

ACSP · Analog Circuits And Signal Processing

Athanasios T. Ramkaj
Marcel J. M. Pelgrom
Michiel S. J. Steyaert
Filip Tavernier

Multi-Gigahertz Nyquist Analog-to-Digital Converters

Architecture and Circuit Innovations
in Deep-Scaled CMOS and FinFET
Technologies

 Springer

Analog Circuits and Signal Processing

Series Editors

Mohammed Ismail, Khalifa University, Dublin, OH, USA

Mohamad Sawan, 18, Shilongshan Road, School of Engineering, Westlake University, Hangzhou, Zhejiang, China

The *Analog Circuits and Signal Processing* book series, formerly known as the *Kluwer International Series in Engineering and Computer Science*, is a high level academic and professional series publishing research on the design and applications of analog integrated circuits and signal processing circuits and systems. Typically per year we publish between 5-15 research monographs, professional books, handbooks, and edited volumes with worldwide distribution to engineers, researchers, educators, and libraries. The book series promotes and expedites the dissemination of new research results and tutorial views in the analog field. There is an exciting and large volume of research activity in the field worldwide. Researchers are striving to bridge the gap between classical analog work and recent advances in very large scale integration (VLSI) technologies with improved analog capabilities. Analog VLSI has been recognized as a major technology for future information processing. Analog work is showing signs of dramatic changes with emphasis on interdisciplinary research efforts combining device/circuit/technology issues. Consequently, new design concepts, strategies and design tools are being unveiled. Topics of interest include: Analog Interface Circuits and Systems; Data converters; Active-RC, switched-capacitor and continuous-time integrated filters; Mixed analog/digital VLSI; Simulation and modeling, mixed-mode simulation; Analog nonlinear and computational circuits and signal processing; Analog Artificial Neural Networks/Artificial Intelligence; Current-mode Signal Processing; Computer-Aided Design (CAD) tools; Analog Design in emerging technologies (Scalable CMOS, BiCMOS, GaAs, heterojunction and floating gate technologies, etc.); Analog Design for Test; Integrated sensors and actuators; Analog Design Automation/Knowledge-based Systems; Analog VLSI cell libraries; Analog product development; RF Front ends, Wireless communications and Microwave Circuits; Analog behavioral modeling, Analog HDL.

Athanasios T. Ramkaj • Marcel J. M. Pelgrom •
Michiel S. J. Steyaert • Filip Tavernier

Multi-Gigahertz Nyquist Analog-to-Digital Converters

Architecture and Circuit Innovations in
Deep-Scaled CMOS and FinFET
Technologies

 Springer

Athanasios T. Ramkaj
Stanford University
Stanford, CA, USA

Marcel J. M. Pelgrom
Helmond, Noord-Brabant
The Netherlands

Michiel S. J. Steyaert
KU Leuven
Leuven, Belgium

Filip Tavernier
KU Leuven
Leuven, Belgium

ISSN 1872-082X ISSN 2197-1854 (electronic)
Analog Circuits and Signal Processing
ISBN 978-3-031-22708-0 ISBN 978-3-031-22709-7 (eBook)
<https://doi.org/10.1007/978-3-031-22709-7>

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2023

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To Theodor and Linda

Preface

The analog-to-digital converter (ADC) is considered the cornerstone of modern electronics due to its fundamental role in virtually any application requiring the transfer of information between the physical (analog) world and the processing (digital) world. This task comes with myriad challenges due to the complex multi-functional ADC nature, further exacerbated when the relevant applications demand stringent performance requirements. Furthermore, bridging the analog and digital worlds fundamentally implies that ADCs must deal with the non-idealities of the former while keeping up with the advancements of the latter.

The rapidly accelerating trend for broader-band signals and software-defined systems has spurred the need for ADCs operating in the multi-GHz sample rate and bandwidth regime. Such converters are highly demanded by applications in the realm of next generation high-speed wireless and wireline communications, automotive radar, and high-end instrumentation, and have attracted a growing attention from both industry and academia. The ever-increasing desire of these systems is to maximize speed, while progressively improving the accuracy and the power efficiency, pushing the performance dimensions to new benchmarks. Meeting these requirements at the multi-GHz regime comes with numerous challenges at the circuit, architecture, and system levels. On top, the constant technology downscaling, dictated by the demand for higher functionality at a reduced power and cost, and the improvement in digital performance, exacerbates these challenges for traditional analog-intensive solutions.

This book follows a holistic approach, from analysis to implementation, to propose innovative circuit, architecture, and system solutions in deep-scaled CMOS and maximize the $accuracy \cdot speed \div power$ of multi-GHz sample rate and bandwidth ADCs. The approach starts by identifying the major error sources of any practical converter's circuits and quantitatively analyzing their significance on the overall performance, establishing the fundamental accuracy-speed-power limits imposed by circuits, and building an understanding as to what may be achievable from a converter's elementary building blocks. The analysis extends to the architecture level, by introducing a mathematical framework to estimate and compare the accuracy-speed-power limits of high-performance architectures, such as flash, SAR, pipeline,

and pipelined-SAR. To gain insight on the system level and peripheral blocks, a framework is introduced to quantitatively compare interleaver architectures, in terms of achievable bandwidth and sampling accuracy. The strength of the newly introduced frameworks is further enhanced by adding technology effects from four deep-scaled CMOS processes: 65 nm, 40 nm, 28 nm, and 16 nm, building insight into both architecture as well as process choices for optimum performance at given specifications.

The validity of the above holistic approach and the feasibility of the proposed solutions are demonstrated by four prototype ICs, realized in 28 nm bulk CMOS and 16 nm FinFET CMOS:

1. An ultrahigh-speed three-stage triple-latch feed-forward dynamic comparator improves the gain and reduces the delay of dynamic comparators across the entire input range. [28 nm CMOS, presented at *ESSCIRC 2019*, and published in *SSC-L 2019* and *TCAS-I 2022*]
2. A high-speed wide-bandwidth medium resolution single-channel SAR ADC maximizes the $accuracy \cdot speed \div power$ ratio with a semi-asynchronous timing, an improved bootstrapped input switch, a triple-tail dynamic comparator, and a Unit-Switch-Plus-Cap DAC. [28 nm CMOS, presented at *ESSCIRC 2017*, and published in *JSSC 2018*]
3. A high-resolution wide-bandwidth 8 \times -interleaved hybrid RF ADC with a bufferless input front end, a 3-stage pipelined-SAR sub-ADC, a low excess jitter clock chain, and co-designed analog–digital calibrations significantly improves the state of the art in RF ADCs. [28 nm CMOS, presented at *ISSCC 2019*, and published in *JSSC 2020*]
4. An ultra-wideband highly linear analog front end with a multi-segment distributed attenuation filter and a hybrid amplifier-buffer extends the bandwidth of next-generation direct RF ADC-based receivers to several tens of GHz, enabling direct RF sampling up to mm-wave frequencies. [16 nm FinFET CMOS, presented at *VLSI 2022*, and two US patents]

Stanford, CA, USA
Helmond, The Netherlands
Leuven, Belgium
Leuven, Belgium
September 2022

Athanasios T. Ramkaj
Marcel J. M. Pelgrom
Michiel S. J. Steyaert
Filip Tavernier

Acknowledgments

We would like to acknowledge several people and organizations for their valuable contributions to the work in this book.

We thank Boris Murmann (Stanford University, CA, USA), Gabriele Manganaro (MediaTek Inc., MA, USA), Marian Verhelst (KU Leuven, Belgium), and Patrick Wambacq (KU Leuven, Belgium) for their feedback on previous versions of this manuscript and for all the engaging and inspiring discussions throughout the last years. We would also like to acknowledge our collaborators' contributions in the articles discussed throughout this book: Adalberto Cantoni (Analog Devices Inc., MA, USA), Siddharth Devarajan (Analog Devices Inc., MA, USA), Juan Carlos Peña Ramos (ICsense, Leuven, Belgium), and Maarten Strackx (MAGICS Instruments, Geel, Belgium).

Finally, the authors thank Nokia Bell Labs, Antwerp, Belgium, and Analog Devices, Wilmington, MA, USA for their financial support.

Contents

1	Introduction	1
1.1	Data Converters in a Digital Era: Need and High-Performance Applications	1
1.2	Challenges in Pushing Performance Boundaries	5
1.2.1	ADC Core and Peripherals Challenges	6
1.2.2	The Good, the Bad, and the Ugly of Deep-Scaled CMOS	7
1.3	Research Goal and Objectives	10
1.4	Structure of This Book	12
2	Analog-to-Digital Conversion Fundamentals	15
2.1	Theoretical Background	15
2.1.1	Sampling	16
2.1.2	Ideal Quantization	21
2.2	Error Sources	26
2.2.1	Noise	26
2.2.2	Non-linearity	30
2.2.3	Calibration	33
2.3	Performance Evaluation	34
2.3.1	Metrics	34
2.3.2	Figures of Merit	36
2.4	Accuracy-Speed-Power Limits	37
2.4.1	Sampler Noise Limit	38
2.4.2	Quantizer Noise Limit	41
2.4.3	Metastability Limit	42
2.4.4	Aperture Jitter Limit	47
2.4.5	Mismatch Limit	48
2.4.6	Heisenberg Uncertainty Principle	51
2.4.7	Putting It All Together	53
2.5	Conclusion	54
	Appendix A: Proper FFT Evaluation Setup	55

3	Architectural Considerations for High-Efficiency GHz-Range ADCs	57
3.1	State of the Art	57
3.2	The Flash Architecture	60
3.2.1	Overview	60
3.2.2	Flash Accuracy-Speed-Power Limits	62
3.2.3	Impact of Scaling	65
3.3	The SAR Architecture	69
3.3.1	Overview	69
3.3.2	The DAC in a SAR	70
3.3.3	SAR Accuracy-Speed-Power Limits	78
3.4	The Pipeline Architecture	81
3.4.1	Overview	81
3.4.2	Pipeline Accuracy-Speed-Power Limits	83
3.5	The Pipelined-SAR: A Powerful Hybrid	88
3.5.1	Overview	88
3.5.2	Pipelined-SAR Accuracy-Speed-Power Limits	93
3.6	Architectural Limits' Comparison	98
3.7	Time-Interleaving	102
3.7.1	Overview	102
3.7.2	Interleaving Errors	104
3.7.3	Interleaver Architectures	109
3.8	Conclusion	118
	Appendix B: Transconductance—Settled RA	119
	Appendix C: Transconductance—Integrator RA	120
4	Ultrahigh-Speed High-Sensitivity Dynamic Comparator	121
4.1	Dynamic Regenerative Comparators	121
4.1.1	Single-Stage Latch-Based Strong-ARM Comparator	122
4.1.2	Two-Stage Double-Tail Latched Comparator	125
4.2	Prototype IC: A 28 nm CMOS Three-Stage Triple-Latch Feed-Forward Comparator	128
4.2.1	Circuit Operation and Analysis	128
4.2.2	Simulation and Comparison with Prior Art	133
4.2.3	On-Chip Delay Evaluation Setup	136
4.3	Experimental Verification	137
4.3.1	Measurement Setup	139
4.3.2	Measurement Results	140
4.3.3	State-of-the-Art Comparison	145
4.4	Conclusion	147
5	High-Speed Wide-Bandwidth Single-Channel SAR ADC	149
5.1	Pushing the SAR Conversion Speed	149
5.1.1	Conventional Synchronous Clocking Scheme	150
5.1.2	Speed-Boosting Techniques	151
5.2	Prototype IC: A 1.25 GS/s 7-bit SAR ADC in 28 nm CMOS	156
5.2.1	High-Level Design	156

- 5.2.2 Semi-asynchronous Processing w/o Logic Delay 157
- 5.2.3 Dual-Loop Bootstrapped Input Switch 159
- 5.2.4 Unit-Switch-Plus-Cap DAC 162
- 5.2.5 Triple-Tail Dynamic Comparator 166
- 5.2.6 Custom SAR Logic 171
- 5.3 Experimental Verification..... 173
 - 5.3.1 Measurement Setup 174
 - 5.3.2 Measurement Results 175
 - 5.3.3 State-of-the-Art Comparison 178
- 5.4 Conclusion 179
- 6 High-Resolution Wide-Bandwidth Time-Interleaved RF ADC 183**
 - 6.1 RF Sampling ADCs: Needs and Challenges 183
 - 6.1.1 The ADC Role in the Receiver 184
 - 6.1.2 ADC Architectural Trade-Offs 185
 - 6.2 Prototype IC: A 5 GS/s 12-bit Hybrid TI-ADC in 28 nm CMOS 188
 - 6.2.1 High-Level Design 188
 - 6.2.2 Interleaving Factor and Sub-ADC Architecture 189
 - 6.2.3 Passive Input Front-End 190
 - 6.2.4 Clock Generation and Distribution 194
 - 6.2.5 Hybrid Sub-ADC Design 199
 - 6.2.6 Digital Calibration 203
 - 6.3 Experimental Verification 204
 - 6.3.1 Measurement Setup 205
 - 6.3.2 Measurement Results 207
 - 6.3.3 State-of-the-Art Comparison 213
 - 6.4 Conclusion 214
- Appendix D: TI ADC Power Estimation with On-Chip Input Buffer 216
- 7 Ultra-Wideband Direct RF Receiver Analog Front-End 217**
 - 7.1 Pushing the Bandwidth Beyond 20 GHz 218
 - 7.1.1 Revisiting the Analog Front-End Problem 218
 - 7.1.2 Increasing Integration and Challenges 221
 - 7.2 Prototype IC: A 30 GHz-Bandwidth <-57 dB-IM3 Front-End in 16 nm FinFET CMOS 222
 - 7.2.1 High-Level Front-End Chain 223
 - 7.2.2 Filter with Distributed ESD and Variable Attenuation 224
 - 7.2.3 Two-Path Push-Pull Hybrid Amplifier 230
 - 7.2.4 Push-Pull Bootstrapped Cascoded Buffer 232
 - 7.3 Experimental Verification 236
 - 7.3.1 Measurement Setup 237
 - 7.3.2 Measurement Results 239
 - 7.3.3 State-of-the-Art Comparison 244
 - 7.4 Conclusion 246

8 Conclusions, Contributions, and Future Work	247
8.1 Overview and General Conclusions	247
8.2 Original Scientific Contributions	251
8.3 Suggestions for Future Work	254
Bibliography	257
Index	267

List of Figures

Fig. 1.1	Illustration of the data conversion as an indispensable bridging function between the real analog world and the digital signal processing world	2
Fig. 1.2	Popular ADC applications and architectures covering them	3
Fig. 1.3	Wideband ADC-based heterodyne (top left) and zero-IF (top right) vs. direct RF sampling receiver (bottom)	4
Fig. 1.4	The three main ADC performance parameters and several factors affecting them on different levels	6
Fig. 1.5	CMOS scaling evolution from planar FET to FinFET to GAAFET. Picture credit: Samsung	8
Fig. 1.6	Theoretical cut-off frequency versus channel length scaling [10]	8
Fig. 1.7	Supply and threshold voltage versus channel length scaling	9
Fig. 1.8	BEOL interconnect comparison between 65 nm (left) and 32 nm (right) CMOS processes [11]	10
Fig. 2.1	Block diagram of an ideal A/D conversion (top) and the resulting waveforms at every part of the chain (bottom)	16
Fig. 2.2	Sampling a continuous-time signal using a Dirac pulse sequence	17
Fig. 2.3	Frequency spectrum of a signal multiplied with a sequence of Dirac pulses	18
Fig. 2.4	(a) Single-tone signals with different frequencies (b) fall in the same frequency location after spectrum processing	18
Fig. 2.5	Dual-sided frequency spectrum highlighting different Nyquist zones	19
Fig. 2.6	(a), (b) Two cases of signals with bands meeting the Nyquist criterion and (c) one scenario where bands are overlapping leading to information loss	19

Fig. 2.7	Anti-aliasing filter on a parasitic tone when (a) sampling at Nyquist rate (slightly oversampled in practice) and (b) oversampling by $M > 1$	20
Fig. 2.8	Conceptual model and transfer characteristic of an ideal quantizer	22
Fig. 2.9	Sawtooth approximation of ϵ_q as a function of time	22
Fig. 2.10	Uniformly distributed PDF of ϵ_q within $\pm\Delta/2$	23
Fig. 2.11	Frequency spectra of an ideally quantized with various resolutions 77 MHz signal sampled at 1 GS/s ($N_{\text{FFT}} = 1024$)	24
Fig. 2.12	Conceptual model of a real converter including error sources from the different blocks	26
Fig. 2.13	(a) Simple model of a sampler and (b) its noise spectrum	27
Fig. 2.14	(a) Simple quantizer model and (b) its allowed operation time	28
Fig. 2.15	(a) Sampler with jitter and (b) time to voltage error translation	29
Fig. 2.16	(a) DNL in transfer characteristic with corresponding curve and (b) INL in transfer characteristic with corresponding curve	31
Fig. 2.17	Simple sampler model with input termination network	39
Fig. 2.18	Fundamental limits due to sampler noise: (a) accuracy-speed and (b) accuracy-power	40
Fig. 2.19	Fundamental limits due to quantizer noise: (a) accuracy-speed and (b) accuracy-power	43
Fig. 2.20	Quantizer output for a valid (gray) and a metastable (black) case	44
Fig. 2.21	Fundamental limits due to metastability of a standalone quantizer: (a) accuracy-speed and (b) accuracy-power	46
Fig. 2.22	Simple model for clock power estimation for a certain jitter	48
Fig. 2.23	Fundamental limits due to aperture jitter: (a) accuracy-speed and (b) accuracy-power	49
Fig. 2.24	Limits imposed by mismatch: (a) accuracy-speed and (b) accuracy-power	51
Fig. 2.25	Fundamental accuracy-speed limit due to Heisenberg	52
Fig. 2.26	Fundamental limit curves from all the error sources analyzed in this chapter: (a) accuracy-speed and (b) accuracy-power	53
Fig. 3.1	State-of-the-art performance of various ADC architectures with data points taken from [36] (a) accuracy-speed and (b) accuracy-energy	58
Fig. 3.2	Block diagram of a B-bit flash ADC (the S/H is optional)	60
Fig. 3.3	Simplified small-signal model of an NMOS transistor (bulk is omitted for simplicity)	66

