better area optimization, reduce parasitics and enable even higher bandwidth modes in the CT $\Sigma\Delta$ ADC. Low-bandwidth options require large passive components, which can increase the ADC's area and introduce additional parasitics for high-bandwidth modes. By reducing the size of these passives, further optimization can be achieved for high-bandwidth performance. Therefore, transitioning more modes to VCO-based operation will unlock the potential for significant improvements in high-end performance. The next-generation Wi-Fi 8 CT $\Sigma\Delta$ ADC is expected to feature most modes implemented using VCO-based methods, reserving only the very highest bandwidth modes for traditional approach.

Kokkuvõte

Pidevaja Sigma-Delta muunduri disain Wi-Fi 7 seadmetele

Käesolev doktoritöö kirjeldab analoog-digitaalmuunduri (ADC) arendusprotsessi, mis on spetsiaalselt mõeldud kasutamiseks juhtmevaba side valdkonnas. Töö käsitleb topoloogia ja arhitektuuri valikut ning disainivariante, eesmärgiga saavutada konkurentsieelis teiste valdkonnas tegutsevate ettevõtete ees. Uurimus käsitleb praktilisi kaalutlusi Wi-Fi 7 vastuvõtja spetsifikatsioonidest tulevatele nõuetele, pakkudes laiaulatuslikku süsteemset vaadet ja erinevate ADC topoloogiate teostatavusuuringuid ning nende võimalikust integreerimisel signaaliahelasse. Teostatavusuuringus tuuakse esile $\Sigma\Delta$ -põhise arhitektuuri eelised ja puudused võrreldes laialdaselt kasutatavate SAR ADC-dega, mis on tööstuses endiselt rohkem levinud.

Projekteerimisprotsess on süsteemset jagatud ADC alamoodulite tasemele, kus iga eraldi funktsioon on üksikasjalikult käsitletud spetsiaalses alapeatükis. Need alapeatükid pakuvad põhjalikke selgitusi ja skeeme, mis illustreerivad projekteerimise väljakutseid. Suuremate ja keerukamate alamoodulite jaoks on esitatud peamised simulatsioonitulemused, keskendudes eelkõige kõige olulisematele aspektidele. Lisaks sisaldab projekteerimise osa mooduleid, mis mõjutavad märkimisväärselt ADC jõudlust, nagu toiteallikas ja ADC-le eelnev muudetava ülekandeturiga võimendi (PGA).

Pärast projekteerimise peatükki esitatakse valideerimise, verifitseerimise ja mõõtmistulemused. Kuigi disaini, valideerimise ja verifitseerimise etappide vahel olid vajalikud mõned muudatused algse skeemidisaini ja tootmiseks saadetud lintfaili vaheliste lahknevuste tõttu, saavutati lõpuks kõik projekti eesmärgid. Lõpptulemus vastas edukalt kõigile seatud eesmärkidele, saavutades nõutud kõrge jõudluse standardid.

Disainitud ADC on loodud tulevikku suunatud lähenemisega, et tulla toime ka ülejäämise põlvkonna väljakutsetega. Loodud toode sisaldab funktsioone, mis tagavad ühilduvuse eeldatavate arengutega tulevases Wi-Fi 8 standardis, vajades võib-olla ainult minimaalseid disainimuudatusi. Mõõtmistulemused näitavad, et ADC tõhusus oma segmendis on võrreldav juhtivatel mikroelektroonika konverentsidel esitletud kõige arenenumate arhitektuurilahedustega. Olles komponent ühest vaid vähem kui kümnest kiibist, millele on antud Wi-Fi 7 sertifikaat, leiab see ADC kindlasti tee ka osade seda lõputööd lugenud inimeste kodudesse.

Jätkuvalt toimub töö selle nimel, et integreerida üha suurema ribalaiusega režiime madalama energiatarbega pingega juhitava ostsillaatori (VCO) režiimi alla, et saavutada parem pindala optimeerimine, vähendada parasiitmahtuvusi ja võimaldada veelgi suuremat jõudlust kõrgema ribalaiuste juures $CT\Sigma\Delta$ ADC-s. Madala ribalaiusega režiimid nõuavad suuri passiivseid komponente, mis omakorda suurendavad ADC pindala ja toovad kaasa täiendavaid parasiitmahtuvusi kõrgema ribalaiusega režiimide juures. Nende passiivsete komponentide suuruse vähendamine võimaldab veelgi paremat optimeerimist kõrge ribalaiusega jõudluse jaoks. Seetõttu avab rohkemate režiimide üleminek VCO-põhisele toimimisele märkimisväärsed võimalusi paranduste saavutamiseks tipptasemel jõudlusel. Arvatavasti on järgmise põlvkonna Wi-Fi 8 $CT\Sigma\Delta$ ADC-l enamik režiime rakendatud VCO-põhiste meetodite abil, reserveerides ainult kõrgeima ribalaiusega režiimid traditsioonilisele $\Sigma\Delta$ lähenemisele.

Appendix

Publication I

Michael Fulde, Gernot Babin, Christoph Duller, Simon Gruenberger, Harun Habibovic, **Vahur Kampus**, Gerhard Knoblinger, Christian Krassnitzer, Franz Kuttner, Edwin Thaller, Davide Ponton, Andreas Santner, "Innovative RFDAC concepts for digital multi-mode transmitter in cellular applications" 9th Conference on Ph.D. Research in Microelectronics and Electronics (PRIME), Jun. 2013.

Innovative RFDAC Concepts for Digital Multi-Mode Transmitter in Cellular Applications

M. Fulde, F. Kuttner, E. Thaller, D. Ponton, V. Kampus, S. Gruenberger, H. Habibovic, A. Santner, C. Duller, C. Krassnitzer, G. Babin, G. Knoblinger

Intel

Digital polar transmitter concepts based on RF-DA-converter recently proved the potential to significantly reduce power consumption. Furthermore, external component count as well as PCB area is minimized since no TX SAW filter is required and a single multimode, multiband power amplifier can be used.

The RFDAC is a key building block for this architecture and needs to provide very high dynamic range (e.g. 17 ENOB for 3G) at high clock frequencies (e.g. 1GHz) with reasonable power efficiency in order to deal with far-off noise, spurious and repetition spectrum issues. Another key requirement for the DAC is the scalability to new CMOS nodes. A novel RFDAC concept is presented and benchmarked; the capability to fulfill the tough linearity/resolution/power specifications even in 28nm technology and below is shown.

A 15bit class-B (=single-ended) current mode DAC with distributed digital mixer has been implemented in 28nm CMOS. To minimize DNL and to meet the stringent quantization noise requirements a 10bit thermometer coded array of 1024 current sources is combined with 5 binary scaled cells. Different calibration techniques are employed to compensate for the dynamic degradation effects inherently connected to the single ended DAC structures and to achieve linearity requirements. A pre-charge mechanism prepares the current source before turning it on. To compensate coupling on the bias node a H2 dummy capacitor is used. The related challenge to match a parasitic capacitor to a dummy MOS capacitor is solved with an integrated regulation loop employing tracking ADC principle to set the correct compensation voltage. The bias voltage of the distributed mixer-cascode structure is also regulated on chip. Preliminary measurements on a 28nm test-chip show that IM3 of better than -45dBc at 6dBm output power is achievable (meeting product specs for 2G/3G/LTE). At 6dBm the RFDAC draws 7mA from 1V (cell array) and 25mA from 2V (trafo/load). The active DAC area is just 0.048mm² (without trafo).

Furthermore a new digital data decoding and mixing scheme is presented. This approach allows splitting the current source array in two independent sub-DACs. In addition the phase of the LO signal can be shifted in the decoder which allows to use a single RFDAC for polar, signed polar and IQ modulation.

Finally an outlook to even more digitized RFDAC based transmitter concepts with integrated PA functionality is given.

Publication II

Vahur Kampus, Toomas Rang, "A smart capless voltage regulator for very high bandwidth A/D and D/A converters in a standard 28nm CMOS process", 15th Biennial Baltic Electronics Conference, Oct. 2016.

A smart capless voltage regulator for very high bandwidth A/D and D/A converters in a standard 28nm CMOS process.

Vahur Kampus

PhD Student, Thomas Johann Seebeck
Department of Electronics
Tallinn University of Technology

Toomas Rang

Professor; Head of the Thomas Johann Seebeck
Department of Electronics
Tallinn University of Technology

Abstract – Analog systems, where good and stable supply plays an important role in total system performance, need highly accurate, clean and reliable low dropout linear regulators (LDOs). High performance and low area are also important in voltage reference circuits for achieving smallest cost and best in the industry Analog-to-Digital (A/D) and Digital-to-Analog (D/A) conversion. This paper describes a smart capless LDO that is designed to be used with very high bandwidth ADCs and DACs in a standard CMOS process. The maximum current, the LDO is able to provide is upwards of 200 mA, while the output voltage range can be programmed from 0.85 V to 1.05V. The intelligence of the LDO allows it to adjust its quiescent current totally based on the load current and can be as small as 80 μ A with no load. Furthermore, the LDO can achieve an efficiency of 85% when used in normal conditions.

Keywords- LDO; CMOS; analog; semiconductor; PSRR; Noise; Power Management; ADC; DAC

I. INTRODUCTION

High quality supply voltage is important to build up top of the line performance A/D and D/A converters. Modern electronics devices require variety of different voltage levels for different blocks. Almost as important is the area and the overall cost. External components increase the Bill of Materials (BOM) and complexity of the application and should be avoided if possible. Therefor capless LDOs play an important role in providing clean supply voltage with low correlated and uncorrelated noise. They are ideally suited for System-on-Chip solutions for their relatively simple nature [1]–[3]. Moreover, the efficiency of such regulators has become vital. Code of Conduct rules for power dissipation are continuously getting tougher to meet and therefor require high level of efficiency in order to have more flexibility for the converters [4]. All this means that quality voltage regulation have to be done on-chip, effectively and with minimum area.

High bandwidth converters also require high bandwidth regulators. In order not to see performance limitations due to supply regulation, the bandwidth of the LDO should be in the same range as the converter. As present day IEEE 802.11 and G.FAST standards already have included channels with 80 MHz and 106 MHz bandwidth respectively and future amendments

could also go up to even 200 MHz of 400 MHz bandwidths, LDO design has to follow suit. From the standpoint of the converter it is more desirable to have a supply with uncorrelated disturbers at the output then correlated. Uncorrelated noise will lower the performance of the converters in the entire band equally, while correlated noise can mix with the clock or carriers and create unwanted tones in band, lowering Signal-to-Noise (SNR) performance unevenly across the wanted bandwidth. Therefor when designing an LDO for a high bandwidth ADC or DAC, it is generally a good idea to trade noise performance for better regulation and Power Supply Rejection Ratio (PSRR). Good PSRR can be a challenge for on-chip voltage regulators [5]. The noise of the LDO is generally less important, however has to still stay within respective limits in order not to limit the performance of the converter itself.

II. STRUCTURE OF THE PROPOSED LDO

Commonly used capless LDO's are designed to work under strict load conditions, what are known and will not change in a very wide range. This will however limit the effectiveness or stability of the LDO since it will either have to be greatly oversized or cannot be used in low load conditions, therefor potentially wasting a lot of power. Capless LDOs are usually designed from two separate voltage rails. One rail for the pass device, where the current is drawn from and one for the error amplifier, what regulates the pass device.

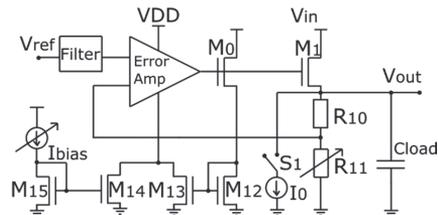


Figure 1. Biasing of the proposed LDO

This kind of voltage rail separation allows the usage of an NMOS pass device, what can be operated if the voltage on the

VDD rail is at least one threshold higher than the required regulated output. An NMOS pass device has several advantages. It does not offer any gain, but since the output impedance is at least one order of magnitude lower than an equivalent PMOS pass device, then the pole created from the NMOS pass device output resistance and the capacitance it has to drive, is also located in a much higher frequency. That allows to place the other poles internally at relatively high frequency even in no load conditions.

The main idea of the proposed LDO is presented in Figure 1. The pass device M1 has a much smaller sense device M0 in parallel, what routes a small portion of the current back to the system for self-use. By doing this, the LDO can track the pole created by the pass device and adjust the compensation created by the internal stages accordingly. This will allow the LDO to still be stable and operational over a wide range of output currents, while itself requiring also far less current to operate in low load conditions, where the performance is less important. M14 sets the constant biasing current and M13 adds the variable part based on the output current.

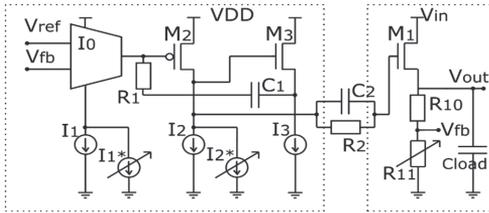


Figure 2: High level diagram of the proposed LDO

The challenge of the design is to keep the phase margin constant over all load conditions. For that, internal compensation schemes have been developed to provide the LDO with some intelligence. The error amplifier itself consists of two stages plus a follower stage for Miller compensation (Figure 2). The first amplifying stage is a common differential pair, what consist of an NMOS input pair and PMOS mirrors on top. The second stage is a little more than a single stage PMOS amplifier M2 with NMOS current source.

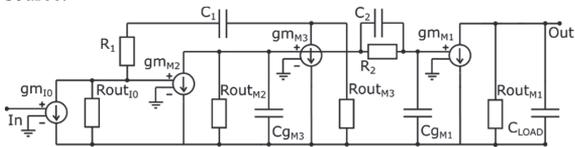


Figure 3: Small signal equivalent circuit of the proposed LDO

The circuit uses a Miller compensation scheme, what is deployed through a buffer composed of M3, what acts like an NMOS follower and a current source I3. The buffer serves two purposes.

First, it level shifts the voltage seen by capacitor C1, making it possible to use much smaller non-linear MOS capacitor in the place of a linear capacitor. Secondly, it decouples the Miller compensation from the signal path. The Miller compensation creates a pole that comes from the output impedance of the first stage and the Miller capacitor. Assuming that the Miller capacitor is much larger than the gate capacitance of M2, a constant gain of one from the NMOS follower composed of M3 and I3 and that the current sources I2 and I2* have a much larger output impedance than M2, the location of the pole is at:

$$P_0 = \frac{1}{2\pi * Rout_{I0} * C_1 + Rout_{M2} * gm_{M2}} \quad (1)$$

The resistor R1 adds a Miller zero to the transfer function. The location of the Miller zero is at:

$$Z_0 = \frac{1}{2\pi * R_1 * C_1} \quad (2)$$

The second stage combined with resistor R2 and capacitor C2 will add another pole and another zero. Assuming that the output resistance of I2 and I2* are much higher than the output impedance of M2, the location of the second pole becomes:

$$P_1 = \frac{1}{2\pi * (Rout_{M2} + R_2) * C_2 * gm_{M1}} \quad (3)$$

The location of the second zero is at:

$$Z_1 = \frac{1}{2\pi * C_2 * R_2} \quad (4)$$

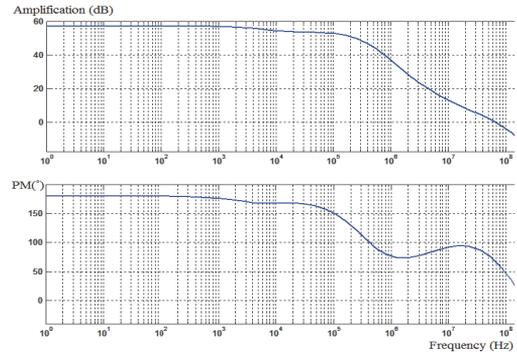


Figure 4: A Bode plot of the LDO with a 20mA Iload

The last pole comes from the pass device. Since the pass device is an NMOS follower, the output impedance of it will be relatively low. However the output capacitance can be – depending on the topology of the ADC or DAC – quite high, spanning up to couple of hundreds of picofarads. The stability is mostly effected in scenarios, where the output current is low. A bigger load capacitor

means that the pole coming from the pass device is shifting to lower frequencies, lowering the phase margin. To combat this, the LDO has been designed in a way, where the load capacitor value can be internally programmed. When the output capacitor is larger a small current source I_0 (Figure 1) is enabled at the output, what limits the minimum current the LDO can supply. This current source is only activated when the current drawn from the pass device is low and will be deactivated automatically with higher currents. Lastly, the pole coming from the pass device is given by the equation:

$$P_2 = \frac{1}{2\pi \cdot R_{out} \cdot M_1 \cdot C_{Load}} \quad (5)$$

It can be seen from the equations, that the two zeroes are constant over load and rely only on the values of passive devices. The poles however depend also on the output impedances of different stages. With the help of carefully designed current feedback, the system is automatically able to change the values output resistances by injection more or less current with the use of two current sources I_1^* and I_2^* . The more current is inserted back to the two stages the further in frequency the poles are shifted, extending therewith the bandwidth of the loop and improving the regulation speed.

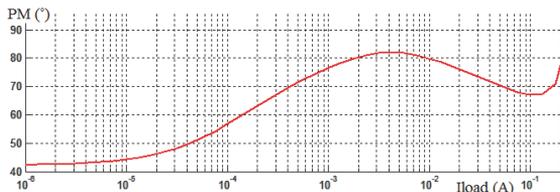


Figure 5: Phase Margin of the LDO over Output load

III. SIMULATED RESULTS

The LDO has been designed in TSMC HPC+ 28nm technology, with 1.8V I/O devices to supply a DAC or an ADC, what are never used simultaneously. The VDD used is 1.8V +/-10%. The target average current for the blocks is from around 20 mA to 30 mA in both transmit and receive modes. A bode plot of the LDO, with such load current is shown in figure 4. The bandwidth of the LDO is around 70 MHz in that case, while consuming around 350 μ A of current from the VDD rail and 100 μ A from Vin rail. The 20 mA to 30 mA use case is however not the only one and therefore it is important that the LDO retains really its stability over all load conditions. The current feedback was designed in a way where the phase margin is kept as constant as possible over PVT (Process, Voltage, Temperature) with a minimum phase margin around 50°, only allowing lower values in no load conditions where the LDO should just be operational, stable and is not intended to be heavily used. The phase margin change over load of a nominal case has been plotted in figure 5. It can be seen that

the phase margin indeed creeps below the wanted 50° in very low load conditions, but is deemed more than sufficient.

One interesting point that can be brought out, is the different movement of poles and zeroes over diverse load conditions. This is plotted in figure 6. While the two zeroes plotted in green and purple stay constant over frequency, all of the three poles shift with load. The relatively highest shift occurs with the pole coming from the first stage plotted in red, while the lowest shift happens with the pole coming from the second stage plotted in blue. The pole coming from the output, what is already in relatively high frequencies with no load, because of low impedance values coming from the NMOS pass device, moves relatively linearly in the entire load range, only saturating with very low load currents, where the current feedback is already the dominant consumer. It can be seen that the dominant pole changes from P_0 to P_1 with higher load. The point of the exchange is at around 10 mA, where the LDO should be used most of the time. This is done intentionally in order to have the best performance and the highest phase margin of the entire range, in the case mentioned.

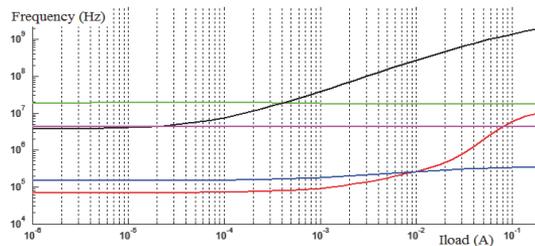


Figure 6: Approximate locations of poles and zeroes over different load conditions. P_0 (Red), P_1 (Blue), P_2 (Black), Z_0 (Purple), Z_1 (Green)

One very important aspect in the design of LDOs for such a systems is the immunity to noise and disturbers coming externally

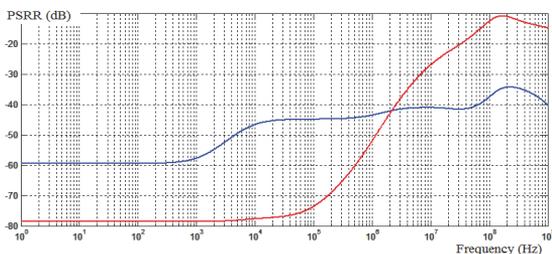


Figure 7: PSRR of the LDO. Blue (VDD), Red (Vin).

from the supplies and references to the LDO itself. Since there are two different supply rails, what are coming from two different DC/DC converters, that can work in different phases or even

frequencies, the rejection is checked separately for both (Figure 7). The red line represents the PSRR from the V_{in} rail, where the current is drawn from, or the pass device supply and the blue line represents the PSRR of the LDO from the voltage rail where the error amplifier is supplied. The two PSRRs intersect at around 1 MHz region, where below that the VDD rail will be the cause of most of the correlated noise and above 1 MHz, V_{in} will cause it. The intersect point at 1 MHz is chosen because most of the commercially available DC/DCs tend to work at or around that frequency.

In addition to the unwanted noise coming from the external supplies, there is also noise coming from the reference of the LDO due to all sorts of noise the reference node picks up on the chip. Because of this, the reference voltage is first filtered with a regular RC filter (Figure 1). In a standard Wi-Fi OFDM modulation, the first wanted subcarrier is already at 312.5 kHz. This means that the integration of the noise has to begin from half of that frequency. Due to PLL inaccuracies, one could easily half even that. So the LDO should provide a clean low noise supply starting from as low as 70 kHz. Therefore the filter is designed to have a cutoff frequency in the tens of kilohertz (Figure 8).

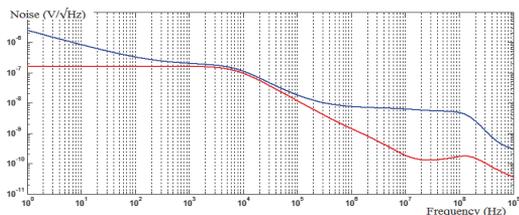


Figure 8: Noise performance of the LDO. Red (Input filter) Blue (Complete LDO)

Although uncorrelated noise is not as important for systems like ADCs or DACs, as correlated noise, the LDO should still provide a decent performance in that category. The designed LDO has a quite flat output noise of less than 7 nV/√Hz starting @ 1 MHz. To check the overall system stability various load jumps were performed from load currents ranging from 1 μ A up to 200 mA (Figure 9). In the simulations performed, no ringing was detected.

IV. CONCLUSION

The LDO has been designed and manufactured in a TSMC 28nm process (Figure 10). The layout measures roughly 200 μ m by 100 μ m and the LDO is capable of supplying up to 200 mA, when connected to a 1.8 V VDD rail and a 1.2 V V_{in} rail. At nominal conditions, with 20 mA load current, the self-current measures about 450 μ A, while having a bandwidth of roughly 70 MHz. In no-load conditions the self-current measures about 80 μ A.

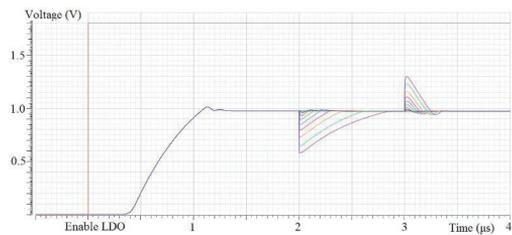


Figure 9: LDO Startup and step response from 1 μ A up to 200 mA.

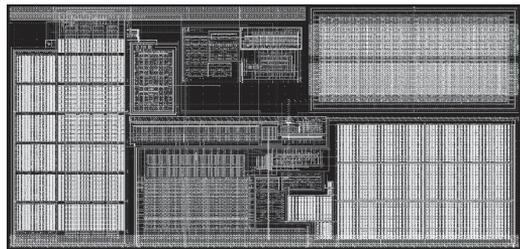


Figure 10: Layout of the proposed LDO

V. ACKNOWLEDGEMENTS

The authors wish to thank Mr. Antonio Di Giandomenico, Mr. Marco Salvucci, Mr. Rein Sabolotny and Mr Eiko Kängsep for the valuable discussions along the preparation of this paper. This research was supported by Intel Austria GmbH, the Horizon 2020 ERA-chair initiative VFP15051 “Cognitive Electronics — COEL”, and by the Estonian Research Council through the Institutional Research Project IUT19-11.

REFERENCES

- [1] G. A. Rincon-Mora and P. E. Allen, “A low-voltage, low quiescent current, low dropout regulator,” *IEEE J. Solid-State Circuits*, vol. 33, no. 1, pp. 36–44, Jan. 1998
- [2] S. K. Singh and G. D. Kanungo, “Ultra-fast cap-less LDO for dual lane USB in 28FDSOI” 28th International Conference on VLSI Design, Jan. 2015
- [3] K. N. Leung and P. K. T. Mok, “A capacitor-free CMOS low-dropout regulator with damping-factor-control frequency compensation,” *IEEE J. Solid-State Circuits*, vol. 38, no. 10, pp. 1691–1702, Oct. 2003.
- [4] European Commission. Joint Research Centre. Institute for Energy. “Code of Conduct on Energy Consumption of Broadband Equipment Version 5.0” Dec. 2013.
- [5] M. El-Nozahi, A. Amer, J. Torres, K. Entesari, and E. Sanchez-Sinencio, “A 25 mA 0.13 μ m CMOS LDO regulator with power supply rejection better than -56 dB,” *IEEE Int. Solid-State Circuits Conf.*, San Francisco Feb. 2009

Publication III

Michael Fulde, Alexander Belitzer, Zdravko Boos, Michael Bruennert, Jonas Fritzin, Hans Geltinger, Marcus Groinig, Daniel Gruber, Simon Grünberger, Thomas Hartig, **Vahur Kampus**, Boris Kapfelsberger, Franz Kuttner, Stephan Leuschner, Thomas Maletz, Andreas Menkhoff, Jose Moreira, Alan Paussa, Davide Ponton, Harald Pretl, Daniel Sira, Ulrich Steinacker, Nenad Stevanovic, "A Digital Multimode Polar Transmitter Supporting 40MHz LTE Carrier Aggregation in 28nm CMOS", IEEE International Solid-State Circuits Conference, Feb. 2017.

13.2 A Digital Multimode Polar Transmitter Supporting 40MHz LTE Carrier Aggregation in 28nm CMOS

Michael Fulde¹, Alexander Belitzer¹, Zdravko Boos¹, Michael Bruennert¹, Jonas Fritzin¹, Hans Geltinger², Marcus Groinig³, Daniel Gruber³, Simon Gruenberger³, Thomas Hartig³, Vahur Kampus³, Boris Kapfelsberger¹, Franz Kuttner³, Stephan Leuschner¹, Thomas Maletz¹, Andreas Menkhoff¹, Jose Moreira¹, Alan Paussa³, Davide Ponton³, Harald Pretl⁴, Daniel Sira¹, Ulrich Steinacker¹, Nenad Stevanovic¹

¹Intel, Neubiberg, Germany

²DPMA, Munich, Germany

³Intel, Villach, Austria

⁴DMCE, Linz, Austria

The evolving trend for increasing data rates in cellular communication systems with limited and fragmented frequency spectrum requires enhanced spectral efficiency of today's and future communication standards. The resulting need for high-order modulation schemes with large peak-to-average power ratio results in stringent requirements on in-channel linearity and SNR. At the same time FDD operation in SAW filter-less designs enforces very demanding limits for out-of-band and spurious emissions. Digital polar TX concepts have demonstrated low power consumption combined with low out-of-band noise in moderate bandwidth applications like WCDMA [1]. However, the implementation of 4G wideband polar systems with 20MHz RF bandwidth and above is very challenging due to the nonlinear conversion from cartesian to polar coordinates extending the effective signal bandwidth even further. A digital quadrature TX concept with up to 80MHz bandwidth, low EVM and moderate power efficiency has been shown in [2]. More than 30% power-added efficiency has been demonstrated in a 20MHz polar TX based on a switched-capacitor digital PA [3]. However, [2,3] suffer from low resolution and limited out-of-band noise performance.

This work introduces a digital polar transmitter architecture based on signed amplitude path and switched-capacitor RFDAC (C-RFDAC) to overcome the conventional bandwidth vs. signal quality limitations in polar TX. It combines a digital concept to limit large phase jumps in the DPLL with a high-speed, high-resolution C-RFDAC and achieves 3.6% EVM in 40MHz LTE Carrier Aggregation mode (2xLTE20) while still fulfilling 3G TX noise requirements of -159dBc/Hz, both at 6dBm output power.

The main limiting factor for increasing the signal bandwidth in a polar TX is typically the restricted dynamic phase/frequency modulation (PM/FM) range of the DCO/DPLL. PM/FM range reduction by clipping of frequency peaks degrades the signal quality in terms of EVM and out-of-band noise performance. The signed polar concept adds another degree of freedom to the signal quality vs. FM range trade-off. The basic concept is shown in Fig. 13.2.1. The main idea is to operate the amplitude (AM) path in signed mode where a change in sign equals a phase-shift of 180° according to $A \cdot \exp[j(\varphi)] = -A \cdot \exp[j(\varphi - \pi)]$. This way large phase jumps can be reduced by 180° in the AM path without sacrificing signal quality.

The core of the AM path is a C-RFDAC as shown in Fig. 13.2.2. The C-RFDAC offers high power efficiency [3] and is inherently wideband: the D-to-A conversion is just done by a digital logic element, gating the LO signal depending on the digital AM input code. A high-resolution, low-noise implementation is presented here. Quantization noise and gain control range requirements lead to 15b resolution. To limit the DNL error a 10b thermometer / 5b binary segmentation is used. The 10b thermometer array is decoded in 2 steps with row/column and local decoder to limit complexity and enable high speed. As in [1] the DAC sampling clock is derived from the phase-modulated LO and can reach up to 2.8GHz in B41 including FM overhead. The unit cell is built in a pseudo-differential structure, both for the local decoder and the LO gating. Differential row/column signals are used to minimize the capacitive coupling from long lines over the array to sensitive nodes like LO or RF out. Differential LO gating keeps the amount of cells connected to VDD and VSS equal all the time avoiding common-mode effects.

The realization of negative amplitudes in the C-RFDAC is only possible with severe power drawbacks, e.g. by using half of full-scale code as virtual zero or negative supply voltage. Therefore the sign change is converted back to 180° phase shift

by inverting the LO in an XOR gate, see Fig. 13.2.2. A re-timing stage is added afterwards to guarantee timing relations and phase-noise performance. The 180° phase shift also changes the effective sampling point of the AM data: since the RF output is only triggered by LO edges, it is moved in time by a half LO period with every sign change, see timing of digital input codes d_1/d_2 and d_3/d_4 in Fig. 13.2.2. The resulting shift of sampling time creates significant distortion in RF output spectrum and is corrected in the digital domain with a simple FIR filter that delays positive and negative signals against each other by a half LO period.

Impedance and thermal noise of the supply network are key parameters for out-of-band emissions. The C-RFDAC samples VDD/VSS at LO rate, thus all disturbances on the supply like ripple or thermal noise are upconverted. To accommodate the requirements on supply quality an on-chip LDO with external blocking cap is implemented. The off-chip cap provides large currents and low impedance at medium frequencies but also introduces resonance peaks due to the inductive wiring parasitics L_{par} , see Fig. 13.2.3. The resonances may occur in the relevant duplex frequencies and degrade RX band noise. To limit the impedance peaks a shunt regulator is added to the supply network, active only for high frequencies, see Fig. 13.2.3. An active compensation circuitry is added to mitigate the problem of current steps/glitches on the supply. Dummy currents are added to smear over steps/glitches and limit the bandwidth of distortion. The result of TX noise optimization is shown exemplified for an LTE20 signal at 1880MHz in Fig. 13.2.3, yielding an improvement of 5dB at the duplex distance of 80MHz. Supply impedance is also critical for LO leakage and EVM at low output power. In order to relax the supply requirements for low output power an extra binary-scaled low-power DAC combined with a dedicated LDO is added.

The proposed transmitter is implemented in 28nm CMOS technology, featuring low-, mid- and high-band path and supporting GSM, EDGE, UMTS/WCDMA, HSPA+, TDSCDMA, LTE and LTE-A (contiguous CA with 2xLTE20), see Fig. 13.2.7. Area for respective C-RFDACs including biasing is 1.3mm². Measured out-of-band noise for different modes/bands is shown in Fig. 13.2.4. -159dBc/Hz is achieved in 3G low- and mid-band at closest duplex distance, in worst case LTE20 scenarios (100RB – B1, 75RB – B7) -157dBc/Hz and -152dBc/Hz, respectively. Besides full RB allocation also single RB use cases with large offset from carrier are critical. Due to the polar modulation no degradation in TX noise or spectral emissions is observed in single RB allocation at the channel edge, CIM3 stays well below -70dBc, see Fig. 13.2.4.

In-band signal quality is summarized in Fig. 13.2.5. Due to the signed AM path operation, EVM does not exceed 3.6% even in 40MHz LTE-CA (2xLTE20) use-case, ACLR is only -47dBc. Due to the additional low-power DAC 3G EVM does not degrade significantly over large dynamic range of 66dB, which is required for 3G gain control, see Fig. 13.2.5.

The complete TX chain, from digital interface via DPLL/RFDAC to RF output, including bias and LO/clock distribution, consumes in low/mid-band 24mA/27mA DG09 weighted current from a 3.6V battery in 3G mode, using an external DC/DC converter with 93% efficiency. The corresponding numbers for LTE10 (50RB) and LTE20 (100RB) are given in Fig. 13.2.6, measured at -18dBm which translates to roughly 0dBm antenna power. Compared to multimode transmitters with similar feature set as presented e.g. in [4] a significant reduction in power consumption is achieved.

Acknowledgement:

The authors would like to thank the whole layout team for the valuable support.

References:

- [1] Z. Boos, et al., "A Fully Digital Multimode Polar Transmitter Employing 17b RF DAC in 3G Mode," *ISSCC*, pp. 376–377, Feb. 2011.
- [2] C. Lu, "A 24.7dBm All-Digital RF Transmitter for Multimode Broadband Applications in 40nm CMOS," *ISSCC*, pp. 332–333, Feb. 2013.
- [3] S. M. Yoo, et al., "A Switched-Capacitor Power Amplifier for EER/Polar Transmitters," *ISSCC*, pp. 428–429, Feb. 2011.
- [4] Th. Georgantas, et al., "A 13mm² 40nm Multiband GSM/EDGE/HSPA+/TDSCDMA/LTE Transceiver," *ISSCC*, pp. 160–161, Feb. 2015.
- [5] P. E. P. Filho, et al., "0.22mm² CMOS Resistive Charge-Based Direct-Launch Digital Transmitter with -159dBc/Hz Out-of-Band Noise," *ISSCC*, pp. 250–251, Feb. 2016.

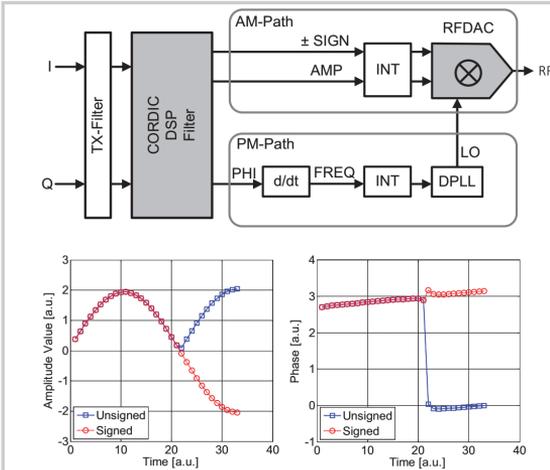


Figure 13.2.1: Concept of digital polar modulator with signed AM path.

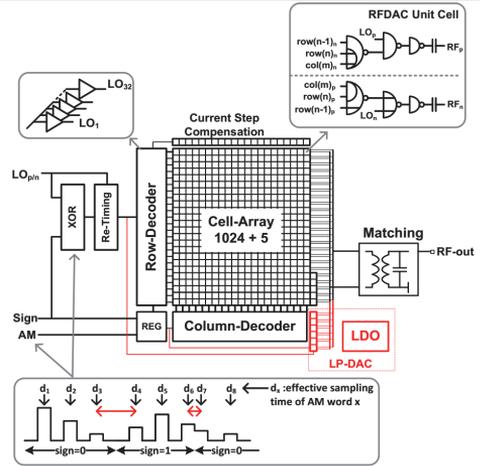


Figure 13.2.2: Schematic of C-RFDAC and sign change details.

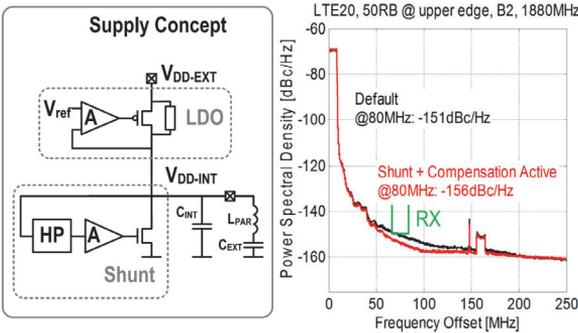


Figure 13.2.3: Supply concept and TX noise optimization (LTE20).

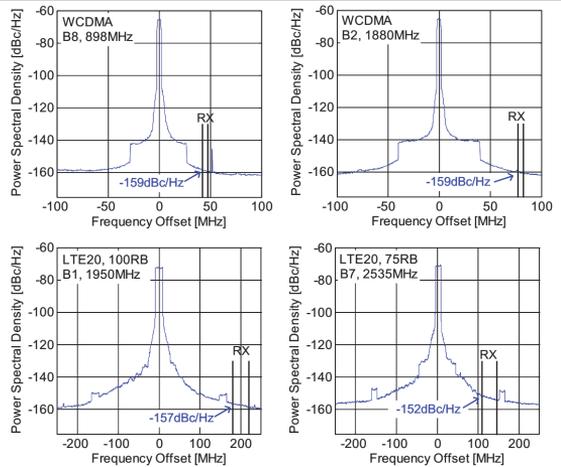


Figure 13.2.4: Measured TX noise into RX band for WCDMA and LTE20.

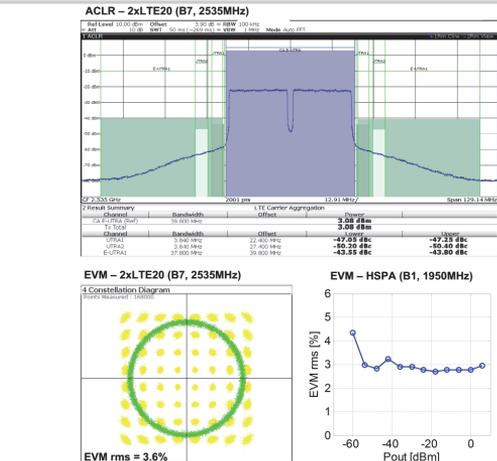


Figure 13.2.5: Measured EVM and ACLR for 2xLTE20 with 64 QAM.

Mode	Parameter	Unit	This Work			[4]		[5]	
			L: 0.7-1.0GHz	M: 1.4-2.0GHz	H: 2.3-2.8GHz	B5/B17	B1	B7	0.9GHz
2G GMSK	Pout	[dBm]	12	12					
	Phase Error RMS	[°]	0.6	0.8					
	Noise @ 20MHz	[dBc/Hz]	-165	-160					
3G WCDMA / UMTS	Pout	[dBm]	6	6		>5	>5		
	EVM RMS	[%]	2.2	2.5		1.2	1.8		
	ACLR1/2	[dB]	-49/-145.3	-47.3/-49.6		-43n.a.	-43n.a.		
4G LTE10 (50RB)	Pout	[dBm]	6	6	6	4	4	4	4
	EVM RMS	[%]	2.4	2.6	2.1	0.9	1.2	1.5	
	UACLR1/2	[dB]	-45.1/-57.6	-44.9/-58.2	-44.3/-56.5	-41n.a.			
4G LTE20 (100RB)	Pout	[dBm]	6	6	6	4	4	4	4
	EVM RMS	[%]	2.9	2.6	2.6	0.9	1.3	1.6	<1.6%
	UACLR1/2	[dB]	-44.4/-49.9	-45.7/-53.6	-44.3/-54.2	-43n.a.	-44.5n.a.	-49/-59	-47/-59
CA - 2xLTE20 (200RB)	Pout	[dBm]	6	6	6	4	4	4	4
	EVM RMS	[%]	2.9	2.6	2.6	0.9	1.3	1.6	<1.6%
	UACLR1/2	[dB]	-44.4/-49.9	-45.7/-53.6	-44.3/-54.2	-43n.a.	-44.5n.a.	-49/-59	-47/-59
Process	Current @ -18dBm	[mA]	29	28	35	n.a.			
	Current @ -1dBm	[mA]	38	39	43	65 ¹⁾	7 ²⁾	15 ³⁾	
Supply	[V]		1.3/1.1	C-RFDAC, 1.0 DIG		n.a.		28	0.9/1.1

¹⁾ not DQ09 weighted, Pout = -50dBm
²⁾ @ 45MHz offset, measured with single-tone BB signal
³⁾ without PLL, Pout = -3.50Bm, value calculated assuming 3.6V battery voltage and 93% DCDC efficiency

Figure 13.2.6: Performance summary and comparison vs prior art.

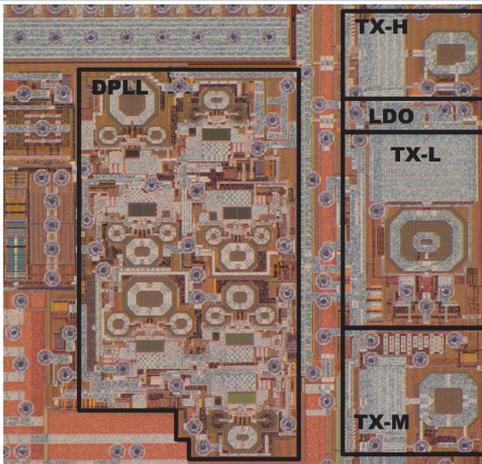


Figure 13.2.7: Die micrograph including DPLL and AM path of TX.

Publication IV

Vahur Kampus, Toomas Rang, Daniel Knaller, Christian Fleischhacker, Markus Korak, Jozef Kiss "A fully differential, 200MHz, programmable gain, level-shifting, hybrid amplifier/power combiner/test buffer, using pre-distortion for enhanced linearity", 14th Conference on Ph.D. Research in Microelectronics and Electronics, Aug. 2018.

A fully differential, 200MHz, programmable gain, level-shifting, hybrid amplifier/power combiner/test buffer, using pre-distortion for enhanced linearity

Vahur Kampus
and Toomas Rang
Department of Electronics
Tallinn University of Technology
Tallinn, Estonia 19086

Daniel Knaller,
Christian Fleischhacker,
Markus Korak,
and Jozef Kiss
Intel Austria GmbH
Villach, Austria 9528

Abstract—With the continuous advancement of standards in high-end telecommunication systems, requirements for analog circuitry have become ever more demanding. The new standards not only want to support better modulation schemes, demanding less noise with higher linearity, but also require increased bandwidth from analog circuitry. This paper describes a fully differential, programmable gain, baseband power combiner with a bandwidth of 200MHz, using pre-distortion for extra linearity, that offers also level-shifting and test buffer functionalities. The pre-distortion allows the improvement of linearity in the amplifiers AB stage, what otherwise would be a limiting factor in rail-to-rail swing operations and lowering the maximum throughput.