Fig. 3.4	f_T vs. g_m/I_D and $f_T \cdot g_m/I_D$ vs. g_m/I_D in four CMOS processes	67
Fig. 3.5	Flash accuracy-speed-power limits: (a) for different f_s in 28 nm and (b) at $f_s = 4$ GHz in the processes under comparison	68
Fig. 3.6	Block diagram of a B -bit SAR ADC	69
Fig. 3.7	(a) Scale equivalent of a binary SA algorithm and (b) waveform operation in the voltage vs. time domain	71
Fig. 3.8	3-bit example of the conventional CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints	72
Fig. 3.9	3-bit example of the split-capacitor CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints	73
Fig. 3.10	3-bit example of the energy-saving CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints	74
Fig. 3.11	3-bit example of the monotonic CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints	75
Fig. 3.12	3-bit example of the MCS CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints	76
Fig. 3.13	Switching energy for the different CDAC switching schemes	77
Fig. 3.14	SAR accuracy-speed-power limits: (a) for different f_s in 28 nm and (b) for $f_s = 500$ MHz in the processes under comparison	80
Fig. 3.15	Block diagram of a B -bit m -stages pipeline ADC	81
Fig. 3.16	Residue plot of stage- s : (a) ideal case with $A_s = 2^{B_s}$, (b) $A_s = 2^{B_s}$ with error and no OR, and (c) $A_s = 2^{B_s-1}$ with error and $2\times$ -OR	82
Fig. 3.17	Basic $g_m - C$ amplifier for modeling the RA gain stage	84
Fig. 3.18	Pipeline with 1,2,3,4-bit/stage effective resolution accuracy-speed-power limits in 28 nm: (a) $f_s = 500$ kHz, (b) $f_s = 500$ MHz, and (c) $f_s = 1.3$ GHz	89
Fig. 3.19	Pipeline accuracy-speed-power limits across different processes at $f_s = 500$ MHz: (a) 1-bit/stage, (b) 2-bit/stage, and (c) 3-bit/stage	90
Fig. 3.20	Block diagram of a B -bit two-stage pipelined-SAR ADC	91
Fig. 3.21	Illustration of conversion energy requirement in (a) a binary SAR ADC and (b) a two-stage pipelined-SAR ADC	92

Fig. 3.22	2,3,4,5-stage pipelined-SAR with accuracy-speed-power limits in 28 nm: (a) $f_s = 500$ kHz, (b) $f_s = 500$ MHz, and (c) $f_s = 1.3$ GHz	96
Fig. 3.23	Pipelined-SAR accuracy-speed-power limits across different processes at $f_s = 500$ MHz: (a) three-stage, (b) four-stage, and (c) five-stage	97
Fig. 3.24	Accuracy-speed-power limits for the different ADC architectures studied at $f_s = 500$ kHz	99
Fig. 3.25	Accuracy-speed-power limits for the different ADC architectures studied at $f_s = 500$ MHz	100
Fig. 3.26	Accuracy-speed-power limits for the different ADC architectures studied at $f_s = 1.3$ GHz	101
Fig. 3.27	(a) High-level block diagram of an N -channel TI-ADC and (b) sampling of a signal using an N -interleaved Dirac pulse sequence	103
Fig. 3.28	Power vs. frequency illustration of a non-TI- and a TI-ADC	103
Fig. 3.29	Illustration of mismatch errors in a four-channel TI-ADC example	105
Fig. 3.30	Graphical illustration of sub-ADC offset mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum	105
Fig. 3.31	Graphical illustration of sub-ADC gain mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum	106
Fig. 3.32	Graphical illustration of sub-ADC timing mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum	107
Fig. 3.33	Graphical illustration of sub-ADC bandwidth mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum	108
Fig. 3.34	Simulated SNDR vs. (a) σ_{OS}/V_{DD} , (b) σ_G/G , (c) $\sigma_{\Delta T}/T_s$, (d) σ_{BW}/BW , and (e) σ_{BW}/BW with separated gain/phase and (f) combined errors	110
Fig. 3.35	Classification tree for different interleaver architectures	110
Fig. 3.36	(a) Direct interleaver architecture and timing diagram for $N = 8$ (b) with 50% duty-cycle clocks and (c) with (1/8)-100% duty-cycle clocks	111
Fig. 3.37	(a) Interleaver architecture with a hierarchical $N = L \times K$ de-multiplexing and (b) timing diagram for $N = 8$ with $L \times K = 2 \times 4$	112
Fig. 3.38	(a) Interleaver architecture with $N = L \times K$ re-sampling hierarchy and (b) timing diagram for $N = 8$ with $L \times K = 2 \times 4$	113

Fig. 3.39	Equivalent RC model for (a) interleaver and (b) simple switch	114
Fig. 3.40	Bandwidth vs. channel count for different interleavers in (a) 65 nm, (b) 40 nm, (c) 28 nm, and (d) 16 nm CMOS	115
Fig. 3.41	Sampling accuracy vs. channel count for the interleavers in 28 nm at (a) $Nf_s = 2.5$ GHz, (b) $Nf_s = 5.0$ GHz, (c) $Nf_s = 7.5$ GHz, and (d) $Nf_s = 10$ GHz	117
Fig. 4.1	Single-stage strong-ARM comparator and its signal waveforms	122
Fig. 4.2	Double-tail comparator and its signal waveforms	126
Fig. 4.3	Proposed three-stage TLFF dynamic comparator	129
Fig. 4.4	LTV representation of the proposed TLFF comparator	130
Fig. 4.5	Simulated timing waveforms of the proposed TLFF comparator	134
Fig. 4.6	Simulated outputs and delay versus ΔV_I for different comparators	135
Fig. 4.7	Top-level diagram of the multiple comparators test chip	137
Fig. 4.8	Die photo of the 28 nm IC with zoomed-in comparator layout views	138
Fig. 4.9	Measurement setup of the multiple comparators test chip with the prototype 13.5 Gb/s TLFF comparator	139
Fig. 4.10	Measured CLK-OUT delays for the TLFF, SAC, and DTC comparators versus (a) ΔV_I , (b) V_{CM} , and (c) V_{DD}	141
Fig. 4.11	Measured noise cumulative distribution and Gaussian distribution fitting curve for (a) the SAC, (b) the DTC, (c) the TLFF, and (d) measured input-referred noise vs. V_{CM}	143
Fig. 4.12	Measured (a) energy consumption of the TLFF, SAC, and DTC versus V_{DD} and (b) energy-delay product versus V_{DD}	144
Fig. 4.13	(a) Measured ΔV_I eye, (b) measured bathtub curve of the TLFF, (c) measured CLK eye, and (d) measured OUT eye	145
Fig. 5.1	Timing sequence illustration of a B -bit SAR with a conventional synchronous clocking scheme	150
Fig. 5.2	Timing sequence illustration of a B -bit SAR with an internally asynchronous clocking scheme	151
Fig. 5.3	Timing sequence illustration of a B -bit SAR with a multi-bit per cycle resolving scheme (2-bit per cycle shown in the example)	153
Fig. 5.4	Timing sequence illustration of a B -bit SAR with multiple comparators loop unrolling clocking scheme	153
Fig. 5.5	Timing sequence illustration of a B -bit SAR with extra cycles and redundancy implemented (one extra cycle shown in the example)	155
Fig. 5.6	Top-level architecture of the proposed ADC and its timing diagram	157

Fig. 5.7	Implemented semi-asynchronous scheme with the logic delay eliminated from the critical path	158
Fig. 5.8	Typical bootstrap circuit with its speed critical loop highlighted	159
Fig. 5.9	Improved dual-loop bootstrap circuit proposed in this work	160
Fig. 5.10	(a) Timing illustration and (b) simulated M_S on-resistance for the typical and the proposed bootstrap circuit	161
Fig. 5.11	DAC topology with a constant V_{CM} and C_H to set the signal range	162
Fig. 5.12	Schematic and simulated settling time of (a) a conventional UC CDAC and (b) the proposed USPC CDAC	164
Fig. 5.13	Single-ended partial layout of the USPC CDAC (the actual implementation is differential)	165
Fig. 5.14	Schematic of the implemented triple-tail dynamic comparator	166
Fig. 5.15	Simulated performance of the triple-tail comparator	167
Fig. 5.16	Simulated outputs of the triple-tail comparator (top) and one of the logic memory latches (bottom)	168
Fig. 5.17	(a) Simulated comparator resolving time and (b) input-referred noise versus $V_{CM,I}$ and (c) resolving time and (d) energy versus ΔV_I	170
Fig. 5.18	Custom SAR logic including the comparator clock, bit phases, and memory elements	171
Fig. 5.19	Schematic of one memory cell with optimized critical path toward the CDAC reference switches (top). Timing diagram and truth table of the memory cell (bottom)	172
Fig. 5.20	Die micrograph of the 28 nm IC with a zoomed-in view of the SAR core occupying an active area of 0.0071 mm^2	173
Fig. 5.21	Measurement setup of the 1.25 GS/s 7-bit SAR prototype	174
Fig. 5.22	Measured static performance with the histogram (code density) test at 1.25 GS/s for a sinusoidal input of 160 kHz: (a) DNL and (b) INL	176
Fig. 5.23	Measured output spectra at 1.25 GS/s for (a) a Nyquist input frequency and (b) an $8\times$ Nyquist input frequency	177
Fig. 5.24	Measured SFDR/SNDR versus (a) input frequency at 1.25 GS/s and (b) sample rate for a 76 MHz input	178
Fig. 5.25	Measured FoM versus (a) input frequency at 1.25 GS/s and (b) sample rate for a 76 MHz input	179
Fig. 6.1	Generic block diagram of a direct RF sampling receiver	184
Fig. 6.2	Sub-ADC and interleaving overhead vs. channel count illustration	186
Fig. 6.3	Major design strategies regarding the choice of the sub-ADC and the interleaving factor	186

Fig. 6.4	Accuracy-speed standings of the ADCs adopting the two design strategies. Points taken from [36]	187
Fig. 6.5	Top-level diagram of the complete 5 GS/s 12-bit TI-ADC architecture (single-ended shown for simplicity)	189
Fig. 6.6	Passive front-end model of this ADC (single-ended shown)	190
Fig. 6.7	(a) Bootstrap circuit employed for S_{IN} and (b) timing waveforms of the important nodes	192
Fig. 6.8	Proposed intertwined input/clock Y-tree structure to minimize the front-end loading	192
Fig. 6.9	Simulated (a) S-parameters and (b) input impedance of this front-end	194
Fig. 6.10	Input current profile of this front-end for (a) 300 MHz and (b) 2.4 GHz input frequencies	195
Fig. 6.11	Simulated HD2 vs. differential input-clock coupling (unbalancing) for a near-2.5 GHz input	195
Fig. 6.12	Timing diagram with the generated clocks of the TI-ADC	196
Fig. 6.13	(a) Simulated SNR vs. σ_{jitter} and (b) SFDR vs. σ_{skew} at Nyquist	196
Fig. 6.14	Block diagram of the proposed clock conditioning chain for this ADC	197
Fig. 6.15	(a) DDT circuit with simulated (b) sampling edge skew and tuning range and (c) capacitance spread of one DDT unit cell	198
Fig. 6.16	Detailed block diagram of the implemented 12-bit three-stage pipelined-SAR sub-ADC	200
Fig. 6.17	Dynamic integrator RA with simulated SAR ₁ - SAR ₂ residue	201
Fig. 6.18	Sub-ADC internal asynchronous timing sequence with re-timing	202
Fig. 6.19	One slice top-level diagram of the 8 \times synthesized correction block	203
Fig. 6.20	Die micrograph of the 28 nm complete IC with a sub-ADC layout view occupying a core area of 0.015 mm ²	205
Fig. 6.21	Measurement setup of the 12-bit 5 GS/s TI ADC prototype (top). Photo of the overall setup (bottom-left). Closer view of the motherboard with the four-layer Rogers 4350 chip board mounted on its center through high-speed Samtec connectors (bottom-right). The bare die is placed in a plated cavity	206
Fig. 6.22	Measured (black solid curve) and simulated (gray dotted curve) ADC transfer characteristic showing a bandwidth in excess of 6 GHz	208
Fig. 6.23	Measured calibrated output spectra at 5 GS/s for (a) 75 MHz, (b) 2.4 GHz, and (c) 4.8 GHz input frequencies	209

Fig. 6.24	Measured SFDR/SNDR versus (a) input frequency at 5 GS/s and (b) sample rate for a 2.4 GHz input	210
Fig. 6.25	Measured static performance at 5 GS/s for a sinusoidal input of 7.4 MHz: (a) DNL and (b) INL	211
Fig. 6.26	Measured output spectrum at 5 GS/s for a -6.1 dBFS two-tone input signal at 74.5 MHz 81.7 MHz	212
Fig. 6.27	Measured output spectrum at 5 GS/s for a -6.4 dBFS two-tone input signal at 1.67 GHz 1.85 GHz	212
Fig. 6.28	Measured power partitioning versus sample rate for a 2.4 GHz	213
Fig. 6.29	FoM _S comparison with relevant SotA RF ADCs [36]	213
Fig. 7.1	Two chip-to-substrate illustrations: (a) flip-chip through copper pillars and (b) wire bonding through gold bondwires	220
Fig. 7.2	Buffered front-end model including on-chip/interface/off-chip contributions (single-ended shown)	220
Fig. 7.3	Simulated S-parameters of the front-end model, gradually adding the contributions: (a) S_{21} and (b) S_{11}	221
Fig. 7.4	Top-level block diagram of the proposed front-end (single-ended shown for simplicity)	223
Fig. 7.5	(a) Ideal ninth-order Chebyshev filter with its component values for two R_T values and (b) simulated S-parameters and group delay	226
Fig. 7.6	Implementation of the proposed filter with the distributed ESD and variable attenuation (single-ended shown)	226
Fig. 7.7	Attenuator cells employed in this work: (a) Π -cell, (b) T-cell, and (c) their resistance values rounded to include multiple matched units	227
Fig. 7.8	Simulated bandwidth (relative) and linearity of the two attenuator cells for the different attenuation settings: (a) 1 dB, (b) 2 dB, and (c) 4 dB	228
Fig. 7.9	Simulated S-parameters of the implemented filter across the different attenuation settings: (a) S_{21} and (b) S_{11}	229
Fig. 7.10	Simulated filter two-tone IM3 vs. frequency for the best (0 dB) and worst (11 dB) attenuation settings	230
Fig. 7.11	Proposed push-pull hybrid CG-CS amplifier with resistive source degeneration and series-shunt peaking	231
Fig. 7.12	Push-pull buffer with two-level bootstrapped cascoding	233
Fig. 7.13	Simulated amplifier-buffer transfer characteristic for a capacitive load and the implemented matched load	234
Fig. 7.14	Simulated two-tone IM3 vs. frequency of the amplifier-buffer chain	234
Fig. 7.15	AC-noise simulation of the amplifier-buffer chain with 300 fF load	235

Fig. 7.16	Simulated two-tone IM3 vs. frequency of the complete front-end chain for three different load cases	236
Fig. 7.17	Die micrograph of the 16 nm FinFET IC with front-end occupying a core area of about 0.15 mm ²	237
Fig. 7.18	Measurement setup of the 30 GHz bandwidth front-end prototype	238
Fig. 7.19	Measured front-end small-signal performance for different attenuation settings and six samples	240
Fig. 7.20	Measured one-tone output spectra for (a) 2.5 GHz and (b) 5 GHz input frequencies and (c) measured HD2/HD3 vs. input frequency	241
Fig. 7.21	Measured two-tone spectra at 5 and 20 GHz for 0 dB (left) and 11 dB (right) attenuation settings	242
Fig. 7.22	Measured two-tone IM3 for six samples vs. (a) carrier frequency and (b) tone spacing	242
Fig. 7.23	Measured constellations and spectra at 5 GHz frequency for (a) 1024-QAM and (b) 2048-QAM signals	243
Fig. 7.24	Measured EVM and ACLR vs. frequency for 1024-QAM and 2048-QAM modulated signals	244

List of Tables

Table 2.1	Comparison between calculated and simulated SQNR for different B	25
Table 2.2	Typical process parameters and comparison with kT	50
Table 3.1	Core supply voltage and simulated minimum capacitance in four deep-scaled CMOS processes	65
Table 3.2	Bit partitioning for different aggregate resolutions in a 2,3,4,5-stage pipelined-SAR including 2×-OR between stages	93
Table 4.1	Summary of the TLFF, SAC, and DTC on the same test chip.....	146
Table 4.2	TLFF comparison with state-of-the-art comparators	146
Table 5.1	Performance summary and comparison with state-of-the-art SAR ADCs	180
Table 6.1	RA gain variation with temperature (typical-typical corner).....	202
Table 6.2	Performance summary and comparison with state-of-the-art wideband TI RF ADCs.....	215
Table 7.1	Performance summary and comparison with state-of-the-art ADC-based receiver front-ends.....	245

Chapter 1

Introduction



Real-world phenomena comprise to their vast majority analog quantities; that is continuous-time and continuous-amplitude signals able to take any value at any particular instant. However, manipulation and storage of data are mainly performed in the digital domain due to several benefits of digital signal processing, such as reduced sensitivity to noise and distortion, increased flexibility and reconfigurability, and continuous performance improvement with technology scaling. As a result, Analog-to-Digital (A/D) conversion performed by an Analog-to-Digital Converter (ADC) and Digital-to-Analog (D/A) conversion performed by a Digital-to-Analog Converter (DAC) are indispensable operations in almost all electronic systems.