Index Terms—CMOS, semiconductor, analog, opamp, power combiner, pre-distortion, class-AB, multi-path, transmit chain

I. INTRODUCTION

Standard OFDM-like signals lose about 3dB of single sub-carrier strength for every doubling of the bandwidth (BW). To benefit from all the effort of increasing the BWs and maintain the same throughput per sub-carrier, analog blocks have to be designed with 3dB better noise and linearity specs for each doubling of the BW. If new standards incorporate within them new modulation schemes, these requirements can be even more challenging [1]. Modern day high-BW transmitter chains tend to consist of digital blocks dedicated to signal processing, a digital-to-analog converter (DAC), analog amplifiers or buffers and either a line driver (for wired systems) or a power amplifier (for wireless systems) (Fig. 1). High performance line drivers and power amplifiers are generally built in different technologies than the rest of the transmit chain – meaning the signal path has to go off-chip. Those external blocks can often be the biggest contributors to the total chain noise figure and non-linearities [2]. Thus lowering the performance and gain requirements for them is both highly welcomed and beneficial. Because the DAC and the external amplifier are often designed for one optimum operating condition, the power combiner can also be the primary means to change the total chain gain. Modern day analog technologies usually offer devices with at least two flavors – one in core voltage domain and one in input/output (I/O) voltage domain. In order to have the most



Fig. 1. Typical signal path in high-BW, high-performance transmitters.

effective level-planning and increase the signal amplitude at the input of the external amplifier, therefore requiring less gain and performance from it, but at the same time taking full advantage of the shrink what the technology node can offer, one would like to design all the digital blocks and the DAC in the core voltage domain and then amplify the signal with the power combiner using I/O devices to the needed levels while also level-shifting the common-mode before going off-chip and arriving at the externals.

II. POWER COMBINER

The proposed power combiner is part of a VLSI system designed in a standard C40 technology with I/O devices allowing max VDD of 2.5V. It is designed to provide programmable gain from -15dB up to +15dB with a level-shifting function to a fixed output common-mode voltage of VDD/2 (Fig. 2a). It can be operated with any input common mode ranging from VSS to VDD for gain configurations less than 5dB. If the configured gain is higher, the input common mode will have a slight range limitation depending a lot on the actual gain used.

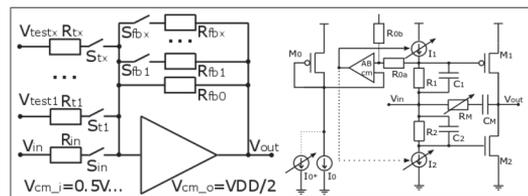


Fig. 2. a) topology of the power combiner. b) topology of a classical push-pull symmetrical high linearity AB stage.

In the buffer mode (0dB), there are no special limitations and the input common mode can actually range from a bit below VSS to a bit above VDD. The combiner has a single optimized main input (V_{in}) connected to the transmitter DAC used in the signal path and several test-mux inputs (V_{test}) that can be used to test the performance of various other IPs across the chip (Fig. 2a). Since there will be a steady DC current flowing through the feedback network, if the common-mode voltages at the input and output differ, the other IPs in test must be able to sink or source the extra current coming from the level shifting functionality, if operating with any other different common-mode. The amount of current needed to be sunk or sourced is proportional to the difference between the two common-mode voltages and inverse proportional to the total resistance of the feedback resistor R_{fb} and the input resistor R_{in} :

$$I_{DC} = 2 * \frac{V_{cm_o} - V_{cm_i}}{R_{fb} + R_{in}} \quad (1)$$

The four main contributors for the non-linearities of the combiner are the finite loop gain of the amplifying circuit, the inherited non-linearity coming from the AB stage, the non-linearities of the signal path switches (S_{fb} , S_{in} , S_t) required for the gain and input programming and the non-linearities associated with the resistors (R_{fb} , R_{in} , R_t). Assuming proper care on the feedback network design – so the switches are sized correctly, that the on-resistance of the switch is non-dominant and the resistance of the resistor is constant with different potential across it – all the dominant non-linearities will come from the amplifier. This is a critical issue in modern high-speed data communications requiring full-rail swing operation for max throughput [3].

III. STRUCTURE OF THE AMPLIFIER

The designed amplifier has a three stage feed-forward compensated structure for the signal path, with one two-stage Miller compensated common-mode loop, one three-stage Miller compensated common-mode loop and one feed-forward compensated three-stage common-mode loop (Fig. 3). The feed-forward structure is used to have a sufficiently high loop-gain in the desired operating BW. The very high DC gain of the multi stage structure is not needed or intentionally designed for, but rather comes as a side benefit, for free. The real value of the feed-forward structure comes from the fact that it allows the designer to place two or more poles in close proximity

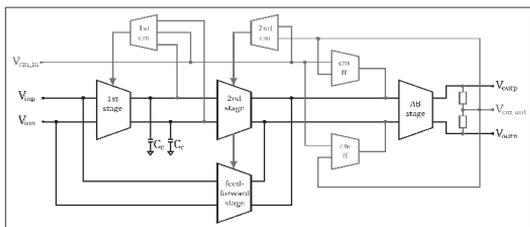


Fig. 3. Topology of the designed multiple feed-forward opamp.

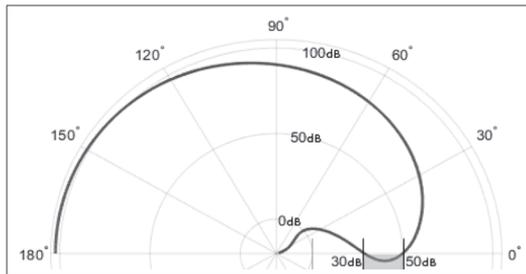


Fig. 4. Nyquist plot of the designed amplifier in buffer (0dB) mode. The amplifier has two clearly visible uncompensated poles between 30dB and 50dB and an optimum PM (~45°) at 0dB crossing.

to have a higher-order (more than 2nd order) roll-off and compensate this with one or more zeros coming from the feed-forward branch before the unity gain frequency (UGF). This can be potentially dangerous if not used properly, because we are creating a conditionally stable system, where the phase can shift easily more 180 degrees before the UGF. Even worse, the location of the UGF changes with the gain configuration of the power combiner. In buffer-mode the amplifier for example has more than 2 clearly visible uncompensated poles between 30dB and 50dB (Fig. 4). However, since it is under the control of the system at all times and cannot be programmed by the user to just about any configuration as it would be with a generic stand-alone amplifier, we can still easily stabilize the system with different predefined biasing programming currents, what shift the singularities for different gains.

IV. AMPLIFIER PROGRAMMING

The basic stability programming of the amplifier is done totally autonomous, with no extra effort from the system, with the help of the basic gain programming. The gain of the amplifier is determined with a binary weighted network of resistive ladders, so the opamp actually has all the information about the gain setting used and can determine itself the compensation needed for best stability. For that reason, the information provided by the programming of the 3MSBs of the resistive feedback is taken and re-used. This means, a stability compensation system, what divides the compensation information into eight sections, has been developed. The amplifier singularities (poles and zeros) are programmed to move together across the frequency range, ensuring in optimum phase margin (PM) for each gain setting (Fig. 4). The lowest gains also result in configuration to lowest open-loop BWs. Doing this, current consumption can be saved in low-gain configurations as shown in Fig. 5.

Setting	7	6	5-3	2	1	0
gain (dB)	-15; -14	-13; -12	...	-2; +2	+3; +5	>+6
I _q (mA)	9 + I _{DC}	10 + I _{DC}	R _M	18 + I _{DC}	21 + I _{DC}	27 + I _{DC}

Fig. 5. Eight modes of compensation for different gains. Switchover of R_M .

In order not to change the noise performance of the amplifier over different programming modes, the operating point of the first stage should not be programmed quite the same way as for other stages. For that reason, the BW of the first stage is programmed rather with programmable capacitive loads C_C , what shift the singularities coming from there without altering the noise performance associated mostly with the first stage in such structures (Fig. 3).

In addition to the autonomous stability programming, the opamp also has one more manual lever to alter performance. The function of the AB stage is not only to provide linear rail-to-rail swing but also to drive the uncertain, sometimes undefined and unknown loads – for example when under mass production test. For that reason an AB boost programming option was added, when the external load is higher than expected, a programmable biasing current I_{0*} only for the AB stage can be increased and the pole coming from the AB stage and output capacitance re-adjusted (Fig. 2b). This allows to change the capacitive driving capabilities of the combiner.

V. AB STAGE

As mentioned before, if the in-band loop gain of the amplifier is sufficiently high, the dominant part of the non-linearities from a rail-to-rail swing tends to come from the AB stage. The designed AB stage itself has a three-stage Miller-compensated common-mode regulation loop (Fig. 2b). This is the only place in the entire amplifier where high DC gain is desired and designed for, because with three gain stages in the common-mode loop, the DC gain should be sufficiently large, that the operating point is set with a high degree accuracy, so that the biasing current for the AB stage would not vary. To improve the linearity of the AB stage, the amplifier has a built in AB stage pre-distort function, what modulates the signal before reaching the AB stage to enhance the gm matching between the pMOS and nMOS output devices.

Inherently one of the most linear push-pull AB stages, where the p-, and n-channel devices have roughly equal amplification and driving strengths is the so called flating battery symmetrical push-pull structure presented in Fig. 2b [4]. The current sources I_1 and I_2 are controlled by a three-stage AB common mode loop, regulated to have an accurate and fixed DC operating point if the DC gain is sufficiently large. The current is regulated to match a reference circuit current composed of M_0 and I_0 . The ratio between M_0 and M_1 will set the current ratio between the reference circuit and the current in the AB stage. The circuit is fully differential, so one reference can regulate the DC current for both the n-, and p-side via a resistive common-mode divider composed of R_{0a} and R_{0b} . The output stage biasing is then done by controlling the potential across resistors R_1 and R_2 . Capacitors C_1 and C_2 are sized to add additional high order zeros after the desired BW but before the UGF to further improve the differential stability.

The AB stage itself incorporates a Miller compensation scheme composing of a Miller capacitor C_M and a Miller resistor R_M . This allows us to control the process variables a

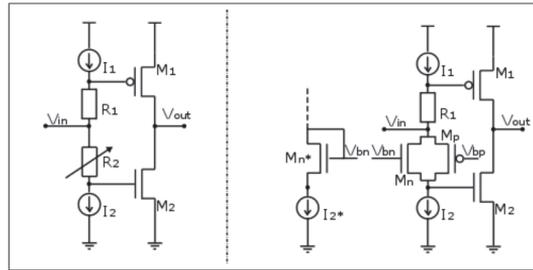


Fig. 6. AB stage pre-distortion. Basic idea and actual implementation.

lot better resulting in a more predictable and better defined UGF, what is important for a conditionally stable system. However controlling the optimum placement of the Miller zero over a large gain range is challenging, so to optimize and simplify matters, the gain range is divided into two subsections, where R_M can be halved with the gain and the Miller zero can be shifted (Fig. 2b). While lowering the frequency of the zero will add extra stability to the lower gain range, it also lowers the inherited performance, so the placement of the switchover is important and should be done with care, as late as possible, for minimum loss of performance (Fig. 5).

VI. AB-STAGE PRE-DISTORTION

In order to have a highly linear push-pull AB stage operation for a large distortion-free range driving capability, one would like to operate both the p-, and n-channel devices of the AB-stage always in similar operating conditions – meaning the gain product coming from one or the other should be similar and matched over the entire output swing. This is very difficult to achieve in full swing, because the last stage of any amplifier structure has only finite gain and its input will therefore already have a considerable amount of signal on top of it. Assuming roughly 20dB of gain from the stage, one will already have a tenth of the total opamp output swing at the input of the last stage, creating a considerable operation point swift during different stages of the swing. This is less apparent in previous gain stages, because with

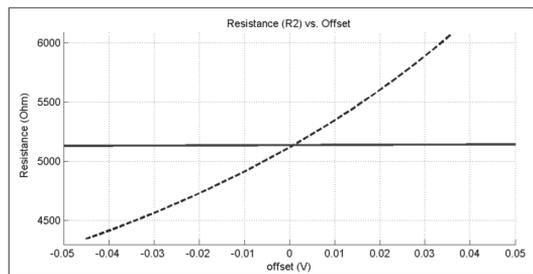


Fig. 7. Modulated resistance of R2 (Dashed line) vs original resistance.

each next stage, the signal residue on top increases. So with near rail-to-rail output swings, every AB structure will start to exhibit some non-linear tendencies, because of the input swing that changes the gms in different directions for the p-, and n-channel devices used for the AB stage. Fortunately, the distinctive and unique structures of AB stages, what are different from all the stages prior, allow us to add some wanted defects to pre-distort the signal seen by the output drivers, to help linearize the signal. For that purpose, a simple modulation scheme has been developed to pre-distort only one of the two signal paths in the AB stage, in order to compensate for the shift of gms over the full rail-to-rail operation. R_1 and R_2 together with a fixed current I_1 and I_2 are typically used in such a structure to bias the drivers M_1 and M_2 . The current sources I_2 are used to create a voltage drop across resistors R_1 and R_2 . The smaller the fixed current and resistances, the smaller the voltage drop across them. This will in return bias the output drivers M_1 and M_2 in wanted operating condition. By keeping R_1 constant and modulating R_2 , the input can be pre-distorted for only the n-side in a beneficial way, so that the total output linearity is actually enhanced (Fig. 6a). The wanted R_2 characteristic can be built with a combination of a pure resistor and output resistances of devices M_n and M_p . Total R_2 is modulated to have a signal dependent resistance by using the output resistances of M_n and M_p by adding another replica circuit composed of current source I_{2s} and M_{ns} , plus a similar topology for M_p , to properly bias the devices according to current I_2 (Fig. 6b). By doing this, any needed slope for the required modulated resistance can be acquired (Fig. 7).

VII. MEASURED RESULTS

The function of the power combiner is to level-shift, amplify and buffer the output of the DAC or to use it as a vehicle to test some other IPs across the die. It has a signal BW of about 200MHz, with a relatively smooth roll-off desired for any blocker scenarios and mask requirements. The internal programmability of the compensation keeps the performance relatively constant over the gain range seen in Fig. 8. The similarities in the slight signal peaking at band-edge over

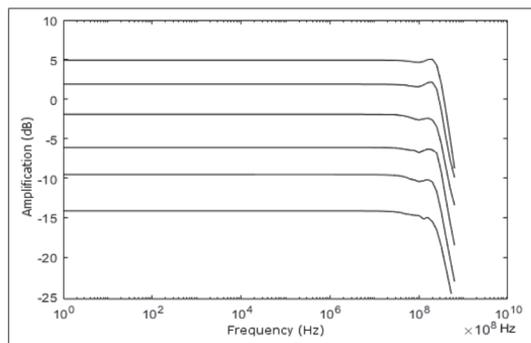


Fig. 8. Measured STF of the designed power combiner over different gains.

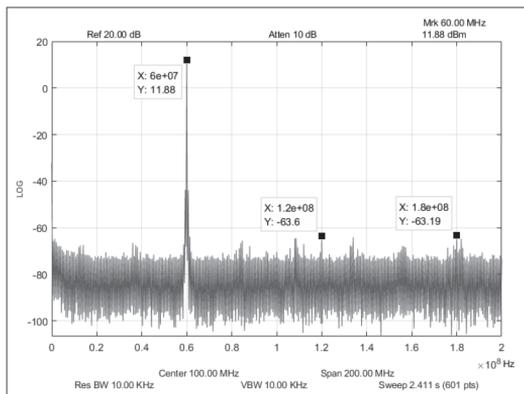


Fig. 9. Measured output spectrum of the power combiner with 5dB gain and with 12dBm output power.

different gains prove the effectiveness of the compensation programming, keeping the loop-gain constant over gain programming. It has a HD2/HD3 performance of <-75 dBc, when driving a 10pF load, with a near full-scale output, with no extra current for the AB stage boosting (Fig. 9). The power consumption of the power combiner is variable, depending from the compensation biasing schemes used, described before, starting from 9mA ($+I_{DC}$ for level-shifting – Eq. 1) for low gains (-15 dB...) and ending with 27mA ($+I_{DC}$) for the highest gain and load configurations from the a 2.5V rail. The capacitive load, the combiner is able to drive, ranges from up to 20pF when boosting capabilities for extra load driving are disabled and up to 200pF when they are enabled. Summary:

Gain	C_{LOAD}	BW	I_Q (mA)	PSRR@2MHz
± 15 dB	< 200 pF	200MHz	9 - 27	< -40 dB

VIII. CONCLUSION

A 200MHz, programmable gain power combiner with pre-modulation for enhanced linearity has been presented. It also offers level-shifting functions to $V_{DD}/2$ of the combiner, to be able to maximize the signal swing at the output. The measured results show a decent HD2/HD3 performance with even a near rail-to-rail operation, when the compensation schemes and power consumption are autonomously self-programmed.

REFERENCES

- [1] A. Mehria, "Optimum DAC resolution for WMAN, WLAN and WPAN OFDM-based standards," ICCE, Jan. 2005.
- [2] H. Pan, Y. Yao, M. Hammad, J. Tan, K. Abdelhalim, E. W. Wang, R. C. J. Hsu, D. Tam and I. Fujimori, "A Full-Duplex Line Driver for Gigabit Ethernet With Rail-to-Rail Class-AB Output Stage in 28 nm CMOS," IEEE J. Solid-State Circuits, vol. 49, no. 12, Dec. 2014.
- [3] S-C. Lee, Y-D. Jeon, J-K. Kwon and J. Kim, "A 10-bit 205-MS/s 1.0-mm² 90-nm CMOS Pipeline ADC for flat Panel Display Applications, IEEE J. Solid-State Circuits, vol. 42, no. 12, Dec. 2007.
- [4] R. Hogervorst, J. P. Tero, R. G. H. Eschauzier and J. H. Huijsing, "A Compact Power-Efficient 3 V CMOS Rail-to-Rail Input/Output Operational Amplifier for VLSI Cell Libraries," IEEE J. Solid-State Circuits, vol. 29, no. 12, Dec. 1994.

Publication V

Vahur Kampus, Martin Trojer, Robert Teschner, “Unleashing the full power of feed-forward opamps: a 200MHz, fully differential, conditionally stable, 36dB gain PGA, using a four-stage multi-path 2.5V amplifier with double feed-forward compensation”, IEEE Nordic Circuits and Systems Conference, Oct. 2018.

Unleashing the full power of feed-forward opamps: a 200MHz, fully differential, conditionally stable, 36dB gain PGA, using a four-stage multi-path 2.5V amplifier with double feed-forward compensation.

Vahur Kampus
Intel Austria GmbH
Villach, Austria 9528

Martin Trojer
Intel Austria GmbH
Villach, Austria 9528

Robert Teschner
IMST GmbH
Kamp-Lintfort, Germany 47475

Abstract—With the continuous advancement of standards in telecommunication systems, requirements for analog circuitry will become ever more demanding. The newer standards not only utilize better modulation schemes, demanding less noise with higher linearity, but also require increased bandwidth from analog circuitry. This paper describes a highly linear, fully differential, feed-forward compensated, 200MHz Programmable Gain Amplifier capable of providing DC gain from -6dB up to 24dB with less than 1dB steps. The Programmable Gain Amplifier also has an option to enable high-pass-like signal transfer characteristics, done by a channel equalizing function, providing up to 36dB of gain at bandedge, to compensate for the losses associated with the propagation medium over frequency.

Index Terms—CMOS, semiconductor, analog, opamp, PGA, multi-path, Receiver chains

I. INTRODUCTION

Standard OFDM-like signals lose about 3dB of single subcarrier strength for every doubling of the bandwidth (BW). To benefit from all the effort of increasing the BWs and maintain the same throughput per sub-carrier, analog blocks have to be designed with 3dB better noise and linearity specs for each doubling of the BW. If new standards incorporate within them new modulation schemes, these requirements can be even more challenging [1]. Programmable Gain Amplifiers (PGA) are important building blocks in modern day receiver chains. They provide the much needed flexibility for the chain to be able to enhance the Analog-to-digital converter (ADC) performance by operating the ADC in optimum back-off with the help of an automatic gain control (AGC) scheme (Fig. 1). Every ADC will have an optimum operating point, where the potential obtainable resolution is the maximum. If the signal at the input is too small, the ADC will be limited by the Signal-to-Noise ratio (SNR) of the converter itself, while when the signal is too large, the Signal-to-Distortion ratio (SDR) will be the biggest limitation. Much of finding the optimum input power comes from the quality of the AGC algorithm what controls the gain and equalization of the PGA, depending on the distance of the transmitter and the quality of the medium. Gain control is usually done in the beginning, so that the signal

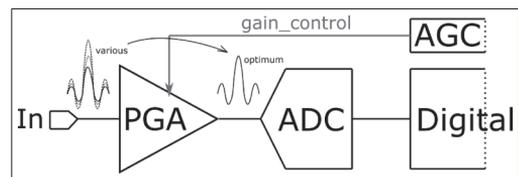


Fig. 1. Typical signal path of high-BW, high-performance receivers, where the AGC controls the gain of the PGA to have an optimum input power for the ADC.

would be amplified to the optimum value as early as possible in the chain and subsequent IPs would not further amplify noise coming from the previous building blocks. That means, the noise requirements placed on the PGA are extremely tight, while at the same time having to provide very high linearity when operating at high gain.

II. THE PGA

The proposed block is designed with 2.5V I/O devices available in standard C40 technology, what have a minimum drawable dimension of 250nm, making it from a device point of view essentially a 0.25 micron design. It is composed of a four-stage opamp, 8-bit binary weighted feedback network ($R_{unit}, R_{2*unit} \dots R_{n*unit}$) and a 6-bit binary weighted channel equalizer ($R_{eq1}C_{eq1} \dots R_{eqn}C_{eqn}$) (Fig. 2).