This introductory chapter starts by briefly outlining the need and applicability of data converters in a digital era, in Sect. 1.1. Key high-performance ADC applications are briefly discussed. Section 1.2 introduces challenges in simultaneously improving the three main ADC performance parameters. These come on a circuit level, an architectural level, a system level, and a technology level, with the last one being particularly important as it affects the other three. The main scope of the research described in this book and its objectives are listed and briefly discussed in Sect. 1.3. Finally, Sect. 1.4 concludes this chapter with the structural organization of this book.

1.1 Data Converters in a Digital Era: Need and High-Performance Applications

Electronic devices undeniably play a crucial role in tremendously improving every aspect of our modern life: from massive communication and transportation infrastructure to personalized entertainment systems and healthcare. To a great extent, this level of accessibility to electronic devices and services owes to the expansion of Digital Signal Processing (DSP), leading to a progressively digital electronic world. The fundamental reason for the DSP advances finds its root

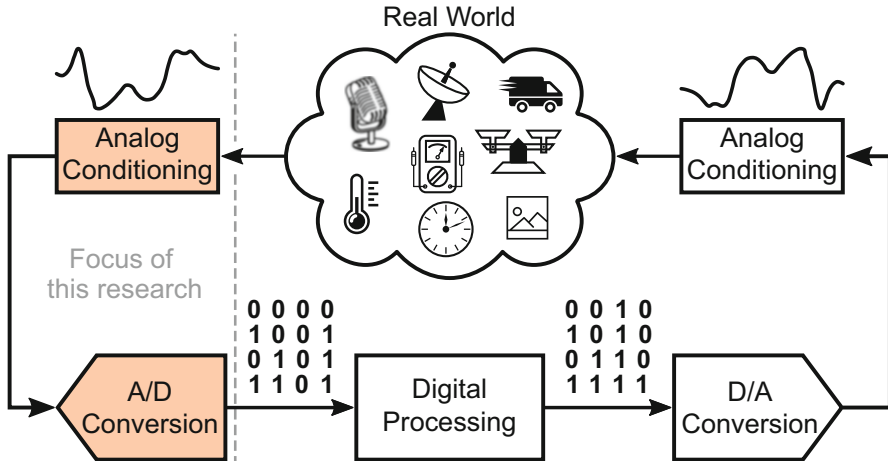


Fig. 1.1 Illustration of the data conversion as an indispensable bridging function between the real analog world and the digital signal processing world

in the down-scaling and integration advantages of Very Large Scale Integration (VLSI) technologies, offering a higher functionality per unit area for a reduced power and cost. Following Gordon E. Moore's law proposed in 1965 [1], the number of devices per chip doubles every roughly 2 years, with an even higher rate recently [2]. Further, digital circuits offer several advantages with respect to their analog equivalents, such as reduced susceptibility to distortion and noise, less dependency on process-voltage-temperature variations, increased flexibility and reconfigurability, and an unprecedented ability to perform complex computations on-demand, making Systems-on-a-Chip (SoC)s the norm.

Yet, real-world phenomena comprise analog quantities, such as velocity, volume, weight, temperature, voltage, current, etc. Therefore, A/D conversion is a vital function in most electronic devices, to provide the translation interface between the physical and electronic worlds, as illustrated in Fig. 1.1. Similarly, D/A conversion enables the interaction of the device with the environment and the humans, who perceive and process data in the form of analog quantities. Even without the human in the loop, the information exchange between two devices still requires conversion of the data from the digital domain to the analog domain and back. At the transmit side, the digital data is converted into the analog domain to travel through a certain medium. Equivalently, at the receive side, the signal comes either in an analog form or in a degraded by the medium digital form; therefore, it is processed as analog information before it is converted into the digital domain. This role of data converters as interfaces between the analog and digital domains puts them in a unique position in the signal processing chain but also poses considerable design challenges since they must deliver an equivalent or better performance than the corresponding digital systems. For the ADC in particular, which is the focus of this work, to maximally leverage the favorable properties of DSP, its function should be

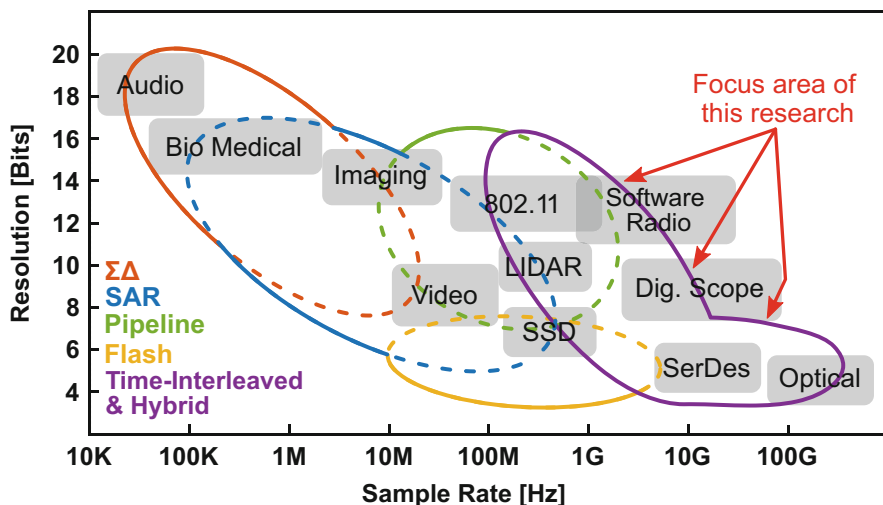


Fig. 1.2 Popular ADC applications and architectures covering them

performed as early as possible in the chain. However, to account for additional and unpredictable signal corruption by the medium, one or more analog conditioning blocks usually precede the converter, whose design is added on the existing set of challenges.

ADCs are employed in a vast and constantly growing number of applications incorporating DSP, either standalone or integrated on the same die or substrate with other blocks composing a larger complex system. More often than not, the performance of the ADC determines the overall performance of the system it is included in. The various applications span a wide range of needs and specifications, including healthcare (diagnostic imaging), consumer electronics (mobile phones, audio, video), automotive (RADAR, LIDAR), next-generation communications (wireline, wireless, optical), and high-end measurement/instrumentation (digital storage oscilloscopes). Some noteworthy applications with established as well as under development specifications are illustrated in Fig. 1.2. To try and meet the demands of such applications, different ADC architectures have been developed. The most widely used up to date are flash, pipeline, Successive Approximation Register (SAR), and sigma-delta ($\Sigma\Delta$). These topologies all have their merits and drawbacks in terms of different specifications, such as sample rate (speed), bandwidth, aggregate resolution, effective resolution (actual), noise, linearity, power consumption, complexity, scalability, etc., making them better tailored for some applications than others. However, as depicted in the conceptual illustration of Fig. 1.2, overlapping target areas exist, such that more than one application can be satisfied by several architectures. To extend the sample rate beyond that of a single converter, time-interleaving has been extensively applied to the above architectures. More recently, hybrid converters have emerged, combining the merits of different

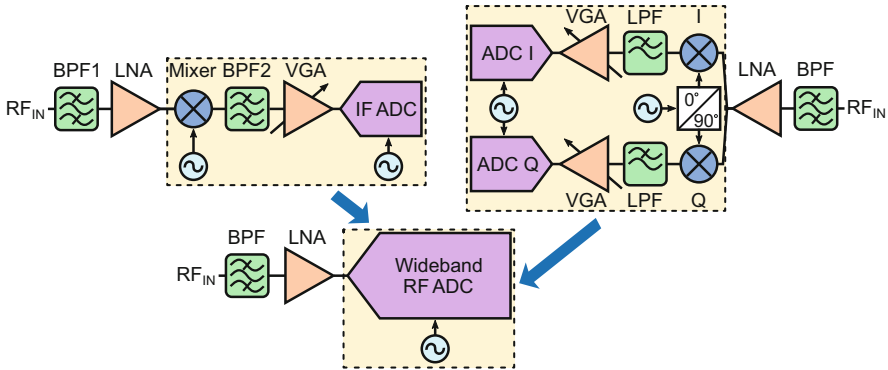


Fig. 1.3 Wideband ADC-based heterodyne (top left) and zero-IF (top right) vs. direct RF sampling receiver (bottom)

architectures,¹ to extend the range of achievable performance and keep up with the rapidly advancing number/demands of future applications.

Undeniably, the field that has established data converters as a hot research topic in both industry and academia for several decades now is that of communications [3]. The constantly increasing demand for higher bandwidth and accuracy in wireline and wireless communication systems is majorly driving the advances in research and development of ADCs, as being key blocks in every receiver. The multi-band requirements of fifth-generation (5G) and future sixth-generation (6G) cellular mobile networks [4] and data over cable networks [5] call for ADCs with high resolution (10–12 bits), multi-GS/s sample rates (5 GS/s or higher), and several GHz of signal bandwidth (half the sample rate or more) while ensuring high linearity (60–70 dB) and low power (preferably below 500 mW). Realizing such ADCs and integrating them with the digital processor in deep-scaled Complementary Metal-Oxide-Semiconductor (CMOS) empower the direct RF sampling receiver topology [6], depicted in Fig. 1.3. This is the closest hardware equivalent to the ideal Software-Defined Radio (SDR) [7]. By leveraging the advanced DSP capabilities in finer CMOS processes, this topology simplifies the analog signal chain and captures multiple bands with a reduced receiver count, lowering area and cost, with improved flexibility and efficiency.

On another high-speed communications front, the rapid emergence of cloud computing and the Internet-of-Things (IoT) has dramatically increased the demand for a higher bandwidth in data center infrastructures. Consequently, the data rates of the transceivers in these systems have reached numbers as high as 112 Gb/s, with future plans to extend to 224 Gb/s and beyond. With Pulse-Amplitude Modulation (PAM-4) currently the prevailing signaling method, for an ADC-based receiver,

¹ One of the proposed prototypes Integrated Circuit (IC)s in this book is a hybrid converter of such kind.

this translates to sample rate requirements of 56 GS/s (112 Gb/s) and a theoretical² 112 GS/s (224 Gb/s) with 6–8-bit resolution. To keep up with these sample rates, suitable ADCs are typically realized with a massive amount of time-interleaving (≥ 64 , rarely ≥ 36) [8, 9]. However, this significantly increases the power and area, on top of deteriorating the signal bandwidth due the large input loading. Further, time-interleaving results in offset, gain, timing, and bandwidth mismatches between the different sub-ADCs, which deteriorate accuracy and require complex calibration schemes, further increasing power and complexity. It is of great interest to develop speed-boosting techniques to increase the sub-ADC sample rate (beyond 1 GS/s), such that the aforementioned interleaving-associated drawbacks are reduced while ensuring negligible accuracy degradation as well as minimizing the power and complexity. Again, to maximally exploit the DSP benefits in advanced CMOS nodes, it is favorable to place the ADC on the same chip; therefore, scalable-friendly solutions are highly desirable.

The above applications require very-high-performance multi-GHz ADCs on multiple specifications and are key drivers for future performance advancements. It is important to note that these specifications are highly desirable to enable next-generation communications, but not entirely achievable.³ As such, they were used to motivate this work in extending the limits beyond what is realizable.

1.2 Challenges in Pushing Performance Boundaries

The previous section briefly discussed some common ADC specifications and requirements for key multi-GHz-range applications. Generally, many different metrics exist, with different importance depending on the target application. However, in a generalized manner, three main parameters encompass the overall ADC performance in a nutshell; these are accuracy, speed, and power. Under accuracy, we include metrics such as aggregate/effective resolution, dynamic range, noise, and linearity. With speed, we denote both sample rate and bandwidth, in the sense that if one increases, the other one must follow. These three main parameters, as illustrated in Fig. 1.4, are bound to each other and influenced by several factors on different levels, such as on a circuit level, an architectural level, a system level, and a technological/process level. These factors present considerable challenges in all the levels, making it non-trivial (even impossible) to simultaneously push all three parameters toward the desired directions, especially at multi-GHz sample rate and bandwidth.

² This is one of the potential future options under consideration with several others being investigated, such as PAM-6 or a different signaling method altogether.

³ At the time of this writing, these applications are undergoing research and prototyping phases to determine viability and long-term reliability.

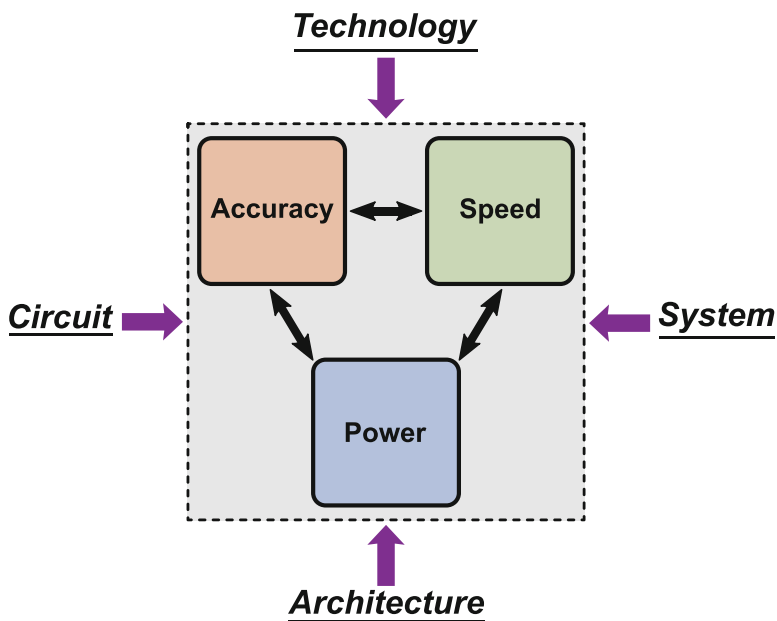


Fig. 1.4 The three main ADC performance parameters and several factors affecting them on different levels

1.2.1 ADC Core and Peripherals Challenges

The challenges start from the already complex ADC core, which, as a mixed-signal IC, contains both analog and digital building blocks. The analog blocks include (but are not limited to) samplers, comparators, amplifiers, and references, while the digital circuits contain logic and control, state and memory, as well as delay cells. On the circuit level, the challenge comes from trying to achieve an optimum accuracy-speed-power set for each of the analog circuits until they are mainly limited by fundamental error sources. That entails first understanding these fundamental sources and quantifying their limits. For the digital circuits, the situation is somewhat easier since their trade-off is between speed and power only. The challenge on the architectural level lies in finding the best combination of the different blocks, such that each one contributes its maximum performance and minimum overhead to the system, in the pursuit of an architecturally optimum accuracy-speed-power set. Achieving this involves understanding existing architectures' strengths and weaknesses in order to choose the appropriate one or combine different ones in innovative hybrid forms. On this level, any interference between the different analog blocks as well as between the analog and digital blocks should be minimized while also considering calibration schemes to correct non-idealities.

The challenges are not limited to the single ADC core. High-performance converters employ time-interleaving, adding challenges in the architectural-system-

level intersection. Further, these converters may need extra preceding analog front-end signal conditioning blocks, such as buffers, amplifiers, and filters (Fig. 1.3). These peripheral blocks may significantly degrade the accuracy and bandwidth of the overall converter while greatly increasing the power. Finally, clock conditioning and distribution circuitry is imperative to ensure high-quality clock pulses and minimize the accuracy degradation due to jitter, but it also adds power. Optimally integrating all these core and peripheral functionalities to yield the best system accuracy-speed-power set brings a significant system-level challenge.⁴ On this level, any critical high-speed interfaces between the IC and the outside world, such as input and clock, should be allocated special attention as they are getting progressively dominant when extending the sample rate and bandwidth to several GHz while maintaining a high accuracy. On a more practical front, putting all the ADC blocks (core and peripherals) together and verifying the combined performance with Computer-Aided Design (CAD) tools can be particularly long due to the multiple characterization methods required. Nevertheless, it is necessary to try and ensure a satisfactory performance, at the expense of adding an extra challenge.

1.2.2 *The Good, the Bad, and the Ugly of Deep-Scaled CMOS*

The aforementioned circuit-, architectural-, and system-level challenges are present in every ADC regardless of the technology or process node it is implemented in. The available technology introduces an extra degree of complexity, which can exacerbate and/or add to these challenges. However, it also presents an extra degree of freedom for potential innovation in multiple directions.

The tremendous advancements in the devices' and materials' field have allowed CMOS technology to progressively scale into finer nodes following Moore's law and beyond, which will continue to happen for at least one more generation. We have come to witness the CMOS scaling evolution from planar Field-Effect Transistor (FET) to FinFET, with the Gate-All-Around (GAAFET) being around the corner, as depicted in Fig. 1.5. Due to its increased integratability and functionality for a reduced cost and power, it is thus the preferred technology for highly integrated digital circuits, containing the lion's share of today's SoCs. On top of these benefits, every next-generation process node offers devices with a higher intrinsic speed by means of their cut-off frequency f_T (Fig. 1.6 [10]). In less than 25 years, the f_T has increased from below 10 GHz in 350 nm (1994) to theoretically above 1 THz in 7 nm (2018), with a starting slope $\propto 1/L^2$, which eventually became $\propto 1/L$ due to velocity saturation limitations, with L the minimum channel length. This clearly

⁴ One peripheral block that is not treated in this work is the supply/reference voltage generation. Instead, external voltages are used in the implemented prototypes, which are sufficiently filtered to ensure high quality and not limit the targeted performance.




Fig. 1.5 CMOS scaling evolution from planar FET to FinFET to GAAFET. Picture credit: Samsung




Fig. 1.6 Theoretical cut-off frequency versus channel length scaling [10]

boosts the performance of the digital processors as well as that of the digital circuits within the ADC core.

However, one major drawback for the performance of analog circuits scaling into finer nodes is the reduction of the supply voltage V_{DD} , shown in Fig. 1.7. This is to preserve reliability when reducing the devices' thin oxide, a natural consequence of scaling down L . Although beneficial for reducing the digital power consumption, a lower V_{DD} limits the available analog signal swing, thereby reducing the dynamic range accordingly since the thermal noise floor remains the same. Traditional analog techniques, such as cascoding and high gain stages with feedback, are steadily succumbing to the aggressive device scaling. On top, the threshold voltage V_{TH}




Fig. 1.7 Supply and threshold voltage versus channel length scaling

does not scale down equally (Fig. 1.7⁵) due to subthreshold conduction issues. To make matters worse, the devices in finer nodes exhibit a higher flicker noise and mismatch⁶ and a lower intrinsic gain. In terms of analog power consumption, the f_T increase partially compensates the V_{DD} drop, enabling a device operation at a lower overdrive voltage, keeping the power consumption to a first-order the same, as will be shown in a later chapter of this book.

Last but not least, a particularly aggravated issue associated with scaling is the increased parasitic contribution due to the Back-End of Line (BEOL) metal interconnect. In every finer node, the metals get closer to each other, to the devices, and to the substrate, while the lower-level ones are also getting thinner. Figure 1.8 shows a BEOL metal comparison between 65 nm and 32 nm [11]. This effect increases the metal parasitic resistance and inter-metal parasitic capacitance as well as the resistance of the vias between the metals, which partially reduces the theoretical f_T improvement of Fig. 1.6. To give an example, the BEOL contribution in 16 nm FinFET results in a speed degradation as much as 35–40% compared to the intrinsic device.

To conclude, the technology brings extra challenges in achieving an optimum accuracy-speed-power set and can stir to a great extent the methodology chosen to collectively address the already existing challenges described in Sect. 1.2.1. To tackle the extra challenges but also explore opportunities deep-scaled CMOS brings

⁵ The plot is constructed based on data from various publications and personal experience with certain flavors of different process nodes.

⁶ This is true if $L = L_{\min}$ and the same W/L ratio are used across different nodes. However, if the area $W \cdot L$ is kept constant, mismatch improves in finer nodes since A_{VT} reduces.




Fig. 1.8 BEOL interconnect comparison between 65 nm (left) and 32 nm (right) CMOS processes [11]

in high-performance ADC design, architectural and circuit innovations should minimize traditional analog functionality and interconnect intensity. These should include dynamic circuit solutions for reducing power and digital assistance to correct imperfections when beneficial. Also, both switches and capacitors improve with scaling, thus should be preferred over current sources.

1.3 Research Goal and Objectives

The challenges discussed in the previous section set the stage for defining the scope of the research described in this book.

The overarching goal of this research is to propose deep-scaled CMOS-friendly architectural and circuit solutions to address the challenges in improving accuracy · speed ÷ power for various aggregate resolutions and realize maximally efficient multi-GHz sample rate and bandwidth ADCs with high spectral purity.

To achieve the milestone set above, the challenges are addressed at the circuit, architectural, and system levels while maximizing technology benefits. The steps below summarize the analytical approach followed to address these challenges:

- **Identify converter circuits' limits.** The major error sources stemming from the circuits of any practical converter are identified, and their significance on the converter's performance is quantitatively analyzed. This analysis leads to establishing the circuit-level fundamental accuracy-speed-power limits. The thermal noise of the sampler and the quantizer and the metastability of the quantizer and aperture jitter limit a converter's performance. Heisenberg imposes the ultimate physics limit.

- ▶ **Establish and compare architectural limits.** The investigation for the optimal ADC architecture starts by examining the state-of-the-art standings of high-performance architectures, such as flash, SAR, pipeline, and pipelined-SAR. Mathematical models are introduced to estimate and compare their accuracy-speed-power limits, using certain assumptions. From these models, the SAR comes out as the optimum architecture for a low-to-medium resolution across a wide range of sample rates. It can either be used standalone or serve as a good base candidate for larger system integration to enhance speed, resolution, or both. Its performance is limited by its single analog block, the comparator. For a medium-to-high resolution, the pipelined-SAR hybrid with more than two stages emerges as a promising candidate that can compete with the traditional pipeline even at GHz sample rates. The analysis reveals that for a similar stage count, the pipelined-SAR is more efficient and potentially faster than the pipeline for an extended range of sample rates and resolutions. Both are limited by the residue amplifiers, which can have different requirements for each architecture depending on their internal operation mechanism.
- ▶ **Enhance analysis with technology effects.** The strength of the newly introduced models is greatly enhanced, and more insight is gained by adding the technology effects. This is done by incorporating in the analysis V_{DD} , f_T , C_{min} , and g_m/I_D from four deep-scaled mainstream CMOS processes: 65 nm, 40 nm, 28 nm, and 16 nm. These inclusions make the power/sample rate (energy) vs. accuracy for the different ADCs to be first limited by the process C_{min} and then noise-limited by the capacitance of one or more blocks depending on the noise allocation and the importance of their role within the ADC. When the sample rate f_s increases for a fixed process f_T , the slopes of the different ADC curves increase above the ones dictated by noise. The internal operation mechanism of each architecture (e.g., how many sequential cycles in a conversion, settling accuracy, etc.) determines the final slope for a given f_s/f_T .
- ▶ **Include system-level considerations.** Any of the above architectures can undergo time-interleaving to boost the sample rate beyond the capabilities of a standalone converter. However, it comes with mismatches between the different sub-ADCs and extra loading to the input and clock distribution chains. Hence, choosing the interleaving factor in conjunction with the ADC and interleaver architectures brings a multidimensional system challenge, affecting accuracy, sample rate, bandwidth, and power altogether. It also dictates to a great extent the analog front-end and clock conditioning and distribution considerations. To this end, interleaving mismatches are modeled and analyzed to determine their individual and combined effects under various circumstances. In addition, a model is introduced to quantitatively compare the main interleaver architectures, namely, direct, de-multiplexing, and re-sampling, in terms of achievable bandwidth and sampling accuracy. This model is also extended across the four deep-scaled CMOS processes, providing insight in determining the optimum interleaver depending on the design and specifications.

To demonstrate the validity of the above analytical approach and the feasibility of the proposed architectural and circuit solutions in this book, four multi-GHz prototype ICs are implemented in two deep-scaled CMOS processes:

- An ultrahigh-speed three-stage triple-latch feed-forward dynamic comparator improves the gain and reduces the delay of dynamic comparators across the entire input range [28 nm CMOS, characterized and published].
- A high-speed wide-bandwidth medium-resolution single-channel SAR ADC maximizes the $accuracy \cdot speed \div power$ ratio with a semi-asynchronous timing, an improved bootstrapped input switch, a triple-tail dynamic comparator, and a Unit-Switch-Plus-Cap DAC [28 nm CMOS, characterized and published].
- A high-resolution wide-bandwidth 8 \times -interleaved hybrid RF ADC with a bufferless input front-end, a three-stage pipelined-SAR sub-ADC, a low excess jitter clock chain, and co-designed analog-digital calibrations significantly improves the state of the art in RF ADCs [28 nm CMOS, characterized and published].
- An ultra-wideband highly linear analog front-end with a multi-segment distributed attenuation filter and a hybrid amplifier-buffer extends the bandwidth of next-generation direct RF ADC-based receivers to several tens of GHz [16 nm FinFET CMOS, characterized and published].

1.4 Structure of This Book

In addition to the current chapter, this book contains six core chapters sequentially linked to the approach and prototypes described in the previous section and a closing chapter. A brief overview of their contents is given below:

Chapter 2 first reviews the two fundamental functions in every A/D conversion: sampling and quantization. It then identifies the error sources stemming from the individual circuit blocks of practical converters and analyzes their performance impact. Based on this analysis, circuit-level fundamental limits in terms of accuracy-speed-power are established.

Chapter 3 extends the fundamental limits' analysis and derivations to the architectural level and compares high-performance architectures, such as flash, SAR, pipeline, and pipelined-SAR, after examining their state-of-the-art standing. This comparison is extended over four deep-scaled CMOS process nodes, enhancing the power of the model and building unique insight into architectural and technological capabilities. Finally, time-interleaving is discussed, with the focus on key aspects, such as interleaving errors and interleaver architectures, which are also compared across the different process nodes.

Chapter 4 (1st IC) covers the analysis and design of ultrahigh-speed dynamic comparators, as key blocks in high-performance mixed-signal systems, with a particular importance in high-speed ADCs. First, the two widely adopted topologies are reviewed, and their delays are analyzed. Next, the proposed prototype comparator is presented, highlighting its improvements and analyzing its delay in detail. Finally,

the fabrication, experimental verification, and a state-of-the-art comparison of the prototype are presented and discussed.

Chapter 5 (2nd IC) discusses techniques and challenges for extending the sample rate of low-to-medium-resolution single-channel SAR ADCs in the GHz range, without compromising their digital nature, excellent efficiency, and simplicity, such that they can be both easily used as standalone blocks and integrated into larger systems. After reviewing the conventional clocking scheme and speed-boosting techniques, the proposed prototype SAR ADC is introduced, and the employed architectural and circuit principles are elaborated. Finally, the experimental verification, including the measurement setup, measured results, and comparison with existing state of the art, are presented and discussed.

Chapter 6 (3rd IC) elaborates on system, architectural, and circuit capabilities to enable ADC resolutions beyond 10 bits while sampling directly at RF frequencies with multi-GHz rate and bandwidth at maximum power efficiency. First, the direct RF sampling receiver is briefly discussed, highlighting the ADC role, and the case for efficient architectural and circuit choices is made. After reviewing prior art architectural choices and their trade-offs, the proposed prototype time-interleaved hybrid RF ADC is introduced, and its performance-enabling principles are detailed. Finally, the experimental verification, including the measurement setup, several measurements, and a state-of-the-art comparison, are presented and discussed. This prototype combines the insight gained from the performed analytical approach to its fullest and significantly advances the state of the art in multiple directions.

Chapter 7 (4th IC) addresses the analog front-end challenges in pushing the sample rate and bandwidth of direct RF ADC-based receivers to several tens of GHz while delivering high spectral purity with low power. These challenges stem from the large ADC input load and the constant pursuit for higher integration in deep-scaled CMOS. After introducing the problem and overviewing some prior art, the proposed highly integrated ultra-wideband front-end solution is presented, and its novel design features are thoroughly discussed. Finally, the experimental verification, including the measurement setup, several key measurements, and a state-of-the-art comparison, are presented and discussed.

Chapter 8 gives the overview and draws the conclusions of the work in this book. Further, it highlights its original scientific contributions and points to some promising future research directions.

Due to the broad range of topics covered in this book, in an attempt to have each chapter technically complete and individually strong, small pieces of information may be partially repeated in various parts. The reader should realize that this is done intentionally and is not a product of careless writing.

Chapter 2

Analog-to-Digital Conversion

Fundamentals



The tremendous popularity but also challenges of data converters as key interface functions between the physical (analog) world and the electronic (digital) world were discussed from a bird's eye view in the previous chapter. Before delving into advanced architectural and design details, this chapter will cover the fundamental A/D conversion principles, some important performance metrics, as well as practical limitations, serving as the foundation for the following chapters.

Section 2.1 serves as a theoretical background by reviewing the two main functions in every A/D conversion: (1) sampling and (2) quantization. The major error sources stemming from the individual blocks of practical converters are identified and analyzed in Sect. 2.2, followed by a review of the most important performance metrics and figures of merit in Sect. 2.3. Section 2.4 derives the impact on the accuracy-speed-power for every major error source. This derivation leads to the establishment of the fundamental limits on a converter's performance, imposed by circuits, by technology, and ultimately by physics. The limits in this chapter form the basis of what may be theoretically achievable and, together with the architectural overheads presented in Chap. 3, serve as guidelines to assist the design choices of the prototypes in Chaps. 4–7. This chapter closes with an overview and conclusions in Sect. 2.5.

2.1 Theoretical Background

As already mentioned, every analog signal is continuous both in time and in amplitude. Therefore, two main processes are essential to obtain the final digital waveform:

1. Sampling (to achieve the time discretization)
2. Quantization (to achieve the amplitude discretization)




Fig. 2.1 Block diagram of an ideal A/D conversion (top) and the resulting waveforms at every part of the chain (bottom)

Figure 2.1 depicts the block diagram of an ideal A/D conversion with its corresponding waveforms. The continuous time and amplitude analog input signal (black waveform) is uniformly sampled with a period of T_s (or at a sample rate¹ of f_s). The resulting time-discrete analog signal (orange waveform) updates its value only at integer multiples of T_s . When the time is equal to an integer multiple of T_s , the sampled signal is equal in value to the analog input at that instant and keeps its value until the next multiple of T_s arrives. Between two consecutive time instants, the sampled signal is held constant and can be further processed down the conversion chain.

Next, the quantization takes place, where the sampled signal is discretized in amplitude and its analog values are mapped onto a set of discrete levels (blue waveform). The digital output, now discrete in both time and amplitude, is an approximation of the initial analog input, with its approximation accuracy limited by the number of the available discrete levels. During both sampling and quantization, there is information loss since an error is introduced on the initial analog signal. This error can be reduced by increasing the number of time samples and/or the number of discrete levels. As we will see in the remainder of this book, guaranteeing simultaneously both can be far from trivial.

2.1.1 Sampling

Sampling is the basic process that transfers a waveform from the continuous time to the discrete time domain. The sampling process can be described mathematically

¹ Throughout this book, the terms sample rate, sampling rate, sampling speed, and/or sampling frequency will all refer to the same quantity f_s .




Fig. 2.2 Sampling a continuous-time signal using a Dirac pulse sequence

by means of the Dirac function $\delta(t)$, whose integral is equal to one at the integration instant and zero elsewhere [12]. The required sampling time frame is determined by a sequence of equidistant in time Dirac pulses, spaced by T_s . The time-discrete signal is a result of the multiplication of the Dirac pulses with the original waveform, with an amplitude equal to the amplitude of the waveform in the sampling instants and undefined elsewhere (Fig. 2.2). The mathematical formula expressing the above is given as

$$V_s(t) = V(t) \cdot \sum_{n=-\infty}^{n=\infty} \delta(t - nT_s) = \sum_{n=-\infty}^{n=\infty} V(nT_s). \quad (2.1)$$

Generally, the transformation of a signal from time domain to frequency domain is done by means of its Fourier Transform (FT). For a time-discrete signal specifically, this transformation in the frequency domain occurs by employing the signal's Discrete Fourier Transform (DFT). Taking into account that a multiplication in time is a convolution in frequency, the spectrum of the time-discrete signal $V_s(t)$ is depicted in Fig. 2.3 and given by

$$V_s(f) = \frac{1}{T_s} \cdot \sum_{n=-\infty}^{n=\infty} V(f - nf_s). \quad (2.2)$$

The dual-sided band around zero with a frequency content within $\pm f_{in}$ is attributed to the original waveform. The replica or alias bands around multiples of f_s result from the multiplication of the original waveform with the repetitive by $T_s = 1/f_s$ Dirac pulse sequence. The signal bands with the same frequency content around any multiple of f_s , after processing the spectrum with a Fast Fourier Transform (FFT) algorithm become indistinguishable around zero. As a numerical example, single-tone signals with 211 MHz, 789 MHz, 1.211 GHz, 1.789 GHz, and 2.211 GHz input




Fig. 2.3 Frequency spectrum of a signal multiplied with a sequence of Dirac pulses




Fig. 2.4 (a) Single-tone signals with different frequencies (b) fall in the same frequency location after spectrum processing

frequencies (Fig. 2.4a) will all end up at the 211 MHz frequency location when sampled at 1 GS/s (Fig. 2.4b).