The feedback network is designed binary weighted controllable with a 8-bit word, to minimize the capacitive load coming from the switches to the virtual ground of the opamp, what can have a major influence at higher frequencies, where it could cause a lot of problems. The drawback of this is, that the gain step is not always uniform or even necessarily monotonic, but this risk was accepted in order to simplify the design of the opamp. When the feedback network is built in a binary way, the gain of the PGA is given by:

$$A = 2^{\frac{n-1}{gain_ctrl}} \quad (1)$$

and the gain step of the PGA is given by:

$$\Delta A = 20 * \log \frac{\text{gain_cntr}}{(\text{gain_cntr} - 1)} \quad (2)$$

where 'n' is the number of bits in the feedback network and 'gain_cntr' is the programming word used to select the on resistors in the feedback network. In the designed working range of the PGA, the minimum gain step ranges from 0.03dB@-6dB to about 1dB@24dB at DC gain.

The channel equalizer also composes of binary weighted capacitors and resistors, with the same corner frequency at around 60MHz, connected in parallel to the input resistor. The functionality of the channel equalizer is built with modulating the resistance of the unity input resistor over frequency (Fig. 3a). If the equalizer is programmed to the maximum, it will modulate the input resistor to only about one fourth, essentially creating an extra 12dBs of amplification at high frequencies (Fig. 3b). The amount of equalizer used is medium dependent and will be determined separately in every use-case in the training phase. This makes the design of the amplifier much more challenging, because not only does one have to guarantee good linearity potentially over a lot more gain, but this also fundamentally changes the unity gain frequency (UGF) location, where the stability must be determined. It does not matter that the gain is required only over high frequencies, because more often than not the DC gain is a byproduct in baseband opamps with no real practical value and the loop gain at bandedge determines the final performance [2].

The four main contributors for the non-linearities of the PGA are the finite loop gain of the amplifying circuit, the inherited non-linearity coming from the AB stage, the non-linearities of the signal path switches (S_{fb} , S_{eq}) required for the gain and input programming and the non-linearities associated with the resistors (R_{fb} , R_{in} , R_{eq}). Assuming proper care on the feedback network design – so the switches and resistors are physically sized correctly, all the dominant non-linearities will come from the amplifier [2]. This is a critical issue in modern high-speed data communications requiring full-rail swing operation for max throughput [3].

III. DESIGN CONSIDERATIONS

The designed amplifier has a four-stage double feed-forward compensated structure for the signal path, one two-stage Miller

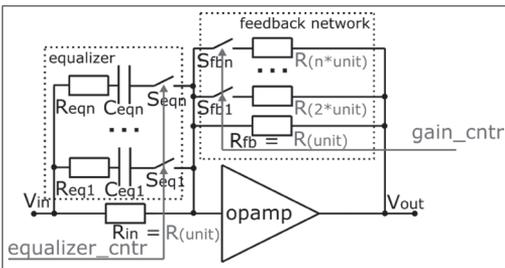


Fig. 2. Topology of the PGA, composing of an opamp, programmable equalizer and feedback network.

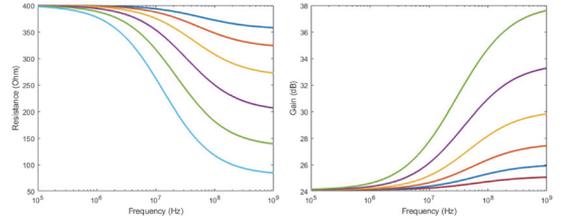


Fig. 3. a) Input resistance over frequency with different equalizer settings. b) Gain of the PGA over frequency with different equalizer settings.

compensated common-mode loop and one four-stage feed-forward compensated common-mode loop (Fig. 4). The feed-forward structure is used to have a sufficiently high loop gain in the whole desired operating BW. The design of opamps with good performance to power ratios often require a paradigm change, because on one hand, gain over the whole signal bandwidth is needed to have good linearity, but on the other hand lower UGF is wanted, to save the power used, because of the placement of non-dominant poles. The UGF in a large part determines the power consumption of the opamp, because the non-dominant poles have to be designed with margin and pushed to higher frequencies. These two aspects of efficient design could be seen as an oxymoron. The real values from the feed-forward structures come from the fact that it allows the designer to place two or more poles in close proximity to have a higher-order (more than 2nd order) roll-off and compensate this with one or more zeros coming from the feedforward branches before the unity gain frequency (UGF). This allows the building of systems where loop gain drops not at all (or very little) in the required signal BW, very fast (more than second order roll-off) between the required BW and just before the UGF and levels to a first-order (or so) roll-off around UGF (Fig. 5c). This can be potentially dangerous if not used properly, because we are creating a conditionally stable system, where the phase will by design shift more 180 degrees before the UGF [2]. Even worse, the location of the UGF changes with the gain and equalizer configurations of the PGA. The biggest intrinsic challenge of the design is to hit the UGF perfectly in all conditions and place the zeros just in front of it, no matter the configuration of the PGA itself.

IV. DESIGN STRATEGY

An n-stage opamp will have at least n high impedance nodes what will all create a pole. A commonly used design strategy, no matter the topology, has always been to place a dominant pole somewhere in low frequencies, a second pole at around three times UGF and push all the other poles out somewhere where they can be considered non-dominant or neglectable. This design strategy will lead to a clearly defined phase margin of around 60 degrees with a perfectly safe first-order roll-off before the UGF (Fig 5a) [4]. A second design strategy has been to use more complex feed-forward structures to compensate in-band poles with zeros at the same location.

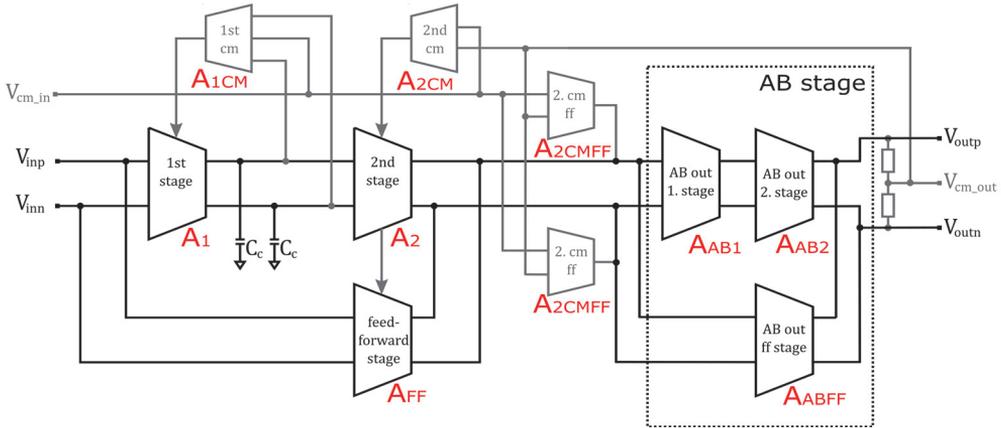


Fig. 4. Topology of the designed multiple feed-forward path opamp, where A_1 , A_2 , A_{FF} , A_{AB1} , A_{AB2} and A_{ABFF} make up the four stages plus two feed-forward stags in the differential path.

This allows to compensate all the poles exactly at the source, while creating pole-zero doublets. Sometimes those doublets do not match creating more unpredictable roll-off scenarios so the PM is less clearly defined. However both of those methods do not fully utilize the potential performance advantages of the opamps based on multi-stage topologies. To get the maximum performance obtainable, one should push the in-band poles just around bandedge, to get the maximum loop gain in-band and leave them uncompensated until just before the UGF, where as many zeros as available should be placed to regain the lost phase (Fig. 5b, 5c) [5].

V. AMPLIFIER DESIGN AND COMPENSATION

A good rule of thumb is to target 40dB of loop gain at signal bandedge for more than 12bit linearity. At the same time the UGF should stay in as low frequencies as possible for power consumption reasons. When designing a first-order roll-off opamp, with 40dB of loop gain at signal BW, the UGF would be at around $100 \cdot BW$, what can be unrealistic for high BW amplifiers, so higher roll-offs are needed. The designed opamp has four gain stages A_1 , A_2 , A_{AB1} and A_{AB2} to achieve high loop gain over the entire signal BW. At DC and low frequencies the gain from these four stages is dominant and provides the gain and linearity. Feed-forward stages A_{FF} and A_{ABFF} can be seen just as load what lower the total

output impedance at the outputs of the second stage and at the output, lowering therefore also somewhat the gain from that stage (Fig. 6). The feed-forward stages become dominant once the gain of the cascaded stages drops below the gain of the single feed-forward stage itself. $A_{FF} > A_1 \cdot A_2$ and $A_{ABFF} > A_{AB1} \cdot A_{AB2}$ (Fig. 4). At frequencies close to UGF, the opamp is basically a two stage amplifier composing of stages A_{FF} and A_{ABFF} . The pole placement becomes extremely critical in these kinds of systems. Poles from A_1 , A_2 , A_{FF} , A_{AB1} and A_{AB2} should be placed just after the signal BW, as close together as possible, to have the maximum positive effect and roll-off. Stages A_2 and A_{FF} are designed to be similar and to have the same corner frequency. This does not have to be the case, but simplifies layout matching considerations and common-mode controlling scenarios, as they now can be controlled with a single stage A_{2CM} (Fig. 7). This means that transition from medium to high frequencies coming from the first and second stage being dominant to the feed-forward stage being dominant, will result in a change from second-order roll-off to a first order roll-off in the first part of the amplifier. In the AB stage however the feed-forward stage A_{ABFF} is designed to have a real non-dominant pole in a much higher frequency, so that it could provide a double zero very close to UGF. When crossing the 0dB gain point, the system can behave much like a single pole amplifier, with the single dominant pole coming from A_{FF} . This means the only current hungry stage is the A_{ABFF} , what in reality will consume close to half of the total current of the amplifier.

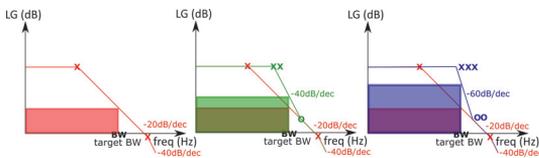


Fig. 5. Potential performance of a similar GBW amplifiers a) single pole b) dual pole and c) triple pole systems

The amplifier also has two common-mode loops. The first one is a two stage Miller compensated loop composing of A_{1CM} and A_1 . The dominant pole in this loop is placed in the first gain stage of the common-mode loop A_{1CM} , so that the second pole coming from the pole-placement described previously would be at least in the distance of $3 \cdot UGF$, much like in the first simple design strategy described earlier, to have

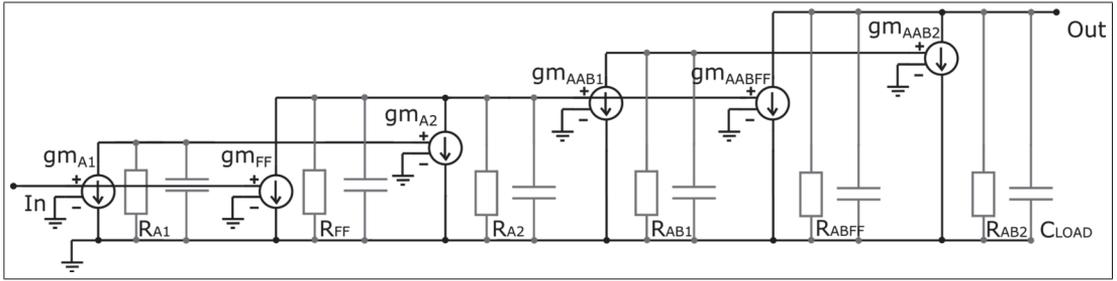


Fig. 6. Small signal equivalent circuit of the designed opamp, where the second stage A_2 shares the output resistance with the feed-forward stage A_{FF}

a PM of better than 60 degrees. The second common-mode loop however is very similar to the differential path with a four-stage double feed-forward compensated loop, where the four gain stages are A_{2CM} , A_2 , A_{AB1} and A_{AB2} and the two feed-forward paths are A_{2CMFF} and A_{ABFF} . Much like in the case of the differential path if A_{2CM} has a pole similar to A_1 and A_{2CMFF} a pole similar to A_{FF} , the system will become a two-stage-like opamp close to UGF, with a similar GBW as in the differential path. Now A_{2CMFF} will act like a dominant pole with A_{ABFF} still being a non-dominant one.

VI. PROGRAMMING OF THE AMPLIFIER

The PGA must be stable and perform with gains from -6dB to 24dB, with any combination of the equalizer enabled. To place the poles and zeros correctly over such a large set of configurations is quite a challenging feat. Not to over-design for one setting, to cover everything with one internal configuration, the amplifier has been done internally configurable with four different modes dedicated from '0' to '3'. These modes set the biasing currents for different stages (Fig. 7). When mode '0' is selected, all current sources (dashed) will be enabled, while with setting '3' only the solid one is selected. This means the opamp has a current dependency on the configuration, meaning that setting '0', what should be used, when the gain is high, will also consume the most current. With such programming the gain range can be divided into four sections what could cover the entire 30dB span. In order not to change the noise performance of the amplifier

over different programming modes, the operating point of the first stage should not be programmed quite the same way as for other stages. For that reason, the BW of the first stage is programmed rather with programmable capacitive loads C_c , what shift the singularities coming from there without altering the noise performance of the opamp associated mostly with the first stage in such feed-forward structures (Fig. 4). The last thing that can also be programmed is the AB stage current. This is done in two scenarios. First, when maximum performance is required with full amplification. Increasing the AB stage current can significantly improve the linearity. Secondly, when the opamp needs to drive a small resistive load, when the load decreases in the resistive feedback. Therefore, in modes '0' and '3', when the opamp is configured either to amplify a lot or very little, the AB stage current is programmed to higher values and in middle codes '1' and '2', where the opamp is configured to operate in middle amplifications, the AB stage current is programmed to lower values (Fig. 8a).

VII. SIMULATIONS AND MEASUREMENTS

The PGA has been designed as a part of a receiver chain, so direct performance measurements are impossible. However, since the total chain measurements are in line with the expected values, it is likely that simulations fit very well to real life. The biggest emphasis has been on retaining good noise performance in all configuration and modes, as any noise created in the beginning of the chain, before bringing it to full scale would degrade the total chain performance a lot. Linearity has been designed to be around the 12bit level, over all gains, with multi-tone patterns over the 200Mhz BW. High BW converter chains generally require good voltage regulation roughly over the interested BW or higher, not to suffer from the regulation imperfections of the regulator [6]. The designed amplifier further enhances the Power insensitivity by having a Power Supply Rejection Ratio (PSRR) that has been designed specially for general purpose voltage converters with switching speeds up to 2MHz, where the PSRR will be guaranteed as below -40dB and regulators with limited BW, so that the PSRR would never creep over -6dB, even at high frequencies.

The partitioning of the compensation over four parts creates four distinguishable zones where the performance and stability make a 'jump' and are non-continuous. It is less obvious

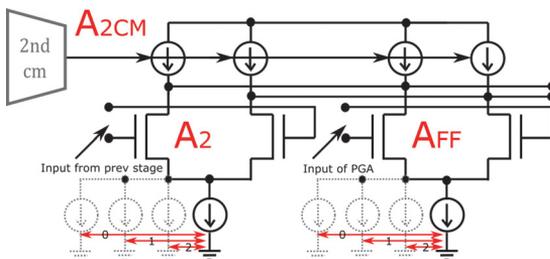


Fig. 7. The second stage and the feed-forward stage can be controlled with the same common mode regulation loop if made similar.

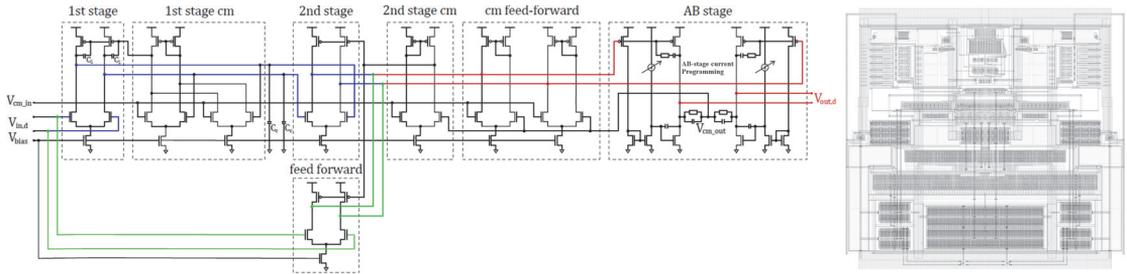


Fig. 8. a) Structure of the designed four-stage multi-path amplifier, with AB-stage programming b) Layout of the designed opamp, built with 0.25um devices.

TABLE I
PERFORMANCE SUMMARY

Mode gain (dB)	I _q (mA)	PSRR @2M (dB)	noise (nV/√Hz)	linearity (ENOB)
(0) 20...●	21	<-45	<1.6	>11.5
(1) 11...19	17	<-45	<1.8	>12.0
(2) 3...10	13	<-40	<2.0	>12.0
(3) ●...2	12	<-40	<2.2	>12.0

Mode gain(dB)	PM (°)	PM _{CM1} (°)	PM _{CM2} (°)
(0) 20...●	>25	>60	>45
(1) 11...19	>25	>60	>45
(2) 3...10	>25	>60	>45
(3) ●...2	>25	>60	>45

in the first common-mode loop, where the PM is designed to exceed 60 degrees for all cases, but more visible in the second common-mode loop and in the differential loop. When designed correctly, stability of the differential loop and the second common-mode loop path track each other in a large part, because the path that the loop follows is common over several sections. However, the common mode stability is usually more vital and BW less critical, then in the differential mode. Therefore it has been designed to have roughly 15 to 20 degrees more PM, as in the differential loop, with the cost of some extra BW. These four zones are built to optimize PM performance, so that the zeroes placed in front of the UGF could have the biggest potential impact. Once the zeros are

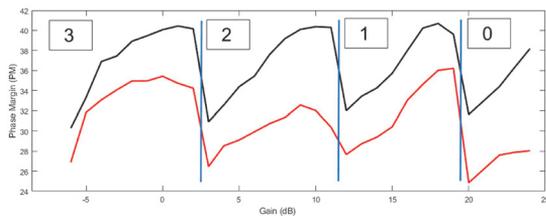


Fig. 9. Nominal and Minimum PM of the differential path of the PGA over different gain settings, with no equalizer used, where the PM was the lowest.

placed too far from the UGF, PM will start to drop compared to the previous gain setting. That's why the PM acts like a half circle over the gain programming (Fig. 9). Changing gains from 2dB gain to 3dB gain, from 11dB gain to 12dB gain and from 19dB gain to 20dB gain will bring with it an internal configuration change that will increase performance, but lower the PM. The same is true for the second common-mode loop, where the PM however is somewhat higher. Minimum performance summary of extractions is presented in Table 1.

VIII. CONCLUSION

A 200MHz BW, four-stage, feed-forward compensated PGA has been designed and presented. It utilizes a design strategy where the all the four primary poles have been placed relatively close together and then later compensated with 2+ zeroes. This allows to have very high loop gain at signal BW without having to push the UGF into extreme frequencies, what in return allows to save power by not having to design many stages with non-dominant poles. To cover a large gain range, the PGA has been designed to have an internally configurable compensation. Measurements show a good correlation with simulations, specially with high gains, where excellent noise performance makes it possible for the ADC to capture even low power signals over long loops.

REFERENCES

- [1] A. Mehmia, "Optimum DAC resolution for WMAN, WLAN and WPAN OFDM-based standards," ICCE, Jan. 2005.
- [2] V. Kampus, T. Rang, D. Knaller, C. Fleischhacker, M. Korak, J. Kiss, "A fully differential, 200MHz, programmable gain, level-shifting, hybrid amplifier/power combiner/test buffer, using pre-distortion for enhanced linearity," 14th Conference on PRIME, Jul. 2018.
- [3] S-C. Lee, Y-D. Jeon, J-K. Kwon and J. Kim, "A 10-bit 205-MS/s 1.0-mm² 90-nm CMOS Pipeline ADC for flat Panel Display Applications," IEEE J. Solid-State Circuits, vol. 42, no. 12, Dec. 2007.
- [4] X. Li, X. Le Cui, B. Wang and C. Lee, "A 100MHz PGA with DC offset cancellation for UWB receiver", 2012 IEEE 11th International Conference on Solid-State and Integrated Circuit Technology, Feb. 2012.
- [5] V. Kampus, T. Rang, "A smart capless voltage regulator for very high bandwidth A/D and D/A converters in a standard 28nm CMOS process," 15th Biennial Baltic Electronics Conference (BEC), Oct 2016.
- [6] M. Fulde, A. Belitzer, Z. Boos, M. Bruennert, J. Fritzin, H. Geltinger, M. Groinig, D. Gruber, S. Gruenberger, T. Hartig, V. Kampus, B. Kapfelsberger, F. Kuttner, S. Leuschner, T. Maletz, A. Menkhoff, J. Moreira, A. Paussa, D. Ponton, H. Pretl, D. Sira, U. Steinacker, N. Stevanovic "A Digital Multimode Polar Transmitter Supporting 40MHz LTE Carrier Aggregation in 28nm CMOS," ISSCC 2017, Feb. 2017

Publication VI

Vahur Kampus, Robert Teschner, Ulrich Gaier, Thomas Linder, Gerhard Nössing, Martin Trojer, “Five-Stage, Power Efficient, Dual Rail, 100MHz, 10dB Programmable Gain Amplifier with Down-Stepping Functions in 28nm CMOS”, IEEE International Symposium on Circuits and Systems, May 2019.

Five-stage, power efficient, dual rail, 100MHz, 10dB Programmable Gain Amplifier with down-stepping functions in 28nm CMOS.

Vahur Kampus	Robert Teschner	Ulrich Gaier	Thomas Linder	Gerhard Nössing	Martin Trojer
<i>Intel Austria GmbH</i>	<i>IMST GmbH</i>	<i>Intel Austria GmbH</i>	<i>Intel Austria GmbH</i>	<i>Intel Austria GmbH</i>	<i>Intel Austria GmbH</i>
Villach	Kamp-Lintfort	Villach	Villach	Villach	Villach
Austria, 9528	Germany, 47475	Austria, 9528	Austria, 9528	Austria, 9528	Austria, 9528

Abstract—With the continuous advancement of standards in telecommunication systems, requirements for analog circuitry become ever more demanding. The newer standards not only utilize better modulation schemes, demanding less noise with higher linearity, but also require increased bandwidth from analog circuitry. Together with performance, one of the most important key parameters is high efficiency. This paper describes a highly efficient, fully differential, five-stage, double feed-forward compensated, down stepping Programmable Gain Amplifier capable of providing DC gain from -12dB up to 10dB with 1dB steps for signal bandwidths of up to 100MHz. It consumes a total of 2.5mW-4mW of power from a split 1.8V/0.9V supply, while delivering a 1.2Vpp into a 1k Ω on-chip load. The PGA can support a maximum differential input swing of up to 5Vpp. The split architecture enables to run the first two gain stages of the amplifier from a 1.8V power rail and the latter three stages from a 0.9V power rail.

Index Terms—CMOS, semiconductor, analog, opamp, down-shifting, dual rail, PGA, multi-path, receiver chain

I. INTRODUCTION

Programmable Gain Amplifiers (PGA) are important building blocks in modern day receiver chains. They provide the much needed flexibility for the chain to be able to maximize the Analog-to-Digital Converter (ADC) dynamic range by operating the ADC in optimum back-off [1]. This is usually done with the help of a smart automatic gain control (AGC) scheme, what detects the power and adjusts the PGA gain accordingly [2]. Receiver chains always need to include some type of ADC for digitization and it is usually desirable to design the converters with core voltage devices in core voltage domains without extra stacking to reach best performance and increase overall efficiency. Every ADC will have an optimum operating point, where the potential obtainable resolution is the maximum. If the signal at the input of the converter is too small, the ADC will be limited by the Signal-to-Noise ratio (SNR) of the chain, while when the signal is too large, the Signal-to-Distortion ratio (SDR) will start to degrade. Input/Output transmission signals are often in domains not suited for core voltage devices and require level-shifting and signal strength optimization before reaching the converter. Analog front-ends in general have to deal with very different input common-modes and signal levels that can be in a higher

voltage domains with a larger common-mode voltage (V_{cm1}) then directly usable for the converter and may therefore require down-shifting to a more useful common-mode voltage (V_{cm2}) before reaching the ADC (Fig. 1). Operating analog front-ends that actually do the down-stepping, like PGAs, from a single power rail is inefficient by design, since the power rail required needs to correspond to the highest voltage domain used and therefore is over-proportionally inefficient [3].

Down-shifting could be done in several ways. One option is to add a level-shifting current mode Digital-to-Analog Converter (DAC) before the PGA (Fig 2a). This can accomplish the task, but needs to have a sufficiently large current in the DAC to have a reasonably good noise figure, before reaching the amplifier, wasting large amounts of power. It also requires the transmitting path to be able to source all the extra current, creating unnecessary origins of distortion. The second option is to capacitively decouple the output of the PGA followed by a buffer done in the ADC domain (Fig 2b). This option needs a very big decoupling capacitor to have low cut-off frequency, what sometimes could be not feasible. It also requires two separate active amplifiers, one in the I/O domain and one in the core domain. The third option is to do the down-shifting directly with the PGA itself, running it completely from the I/O domain (Fig 2c). This option needs the functionality to sink some extra current from the AB-stage and source the same amount from the transmitting circuit [4].

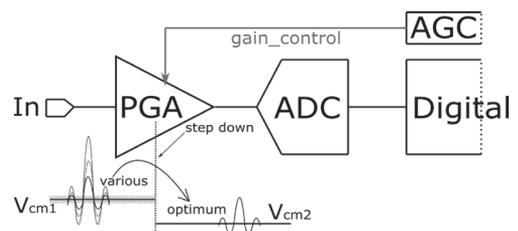


Fig. 1. Typical signal path of high-performance receivers, where the AGC controls the gain of the PGA to have an optimum input power for the ADC.

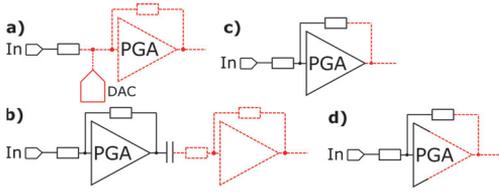


Fig. 2. Possible solutions for down-stepping the signal levels. Core voltage domain is drawn in dashed red and I/O voltage domain in solid black.

II. ARCHITECTURE OF THE PGA

All the solutions discussed waste substantial amounts of current to do the level-shifting, making them potentially very inefficient. To improve efficiency while at the same time keeping the possible origin of distortions to a minimum, a better solution in a form of a five-stage, dual rail, feed-forward compensated PGA, that does the down-stepping is proposed and designed in a standard 28nm CMOS process (Fig 2d). The value from a feed-forward structure comes from the fact that it allows the designer to place two or more poles in close proximity to have a higher-order (more than 2nd order) roll-off and compensate this with one or more zeros coming from the feed-forward branches before the unity gain frequency (UGF). This allows the building of systems where loop gain drops not much in the required signal BW, very fast (more than second order roll-off) between the required BW and just before the UGF and levels to around first-order roll-off at UGF [5].

A. Design Considerations

The PGA utilizes a thick-oxide device option provided by most modern technologies for the first two stages and thin-oxide core devices for the latter three stages and it has a built-in double common-mode down-stepping function. The first down-step is at the edge of the switchover from the thick-oxide devices to the core voltage devices (Fig. 3a ●2) and the second down-step is from the input of the thin-oxide devices to the output of the amplifier (Fig. 3a ●5). To do the step-down in a single step from the input of the amplifier to the device switch-over turned out to be too large and not doable, so a two-step solution was chosen. Signal path stages A_1 and A_2 plus common-mode regulation stage A_{1CM} are built with 1.8V thick-oxide devices and powered from a 1.8V rail. The first

stage A_1 uses self-biasing for the output. This is done because the common-mode voltage at the input of the first stage is very dependent on the gain setting to what the amplifier is configured to. The common-mode voltage seen by the first stage is given with:

$$V_{cm0} = \frac{V_{cmi} - V_{cmo}}{1 + \frac{1}{10^{20}A}} + V_{cmo} \quad (1)$$

where V_{cm0} is the common-mode at the input of the first stage, V_{cmi} the common mode of the PGA at the input, V_{cmo} the common mode of the PGA at the output and A the gain configuration of the PGA. Assuming half-VDD common-mode voltage at the output ($\sim 0.45V$) and something between 1V and 1.8V at the input of the PGA, different gains can bring about common-mode voltages between 0.6V and 1.45V at the input of the amplifier. The first self-biasing amplifying stage therefore also provides a much needed leveling function inside the amplifier, where it takes a variable common-mode voltage at the input of the amplifier and translates this into a more equal voltage of around 1V (depending on the process corner) at the input of the second stage (Fig. 3a ●1). Having the first stage self-biased also allows us to simplify the first common-mode regulation loop to an only two-stage amplifier composing of A_{1CM} and A_2 . The common-mode is regulated at this point to 0.65V (Fig. 3a ●2) at the input of the third stage, from where the rest of the amplifier is built with thin-oxide devices powered from a 0.9V rail. Stages A_3 and A_{FF} are designed to be similar and to have the same corner frequency. This does not have to be the case, but simplifies layout matching considerations and common-mode controlling techniques, as they now can be controlled with a single stage A_{2CM} (Fig. 3b) [5]. This means that the input of the feed-forward stage A_{FF} should also be built with thin-oxide devices and needs to be capacitively decoupled and set to a usable common-mode voltage, as the input common-mode voltage of the amplifier would be far too high. The common-mode of the feed-forward stage is set after the decoupling via resistors R_d to the same common-mode (0.65V) as the input of the third stage of the amplifier is regulated to (Fig. 3a ●2a). The RC-constant of the decoupling capacitor C_d and common-mode resistor R_d needs to be high enough not to interfere in-band of the required BW but low enough to provide a zero before UGF coming from the feed-forward stage A_{FF} itself.

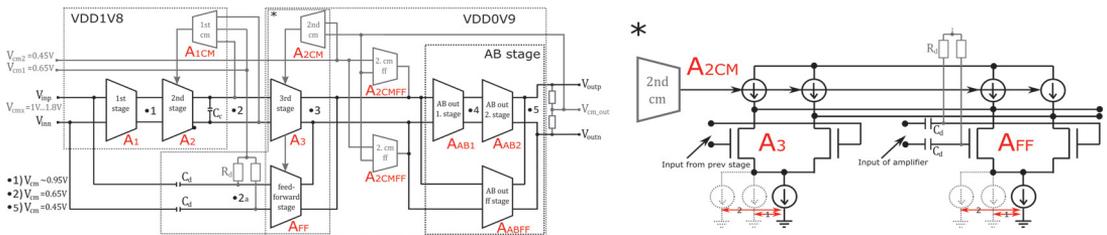


Fig. 3. a) Topology of the five-stage feed-forward amplifier. b) Shared common-mode regulation loop of Stages A_3 and A_{FF} what requires the decoupling.

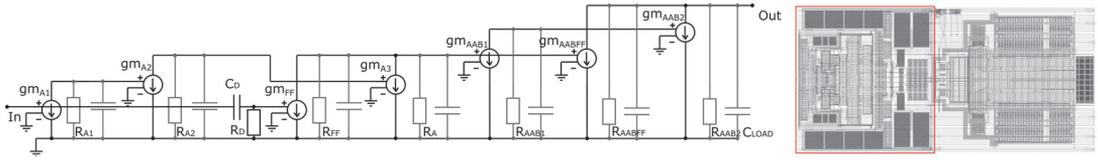


Fig. 4. a) Small signal equivalent circuit of the designed five-stage opamp, where the feed-forward stage A_{FF} is capacitively decoupled from the input and the third stage A_3 shares the output resistance with the feed-forward stage A_{FF} . b) Layout of the amplifier (marked in red) with-in the entire PGA.

B. Noise and Stability Considerations

Using such five-stage topology allows the feed-forward stage A_{FF} theoretically provide a double-zero for the amplifier. If stages A_1 , A_2 and A_3 have dominant poles, one could create a third-order roll-off in-band and compensate this with a double-zero coming from A_{FF} [5]. This was not implemented however, since the noise performance of the amplifier follows the dominant gain path, meaning the dominant noise source comes from the first stage of the signal path where the gain is the highest. From low frequencies up to signal BW, the dominant noise source is stage A_1 , while at frequencies above the signal BW (where the gain of the feed-forward path A_{FF} is higher than the total gain of A_1 , A_2 and A_3) the feed-forward stage A_{FF} will become dominant. If good noise performance is wanted in-band, stage A_1 has to be designed with a relatively large biasing current, pushing the pole frequency to a non-dominant range, reducing the maximum achievable roll-off from the first three stages to a second order, coming from stages A_2 and A_3 (Fig. 5). As also the very large gain of A_1 , A_2 and A_3 becomes a problem, the pole from A_2 has been made very dominant and pushed to around 1MHz with a loadcap C_C (Fig. 3a). However, also with that, getting rid of all the extra gain from five stages is not possible, when placing all the poles just out of band, so poles from A_3 and A_4 have been brought a decade in-band and not on the band-edge as ideally wished (Fig. 5). Out-of-band noise performance is dominated by the feed-forward stage A_{FF} and the biasing resistor R_d . This can be potentially important and should be taken into account for non-oversampling converters where the noise could be folded back in-band.

C. The Five-stage Amplifier

The designed opamp has five gain stages A_1 , A_2 , A_3 , A_{AB1} and A_{AB2} to achieve high loop gain over the entire signal BW and do the down-stepping. All stages offer around 20dB of gain while having singularities strategically placed to achieve the performance required and maximum stability (Fig. 5). At DC and low frequencies the gain from these five stages is dominant and provides the gain and linearity. Feed-forward stages A_{FF} and A_{ABFF} can be seen just as load what lower the total output impedance at the outputs of the third stage and at the output, lowering therefore also somewhat the gain from that stage (Fig. 4a). The feed-forward stages become dominant once the gain of the cascaded stages drops below the gain of the single feed-forward stage itself: $A_{FF} > A_1 * A_2 * A_3$ and $A_{ABFF} > A_{AB1} * A_{AB2}$ (Fig. 3a). At frequencies close to UGF,

the opamp is basically a two stage amplifier composing of stages A_{FF} and A_{ABFF} . The pole placement becomes extremely critical in these kinds of systems. Poles from A_2 , A_3 , A_{FF} , A_{AB1} and A_{AB2} should be placed just after the signal BW (or in this case just a little bit inside), as close together as possible, to have the maximum positive effect and roll-off while the pole from A_1 needs to stay non-dominant to achieve good noise performance in-band. Since stages A_3 and A_{FF} are designed to be similar and to have the same corner frequency, the transition from medium to high frequencies coming from the first, second and third stage being dominant to the feed-forward stage being dominant, will result in a change from second-order roll-off to a first order roll-off in the first part of the amplifier. In the AB stage however the feed-forward stage A_{ABFF} is designed to have a real non-dominant pole in a much higher frequency, so that it could provide a double zero very close to UGF. When crossing the 0dB gain point, the system can behave much like a single pole amplifier, with the single dominant pole coming from A_{FF} . This means, the only current hungry non-dominant pole stage of the amplifier is A_{ABFF} , what in reality will consume close to half of the total current. Compensating higher roll-off amplifiers with feed-forward stages is always challenging as the UGF location changes with the gain configuration. A three mode internal compensation has therefore been used, where the pole and zero locations are shifted automatically with the stage basing based on the PGA gain (Fig. 3b) [5]. This means also that in lower gain modes the PGA will use less power.

The amplifier also has two common-mode loops. The first one is a two stage Miller compensated loop composing of stages A_{1CM} and A_2 . The dominant pole in this loop is placed in the first gain stage of the common-mode loop A_{1CM} , so that the second pole coming from the pole-placement described

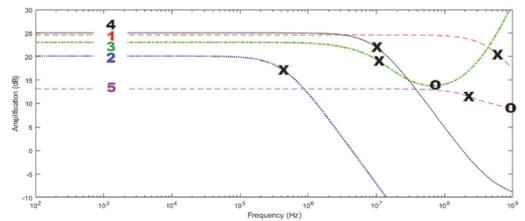


Fig. 5. Amplification of signal over frequency as seen by different stages. X mark the location of poles and O the approx location of feed-forward zeroes.

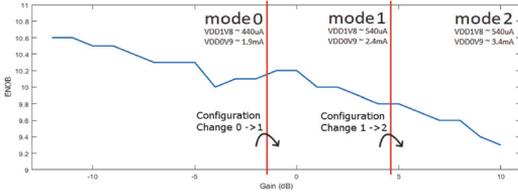


Fig. 6. SNR of the designed PGA over different gains, with the conf. changes.

previously in the signal path, would be at least around UGF, to have a Phase Margin of better than 45 degrees. Ideally one would like to push A_{1CM} to such low frequencies, that the pole coming from A_2 would be at $3 \cdot UGF$ [5], but since the pole coming from A_2 was intentionally pushed to be very dominant in the signal path, the phase requirements for the first common-mode loop were relaxed. The second common-mode loop is very similar to the differential path with a four-stage double feed-forward compensated loop, where the four gain stages are A_{2CM} , A_3 , A_{AB1} and A_{AB2} and the two feed-forward paths are A_{2CMFF} and A_{ABFF} . Much like in the case of the differential path if A_{2CM} has a pole similar to A_2 and A_{2CMFF} a pole similar to A_{FF} , the system will become a two-stage-like opamp close to UGF, with a similar gain bandwidth as in the differential path. Now A_{2CMFF} will act like a dominant pole and A_{ABFF} will still end up being non-dominant.

III. MEASUREMENTS AND GENERAL PERFORMANCE

The PGA has been manufactured and tested over 11 different skew lots with a maximum 2-sigma target offset. All silicon flavors show very similar performance. The SNR of the designed PGA over different gains is presented in Fig. 6. The dominant noise sources of the PGA however are the input and feedback resistors, what are relatively large to save power. The amplifier itself has a much better noise figure, what is a deviation from the normal approach, where the amplifier is designed as the dominant noise source and the feedback network is over-designed. The noise performance of the amplifier itself without the feedback network would be in the 11ENOB range. This also means that the distortion of

TABLE I
PERFORMANCE SUMMARY

Mode gain (dB)	$I_q@1.8V$ (uA)	$I_q@0.9V$ (mA)	SNR (ENOB)	linearity (ENOB)	PSRR @>2M (dB)
(2) 5..10	540	3.40	>9.30	>11.5	>10
(1) -1..4	540	2.40	>9.80	>12.0	>10
(0) ●..-2	440	1.90	>10.3	>12.0	>10
Mode gain (dB)	PSRR@1.8V @<2M (dB)	PSRR@0.9V @<2M (dB)	PM (°)	PM _{CM1} (°)	PM _{CM2} (°)
(2) 5..10	>80	>50	>45	>40	>45
(1) -1..4	>80	>50	>45	>45	>45
(0) ●..-2	>80	>50	>45	>45	>45

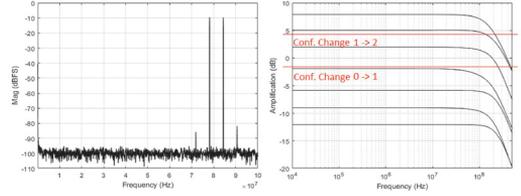


Fig. 7. a) Output of the PGA with a dual-tone input. b) STF of over gain.

TABLE II
COMPARISON TABLE

Specification	ESSCIRC'12 [3]	NORCAS'18 [5]	This Work
Technology	65nm	40nm	28nm
BandWidth	30MHz	200MHz	100MHz
Power Cons.	25mW	25mW-50mW	2.5mW-4mW
CM Shifting	No	No	Yes
Input ref. SNR	14ENOB	12ENOB	10ENOB
Max. Gain	27dB	36dB	10dB

the amplifier is not dominant ($>12ENOB$) and neglectable in overall SNDR (Fig. 7a). One of the most important aspects of any system with different domains is the Power Supply Rejection Ratio (PSRR) from both rails. Since the target the PGA was to be directly supplied from a dirty source (DCDC), a good PSRR rejection was a major requirement. The PSRR up to 2MHz is designed to be better than 80dB from the 1.8V power rail and better than 50dB from the 0.9V power rail and better than 10dB in both cases for frequencies higher than 2MHz. A short performance summary is presented in Table I.

IV. CONCLUSION

A 100MHz BW, five-stage, down-stepping, feed-forward compensated PGA has been designed and manufactured. It utilizes a design strategy where the primary poles have been placed relatively close together and then later compensated with 2+ zeros. This allows to have very high loop gain at signal BW without having to push the UGF into extreme frequencies, what in return allows to save power by not having to design many stages with non-dominant poles. The first two stages of the amplifier are designed with thick-oxide devices and the other three with core devices. To use a multi-path compensation in the signal path, the feed-forward stage has been capacitively decoupled and re-set to a new common-mode to be able to use the same common-mode controlling techniques to both feed-forward and third stages. To cover a 20dB gain range with systems that can have two or more uncompensated poles, the PGA has been designed to have an internally adjustable compensation with three different modes what are automatically selected with different PGA gains.

REFERENCES

- [1] C. C. Hsu, J. T. Wu, "A highly linear 125-MHz CMOS switched-resistor programmable-gain amplifier," IEEE Journal of Solid-State Circuits, Vol. 38, No. 10, pp. 1663 - 1670, Oct. 2003.

- [2] Xiaoman Wang, Baoyong Chi, Zhihua Wang, "A Low-Power High-Data-Rate ASK IF Receiver With a Digital-Control AGC Loop," IEEE Transactions on Circuits and Systems, pp. 617 - 621, Aug. 2010.
- [3] D. Gruber, M. Clara, B. Seger, "A High-Input-Swing Common-Mode-Robust Programmable Gain Amplifier in 65nm CMOS," 2012 Proceedings of the ESSCIRC, Sep. 2012.
- [4] V. Kampus, T. Rang, D. Knaller, C. Fleischhacker, M. Korak, J. Kiss, "A fully differential, 200MHz, programmable gain, level-shifting, hybrid amplifier/power combiner/test buffer, using pre-distortion for enhanced linearity," 14th Conference on PRIME, Jul. 2018.
- [5] V. Kampus, M. Trojer, R. Teschner, "Unleashing the full power of feed-forward opamps: a 200MHz, fully differential, conditionally stable, 36dB gain PGA, using a four-stage multi-path 2.5V amplifier with double feed-forward compensation," 2018 IEEE NorCAS, Oct. 2018.

Publication VII

Leidy Mabel Alvero-Gonzalez, Eric Gutierrez, Luis Hernandez, **Vahur Kampus**, Viktor Medina, Susana Paton, "Ring-Oscillator with Multiple Transconductors for Linear Analog-to-Digital Conversion", MDPI Electronics Journal, Jun. 2021.



Article

Ring-Oscillator with Multiple Transconductors for Linear Analog-to-Digital Conversion

Leidy Mabel Alvero-Gonzalez ^{1,*}, Victor Medina ¹, Vahur Kampus ^{2,3}, Susana Paton ¹ and Luis Hernandez ¹
and Eric Gutierrez ^{1,†}

¹ Electronics Technology Department, Carlos III University, 28903 Madrid, Spain; vmedina@ing.uc3m.es (V.M.); spaton@ing.uc3m.es (S.P.); luish@ing.uc3m.es (L.H.); eric.gutierrez@uc3m.es (E.G.)

² MaxLinear Austria, 9524 Villach, Austria; vkampus@maxlinear.com

³ Thomas Johann Seebeck Department of Electronics, Tallinn University of Technology, 12616 Tallinn, Estonia

* Correspondence: lalvero@ing.uc3m.es

† These authors contributed equally to this work.