If the band of the original waveform increases in width, so will its alias bands. This will eventually lead to the bands overlapping, causing mixing of information between them and making it impossible to isolate the information from each band correctly. This irreversible situation is described as aliasing. In order to prevent information loss due to aliasing and yield the sampling process reversible, the following condition between the instantaneous signal bandwidth $f_{in,bw}$ and the sample rate f_s must be obeyed:




Fig. 2.5 Dual-sided frequency spectrum highlighting different Nyquist zones




Fig. 2.6 (a), (b) Two cases of signals with bands meeting the Nyquist criterion and (c) one scenario where bands are overlapping leading to information loss

$$\frac{f_s}{2} > f_{in,bw}. \tag{2.3}$$

Known as the sampling theorem or Nyquist sampling criterion [13, 14], the above expression can be translated to

A band-limited continuous-time signal can be sampled and perfectly reconstructed if the sample rate is more than twice the signal’s instantaneous bandwidth. The frequency band between zero and $f_s/2$ is defined as the Nyquist bandwidth or the 1st Nyquist zone. The total spectrum comprises an infinite number of Nyquist zones, each with a width of $f_s/2$. Figure 2.5 shows the first four Nyquist zones in the spectrum, indicating their frequency allocation and width. For signals originally residing in the odd-order zones, their bands after sampling are copied to the 1st Nyquist zone as they are, while the bands of even-order zones are mirrored. Under the assumption that Eq. (2.3) holds (Fig. 2.6a, b), the original signal can be accurately reconstructed by a reconstruction filter. However, a violation

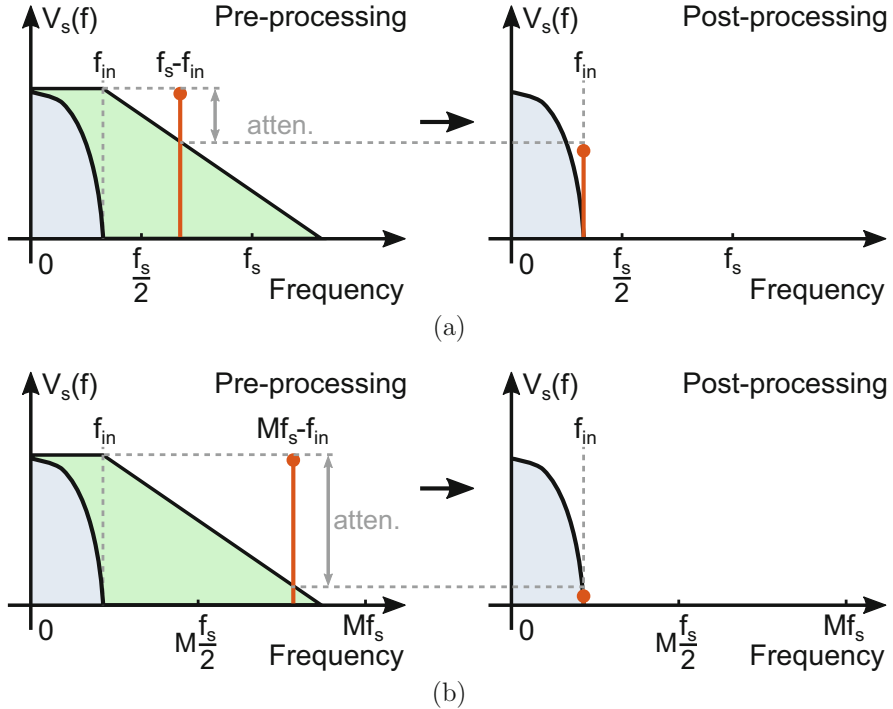


Fig. 2.7 Anti-aliasing filter on a parasitic tone when (a) sampling at Nyquist rate (slightly oversampled in practice) and (b) oversampling by $M > 1$

of the Nyquist criterion (Fig. 2.6c) will result in aliasing and render an accurate reconstruction of the original signal impossible.

Even if the useful signal resides within the Nyquist bandwidth, different types of undesired signals or interferers may appear at higher Nyquist zones, mixing up with the useful signal after sampling in the 1st Nyquist zone. Examples of such undesired signals are harmonic-related products of the main signal and/or interferers/noise from parts in the signal chain preceding the sampling. To prevent these unwanted signals from limiting the Dynamic Range (DR) of the chain, an anti-aliasing filter is typically employed prior to sampling to remove any component outside the Nyquist bandwidth. The specifications of this filter, whose implementation may include active and/or passive components, heavily depend on how much attenuation it needs to provide at which frequency distance with respect to $f_s/2$. Given that typical filters provide an attenuation of 20 dB/decade per order, a multi-order robust filter design becomes increasingly challenging and expensive as the frequency band of interest approaches $f_s/2$. Figure 2.7a illustrates the case of attenuating a parasitic tone by a finite-order anti-aliasing filter for a signal with $f_{in} < f_{in,bw}$ sampled at the Nyquist rate.

One way to improve the filter attenuation for a certain order or relax the filter order for a certain attenuation is to sample faster than the Nyquist criterion imposes. Increasing the sample rate (oversampling) provides a trade-off between parasitic tone attenuation and clock speed to sample and process data for a certain filter order [15]. Figure 2.7b illustrates how oversampling by a factor of M significantly improves the parasitic tone attenuation for the same filter order. However, for very wideband signals, generating the clock for a certain oversampling becomes equivalently challenging as increasing the filter order.

As a final note on sampling, it is worth mentioning that the Nyquist criterion is still satisfied and aliasing is not an issue for a signal residing in any of the Nyquist zones, as long as it is band-limited within one. In fact, this sampling property is utilized in the increasingly popular sub-sampling ADCs in communication systems. Directly sampling Intermediate Frequency (IF)/Radio Frequency (RF) signals in higher Nyquist zones and processing them digitally allow simplification of the signal chain by eliminating several frequency down-conversion blocks, such as a mixer, an IF amplifier, and filters. However, this increases the sub-sampling ADC's bandwidth and spectral purity requirements at higher Nyquist zones. Chapter 6 of this book introduces circuit and architecture techniques for efficiently realizing wideband RF sampling ADCs.

2.1.2 Ideal Quantization

An ideal quantizer is a memoryless non-linear block, which uses B bits to translate the sampled signal to a digital word of binary format (0s and 1s). B represents the aggregate resolution with which the digital output resembles the analog input. Figure 2.8 shows the conceptual model and transfer characteristic of an ideal B -bit quantizer. Each signal value is compared against 2^B discrete levels, and its amplitude is rounded to the nearest level. The output Encoding Logic (ENC) decides how the rounding is done. The maximum input amplitude is defined as the Full-Scale (FS), and the difference between two adjacent transition levels (a.k.a. the step width), Δ , is quantified in the analog domain as the Least Significant Bit (LSB) such that $\Delta = \text{FS}/2^B$.

The digital word can be back-converted to a discrete amplitude analog signal V_q by multiplying each bit with its assigned binary weight, provided that the analog value of Δ is known

$$V_q = \Delta * \left(\sum_{i=0}^{B-1} bit_i * 2^i + bit_1 * 2^1 + bit_2 * 2^2 + \dots + bit_{B-1} * 2^{B-1} \right). \quad (2.4)$$

Due to the rounding process, there is a quantization error ϵ_q added to the original signal V_{in} , with a value ideally within $\pm\Delta/2$ for signals inside FS, while growing out of bounds outside FS (Fig. 2.8). The minimum error power is achieved for uniformly

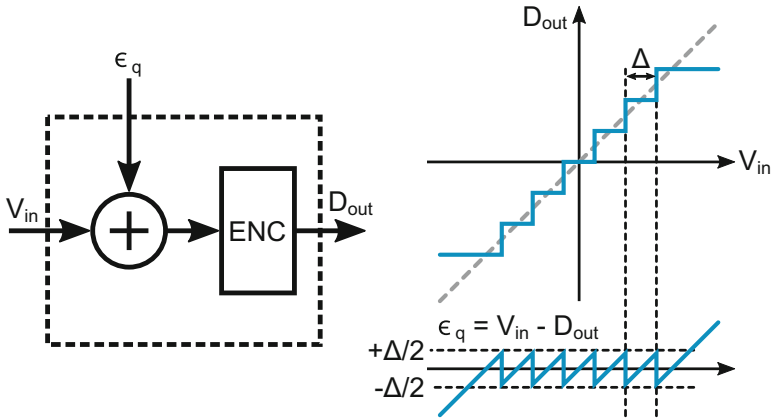



Fig. 2.8 Conceptual model and transfer characteristic of an ideal quantizer

Fig. 2.9 Sawtooth approximation of ϵ_q as a function of time



spaced discrete levels [16]. The back-converted signal relation with the original signal is expressed as


$$V_q = V_{in} + \bar{\epsilon}_q. \tag{2.5}$$

Strictly speaking, ϵ_q is a deterministic quantity, heavily depending on the properties of the signal at hand. For a linear ramp signal that contains several LSBs, ϵ_q can be approximated in time domain by a sawtooth waveform with a peak-to-peak amplitude of Δ , as shown in Fig. 2.9

$$\epsilon_q(t) = slope \cdot t, \quad -\frac{\Delta}{2} \leq slope \cdot t \leq \frac{\Delta}{2}. \tag{2.6}$$

Due to the signal periodicity, an integration over a single period of T_p suffices to calculate the Root-Mean-Square (RMS) value of the error

Fig. 2.10 Uniformly distributed PDF of ϵ_q within $\pm\Delta/2$



$$\bar{\epsilon}_q^2 = \frac{1}{T_p} \int_{-\frac{T_p}{2}}^{\frac{T_p}{2}} \epsilon_q^2(t) dt = \frac{\text{slope}}{\Delta} \int_{-\frac{T_p}{2}}^{\frac{T_p}{2}} \left(\frac{\Delta}{T_p}\right)^2 t dt = \frac{\Delta^2}{12} \Rightarrow \bar{\epsilon}_q = \frac{\Delta}{\sqrt{12}}. \quad (2.7)$$

If a more statistical approach is followed, considering that over a long time span all values within $\pm\Delta/2$ will show up with the same probability, ϵ_q assumes a uniform Probability Density Function (PDF) within that same region as is illustrated in Fig. 2.10. The necessary conditions for the validity of this approach are:

- The signal is sufficiently large or the quantizer resolution is large, such as to cover an adequate amount of levels
- The input is uncorrelated with the quantization error or the input frequency is not harmonically linked to the sample rate
- The signal is limited to FS, such that there is no quantizer overloading

If the above conditions hold, ϵ_q may be allocated a zero mean μ_{ϵ_q} and a variance $\sigma_{\epsilon_q}^2$ that can be calculated as in [17]

$$\sigma_{\epsilon_q}^2 = \bar{\epsilon}_q^2 = \frac{1}{\Delta} \int_{-\frac{\Delta}{2}}^{\frac{\Delta}{2}} \epsilon_q^2 d\epsilon_q = \frac{\Delta^2}{12}, \quad (2.8)$$

which matches the result of Eq. (2.7). As pointed out in [17], this quantization “noise” upon sampling shows a uniform spread across the entire Nyquist bandwidth. In case the input frequency is harmonically linked to the sample rate, there exists a relation between the input and ϵ_q resulting in the energy being accumulated in the harmonics of the signal. When performing a spectral analysis through FFT, this correlation can be avoided by choosing an integer number of signal periods (coherent sampling) and relatively prime number of periods and points [18]. Appendix A describes such an FFT setup.

As the quantizer resolution decreases, the non-linear nature of the quantization process dominates over its noise-like approximation, resulting in a distortion dominated spectrum rather than the flat noise-like. Figure 2.11 plots the spectra of an ideally quantized 77 MHz input signal coherently sampled at 1 GS/s for various resolutions. A reduction of about 8–9 dB per added bit is seen in the odd harmonic

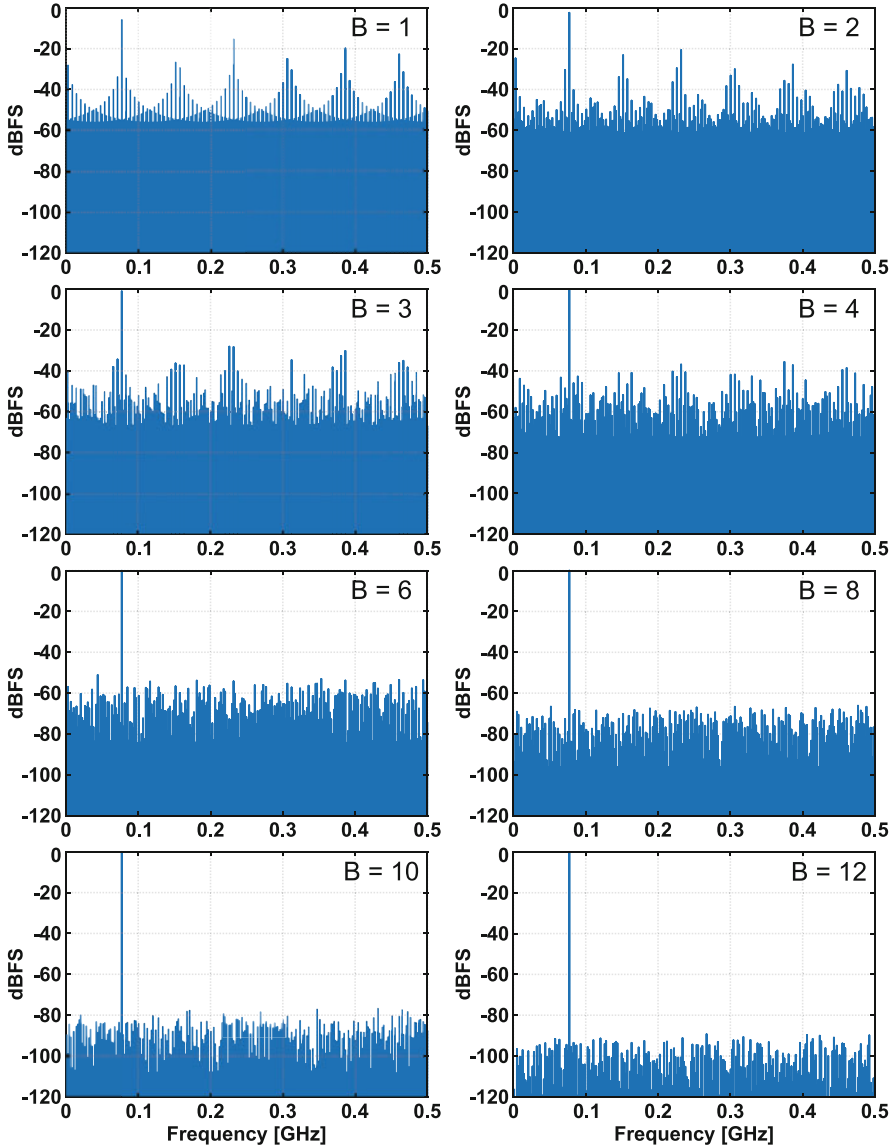


Fig. 2.11 Frequency spectra of an ideally quantized with various resolutions 77 MHz signal sampled at 1 GS/s ($N_{\text{FFT}} = 1024$)

spurs [19]. This is understood by the fact that for every added bit $\Delta^2/12$ reduces by 6 dB, while the additional 3 dB results from preserving the same total harmonic energy with twice the number of harmonics.

Table 2.1 Comparison between calculated and simulated SQNR for different B

Number of bits [B]	Calculated SQNR	Simulated SQNR
1	7.78 dB	6.31 dB
2	13.80 dB	13.30 dB
3	19.82 dB	19.53 dB
4	25.84 dB	25.61 dB
5	31.86 dB	31.66 dB
6	37.88 dB	37.73 dB
7	43.90 dB	43.84 dB
8	49.92 dB	49.84 dB
10	61.96 dB	61.92 dB
12	74.00 dB	73.98 dB

Having determined the conditions under which ϵ_q is considered white noise, the Signal-to-Quantization-Noise Ratio (SQNR) within the Nyquist bandwidth can be computed for a FS input sinusoid with a peak-to-peak amplitude of V_{FS}

$$\begin{aligned} SQNR &= 10 \log \left[\frac{(\frac{V_{FS}}{2\sqrt{2}})^2}{(\frac{V_{FS}}{\sqrt{12 \cdot 2^B}})^2} \right] = 10 \log(1.5 \cdot 2^{2B}) \quad [\text{dB}] \quad (2.9) \\ &= 6.02 \cdot B + 1.76. \end{aligned}$$

As anticipated, due to the non-linear nature of ϵ_q , the validity of the above expression may be questionable as the resolution decreases or for a signal that doesn't uniformly occupy a sufficient range [12]. Table 2.1 compares the calculated ideal SQNR against the simulated value for different resolutions. The noise approximation leading to Eq. (2.9) provides an overestimation, which reduces as the resolution increases, eventually matching the simulated value.

Finally, if the utilized signal bandwidth $f_{in,bw}$ does not include the complete Nyquist band, such that the sampling happens at a higher rate than Nyquist, there is an improvement in SQNR equivalent to the oversampling ratio $f_s/(2 \cdot f_{in,bw})$. In this case, an extra term known as the processing gain needs to be included in Eq. (2.9), which now becomes

$$SQNR = 6.02 \cdot B + 1.76 + 10 \log \left[\frac{f_s}{2 \cdot f_{in,bw}} \right]. \quad [\text{dB}] \quad (2.10)$$

Oversampling combined with quantization error shaping and digital filtering to remove out-of-band noise are fundamental concepts in $\Delta\Sigma$ converters [20].