Abstract: This paper proposes a new circuit-based approach to mitigate nonlinearity in open-loop ring-oscillator-based analog-to-digital converters (ADCs). The approach consists of driving a current-controlled oscillator (CCO) with several transconductors connected in parallel with different bias conditions. The current injected into the oscillator can then be properly sized to linearize the oscillator, performing the inverse current-to-frequency function. To evaluate the approach, a circuit example has been designed in a 65-nm CMOS process, leading to a more than 3-ENOB enhancement in simulation for a high-swing differential input voltage signal of 800-mV_{pp}, with considerable less complex design and lower power and expected area in comparison to state-of-the-art circuit based solutions. The architecture has also been checked against PVT and mismatch variations, proving to be highly robust, requiring only very simple calibration techniques. The solution is especially suitable for high-bandwidth (tens of MHz) medium-resolution applications (10–12 ENOBs), such as 5G or Internet-of-Things (IoT) devices.

Keywords: data conversion; frequency modulation; voltage-controlled oscillator; linearization techniques; 5G; IoT



Citation: Alvero-Gonzalez, L.M.; Medina, V.; Kampus, V.; Paton, S.; Hernandez, L.; Gutierrez, E. Ring-Oscillator with Multiple Transconductors for Linear Analog-to-Digital Conversion. *Electronics* **2021**, *10*, 1408. <https://doi.org/10.3390/electronics10121408>

Academic Editor: Lodovico Ratti

Received: 29 April 2021

Accepted: 10 June 2021

Published: 11 June 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The scaling down of CMOS processes has posed new challenges for analog-to-digital conversion, especially from the analog design perspective. Digital logic consumes less power, occupies less area, and works faster as design processes get smaller. Nevertheless, analog designs have become highly complex due to the low voltage supply and limited devices' gain, higher noise impact, mismatch, and parasitic effects [1]. Thus, the current trend is towards mostly digital implementations.

Energy-efficient wide-band ADCs are essential for applications such as portable battery-powered devices or radio-receivers. Flash ADCs are usually implemented for high-speed analog-to-digital conversion [2]. However, the power consumption increases exponentially with the number of bits, making them less power-efficient for more than 10 ENOBs [3]. Additionally, they are extremely sensitive to mismatch phenomena when implemented with minimum-size devices and not scalable. To minimize this impact, devices' dimensions are often increased at the cost of larger occupied area and higher power consumption [4,5].

SAR (Successive Approximation Register) ADCs are low-power, simple, and friendly-digital architectures. They are more energy-efficient than Flash-based architectures due to their mostly-digital implementation suitable with modern processes nodes [6,7]. Nevertheless, their sampling rate is limited by the need of a high-speed clock for the control

logic [2], limiting their use for medium bandwidth and leading to dramatically high-power solutions when used for MHz-bandwidths [8–11].

Pipeline ADCs divide the analog-to-digital high-resolution conversion operation into several low-resolution conversion stages operating sequentially, showing a proper balance between accuracy and operation speed. The drawback of this kind of ADCs is the high-performance operational amplifier (opamp) mandated by a multiplying digital-to-analog converter (MDAC) [2]. This block consumes a large amount of power, reducing the power efficiency, and entails challenging designs in deep submicron CMOS nodes [12,13].

Time-interleaved (TI) architectures combine low-speed ADCs connected in parallel and sampled with uniformly distributed different phases from a single clock signal to achieve a high sampling rate and high energy efficiency [2]. SAR, Flash, and pipeline ADCs are commonly used in TI-based structures. However, mismatch phenomena among the channels significantly degrade the resolution. To alleviate this issue, calibration circuits are needed, which require extra power consumption and increase the system complexity [14,15].

Continuous-Time Delta-Sigma ($\Delta\Sigma$) modulator (CTSDM) ADCs have also been reported for wide-band analog-to-digital conversion [16], but some key points such as the analog nature of the filter loop, the excess loop delay, and the clock jitter increase the complexity design especially with narrow processes nodes. Noise-Shaping SAR (NS-SAR) structures combined with time-interleaving techniques remove the need for analog-circuits, relaxing technology scaling requirements and enabling high-bandwidth conversion with lower power consumption [17,18].

Voltage-controlled oscillator based analog-to-digital converters (VCO-based ADCs) have emerged as a promising solution due to their highly digital nature adequate for very low voltage supplies. Diverse hybrid structures with VCOs have been published: with VCOs as integrators/quantizers in CTSDMs [19–23], placing it into SAR-based structures for pipeline ADCs [24–27] or in multi-stage noise shaping (MASH) architectures [28–31]. If we look for simplicity, they are of special interest if a ring-oscillator is used in an open-loop configuration [32–34]. Apart from its simple digital architecture, composed of CMOS logic gates (NOT gates), the spectral properties of pulse frequency modulation enable first-order noise-shaped output data, with a performance similar to CTSDMs [35].

The main limitation of the ring-oscillator is the nonlinear voltage/current-to-frequency response, which translates into distortion and limits the dynamic range of the whole ADC. Several ways of correcting the ring-oscillators' nonlinearity have been proposed in the literature. One of the most explored is digital calibration that requires large occupied area and wastes a lot of power [32,36]. Circuit-level solutions have also been introduced at the cost of a much lower oscillator gain [37] or more complex designs [38], being in many cases in applications for low-bandwidth ADCs. Another alternative consists of a VCO-based quantizer within a CTSDM [19,22], where the nonlinearity is corrected by the gain of the loop at the expense of more complex and non-scalable structures.

In this work, we propose a new way to mitigate the distortion generated by the nonlinearity of the ring-oscillator, exploiting a circuit design with significantly lower power consumption and area comparing to prior art. The idea makes use of several transconductors with different bias conditions connected in parallel to inject the current into a current-controlled oscillator (CCO) architecture. By means of selecting a proper distribution of those bias conditions, we can implement a nonlinear voltage-to-current function that approximates the inverse nonlinear current-to-frequency CCO relation. Figure 1 depicts a scheme that summarizes the idea described above. The VCO is composed of an input stage that converts an input voltage into a current, and a CCO that makes a current-to-frequency translation (Figure 1a). The transfer functions of both the front-end circuit and the ring-oscillator affect the whole response of the VCO (Figure 1b). To compensate for the nonlinearity of the structure, the inverse nonlinear function of the oscillator is artificially performed by the front-end circuit, thus canceling both nonlinear effects (Figure 1c). Something similar was already proposed in [39] with a resistive divider as the oscillator

front-end circuit. The main disadvantage is that the input signal attenuation directly entails a lower dynamic range.

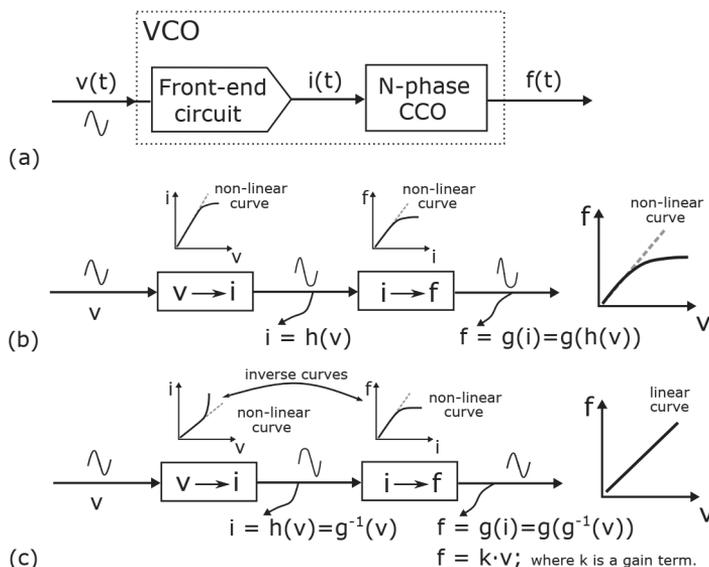


Figure 1. (a) General scheme of a VCO; (b) nonlinear VCO operation; and (c) proposed linearization technique.

The possibility of going towards a fully-synthesizable architecture might be considered in the future, but it is currently out of the scope of the present manuscript, mainly due to the analog nature of the transconductors. All the digital parts can be implemented by using digital synthesis [3,6,7,40–45]. Tools for automated analog design are on-the-spot now and may become of application for the current architecture [46–49].

Our proposal is evaluated by transient simulation with a low power (LP) TSMC 65-nm CMOS process, showing excellent performance in power and area, especially compared to digital calibration techniques [32], significantly reducing the total harmonic distortion (THD) power. In addition, PVT variation and Monte Carlo simulations show that the solution is robust against those variations, requiring simple calibration to achieve substantial distortion enhancement and proper resolution improvement for high-swing input signals. Finally, the proposed circuit is completely scalable as the calibration is also performed digitally, making it suitable for PVT/mismatch compensation in very deep submicron CMOS nodes such as 16-nm or 7-nm.

2. Materials and Methods

2.1. Nonlinearity in Ring-Oscillators

The conventional structure for a pseudo-differential open-loop VCO-based ADC is depicted in Figure 2a, built with one ring-oscillator per branch followed by some digital logic that samples and computes the first difference of the output phases (Figure 2b [35]). The VCO is composed of a transconductor-based front-end stage (g_m) and an N-phase CCO. The voltage-to-current conversion could be performed with a single NMOS transistor. Here, the instantaneous oscillation frequency ($f_{osc}(t)$) of the ring-oscillator follows the expression:

$$f_{osc}(t) = f_o + K_{VCO} \cdot g_m \cdot x(t), \tag{1}$$

where f_o is the rest oscillation frequency, K_{VCO} is the ring-oscillator gain, g_m is the transconductance of the front-end circuit (the single NMOS device in Figure 2b), and $x(t)$ is the

input voltage centered at zero. Level shifters are required to saturate amplitude-modulated oscillator output signals before being sampled and operated in discrete time.

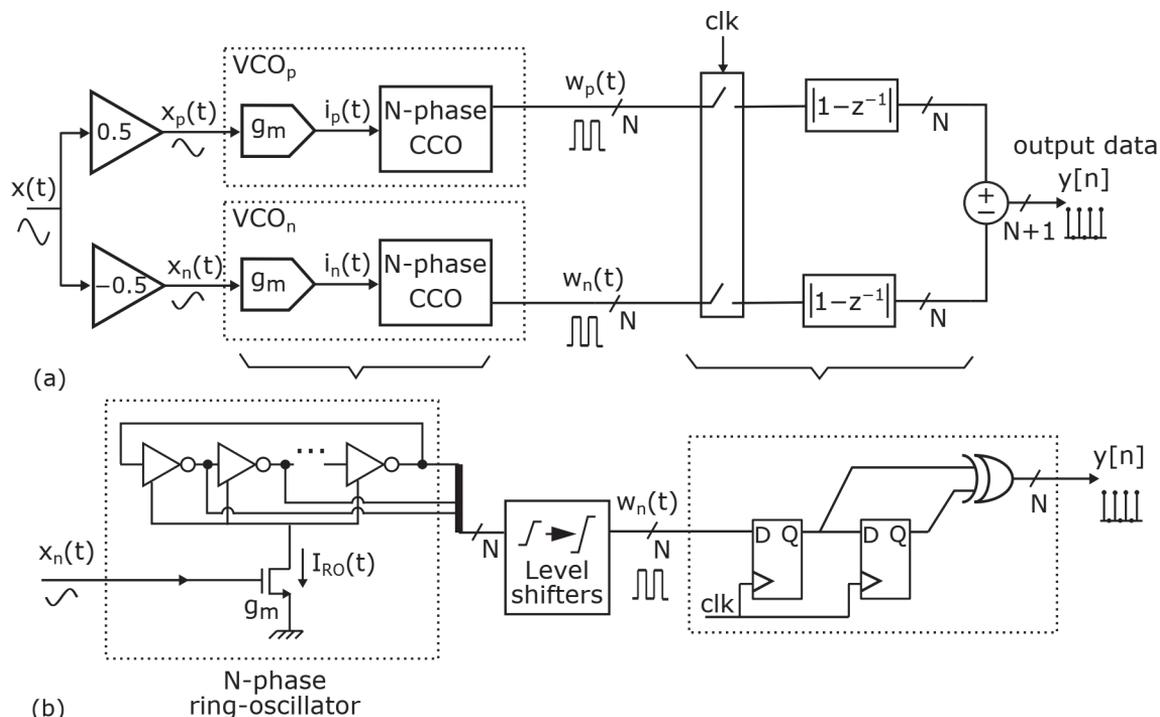


Figure 2. (a) Pseudo-differential configuration for an open-loop VCO-based ADC; (b) circuit built with an N-phase ring-oscillator, level shifters, flip-flops, and XOR gates.

As open-loop configuration corresponds to a highly technology-scalable architecture; this ring-oscillator is implemented with conventional CMOS logic gates and the circuitry afterward is composed of registers and XOR gates. However, the voltage-to-frequency response of the VCO is not linear due to the nonlinear time-delay dependence of the logic gates of the oscillator with respect to the flowing current [50], and the nonlinear voltage-to-current relation in the transconductance g_m .

To have some intuition about the ring-oscillator nonlinear characteristic and its effect on the output data, a behavioral model of Figure 2 was designed in a 65-nm CMOS technology. The model included a 45-phase ring-oscillator with the rest of the blocks modeled with Verilog-A. The sampling frequency (f_s) was 2 GHz with an analog bandwidth (ABW) of 50 MHz. The oscillation parameters were $f_o = 450$ MHz and $K_{VCO} \cdot g_m = 1$ GHz/V, for a 3 MHz sinusoidal differential input signal of 800 mV_{pp}. The nominal voltage supply was 1.2 V.

Figure 3 depicts the spectrum of the output data $y[n]$ resulting from a transient simulation. The harmonic distortion is clearly visible, with the third and fifth harmonic distortion terms (HD3 and HD5) equal to -44 dBc and -68 dBc, respectively (the second harmonic distortion term, HD2, term was -39 dBc in a single-ended configuration). The signal-to-noise ratio (SNR) is 63 dB and the signal-to-noise-distortion ratio (SNDR) is 44 dB, which means a degradation of approximately three ENOBs due to distortion. As seen, a peak SNDR higher than 50 dB becomes impossible, making this architecture unsuitable for next generation WLAN standard or 5G, where a peak SNDR higher than 65 dB over wide bandwidth is required [51].

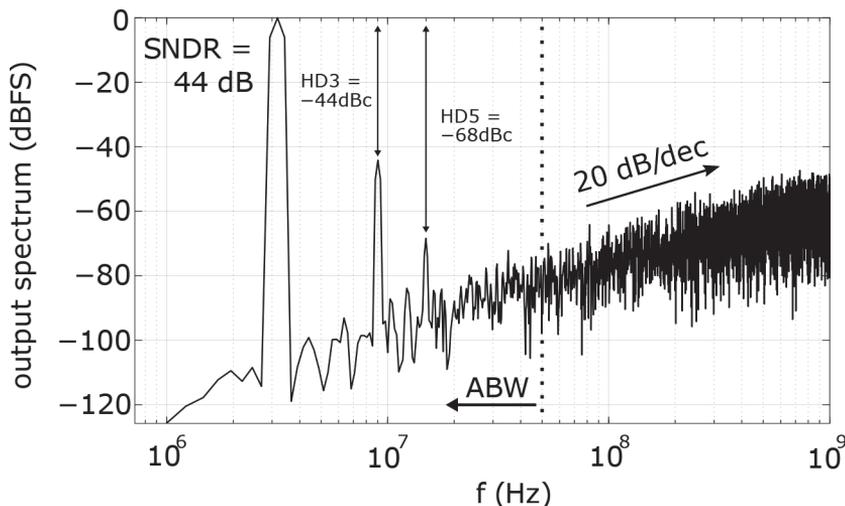


Figure 3. Distortion due to a nonlinear ring-oscillator in analog-to-digital conversion.

2.2. Proposed Multiple-Transconductor Ring-Oscillator

Figure 4 displays the voltage-to-frequency response of the oscillator simulated above. The ideal linear response has also been plotted. Assuming the ring-oscillator configuration of Figure 2b, we may distinguish between three different regions for the NMOS transistor: the saturated region, where the ring-oscillator shows a behavior approximately linear; and the sub-threshold and ohmic regions, where the nonlinearity is clearly visible. When having the NMOS working at such regions, the harmonic distortion will increase due to the joint action of both the nonlinear voltage-to-current conversion and the nonlinear time delay dependence of the ring-oscillator. This restricts most of the linear operating region to a small input voltage range, which might be suitable for low-swing input voltage applications [37], but not for high-swing ones. To increase the linear voltage range, we can inject more current into the oscillator in the ohmic region to move the curve up (see the arrows in Figure 4), and drain current in the sub-threshold region to move the curve down.

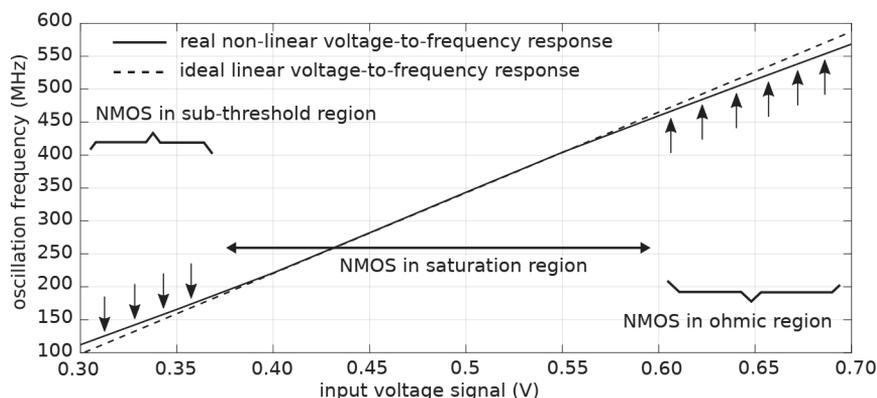


Figure 4. Voltage-to-frequency conversion function of a conventional ring-oscillator (solid line) and an ideal linear curve estimation (dashed line).

With that purpose in mind, we propose a circuit-level solution to extend the linear response of the ring-oscillator. The solution is based on injecting current into the ring-oscillator with several transconductors, instead of a single one, shifting the saturation

regions of them throughout the whole desired input voltage range. On the one hand, we assume that we have at least one device working in saturation at any point of the input voltage range, getting an approximately linear voltage-to-frequency response. On the other hand, we have more flexibility to control the injected current into the ring-oscillator and perform a voltage-to-current conversion that mitigates the distortion due to the nonlinear time dependence of the logic gates. Looking at Figure 4, two current-based operations are needed: injection and draining, requiring respectively both NMOS and PMOS-based devices in the front-end circuit. Thus, the current flowing through the ring-oscillator $I_{RO}(t)$ can be expressed as follows:

$$I_{RO}(t) = \sum_{i=1}^M I_{N,i}(t) - \sum_{j=1}^N I_{P,j}(t), \tag{2}$$

where $I_{N,i}(t)$ is the current provided by the i -th NMOS device, and $I_{P,j}(t)$ is the current drained by the j -th PMOS device.

Using both PMOS and NMOS devices may suppose issues arising from matching and more complexity in making the calibration of the circuit. This is why, for a proof of concept, it was decided to correct and extend only the ohmic region by means of NMOS devices. Several NMOS devices are then connected in parallel. Each of these devices has its own offset component to control the point when they get into the saturation region. Figure 5 depicts a diagram of the proposed solution with a ring-oscillator whose input current is provided by M NMOS devices, where $x_{off,i}$ represents the offset voltage for each of the transconductors. These offset values must be allocated throughout the desired input voltage range to feed the proper current that approximates the inverse current–frequency relation of the ring-oscillator. Linearity is kept mainly because there is always at least one transconductor which is providing sufficient transconductance. Although the rest of transconductors are still providing some current gain, these values are negligible and do not affect the approximation made.

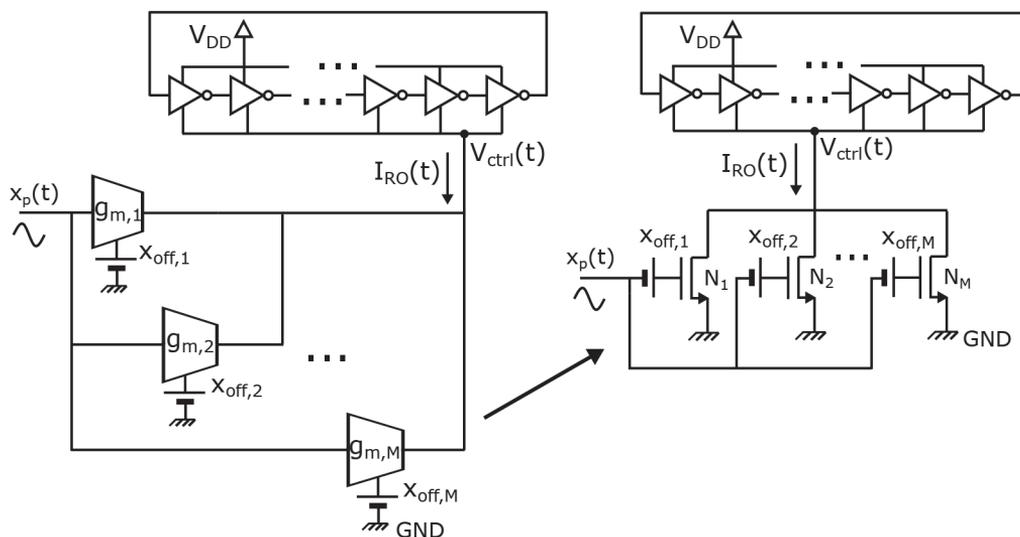


Figure 5. Proposed multiple-transconductor circuit for a linear ring-oscillator-based analog-to-digital conversion.

Apart from the mitigation of the nonlinearity, the proposed structure could also be used to enhance the oscillator gain K_{VCO} . Note that, for the conventional approach (Figure 4), the oscillator gain is limited because the transconductor g_m drops into the ohmic region and is not able to provide sufficient current to keep increasing the oscillation fre-

quency with the same slope as in the saturation region. With the proposed solution, at any point of the selected input range, there is at least one transistor providing enough current to keep a linear relation, so that the highest achievable oscillation frequency increases.

Finally, jitter requirements are alleviated due to the open-loop structure, in comparison to closed-loop architectures [21].

3. Results

3.1. Circuit Validation

To evaluate the proposed architecture, a version of the circuit of Figure 5 was designed in a 65-nm CMOS process with a pseudo-differential configuration (Figure 2a) and 45 phases in the ring-oscillator. The front-end circuit of the ring-oscillator, together with the circuit used to trim the offset component of each transconductor, is shown in Figure 6. The offset component of the input signal is one of the offset references needed, and the other voltage references are generated by means of fixed bias currents and diode-connected transistors M_1 and M_2 . Due to the limited headroom voltage, M_2 operates in the sub-threshold region. In our case, three offset references were needed, with the values displayed in Figure 6. The size distribution of the transconductors was chosen according to a static behavior, performing a sweep of several transient simulations with different DC input values to get the voltage-to-frequency relation and quantify the nonlinear coefficients.

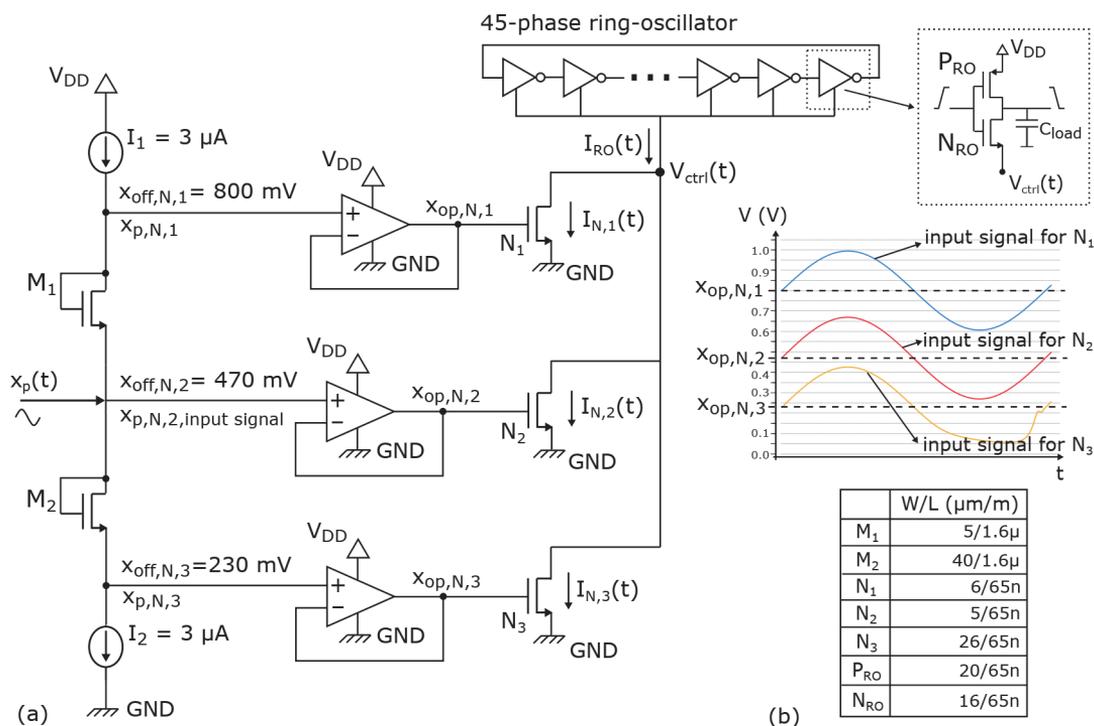


Figure 6. Trimming circuit used to generate the offset references of the ring-oscillator, (a) schematic and (b) device sizes and chronograms.

Low-power opamps connected in a buffer configuration were used to avoid kick-back noise. The opamp does not require high DC gain, hence its design does not lead to a technology scaling issue. The gain-bandwidth product value was 200 MHz and the margin phase was 55° , ensuring the proper operation of the system. Figure 6b illustrates the time response of $x_{op,N,1}$, $x_{op,N,2}$, and $x_{op,N,3}$ from a transient simulation where $x_p(t)$ is a sinusoidal signal. The unitary gain configuration is accomplished totally for $x_{op,N,1}$, $x_{op,N,2}$,

and partially for $x_{op,N,3}$, which gets distorted for voltage values lower than the threshold voltage (around 0.25 V). This is not relevant because, for gate voltage values lower than the threshold voltage, N_3 operates in a sub-threshold, and the current provided is negligible and does not significantly modify the approximation to the inverse oscillator’s current-to-frequency function. Opamp’s circuit is depicted in Figure 7, which consists of a two-stage Miller-compensated opamp with low offset (3 mV simulated, which is sufficiently low for proper performance).

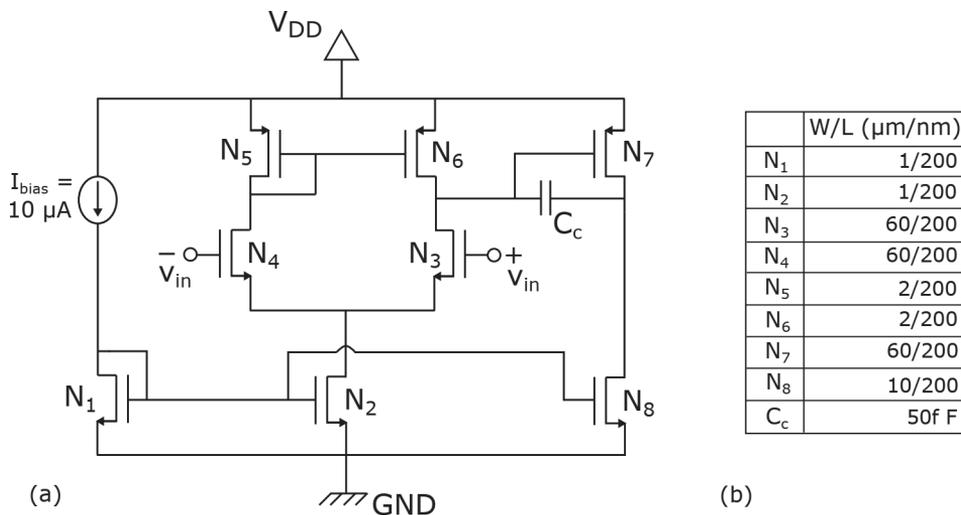


Figure 7. Two-stage Miller-compensated opamp of Figure 6, (a) schematic and (b) device sizes.

Transient simulations were used to check the nonlinearity mitigation. Digital demodulation and sampling logic were modeled in VerilogA language. 10-fF capacitors were placed at the outputs of the ring-oscillator to simulate the input capacitance of the digital logic. Oscillation parameters were kept similar to Figure 3, except $K_{VCO} \cdot g_m$, which was equal to 1.25 GHz/V, slightly higher than before due to linearity compensation. Figure 8 depicts the result of the nominal transient simulation. The HD3 and HD5 were equal to -67 and -75 dBc, respectively, leading to an SNDR of 63 dB (the HD2 value observed in a single-ended configuration was -63 dBc). The distortion is substantially mitigated in comparison to Figure 3. The same simulation was performed with output capacitors of 30 and 50 fF, achieving SNDR values of 62.8 dB and 61.7 dB, respectively, due to the slight decrease in the oscillator gain. The proposed solution achieved a THD value of -66.4 dBc for a differential input of 800 mV_{pp}, [32] reported a THD chip measurement of -63.6 dBc for a differential input of 566 mV_{pp}, and [52] reported a THD of -65.3 dBc for a differential input of 400 mV_{pp} in simulation.

Integral nonlinearity (INL) performance, as a static characterization, was also analyzed. The results of a linear ramp test, using the best-fit straight-line approximation, are shown in Figure 9. The single-ended peak-to-peak INL error was $[-0.44 \ 0.26]$, within 1 LSB. In [40], a fully-synthesizable VCO-ADC is presented, where the INL achieved after digital correction was in the range $(-1.4 \ 1.49)$ LSBs and $(-1.9 \ 1.6)$ LSBs for simulation and measurement results, respectively.

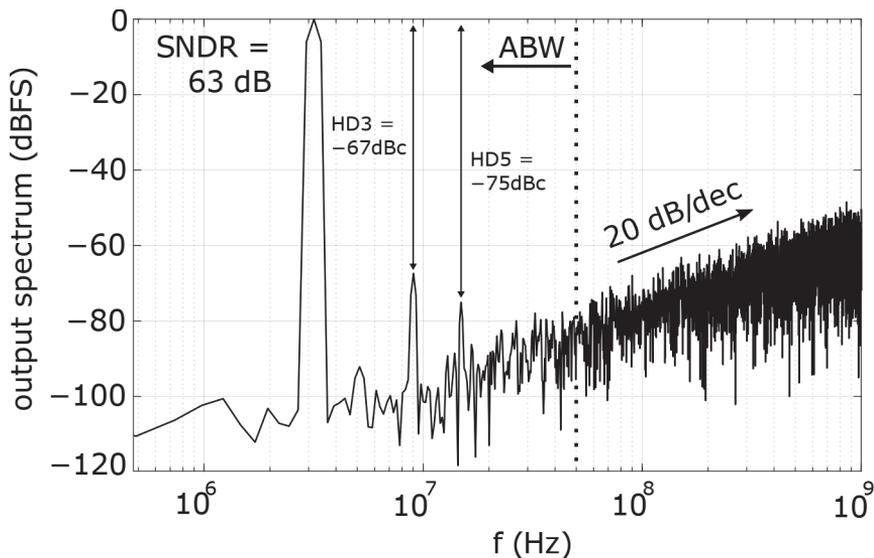


Figure 8. Output spectrum of transient simulation of the circuit proposed in Figure 6.

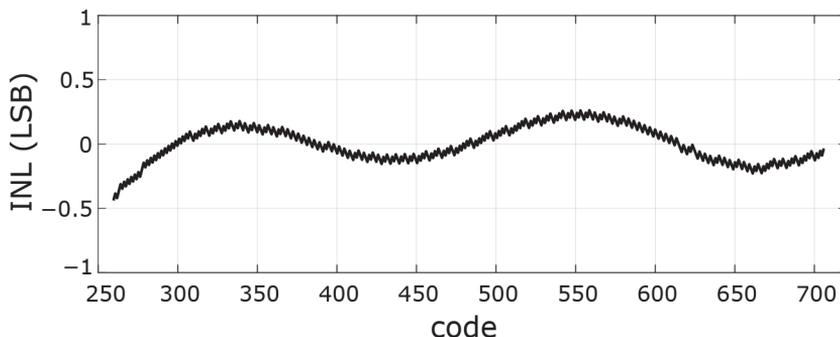


Figure 9. INL performance of the system with the circuit proposed in Figure 6.

3.2. Circuit-Level Impairments

To achieve a correct performance from the proposed solution, proper offset references for each of the transconductors need to be selected. The ring-oscillator might be affected by mismatch effects and PVT variations. This will result in variations in the oscillator current-to-frequency relation and in the required voltage-to-current function. The trimming circuit will be affected by these effects as well, modifying the transconductors’ bias points. To reduce mismatch effects between devices, the diode-connected transistors (M_1 and M_2) were designed to be large. To analyze the variation of $x_{off,N,1}$, $x_{off,N,2}$ and $x_{off,N,3}$; and the oscillation frequency due to mismatching a set of 300 runs of Monte Carlo simulations was performed. The SNDR was normally distributed with a mean value of 63 dB and a standard deviation of 1.8 dB, with a worst-case SNDR value, was equal to 59.5 dB.

Apart from the mismatch verification, the impact of PVT variations in the VCO linearity was also checked. We noticed that the function of Figure 4 is mainly shifted throughout the horizontal axis, but its shape did not vary significantly for different PVT conditions. Consequently, centering the input offset in the linear region is the easiest requirement to keep the oscillator working in a linear manner. Figure 10 shows the SNDR values achieved for different PVT cases with the offset of the input signal ($x_{off,N,2}$) correctly tuned. SNDR degradation can be mainly observed for high temperatures. This degradation occurs when all the transconductors drop into the ohmic region due to the limited V_{ctrl}

(Figure 6), which is of special relevance for the FF case where the working linear region is dramatically reduced. To improve the linearity of the proposed solution in these cases, the voltage supply could be increased at the expense of increasing the power consumption or resizing the ring-oscillator at the expense of reducing the gain. All of the simulation results depicted in Figure 10 were obtained with a nominal voltage supply of 1.2 V. The same simulations were repeated including a variation of ± 50 mV in the voltage supply. The resulting SNDR values did not differ from the ones shown in Figure 8, which means strong robustness against voltage supply variations.

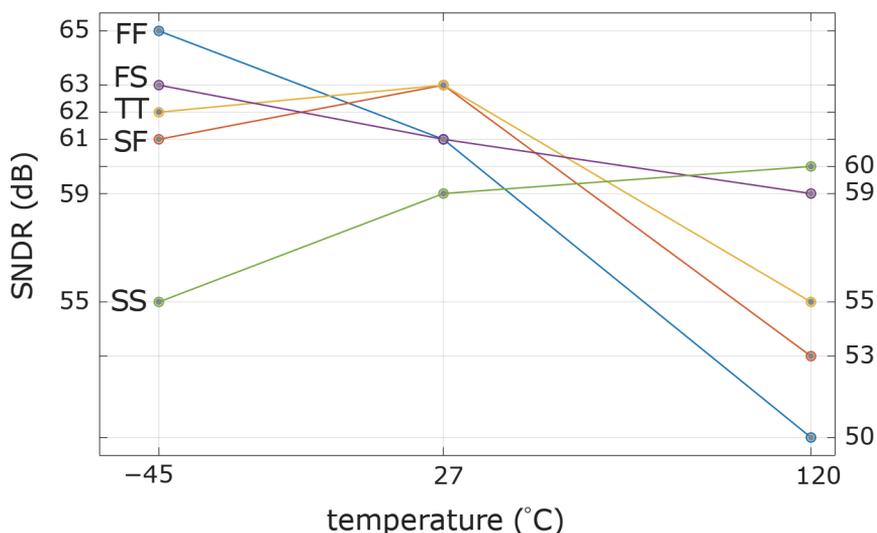


Figure 10. SNDR variation of Figure 8 due to different PVT conditions. Capital letters indicate devices' process: 'S' means "slow", 'F' means "fast", and 'T' means typical. The first letter refers to NMOS devices and the second one refers to PMOS devices. The nominal voltage supply (1.2 V) is the same for all the cases.

Additionally, periodic steady-state and phase-noise analyses were performed to evaluate the limitation that the oscillator imposes to the system in terms of phase noise [53]. The estimated value of SNDR regarding the oscillator phase noise was 71 dB, much lower than the quantization noise based SNDR limit observed in Figure 8.

3.3. Calibration Circuit

A calibration circuit is required to correctly select the offset of the input signal $x_{\text{off},N,2}$, approximate the inverse current-to-frequency ring oscillator's function, and keep proper linearity. With that purpose in mind, we propose to use a digital foreground calibration circuit enabled periodically. Therefore, static power is not increased. The circuit is depicted in Figure 11. The idea is measuring the rest oscillation frequency f_o , identifying the PVT operating point of the prototype and correctly selecting the best input signal offset, similar to what was done for Figure 10 elaboration. The rest oscillation frequency is measured by means of a digital delay chain (NAND gates) that measures the semiperiod of the oscillating signal coming from one of the phases of the ring-oscillator. The outputs of the NAND gates are stored with flip-flops, whose outputs are thermometrically encoded and represent an estimation of the semiperiod of f_o . This digital estimation gets into a Look-Up-Table (LUT) with previously defined values that select the best offset of the input signal. This offset value is finally controlled by means of an opamp whose output offset component can be digitally tuned.

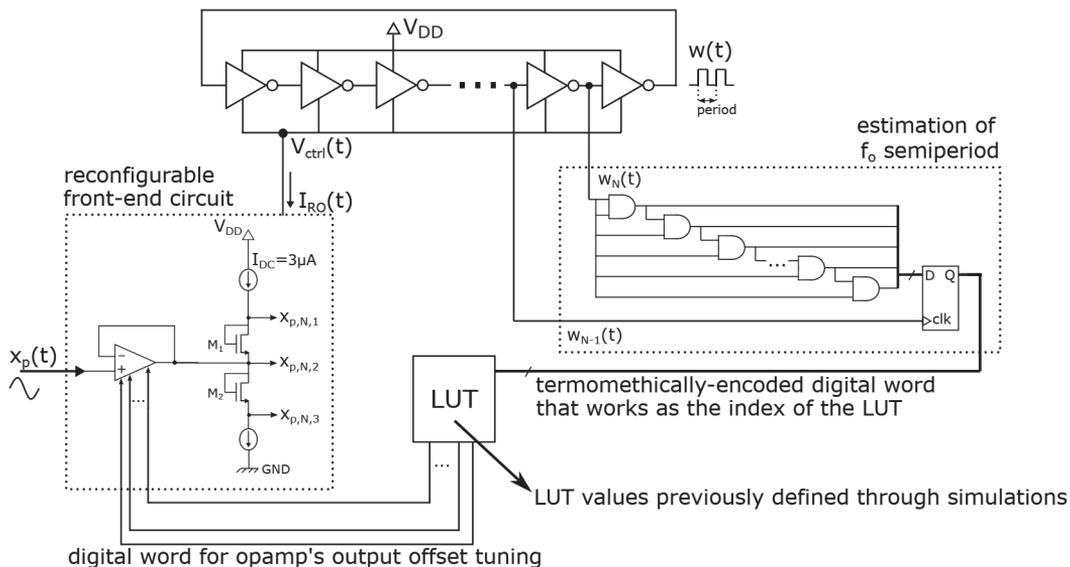


Figure 11. Scheme of the proposed calibration technique.

4. Discussion

The proposed architecture does not add extra relevant power consumption. The main contributor to the power breakdown is the oscillator (1.5 mW per oscillator—in line with other previous publications). The power consumption of each opamp is of 115 μ W, and the branch needed to generate the offset references consumes 3.6 μ W. As a consequence, the power consumption associated with the extra circuitry for linearization is less than 19% of the total required power, while, in [32], it is 60%. Although this power estimation will no doubt increase in an experimental prototype, it is expected that the power ratio between these elements remains. In relation to the area, it is not expected that our solution adds extra area growth in comparison to other digital calibration techniques.

The architecture in [32] introduces a VCO-based nonuniform sampling (NUS) ADC, which involves on-chip nonlinearity estimation and off-chip nonlinearity correction embedded within a non-uniform digital signal processing (DSP). The area of the set of structures dedicated to the calibration occupies almost the same of the ADC itself: 0.13 mm² and 0.14 mm², respectively. In [36], a reconfigurable CTSDM for analog-to-digital conversion is presented with an on-chip digital background calibration and self-canceling dither techniques. The calibration unit occupies 64% of the area of the whole chip and the voltage-to-current converter and the ring-oscillator consume less than a fifth part of the total power dissipation [54]. Our solution for nonlinearity mitigation does not involve background continuous intensive digital operations and is performed in the analog domain, reducing the required power consumption and also the occupied area.

The VCO-based open-loop configuration ADC in [37] presents a resistive network to tune the ring-oscillator voltage-to-frequency function and reduce harmonic distortion. The principle of the resistive divider scheme was originally introduced in [39], and compromises ring-oscillator's gain by attenuating the input amplitude. Our solution does not restrict oscillator's gain.

The structure explored in [19] contains a VCO as a quantizer within a $\Delta\Sigma$ loop. Nonlinearity is mitigated through a high-gain loop filter. The VCO-based quantizer only occupies 3.7% of the active area and consumes 13% of the total power, but it is hardly extended to low voltage supply environments. While enough voltage supply is granted to allow digital switching in the ring-oscillator, our solution can be implemented for lower voltage supply applications.

Table 1 summarizes the performance of the proposed design and provides a comparison to prior works. Different VCO-based ADCs structures (a fully-synthesizable design [40] is also included), hybrid SAR-VCOs, and traditional Flash and SAR architectures are characterized as competitive alternatives. To achieve a fair comparison between the reference solutions and the proposed design, area, power, and Figure-of-Merit (FoM) values only include the ring-oscillator and the associated linearization blocks for the cases where data are available (*).

Table 1. Comparison to state-of-the-art.

Parameter	[3]	[9]	[19]	[23]	[25]	[32] *	[36]	[37] *	[38]	[40]	This Work
Meas./Sim. Results	MR	MR	MR	MR	SR	MR	MR	MR	SR	MR	SR
Supply [V]	1.2	1	1.2	0.9	1	1	1.2	0.2	1	0.6	1.2
Process [nm]	90	28	130	40	65	65	65	28	28	65	65
BW [MHz]	105	50	20	40	5	200	37.5	0.061	10	25.6	50
SNDR [dB]	35.89	67	67	59.5	75.7	57	70	68	62	50.3	63
THD [dBc]	–	–74.4	–67.7	–65.7	–80.7	–63.6	–76	–72.5	–	–53.9	–66.4
INL [LSB] ^a	–2/0.44	–	–	–	–	–	–	–	–	–1.9/1.6	–0.44/0.26
ENOB	5.67	10.85	10.84	9.59	12.3	9.18	11.34	11	10	8.06	10.17
Diff. Input Range [mV]	280	2000	180	715	1800	566	800	355	800	600 ^b	800
Power [mW]	34.8	8	40	2.57	0.51	35.4	39	0.0065	0.23	3.3 ^b	3.69
Area [mm ²]	0.18	0.1	0.42	0.017	–	0.1557	0.11	0.07	–	0.026 ^c	–
FoM [dB] ¹	130.7	165	154	160.9	173.9	154.5	159.8	167.7	166.4	149.2	164.3
FoM [fj/c-s] ²	3256	43.2	500	42	14.9	153	201.2	26	14	235	32
Linearization technique	Inv. Gauss.	Digital calib.	$\Delta\Sigma$ loop	Two-step VCO-ADC	Digital f-calib.	Digital f-calib.	Digital b-calib.	Resist. netw.	Bulk-driven	Digital f-calib	Multiple transc.