2.2 Error Sources

Although ideally ϵ_q sets the theoretical single conversion error source, imperfections of electronic components utilized in a real A/D conversion introduce several noise and distortion sources to the signal. The sampling network comes with thermal noise, non-linear distortion, and aperture jitter. The actual quantizer introduces further thermal noise and both integral and differential non-linearity on top of its existing quantization noise. For very wide bandwidth, if an additional analog front-end needs to be utilized, it adds extra thermal noise and non-linear distortion. Figure 2.12 illustrates the model of a real converter including the aforementioned error sources.

2.2.1 Noise

The wideband internal circuits in a converter produce a certain amount of thermal noise due to Brownian motion of charges. Although the instantaneous value of noise cannot be predicted, its Gaussian nature allows for the construction of a statistical model by means of a distribution. To measure its RMS value a large number of output samples are collected and plotted as a histogram, from where the mean μ and the standard deviation σ (or variance σ^2) can be calculated² [21]. The RMS noise voltage is equal to σ and can be expressed either with respect to an LSB or as an RMS absolute voltage.

Three main noise sources can be identified in a converter chain (Fig. 2.12), namely, thermal noise from the sampling network; thermal noise due to the quantizer; and aperture jitter during the sampling instants.

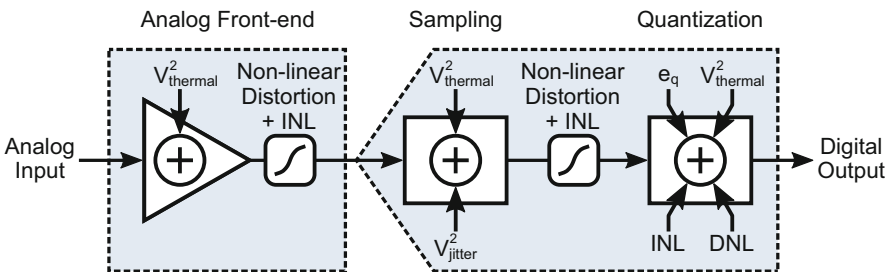


Fig. 2.12 Conceptual model of a real converter including error sources from the different blocks

² In the subsequent calculations, the noise variance will be expressed as voltage squared.




Fig. 2.13 (a) Simple model of a sampler and (b) its noise spectrum

Sampler Thermal Noise

The simplest implementation of a sampler comprises a switch S (Metal-Oxide-Semiconductor (MOS) device) and a capacitor C_S , as illustrated in Fig. 2.13a. When S is turned on, the MOS device is operating in triode region; therefore, it exhibits an on-resistance R_S . R_S produces white noise with a spectral density (single-sided) of

$$\overline{V_{R_S}^2} = 4kTR_S, \quad [\text{V}^2/\text{Hz}] \quad (2.11)$$

where $k = 1.38 \cdot 10^{-23} \text{ J/K}$ is the Boltzmann constant and T is the absolute temperature.³ The RC network of the sampler shows a first-order low-pass characteristic with a cut-off frequency of

$$f_{-3\text{dB}} = \frac{1}{2\pi R_S C_S}, \quad [\text{Hz}] \quad (2.12)$$

which shapes the noise spectrum of R_S as shown in Fig. 2.13b. The sampler noise power can then be calculated by integrating $\overline{V_{R_S}^2}$ over the entire noise bandwidth

$$\overline{V_{n,\text{samp}}^2} = \alpha_{\text{FE}} \int_0^\infty \frac{4kTR_S}{(2\pi f R_S C_S)^2 + 1} df = \alpha_{\text{FE}} \frac{kT}{C_S}, \quad \alpha_{\text{FE}} \geq 1, [\text{V}^2] \quad (2.13)$$

where α_{FE} accounts for any excess noise in the presence of an analog front-end.

Quantizer Thermal Noise

A typical 1-bit quantizer employs a dynamic latch-based comparator (see Chap. 4) in some form and combination. To provide a simple expression as a basis for the noise of the quantizer, we construct the model shown in Fig. 2.14a. It assumes a two-stage comparator with a $g_{m,L}$ latch output and a $g_{m,I}$ integrator input [22] to provide some gain prior to regeneration and lower the noise of the latch. Ignoring large signal behavior and considering the latch as a settling stage with a $g_{m,L}$ noise

³ Throughout this book, T is set to 323 K (50 °C), unless explicitly stated otherwise.




Fig. 2.14 (a) Simple quantizer model and (b) its allowed operation time

contribution equivalent to an effective resistor of $1/g_{m,L}$ [23], the latch noise power at V_I is given by

$$\overline{V_{g_{m,L}}^2} = \frac{4kT\gamma}{g_{m,L}} \cdot \frac{g_{m,L}}{4C_L} = \frac{kT}{C_L}, \quad [\text{V}^2] \quad (2.14)$$

where γ is the thermal noise excess factor.⁴ The input stage integrates its own noise over a noise bandwidth proportional to $1/2T_I$, where T_I is the integration time allowed for the quantizer (Fig. 2.14b). Its input noise power can be calculated similarly to [25] and is given by

$$\overline{V_{g_{m,I}}^2} = \frac{4kT}{g_{m,I}} \cdot \frac{1}{2T_I} = \kappa \frac{kT}{AC_I}, \quad [\text{V}^2] \quad (2.15)$$

where κ depends on the integration time, the integration voltage on V_I , and the relative biasing of the input devices. Assuming for simplicity equal values for C_I and C_L , the total input-referred noise power can be approximated as

$$\overline{V_{n,\text{quant}}^2} = \overline{V_{g_{m,I}}^2} + \frac{1}{A^2} \overline{V_{g_{m,L}}^2} \approx \frac{kT}{AC_I}, \quad [\text{V}^2] \quad (2.16)$$

where in the last step we substituted $\kappa = 1$ and $A = 4$ for the input stage.⁵

⁴ In literature, values of γ for short-channel devices span between 0.7 and 2.9 [24]. In this book, the value of 1 will be used unless otherwise stated.

⁵ The maximum gain for a g_m - C integrator cannot exceed the $g_m R_o$ of a differential pair, which in a 28 nm bulk CMOS process can reach values of 4 (12 dB) at GHz operation.




Fig. 2.15 (a) Sampler with jitter and (b) time to voltage error translation

Aperture Jitter

During ideal sampling (Sect. 2.1.1), the continuous-time input signal is sampled precisely at instants equally spaced by T_s . However, noise and mismatch in the devices of a real sampling network result in random variations in the clock edge (Fig. 2.15a), leading to sampling uncertainty, known as aperture uncertainty or aperture jitter. It is generally measured in picoseconds RMS. Jitter in time (Δt) translates into an output voltage error (ΔV), whose value is strongly related to the slope of the input signal, as illustrated in Fig. 2.15b. It is worth mentioning that jitter on the sampling clock or on the analog input produce exactly the same type of error. In fact, assuming that the sources are uncorrelated, they simply add in a Root-Sum-Square (RSS) fashion to yield the total error at the output.

The voltage error due to jitter can be easily calculated for a sinusoidal input of $V_{in}(t) = 0.5V_{FS}\sin(2\pi f_{in}t)$.⁶ Since this error depends on the slope of the signal, it is maximum at the zero crossings

$$\begin{aligned}\Delta V_{\max} &= \left. \frac{d}{dt} V_{in}(t) \cdot \Delta t \right|_{t=0} = 2\pi f_{in} \frac{V_{FS}}{2} \cos(2\pi f_{in}t) \cdot \Delta t \Big|_{t=0} \\ &= \pi f_{in} V_{FS} \cdot \Delta t.\end{aligned}\quad (2.17)$$

Since Δt is assumed to be random with a standard deviation of t_{jit} , the integrated error noise power can be approximated as

$$\begin{aligned}\overline{V_{n,jitter}^2} &= \frac{1}{T_{sig}} \int_0^{T_{sig}} \left(\frac{d}{dt} V_{in}(t) \right)^2 dt \cdot t_{jit}^2 \\ &= \frac{1}{2} (\pi f_{in} V_{FS})^2 \cdot t_{jit}^2,\end{aligned}\quad [V^2] \quad (2.18)$$

⁶ The calculation is done with respect to a peak-to-peak signal amplitude to preserve consistency with all our subsequent calculations.

where T_{sig} is the integration period, which for a sinusoid can be chosen as the signal period.

As a final note on jitter, special care must be taken across the entire input and clock chains to minimize the accumulative contribution of every added block. In Chap. 6, we will present an ultra-low jitter clock chain that shows how such a minimization can be achieved.

Now that we derived all the major noise contributions referred to the residue node (quantizer input), they can be summed and added to $\bar{\epsilon}_q$ to yield a first-order total quantization and noise power (single-ended)

$$\overline{V_{\epsilon_q+n,\text{total}}^2} = \frac{\Delta^2}{12} + \alpha_{\text{FE}} \frac{kT}{C_S} + \frac{kT}{AC_1} + \frac{1}{2} (\pi f_{\text{in}} V_{\text{FS}})^2 \cdot t_{\text{jitter}}^2 \cdot [V^2] \quad (2.19)$$

One quick observation arising from the above expression is that $\overline{V_{n,\text{jitter}}^2}$ increases with f_{in} , whereas both $\overline{V_{n,\text{samp}}^2}$ and $\overline{V_{n,\text{quant}}^2}$ are to a first-order input frequency independent. Additionally, to reduce both $\overline{V_{n,\text{samp}}^2}$ and $\overline{V_{n,\text{quant}}^2}$ the capacitors at the corresponding band-limiting nodes must increase, adversely affecting the bandwidth. Section 2.4 analyzes the accuracy degradation of a converter due to the above noise sources and establishes some fundamental accuracy-speed-power limits.

2.2.2 Non-linearity

The non-linearity of the circuit elements utilized in a real converter will make its transfer characteristic deviate from an ideal equal step width linear curve. As illustrated in Fig. 2.16, these deviations manifest themselves both locally in each step (Fig. 2.16a) and globally across the entire characteristic (Fig. 2.16b). The two main types of non-linearity encountered in a real converter are characterized by the Differential Non-Linearity (DNL) and the Integral Non-Linearity (INL)

DNL quantifies the individual deviation of each step's width from the ideal value Δ (1 LSB) according to the following expression:

$$DNL_i = \frac{(V_{i+1} - V_i) - \Delta}{\Delta}, \quad \forall i = 0 \dots (2^B - 2). \quad (2.20)$$

For each step, the relative deviation of its width from Δ is uncorrelated with the equivalent deviation of the previous and next steps. Positive or negative DNL implies a larger or smaller step compared to Δ , respectively. A value of -1 LSB is the smallest possible and indicates that a step was completely skipped, a situation described as a missing code (Fig. 2.16a). In the presence of a noisy signal, such that the transition levels carry noise comparable to Δ , this noise can affect the

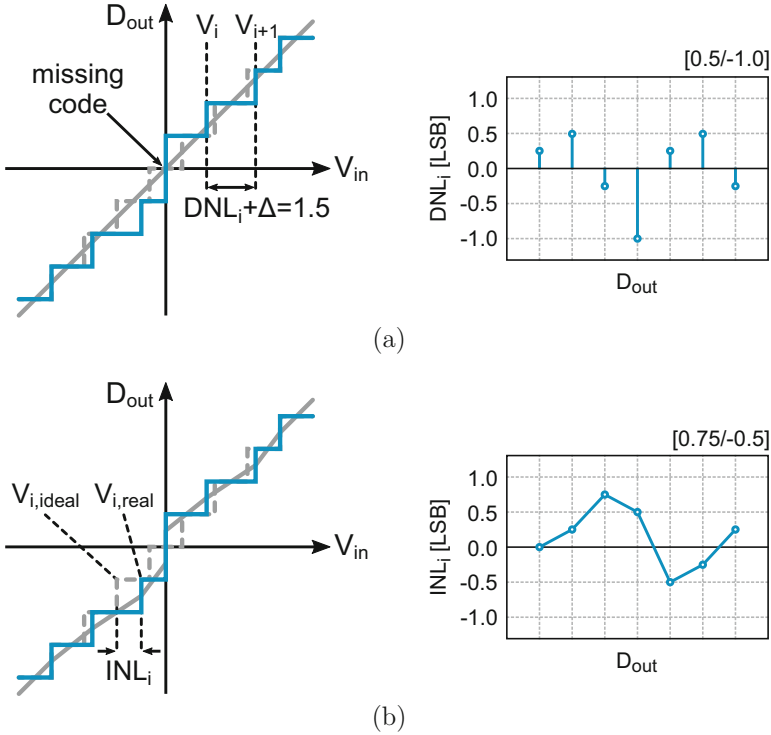


Fig. 2.16 (a) DNL in transfer characteristic with corresponding curve and (b) INL in transfer characteristic with corresponding curve

DNL true value and potentially hide missing codes [26, 27]. Therefore, its value alone should not be trusted blindly. DNL is due exclusively to the quantizer, and the ENC (Fig. 2.8) determines how its errors spread across the transfer curve. Strictly speaking, these errors result in distortion products at the converter output, which depend both on the amplitude of the signal and on their relative position along the transfer curve. However, similar to ϵ_q , under the assumption of a uniform DNL spread across the FS, its effect can be seen more as random noise rather than distortion. In that case, the degradation in SQNR can be estimated if a DNL within $\pm d$ is added to the signal, resulting in a worst-case total quantization + DNL error within $\pm 1/2(\Delta + d)$. Adding this to Eq. (2.9) results in the Signal-to-Quantization-and-DNL-Noise Ratio (SQDNR)

$$\begin{aligned}
 SQDNR &= 10 \log \left[\frac{\left(\frac{V_{FS}}{2\sqrt{2}}\right)^2}{\left(\frac{(1+d)V_{FS}}{\sqrt{12} \cdot 2^B}\right)^2} \right] = 10 \log \left(1.5 \cdot \frac{2^{2B}}{(1+d)^2} \right) [\text{dB}] \quad (2.21) \\
 &= 6.02 \cdot B - 1.76, \quad \text{if } d = 0.5.
 \end{aligned}$$

INL quantifies the overall deviation of the actual converter transfer characteristic from a straight line passing through the first and last transitions. Alternatively, if we draw a line passing through all the real transitions (Fig. 2.16b), its deviation from the ideal straight line (Fig. 2.16a) reveals INL. In each step, INL can be calculated as follows:

$$INL_i = \frac{(V_{i,real} - V_{i,ideal})}{\Delta}, \quad \forall i = 0 \dots (2^B - 1). \quad (2.22)$$

In contrast to the DNL, INL has a cumulative nature adding up errors from the consecutive steps to move the transfer curve with respect to the straight line, therefore resulting in an integral error. As such, its “purity” is affected less than the DNL in the presence of noise, making its value more trustworthy. It can be shown that INL from the quantizer only in each step can be calculated by a cumulative summation of the individual DNLs up to the previous step by the following expression:

$$INL_j = \sum_{i=0}^{j-1} DNL_i. \quad (2.23)$$

The total converter INL is a summation in RSS of different contributions from all the blocks in the chain that generate distortion, including the sampling network and the analog front-end (if utilized) (Fig. 2.12). It is not exclusively a quantizer property like DNL. Overall, INL results in input signal-dependent distortion products at the converter output, making it hard sometimes to identify which one of the individual contributors is dominant.

Non-monotonicity describes a special situation, where an increasing/decreasing input signal results in a decreasing/increasing step in the transfer curve, making the width of that step (hence its DNL) “ill-defined” [26]. This situation is especially important for converters used in closed-loop configurations; therefore, it should be avoided by design. It can be shown that a sufficient but not necessary condition for INL to prevent non-monotonicity is given below

$$|INL_i| \leq 0.5 \text{ LSB}, \quad \forall i, \quad (2.24)$$

which then results in an equivalent condition for DNL as follows:

$$|DNL_i| \leq 1 \text{ LSB}, \quad \forall i. \quad (2.25)$$

2.2.3 Calibration

Generally speaking, any type of non-linearity error, including DNL and INL, originate from circuit imperfections (mismatch [28], leakage, incomplete settling, voltage-temperature variations, etc.) and/or technology limitations to achieve a required performance. Their contribution can be minimized by proper design (e.g., device up-scaling) and/or architectural choices, which often increase the power consumption and area while compromising speed.

Alternatively, deterministic errors that are not associated with random noise but stem from circuit or technology imperfections may be compensated by means of calibration techniques. Such techniques can potentially yield a better overall performance with a reduced impact on the power consumption. The compensation process primarily comprises the following steps:

1. Error detection by measuring circuits' parameters that are considered for modification
2. Error correction by modifying the parameters to desired values by the correction circuitry, such that the errors are minimized or eliminated

The error detection can be implemented either in the analog or in the digital domain. The optimal implementation depends on the type and magnitude of errors as well as the application, performance, and technology at hand. Additional circuits and test signals are often necessary to perform the detection; however, it can be also performed by a statistical analysis without requiring extra hardware or modifications to the core circuitry. The error correction can be also performed either in the analog or in the digital domain (or a combination of both), with the two having distinct differences regarding the end result of the calibration and circuitry used. For example, if correction is performed in the analog domain, modifications in the core circuits are necessary in order to re-adjust the parameters (e.g., by changing biasing voltages/currents or adding/subtracting tunable loads) and eliminate the error. The loading effects of such modifications on the core circuits' performance must then be taken into account. If digital correction is performed, the core circuits are left untouched, and the inverse of the error function is digitally created and applied to the digital output to reduce the error. In this case, the calibration accuracy may be somewhat inferior due to rounding effects but with increasing power and speed benefits moving into finer CMOS processes.

A final difference lies with how often the calibration is performed and how disruptive it is to the normal operation. In case of the so-called "foreground" calibration, the converter operation is halted, and once the calibration is performed, it becomes available again to continue its operation. In the case of "background" calibration, the converter errors are corrected simultaneously to its normal operation, and the calibration is integrated ideally seamlessly into the core functionality. As expected, both methods have advantages and drawbacks in terms of hardware, signal range utilization, correction accuracy, and error tractability. Therefore, the optimal

choice depends on the nature of errors and the specific application requirements and tolerances.

2.3 Performance Evaluation

A converter's achievable performance can be evaluated in the time domain and in the frequency domain [3, 12, 15], and several metrics exist for such evaluations. Below, we will limit ourselves to the frequency domain evaluation by means of an FFT [29] and define the metrics that will be used in the following chapters.

2.3.1 Metrics

Nth-order Harmonic Distortion (HD_n) is normally specified in dBc (decibels below carrier) and is the reciprocal of the ratio between the RMS value of the fundamental signal and the RMS value of its nth-order harmonic. The harmonics of the input signal can be distinguished from other distortion products because of their location in the frequency spectrum at integer multiples of the input frequency. HD_n is generally specified for input signals near FS since for much smaller signals, there may be other error mechanisms that dominate.

Total Harmonic Distortion (THD) is the inverse ratio of the RMS value of the fundamental signal to the mean RSS value of its harmonics. Depending on the specific design and application, the first five to seven harmonics are considered significant. For a FS input sinusoid with a peak-to-peak amplitude of V_{FS} and harmonics' amplitude of $\overline{V_{\text{harm},n}}$, $n = 2, 3, \dots, 7$, THD is evaluated by the following expression:

$$THD = -10 \log \left[\frac{\left(\frac{V_{FS}}{2\sqrt{2}}\right)^2}{\sqrt{V_{\text{harm},2}^2 + \dots + V_{\text{harm},7}^2}} \right]. \quad [\text{dB}] \quad (2.26)$$

Signal-to-Noise Ratio (SNR) is the ratio of the RMS signal amplitude to the mean RSS value of all noise-related spectral components, including quantization (plus DNL), thermal, and jitter. For a FS input sinusoid with a peak-to-peak amplitude of V_{FS} , its value is evaluated as

$$SNR = 10 \log \left[\frac{\left(\frac{V_{FS}}{2\sqrt{2}}\right)^2}{\sqrt{\epsilon_{q+DNL}^2 + V_{thermal}^2 + V_{jitter}^2}} \right]. \quad [\text{dB}] \quad (2.27)$$

Signal-to-Noise-and-Distortion Ratio (SNDR) or SINAD is the ratio of the RMS signal amplitude to the mean RSS value of all spectral components, including quantization error, noise, and harmonics. Again, for a FS input sinusoid with a peak-to-peak amplitude of V_{FS} , the following expression evaluates SNDR:

$$SNDR = 10 \log \left[\frac{\left(\frac{V_{FS}}{2\sqrt{2}}\right)^2}{\sqrt{V_{noise}^2 + V_{harmonics}^2}} \right]. \quad [\text{dB}] \quad (2.28)$$

There exists a relation between THD, SNR, and SNDR provided all of them are characterized under the same input signal conditions (amplitude and frequency) [30]. This relation is summarized with the equations below

$$THD = -10 \log \left[10^{-(SNDR/10)} - 10^{-(SNR/10)} \right], \quad [\text{dB}] \quad (2.29)$$

$$SNR = -10 \log \left[10^{-(SNDR/10)} - 10^{-(THD/10)} \right], \quad [\text{dB}] \quad (2.30)$$

$$SNDR = -10 \log \left[10^{-(SNR/10)} + 10^{-(THD/10)} \right]. \quad [\text{dB}] \quad (2.31)$$

Effective Number of Bits (ENOB) is the actual converter accuracy after adding up all error sources. It can be calculated by using Eq. (2.9) and solving for B after substituting SNDR for SQNR

$$ENOB = \frac{SNDR - 1.76}{6.02}. \quad (2.32)$$

Spurious Free Dynamic Range (SFDR) is one of the most important specifications in ADCs for communications applications. It is quantified as the ratio of the RMS value of the fundamental signal to the RMS value of the largest undesired spectral content. It may be specified either in dBc or in dBFS (decibels below FS). For input signals near FS, it typically coincides with the largest HDn. There might be cases though, where some other distortion product determines SFDR (e.g., an error tone due to interleaving; see Sect. 3.7 from the next chapter).

Analog Bandwidth (BW) is defined as the frequency at which the output power of the reconstructed fundamental drops by 3 dB below its low-frequency value. It does not contain any useful information regarding the spectral purity of the converter at that frequency.

Effective Resolution Bandwidth (ERBW) is defined as the frequency at which there is a 3 dB drop in SNDR (or a 0.5 bit drop in ENOB) compared to its low-frequency value. For reasons that will become obvious in the following chapter, it is highly desirable (but not always easily achievable) that both the analog BW and the ERBW are above the Nyquist frequency.

Noise Spectral Density (NSD) is another important frequency domain metric that measures the noise per unit bandwidth at a given frequency. It may be specified either in V^2/Hz or in dB/Hz . Assuming a flat NSD over a certain band, the SNR within this bandwidth is linked with the NSD via the expression

$$NSD = -SNR - 10 \log(BW). \quad [\text{dB}/\text{Hz}] \quad (2.33)$$

N^{th} -order Intermodulation Distortion (IM n) is the equivalent HD n when applying two closely spaced sinusoidal inputs at frequencies f_1 and f_2 . The amplitude of each tone is backed off by at least 6 dB compared to a one-tone to avoid clipping upon in-phase addition of the two tones. The second-order and third-order products are usually the dominant ones. The second-order products are located at $f_2 \pm f_1$ and can be removed by filtering. The third-order products contain two pairs located at $2f_1 \pm f_2$ and $2f_2 \pm f_1$, respectively. The ones at $2f_1 - f_2$ and $2f_2 - f_1$ are of special interest since they fall close to the two fundamentals and properly characterize the converter's spectral purity.

Multi-Tone Power Ratio (MTPR) can be seen as an evaluation metric for the in-band SFDR when multiple sinusoidal inputs are applied. This metric is particularly useful in multi-channel communication systems such as Orthogonal Frequency Division Multiplexing (OFDM) [31]. A large number of tones equal in amplitude and in frequency spacing are applied, and one of them is eliminated from the input signal leaving an empty bin [32]. However, due to the converter's distortion, a small signal appears in that bin. The ratio between the RMS value of one of the fundamental signals and the RMS value of the undesired spectral content in the empty bin yields the MTPR.

2.3.2 Figures of Merit

Some of the metrics described in Sect. 2.3.1 can be used in different combinations and ratios in order to compare the performance of different converters covering similar applications. For this reason, the Figure-of-Merit (FoM) concept has been introduced, serving to measure the power efficiency of a converter with respect

to other specifications, with speed (sample rate) and accuracy the dominant ones. Although many different FoMs exist, two are extensively used in literature and will be summarized below.

Walden's FoM Originally proposed in [33] for Nyquist converters and later adjusted to also cover oversampled converters [34], FoM_W is defined as

$$FoM_W = \frac{Power}{2^{ENOB} \cdot \min\{2BW, f_s\}} \quad [\text{J/conv.-step}] \quad (2.34)$$

and quantifies the energy spent by a converter to achieve a certain accuracy while performing the conversion at a certain speed. Its units are energy (in J) per conversion step. As Eq. (2.34) suggests, for every extra bit of ENOB, power increases by $2\times$. This trend is not obeyed by noise-limited converters, whose power would need to increase by $4\times$ (see Sect. 2.4), which is an important limitation of this FoM.

Schreier's FoM To alleviate the limitation regarding noise-limited converters, FoM_S was proposed, initially ignoring distortion [20] and later adjusted to include both noise and distortion [35]. It is defined as

$$FoM_S = SNDR + 10 \log \left[\frac{\min\{BW, f_s/2\}}{Power} \right]. \quad [\text{dB}] \quad (2.35)$$

Its units are accuracy (in dB) and it depicts more correctly the $4\times$ higher energy per 6 dB of SNDR increase, which is the prevailing trend in the highest-performance designs of recent years. An extensive ADC performance survey by gathering data from works published at the major scientific venues for more than 20 years has been carried out by Prof. Boris Murmann of Stanford University and can be found in [36].

2.4 Accuracy-Speed-Power Limits

In Sect. 2.3.2, it was argued that a converter's performance is a trade-off between accuracy, speed,⁷ and power. The key challenge lies in maximizing the product with accuracy and bandwidth in the numerator and power in the denominator or minimizing its reciprocal by simultaneously pushing all the three parameters as far as possible toward the desired directions.

⁷ It is assumed that for a certain sample rate (speed), the converter needs to achieve the required accuracy for a bandwidth of at least half of that sample rate, and this assumption is used in the equations and plots to follow.

$$\uparrow \left[\frac{\uparrow \text{Accuracy} \cdot \text{Speed} \uparrow}{\text{Power} \downarrow} \right] \iff \left[\frac{\text{Power} \downarrow}{\uparrow \text{Accuracy} \cdot \text{Speed} \uparrow} \right] \downarrow. \quad (2.36)$$

Several error sources were identified in Sect. 2.2, which degrade the accuracy of a real converter below the ideal quantization error. As discussed in the previous section, errors that are associated with mismatch⁸ or non-linearity can be compensated either by design or by calibration with a small overhead on the other two parameters. On the other hand, errors stemming from thermal noise introduce a more fundamental trade-off on Eq. (2.36); improving one of its parameters will most likely result in an analogous degradation of the other two. The significance of such errors on the accuracy-speed-power are analyzed, and some fundamental limits on a converter's performance are established.

2.4.1 Sampler Noise Limit

In Sect. 2.2, Eq. (2.13) was derived for the single-ended sampler thermal noise. We repeat this expression here for a differential configuration,⁹ which is the start for our derivations, assuming an ideal noiseless front-end ($\alpha_{FE} = 1$)

$$\overline{V_{n,\text{samp}}^2} = \frac{2kT}{C_S}, \quad [\text{V}^2] \quad (2.37)$$


The accuracy degradation due to $\overline{V_{n,\text{samp}}^2}$ can be calculated by combining Eqs. 2.27 and 2.32 and considering a differential peak-to-peak signal swing of $V_{\text{FS-diff}}$

$$\begin{aligned} ENOB_{n,\text{samp}} &= \frac{1}{6.02} \cdot \left[10 \log \left(\frac{1}{8} \frac{V_{\text{FS-diff}}^2}{\epsilon_q^2 + \overline{V_{n,\text{samp}}^2}} \right) - 1.76 \right] \\ &= \frac{1}{6.02} \cdot \left[10 \log \left(\frac{1}{8} \frac{V_{\text{FS-diff}}^2}{\epsilon_q^2} \cdot \frac{1}{1 + \frac{\overline{V_{n,\text{samp}}^2}}{\epsilon_q^2}} \right) - 1.76 \right] \quad (2.38) \\ &= B - \frac{1}{6.02} \cdot 10 \log \left(1 + \frac{24}{2^{2B}} \frac{kT}{C_S} \right). \end{aligned}$$

⁸ A comprehensive analysis on the implications of mismatch in the design of analog circuits can be found in [37].

⁹ For differential signaling, the signal power increases by 4x, while the noise increases by 2x, leading to a 3 dB SNR improvement. Furthermore, the even-order harmonics are ideally fully suppressed, leading to an SFDR boost. On the downside, the power increases by 2x.

Fig. 2.17 Simple sampler model with input termination network



The minimum capacitance for a tolerable ENOB reduction can then be obtained for a certain input swing. It is evident from the above expression that to minimize the accuracy degradation due to $\overline{V_{n,samp}^2}$, C_S must be maximized. On the other hand, Eq. (2.12) implies that in order to maximize the bandwidth, C_S must be minimized (for a fixed R_S). To quantify this fundamental trade-off more completely, we add in the simple sampler model (Fig. 2.13) the basic input termination network, as shown in Fig. 2.17, which in some form is given in every converter measurement system. C_S can then be written as

$$C_S = \frac{1}{2\pi[(R_{i,src}/R_{i,int}) + R_S]f_{in}} = \frac{1}{\pi(0.5R_{i,int} + R_S)f_s}, \quad [F] \quad (2.39)$$

where $R_{i,src} = R_{i,int}$ and represent the external source resistance and the internal termination, respectively. Employing Eq. (2.28) with $\overline{V_{n,samp}^2}$ the sole noise contribution, and combining Eqs. (2.37) and (2.39), we reach to the final accuracy-speed limit

$$SNDR_{n,samp} = 10 \log \left[\frac{V_{FS-diff}^2}{8\pi kT(0.5R_{i,int} + R_S)f_s} \right]. \quad [dB] \quad (2.40)$$

The outcome of the above expression is that for a fixed termination network and C_S value, the only optimization “knob” in preserving the Nyquist $SNDR_{samp}$ as the sample rate increases is to reduce R_S accordingly. In Chap. 5, a sampling circuit that outperforms existing circuits in minimizing R_S will be presented.

The absolute minimum power required to charge C_S can be calculated in a similar fashion as in [38]. We assume that the charging occurs within half a period of f_s and the signal utilizes an input swing V_{FS} equal to the supply voltage V_{DD} . Keeping the $SNDR_{samp}$ as a measure of accuracy, the minimum power to achieve a certain accuracy dictated by the sampler noise is given by

$$\begin{aligned} P_{n,samp} &= V_{DD} \cdot I_{samp} = 2 \cdot 8 \cdot \overline{V_{FS}^2} \cdot f_s \cdot C_S \\ &= 16kTf_s \cdot SNDR_{n,samp}, \end{aligned} \quad [W] \quad (2.41)$$

where we substitute $SNDR_{n,samp} = \overline{V_{FS}^2}/\overline{V_{n,samp}^2}$. The above expression gives the accuracy-power limit due to the sampler noise. We can obtain the same result by allocating a full quantization noise contribution to the sampler and substituting C_S

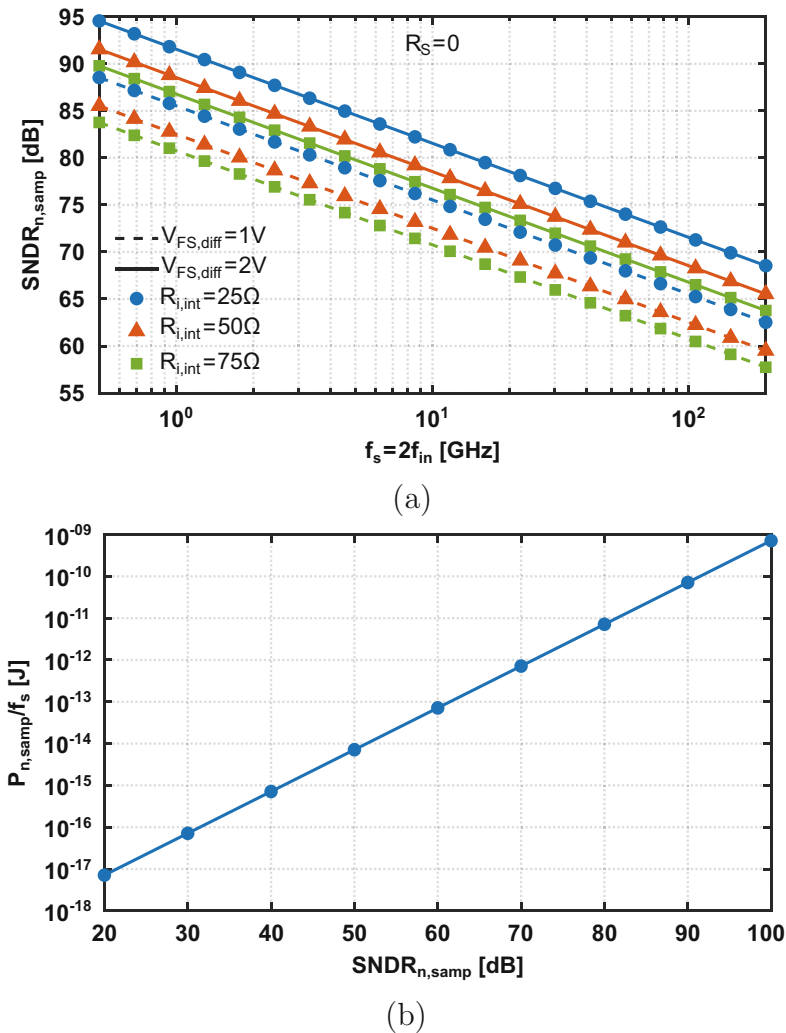


Fig. 2.18 Fundamental limits due to sampler noise: (a) accuracy-speed and (b) accuracy-power

in the above expression. The fundamental limits described by Eqs. 2.40 and 2.41 are plotted in Fig. 2.18 sweeping different parameters.