FoM [dB]¹ = SNDR + 10log₁₀(BW/Power). FoM [fj/c-s]² = Power/(2·BW·2^{(SNDR–1.76)/6.02}).^a Referred to single-ended mode.
^b These area and power consumption values are only for the ADC core, the digital correction block is not included. ^c This input voltage swing is for a single-ended VCO-ADC architecture. Inv. Gaus. CDF = The inverse of Gaussian Cumulative Distribution Function (CDF). Digital f-calib./b-calib. = Digital foreground-calibration/background-calibration. Resist. netw. = Resistive network. Multiple transc. = Multiple transconductors.

5. Conclusions

A circuit-level solution to linearize ring-oscillators-based ADCs is proposed. The solution is based on making use of several transconductors connected in parallel with different bias conditions to implement a voltage-to-current function that approximates the inverse nonlinear current-to-frequency function of the ring-oscillator. To evaluate the approach, a ring-oscillator-based ADC with the proposed circuit was designed and simulated in 65-nm. Nonlinearity was strongly reduced resulting in an ENOB enhancement of more than three bits for high-swing inputs. Additionally, mismatch effects, PVT variations, and noise impact were assessed. The proposal exhibited great robustness without resorting to a complex circuit design and just requiring simple foreground digital calibration. The new VCO-based ADC structure benefits from important power savings in comparison to state-of-the-art digital calibration circuits conventionally used to mitigate distortion. It is also expected to achieve area savings, but this needs to be confirmed through experimental prototypes. The proposal is particularly intended for high-bandwidth and medium-resolution applications, such as 5G or IoT (Internet-of-Things) modules.

Author Contributions: Conceptualization, L.M.A.-G., V.M., V.K., S.P., L.H. and E.G.; methodology, L.M.A.-G., V.M., V.K., S.P., L.H. and E.G.; software, L.M.A.-G. and E.G.; validation, L.M.A.-G. and E.G.; formal analysis, L.M.A.-G. and E.G.; investigation, E.G.; resources, S.P., L.H. and E.G.; data curation, L.M.A.-G. and E.G.; writing—original draft preparation, L.M.A.-G. and E.G.; writing—review and editing, L.M.A.-G., V.M., V.K., S.P., L.H. and E.G.; visualization, L.M.A.-G. and E.G.; supervision, L.H. and E.G.; project administration, L.H. and E.G.; funding acquisition, S.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Project TEC2017-82653-R, Spain.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

ADC	Analog-to-digital converter
CCO	Current-controlled oscillator
ENOB	Effective number of bits
IoT	Internet-of-things
VCO	Voltage-controlled oscillator
SAR	Successive approximation register
Opamp	Operational amplifier
TI	Time-Interleaved
NS-SAR	Noise-shaping SAR
TINS-SAR	Time interleaving noise-shaping SAR
$\Delta\Sigma$	Delta-Sigma
CTSDM	Continuous-time $\Sigma\Delta$ modulator
VCO-based ADC	Voltage-controlled oscillator-based analog-to-digital converter
MASH	Multi-stage noise shaping
THD	Total harmonic distortion
PVT	Process, voltage and temperature
ABW	Analog bandwidth
HD3	Third harmonic distortion term
HD5	Fifth harmonic distortion term
HD2	Second harmonic distortion term
SNDR	Signal-to-noise-distortion ratio
SNR	Signal-to-noise ratio
INL	Integral nonlinearity
LSB	Least significant bit
NUS	Nonuniform sampling
DSP	Digital signal processing
CDF	Cumulative distribution function

References

- Lewyn, L.L.; Ytterdal, T.; Wulff, C.; Martin, K. Analog Circuit Design in Nanoscale CMOS Technologies. *Proc. IEEE* **2009**, *97*, 1687–1714. [[CrossRef](#)]
- Zahrai, S.A.; Onabajo, M. Review of Analog-To-Digital Conversion Characteristics and Design Considerations for the Creation of Power-Efficient Hybrid Data Converters. *J. Low Power Electron. Appl.* **2018**, *8*, 12. [[CrossRef](#)]
- Weaver, S.; Hershberg, B.; Moon, U. Digitally Synthesized Stochastic Flash ADC Using Only Standard Digital Cells. *IEEE Trans. Circuits Syst. Regul. Pap.* **2014**, *61*, 84–91. [[CrossRef](#)]
- Fahmy, A.; Liu, J.; Kim, T.; Maghari, N. An All-Digital Scalable and Reconfigurable Wide-Input Range Stochastic ADC Using Only Standard Cells. *IEEE Trans. Circuits Syst. II Express Briefs* **2015**, *62*, 731–735. [[CrossRef](#)]
- Jeon, M.; Yoo, W.; Kim, C.; Yoo, C. A Stochastic Flash Analog-to-Digital Converter Linearized by Reference Swapping. *IEEE Access* **2017**, *5*, 23046–23051. [[CrossRef](#)]
- Aiello, O.; Crovetti, P.; Alioto, M. Fully Synthesizable Low-Area Analogue-to-Digital Converters With Minimal Design Effort Based on the Dyadic Digital Pulse Modulation. *IEEE Access* **2020**, *8*, 70890–70899. [[CrossRef](#)]
- Park, J.; Hwang, Y.; Jeong, D. A 0.5-V Fully Synthesizable SAR ADC for On-Chip Distributed Waveform Monitors. *IEEE Access* **2019**, *7*, 63686–63697. [[CrossRef](#)]
- Cao, Z.; Yan, S.; Li, Y. A 32 mW 1.25 GS/s 6b 2b/Step SAR ADC in 0.13 μm CMOS. *IEEE J. Solid-State Circuits* **2009**, *44*, 862–873. [[CrossRef](#)]

9. Inerfield, M.; Kamath, A.; Su, F.; Hu, J.; Yu, X.; Fong, V.; Alnaggar, O.; Lin, F.; Kwan, T. An 11.5-ENOB 100-MS/s 8mW dual-reference SAR ADC in 28nm CMOS. In Proceedings of the 2014 Symposium on VLSI Circuits Digest of Technical Papers, Honolulu, HI, USA, 10–13 June 2014; pp. 1–2.
10. Kapusta, R.; Shen, J.; Decker, S.; Li, H.; Ibaragi, E.; Zhu, H. A 14b 80 MS/s SAR ADC With 73.6 dB SNDR in 65 nm CMOS. *IEEE J. Solid-State Circuits* **2013**, *48*, 3059–3066. [[CrossRef](#)]
11. Morie, T.; Miki, T.; Matsukawa, K.; Bando, Y.; Okumoto, T.; Obata, K.; Sakiyama, S.; Doshio, S. A 71dB-SNDR 50 MS/s 4.2 mW CMOS SAR ADC by SNR enhancement techniques utilizing noise. In Proceedings of the 2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers, San Francisco, CA, USA, 17–21 February 2013; pp. 272–273.
12. Chang, D.; Moon, U. A 1.4-V 10-bit 25-MS/s pipelined ADC using opamp-reset switching technique. *IEEE J. Solid-State Circuits* **2003**, *38*, 1401–1404. [[CrossRef](#)]
13. Shibata, H.; Taylor, G.; Schell, B.; Kozlov, V.; Patil, S.; Paterson, D.; Ganesan, A.; Dong, Y.; Yang, W.; Yin, Y.; et al. 16.6 An 800MHz-BW VCO-Based Continuous-Time Pipelined ADC with Inherent Anti-Aliasing and On-Chip Digital Reconstruction Filter. In Proceedings of the 2020 IEEE International Solid-State Circuits Conference—(ISSCC), San Francisco, CA, USA, 16–20 February 2020; pp. 260–262.
14. Dortz, N.L.; Blanc, J.; Simon, T.; Verhaeren, S.; Rouat, E.; Urard, P.; Tu, S.L.; Goguet, D.; Lelandais-Perrault, C.; Benabes, P. 22.5 A 1.62GS/s time-interleaved SAR ADC with digital background mismatch calibration achieving interleaving spurs below 70dBFS. In Proceedings of the 2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC), San Francisco, CA, USA, 9–13 February 2014; pp. 386–388.
15. Baert, M.; Dehaene, W. A 5-GS/s 7.2-ENOB Time-Interleaved VCO-Based ADC Achieving 30.5 fJ/cs. *IEEE J. Solid-State Circuits* **2020**, *55*, 1577–1587. [[CrossRef](#)]
16. Paton, S.; Giandomenico, A.D.; Hernandez, L.; Wiesbauer, A.; Potscher, T.; Clara, M. A 70-mW 300-MHz CMOS continuous-time $\Sigma\Delta$ ADC with 15-MHz bandwidth and 11 bits of resolution. *IEEE J. Solid-State Circuits* **2004**, *39*, 1056–1063. [[CrossRef](#)]
17. Jie, L.; Zheng, B.; Flynn, M.P. 20.3 A 50MHz-Bandwidth 70.4dB-SNDR Calibration-Free Time-Interleaved 4th-Order Noise-Shaping SAR ADC. In Proceedings of the 2014 2019 IEEE International Solid-State Circuits Conference—(ISSCC), San Francisco, CA, USA, 17–21 February 2019; pp. 332–334.
18. Lin, C.-Y.; Lin, Y.-Z.; Tsai, C.-H.; Lu, C.-H. 27.5 An 80MHz-BW 640MS/s Time-Interleaved Passive Noise-Shaping SAR ADC in 22nm FDSOI Process. In Proceedings of the 2021 IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, USA, 13–22 February 2021; pp. 378–380.
19. Straayer, M.Z.; Perrott, M.H. A 12-Bit, 10-MHz Bandwidth, Continuous-Time $\Sigma\Delta$ ADC With a 5-Bit, 950-MS/s VCO-Based Quantizer. *IEEE J. Solid-State Circuits* **2008**, *43*, 805–814. [[CrossRef](#)]
20. Cardes, F.; Gutierrez, E.; Quintero, A.; Buffa, C.; Wiesbauer, A.; Hernandez, L. 0.04-mm² 103-dB-A Dynamic Range Second-Order VCO-Based Audio $\Sigma\Delta$ ADC in 0.13- μ m CMOS. *IEEE J. Solid-State Circuits* **2018**, *53*, 1731–1742. [[CrossRef](#)]
21. Huang, S.; Egan, N.; Kesharwani, D.; Opteynde, F.; Ashburn, M. 28.3 A 125MHz-BW 71.9dB-SNDR VCO-based CT $\Delta\Sigma$ ADC with segmented phase-domain ELD compensation in 16 nm CMOS. In Proceedings of the 2017 IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, USA, 5–9 February 2017; pp. 470–471.
22. Zhong, Y.; Tang, X.; Liu, J.; Zhao, W.; Li, S.; Sun, N. An 81.5dB-DR 1.25MHz-BW VCO-Based CT $\Delta\Sigma$ ADC with Double-PFD Quantizer. In Proceedings of the 2021 IEEE Custom Integrated Circuits Conference (CICC), Austin, TX, USA, 25–30 April 2021; pp. 1–2.
23. Xing, X.; Gielen, G.G.E. A 42 fJ/Step-FoM Two-Step VCO-Based Delta-Sigma ADC in 40 nm CMOS. *IEEE J. Solid-State Circuits* **2015**, *50*, 714–723. [[CrossRef](#)]
24. Dey, S.; Mayaram, K.; Fiez, T. A 12 MHz BW, 80 dB SNDR, 83 dB DR, 4th order CT- $\Delta\Sigma$ modulator with 2nd order noise-shaping and pipelined SAR-VCO based quantizer. In Proceedings of the 2019 IEEE Custom Integrated Circuits Conference (CICC), Austin, TX, USA, 14–17 April 2019; pp. 1–4.
25. Hou, Y.; Chen, Z.; Miyahara, M.; Matsuzawa, A. An Op-amp free SAR-VCO hybrid ADC with second-order noise shaping. In Proceedings of the 2016 IEEE International Conference on Electron Devices and Solid-State Circuits (EDSSC), Hong Kong, China, 3–5 August 2016.
26. Sanyal, A.; Ragab, K.; Chen, L.; Viswanathan, T.R.; Yan, S.; Sun, N. A hybrid SAR-VCO $\Delta\Sigma$ ADC with first-order noise shaping. In Proceedings of the IEEE 2014 Custom Integrated Circuits Conference, San Jose, CA, USA, 15–17 September 2014; pp. 1–4.
27. Xie, Y.; Liang, Y.; Liu, M.; Liu, S.; Zhu, Z. A 10-Bit 5 MS/s VCO-SAR ADC in 0.18- μ m CMOS. *IEEE Trans. Circuits Syst. II Express Briefs* **2019**, *66*, 26–30. [[CrossRef](#)]
28. Ragab, K.; Sun, N. A 12-b ENOB 2.5-MHz BW VCO-Based 0-1 MASH ADC With Direct Digital Background Calibration. *IEEE J. Solid-State Circuits* **2017**, *52*, 433–447. [[CrossRef](#)]
29. Gutierrez, E.; Hernandez, L.; Cardes, F. VCO-based sturdy MASH ADC architecture. *Electron. Lett.* **2017**, *53*, 14–16. [[CrossRef](#)]
30. Sacco, E.; Vergauwen, J.; Gielen, G. A 16.1-bit Resolution 0.064-mm² Compact Highly Digital Closed-Loop Single-VCO-Based 1-1 Sturdy-MASH Resistance-to-Digital Converter With High Robustness in 180-nm CMOS. *IEEE J. Solid-State Circuits* **2020**, *55*, 2456–2467. [[CrossRef](#)]
31. Maghami, H.; Payandehnia, P.; Mirzaie, H.; Zanbaghi, R.; Zareie, H.; Goins, J.; Dey, S.; Mayaram, K.; Fiez, T.S. A Highly Linear OTA-Less 1-1 MASH VCO-Based $\Delta\Sigma$ ADC With an Efficient Phase Quantization Noise Extraction Technique. *IEEE J. Solid-State Circuits* **2020**, *55*, 706–718. [[CrossRef](#)]

32. Wu, T.; Chen, M.S. A Noise-Shaped VCO-Based Nonuniform Sampling ADC With Phase-Domain Level Crossing. *IEEE J. Solid-State Circuits* **2019**, *54*, 623–635. [[CrossRef](#)]
33. Gutierrez, E.; Perez, C.; Hernandez, L.; Cardes, F.; Petrescu, V.; Walter, S.; Gaier, U. A Pulse Frequency Modulation VCO-ADC in 40 nm. *IEEE Trans. Circuits Syst. II Express Briefs* **2019**, *66*, 51–55. [[CrossRef](#)]
34. Perez, C.; Quintero, A.; Amaral, P.; Wiesbauer, A.; Hernandez, L. A 73dB-A Audio VCO-ADC based on a Maximum Length Sequence Generator in 130nm CMOS. *IEEE Trans. Circuits Syst. II Express Briefs* **2021**. [[CrossRef](#)]
35. Gutierrez, E.; Hernandez, L.; Cardes, F.; Rombouts, P. A Pulse Frequency Modulation Interpretation of VCOs Enabling VCO-ADC Architectures With Extended Noise Shaping. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2018**, *65*, 444–457. [[CrossRef](#)]
36. Taylor, G.; Galton, I. A Reconfigurable Mostly-Digital Delta-Sigma ADC With a Worst-Case FOM of 160 dB. *IEEE J. Solid-State Circuits* **2013**, *48*, 983–995. [[CrossRef](#)]
37. Nguyen, V.; Schembari, F.; Staszewski, R.B. A 0.2-V 30-MS/s 11b-ENOB Open-Loop VCO-Based ADC in 28-nm CMOS. *IEEE J. Solid-State Circuits Lett.* **2018**, *1*, 190–193. [[CrossRef](#)]
38. Ahmadi-Farsani, J.; Rosa, J.M.d. Bulk-Input VCO-Based Sigma-Delta ADCs with Enhanced Linearity in 28-nm FD-SOI CMOS. In Proceedings of the 2019 IEEE International Symposium on Circuits and Systems (ISCAS), Sapporo, Japan, 26–29 May 2019; pp. 1–5.
39. Fishani, A.B.; Rombouts, P. Highly linear VCO for use in VCO-ADCs. *Electron. Lett.* **2016**, *52*, 268–270. [[CrossRef](#)]
40. Unnikrishnan, V.; Vesterbacka, M. Time-Mode Analog-to-Digital Conversion Using Standard Cells. *IEEE Trans. Circuits Syst. Regul. Pap.* **2014**, *61*, 3348–3357. [[CrossRef](#)]
41. Wulff, C.; Ytterdal, T. A Compiled 9-bit 20-MS/s 3.5-fJ/conv.step SAR ADC in 28-nm FDSOI for Bluetooth Low Energy Receivers. *IEEE J. Solid-State Circuits* **2017**, *52*, 1915–1926. [[CrossRef](#)]
42. Weaver, S.; Hershberg, B.; Maghari, N.; Moon, U. Domino-Logic-Based ADC for Digital Synthesis. *IEEE Trans. Circuits Syst. II Express Briefs* **2011**, *58*, 744–747. [[CrossRef](#)]
43. Ding, M.; Harpe, P.; Chen, G.; Busze, B.; Liu, Y.; Bachmann, C.; Philips, K.; van Roermund, A. A Hybrid Design Automation Tool for SAR ADCs in IoT. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **2018**, *26*, 2853–2862. [[CrossRef](#)]
44. Waters, A.; Moon, U.-K. A fully automated verilog-to-layout synthesized ADC demonstrating 56dB-SNDR with 2MHz-BW. In Proceedings of the 2015 IEEE Asian Solid-State Circuits Conference (A-SSCC), Xiamen, China, 9–11 November 2015; pp. 1–4.
45. Aiello, O.; Crovetti, P.; Sharma, A.; Alioto, M. Fully-Synthesizable Current-Input ADCs for Ultra-Low Area and Minimal Design Effort. In Proceedings of the 2019 26th IEEE International Conference on Electronics, Circuits and Systems (ICECS), Genova, Italy, 27–29 November 2019; pp. 715–718.
46. McConaghy, T.; Palmers, P.; Steyaert, M.; Gielen, G.G.E. Trustworthy Genetic Programming-Based Synthesis of Analog Circuit Topologies Using Hierarchical Domain-Specific Building Blocks. *IEEE Trans. Evol. Comput.* **2011**, *15*, 557–570. [[CrossRef](#)]
47. Gielen, G.; Eeckelaert, T.; Martens, E.; McConaghy, T. Automated synthesis of complex analog circuits. In Proceedings of the 2007 18th European Conference on Circuit Theory and Design, Sevilla, Spain, 26–30 August 2007; pp. 20–23.
48. Rutenbar, R.A.; Gielen, G.G.E.; Roychowdhury, J. Hierarchical Modeling, Optimization, and Synthesis for System-Level Analog and RF Designs. *Proc. IEEE* **2007**, *95*, 640–669. [[CrossRef](#)]
49. der Plas, G.V.; Vandebussche, J.; Gielen, G.G.E.; Sansen, W. A layout synthesis methodology for array-type analog blocks. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **2002**, *21*, 645–661. [[CrossRef](#)]
50. Rabaey, J.M. *Digital Integrated Circuits: A Design Perspective*; Prentice-Hall, Inc.: Upper Saddle River, NJ, USA, 1996.
51. Wang, C.; Tsai, J.; Su, S.; Tsai, J.; Chen, J.; Lou, C. 20.6 An 80MHz-BW 31.9fJ/conv-step Filtering $\Delta\Sigma$ ADC with a Built-In DAC-Segmentation/ELD-Compensation 6b 960MS/s SAR-Quantizer in 28nm LP for 802.11ax Applications. In Proceedings of the 2019 IEEE International Solid-State Circuits Conference—(ISSCC), San Francisco, CA, USA, 17–21 February 2019; pp. 338–340.
52. Borgmans, J.; Rombouts, P. Enhanced circuit for linear ring VCO-ADCs. *Electron. Lett.* **2019**, *55*, 583–585. [[CrossRef](#)]
53. Tu, C.; Wang, Y.; Lin, T. A Low-Noise Area-Efficient Chopped VCO-Based CTDSM for Sensor Applications in 40-nm CMOS. *IEEE J. Solid-State Circuits* **2017**, *52*, 2523–2532. [[CrossRef](#)]
54. Taylor, G. Mostly Digital ADCs for Highly-Scaled CMOS Processes. Ph.D. Thesis, University of California, Oakland, CA, USA, 2011.

Curriculum Vitae

Personal data

Name: Vahur Kampus
Date of birth: 17.12.1984
Place of birth: Viljandi, Estonia
Citizenship: Estonian

Contact data

E-mail: vahur.kampus@taltech.ee

Education

2011–2025 Tallinn University of Technology, PhD
2004–2009 Tallinn University of Technology, MSC

Language competence

English Native
Estonian Native
German Fluent
Japanese Basic Knowledge

Professional employment

2020–2024 **Maxlinear**
Senior Principal RF/MS IC Design Engineer
2012–2020 **Intel Corporation**
Senior Staff Mixed-Signal IC Design Engineer
2008–2012 **Texas Instruments Incorporated**
Senior Analog IC Design Engineer
2007–2008 **Leab Group**
Test Engineer

Elulookirjeldus

Isikuandmed

Nimi: Vahur Kampus
Sünniaeg: 17.12.1984
Sünnikoht: Viljandi, Eesti
Kodakondsus: Eesti

Kontaktandmed

E-post: vahur.kampus@taltech.ee

Hariduskäik

2011–2025 Tallinna Tehnikaülikool, PhD
2004–2009 Tallinna Tehnikaülikool, MSC

Keelteoskus

Inglise keel kõrgtase
Eesti keel kõrgtase
Saksa keel kõrgtase
Jaapani keel algtase

Teenistuskäik

2020–2024 **Maxlinear**
Segasignaalkiipide vanem peadisainiinsener
2012–2020 **Intel Corporation**
Segasignaalkiipide vanem kaadri disainiinsener
2008–2012 **Texas Instruments Incorporated**
Vanem kaadri disainiinsener
2007–2008 **Leab Group**
Testiinsener

ISSN 2585-6901 (PDF)
ISBN 978-9916-80-327-1 (PDF)