It is worth mentioning that recently published works [39–41] have shown progress in attempting to “break” the $V_{n,samp}^2$ fundamental limits described above. The underlying principle is to either decouple the generating noise source from the sampling bandwidth or sample the noise and then somehow cancel it. As such, these techniques necessitate additional components (resistors, capacitors, switches, amplifiers) in either open-loop or closed-loop configurations. When going at very high sample rates (>GHz), achieving the necessary amplification

and/or generating extra clocks (including associated routing overhead) for complex switching schemes, to bring down the noise, might take away some or all of the power, bandwidth, and area benefits of scaling down C_S . These might explain why such designs have yet to achieve sample rates beyond several MS/s.

2.4.2 Quantizer Noise Limit

The quantizer thermal noise introduces a second fundamental converter accuracy-speed-power limit. It is mainly defined by the input integrator stage preceding the final latch, as we also derived for our simple model of Fig. 2.14. This also makes the quantizer analysis easier, separating the noise critical input from the bandwidth critical latch (see Sect. 2.4.3). The two stages will be analyzed separately as they both impose different limits, and their contributions will be quantified. The noise power with all the assumptions from our basic model is written here in its differential form to start our derivations and given by

$$\overline{V_{n,\text{quant}}^2} = \frac{2kT}{AC_I}. \quad [\text{V}^2] \quad (2.42)$$

The accuracy reduction due to $\overline{V_{n,\text{quant}}^2}$ can be calculated by combining Eqs. 2.27 and 2.32 and considering a differential peak-to-peak signal swing of $V_{\text{FS-diff}}$

$$\begin{aligned} ENOB_{n,\text{quant}} &= \frac{1}{6.02} \cdot \left[10 \log \left(\frac{1}{8} \frac{V_{\text{FS-diff}}^2}{\epsilon_q^2 + \overline{V_{n,\text{quant}}^2}} \right) - 1.76 \right] \\ &= \frac{1}{6.02} \cdot \left[10 \log \left(\frac{1}{8} \frac{V_{\text{FS-diff}}^2}{\epsilon_q^2} \cdot \frac{1}{1 + \frac{\overline{V_{n,\text{quant}}^2}}{\epsilon_q^2}} \right) - 1.76 \right] \\ &= B - \frac{1}{6.02} \cdot 10 \log \left(1 + \frac{24 \frac{kT}{C_I}}{\frac{V_{\text{FS-diff}}^2}{2^B}} \right), \end{aligned} \quad (2.43)$$

which yields the minimum capacitance at the integrator output for a targeted reduction in ENOB and a given signal swing. To minimize this reduction, C_I ¹⁰ must be maximized, which adversely affects the input integrator's operating frequency, expressed as

¹⁰ Our model assumed $C_I = C_L$, which is not far from a realistic design scenario in 28 nm CMOS (see Chap. 4).

$$f_I = \frac{I_I}{C_I \Delta V_I} = \frac{g_{m,I} V_{GT,I}}{2C_I \Delta V_I}. \quad [\text{Hz}] \quad (2.44)$$

ΔV_I is the common-mode voltage rise/fall at the integrator output to build a certain gain, and I_I follows the basic MOS equation [42]

$$\frac{g_m}{I_D} = \frac{2}{V_{GT}}, \quad V_{GT} = \begin{cases} 2nkT/q \approx 60 - 80 \text{ mV}, & \text{Weak - Inversion} \\ V_{GS} - V_{TH}, & \text{Strong - Inversion} \\ 2(V_{GS} - V_{TH}), & \text{Velocity - Saturation} \end{cases}. \quad (2.45)$$

As with the sampler, we allocate half a period of f_s to the quantizer; thus, this is the maximum available time for the integrator. Combining Eqs. (2.42) and (2.44) and employing Eq. (2.28) with $\overline{V_{n,\text{quant}}^2}$ its only noise contribution, the accuracy-speed limit is derived

$$SNDR_{n,\text{quant}} = 10 \log \left[\frac{g_{m,I} V_{GT,I} V_{FS-\text{diff}}^2}{32kT \Delta V_I f_s} \right]. \quad [\text{dB}] \quad (2.46)$$

The minimum necessary power to charge C_I can be calculated with a similar method as for Eq. (2.41), following the same assumptions about the input signal. Additionally, by allocating a maximum value of $\Delta^2/12$ to $\overline{V_{n,\text{quant}}^2}$ for convenience,¹¹ the minimum power to achieve a certain accuracy dictated by the quantizer noise (accuracy-power limit) can be found as


$$\begin{aligned} P_{n,\text{quant}} &= V_{DD} \cdot I_I = 2 \cdot V_{FS} \cdot f_s \cdot C_I \cdot \Delta V_I \\ &= 4 \cdot V_{FS} \cdot f_s \cdot \frac{12kT}{V_{FS}^2} \cdot 2^{2ENOB_{n,\text{quant}}} \cdot \frac{V_{FS}}{2} \\ &= 24kT f_s \cdot 2^{\frac{SNDR_{n,\text{quant}} - 1.76}{6.02}}, \end{aligned} \quad [\text{W}] \quad (2.47)$$

where Eq. (2.32) is used, V_{DD} is assumed to be equal to V_{FS} , and ΔV_I is assumed to be half V_{FS} at the end of the integration. The fundamental limits described by Eqs. (2.46) and (2.47) are plotted in Fig. 2.19 sweeping different parameters. In Sect. 2.4.7, all limits will be plotted together for comparison.


2.4.3 Metastability Limit

In addition to the noise, metastability is another fundamental error source associated with the output latch stage of the quantizer. The latch regenerates exponentially on

¹¹ In high-speed converters, it is general practice to design the various thermal noise sources in the same order as the quantization noise.



(a)



(b)

Fig. 2.19 Fundamental limits due to quantizer noise: (a) accuracy-speed and (b) accuracy-power

an input according to the following expression:

$$V_{\text{out}} = A V_{\text{in}} \cdot e^{-\frac{T_L}{\tau}} = A V_{\text{in}} \cdot e^{-\frac{g_{m,L} T_L}{C_L}}, \quad [\text{V}] \quad (2.48)$$

where A is the integrator's gain (see Sect. 2.2.1), while the time constant $\tau = C_L/g_{m,L}$ is a measure of the latch's bandwidth. Metastability refers to the situation where the quantizer differential input is so small (e.g., a fraction of an LSB), such that for the allowed operation time, the latch of the quantizer cannot produce a sufficiently large differential output for the following circuitry to unambiguously




Fig. 2.20 Quantizer output for a valid (gray) and a metastable (black) case

perceive it as a clear logical level. This scenario, portrayed in Fig. 2.20, results in a conversion error, therefore leading to accuracy degradation. For a certain input voltage and a fixed gain A , this error can be reduced either by allowing more time to the quantizer to produce a sufficiently large output difference or by minimizing τ .

The error due to metastability may be interpreted as an increased quantization noise floor with a variance $\overline{\epsilon_q^2}$ multiplied by a certain probability of occurrence $PR(meta)$ [43]. The total converter noise may be then written as

$$\overline{V_{q+meta}^2} = \overline{\epsilon_q^2} \cdot [1 + PR(meta)]. \quad [V^2] \quad (2.49)$$

The second term inside the square brackets denotes the excess noise due to metastability. If we consider a differential input signal uniformly distributed within $\pm V_{FS-diff}/2$, then the probability of a metastable occurrence, otherwise known as Bit Error Rate (BER), can be seen as the ratio of the smallest input the latch can correctly regenerate on its given time divided by the full input range. For a B -bit quantizer with an equal probability of showing metastability in any of the 2^B steps, utilizing Eq. (2.48), $PR(meta)$ can be expressed as

$$PR(meta) = BER \cdot 2^{B_{meta}} = \frac{2^{B_{meta}} V_{in,min}}{\frac{V_{FS-diff}}{2^{B_{meta}+1}}} = \frac{2^{2B_{meta}} \cdot e^{-T_L/\tau}}{A}, \quad (2.50)$$

where it is assumed that the quantizer latch regenerates to V_{FS} and $B = B_{meta}$. $PR(meta)$ has an exponential dependency on τ ; therefore, minimizing it is extremely desirable. Further, if we re-write τ lumping the total capacitance at the quantizer output, we can see that the technology ultimately dictates the minimum achievable value

$$\tau \approx \frac{C_{gg}}{g_{m,L}} \approx \frac{1}{2\pi f_T}, \quad (2.51)$$

where f_T is the cut-off frequency for which the current gain is unity. In order to take into account practical limitations (e.g., layout parasitics), a more realistic value of $1/\pi f_T$ is adopted for τ , in all the subsequent analysis. Substituting Eqs. (2.50) and (2.51) into (2.49), we have

$$\overline{V_{q+meta}^2} = \overline{\epsilon_q^2} \cdot \left[1 + \frac{2^{2B_{meta}} \cdot e^{-\pi f_T/2f_s}}{A} \right], \quad [V^2] \quad (2.52)$$

where half a period of f_s is allocated for latch regeneration.¹² By allocating a certain small LSB fraction $a_{er} < 1$ to the error due to the excess noise in the above expression, and employing Eq. (2.32), the accuracy-speed limit imposed by metastability can be derived for various f_T values

$$SNDR_{meta} = 6.02 \left[\frac{\log_2(a_{er}A)}{2} + \frac{\pi f_T}{4 f_s \ln 2} \right] + 1.76. \quad [dB] \quad (2.53)$$

The take from the above expression is that if the quantizer resolution increases while preserving the same f_s and f_T , there is an increased excess noise due to metastability on the total quantization noise.

It is important to clarify that the above limit is derived under the assumption of f_s being the sample rate of a standalone non-pipelined non-interleaved quantizer. As such, it is the reciprocal of the standalone quantizer's latch delay T_L to achieve a certain resolution. Pipelining can improve this limit by reducing the quantizer resolution per pipeline stage, therefore increasing the overall resolution for the same total f_s or increasing the total f_s for the same overall resolution. Interleaving can also improve this limit, as discussed in the next chapter. By multiplexing several quantizers in time, each running at a lower standalone f_s , the aggregate f_s can be increased by the interleaving factor while also preserving the resolution.

In order to estimate the minimum power required by the latch to resolve within half a period of f_s a certain small input $A V_{FS-diff}/2^{B_{meta}+1}$ ($A = 4 = 2^2$) and regenerate to V_{FS} , we start the derivation by substituting this value in Eq. (2.48) and solve for $g_{m,L}$

$$g_{m,L} = 2(B_{meta} - 2) \cdot \ln 2 \cdot f_s \cdot C_L. \quad [S] \quad (2.54)$$

This $g_{m,L}$ will require a minimum current I_L , and these two are related through the basic MOS Eq. (2.45). Before we reach to the final expression for the power, we need to substitute C_L from the latch noise Eq. (2.14) and assume that the input-referred latch noise voltage is at least 4× smaller than the input that leads to metastability. This assumption aligns well with our two-stage quantizer model

¹² In a practical design, the input stage and the latch will each occupy a portion of the quantizer allocated time. Our simplification will affect our derivations by about 2×, which is tolerable for first-order generic derivations.

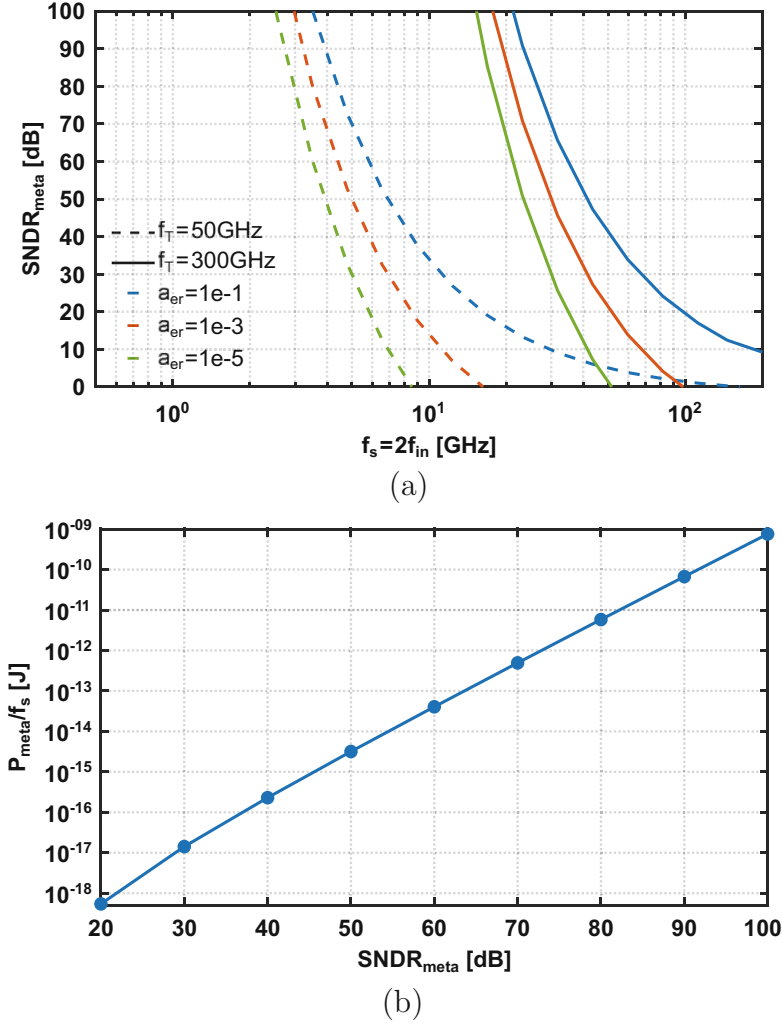


Fig. 2.21 Fundamental limits due to metastability of a standalone quantizer: (a) accuracy-speed and (b) accuracy-power

and allows to a first-order a proper metastability assessment. Finally, utilizing a supply voltage $V_{DD} = 1$ V equal to V_{FS} , we obtain the minimum power dictated by metastability, translating to the accuracy-power limit

$$\begin{aligned}
 P_{\text{meta}} &= V_{DD} \cdot I_L \\
 &= 24(B_{\text{meta}} - 2) \cdot \ln 2 \cdot f_s \cdot kT \cdot \frac{V_{GT}}{V_{FS}} \cdot 2^{2B_{\text{meta}}}. \quad [\text{W}] \quad (2.55)
 \end{aligned}$$

The fundamental metastability limits described by Eqs. (2.53) and (2.55) are plotted in Fig. 2.21 for different values of f_T and a_{er} .

2.4.4 Aperture Jitter Limit

Equation (2.18), from which the error noise power for a sinusoidal signal was obtained, can be adjusted to yield the differential jitter noise power for a differential peak-to-peak signal swing of $V_{\text{FS-diff}}$

$$\overline{V_{\text{n,jitter}}^2} = \frac{1}{2} (\pi f_{\text{in}} V_{\text{FS-diff}})^2 \cdot t_{\text{jitter}}^2. \quad [\text{V}^2] \quad (2.56)$$

The accuracy reduction due to $\overline{V_{\text{n,jitter}}^2}$ can be calculated in a similar way as in Eqs. (2.38) and (2.43)

$$\begin{aligned} ENOB_{\text{n,jitter}} &= \frac{1}{6.02} \cdot \left[10 \log \left(\frac{1}{8} \frac{V_{\text{FS,diff}}^2}{\epsilon_{\text{q}}^2 + \overline{V_{\text{n,jitter}}^2}} \right) - 1.76 \right] \\ &= B - \frac{1}{6.02} \cdot 10 \log \left(1 + 2^{2B} \cdot 6(\pi f_{\text{in}})^2 \cdot t_{\text{jitter}}^2 \right), \end{aligned} \quad (2.57)$$

from which the jitter value is obtained for a tolerable ENOB degradation and at a certain input frequency. The voltage error due to jitter is an increasing function of the frequency. This can be intuitively understood by the fact that a fixed error in time results in a larger voltage error when reflected to a signal with a faster slope compared to a slower slope signal. If we substitute Eq. (2.56) in the SNDR expression (Eq. (2.28)) and consider $\overline{V_{\text{n,jitter}}^2}$ the only noise source, the accuracy-speed limit due to jitter can be obtained

$$SNDR_{\text{n,jitter}} = 10 \log \left[\frac{1}{4(\pi f_{\text{in}})^2 \cdot t_{\text{jitter}}^2} \right], \quad [\text{dB}] \quad (2.58)$$


which is an already known expression [26], re-verified here by our analysis.

The minimum power to achieve a certain accuracy imposed by jitter noise is not entirely straightforward because strictly speaking, this power is not dissipated in the core converter parts (sampler and quantizer) but in the clock generation. Nevertheless, since the clock is an imperative part in any converter,¹³ we are including it in our fundamental limits for a comparison point.

To provide a first-order estimation of the clock power for a certain jitter, we model the clock generation as a single $g_{\text{m,CK}}$ unity gain buffer (Fig. 2.22) and assume linear operation for the entire clock swing, which is equal to V_{DD} . To simplify the analysis, we also assume that the dominant source leading to jitter is the

¹³ In the case of continuous-time ADCs [44, 45], although the input is not sampled, sampling is still performed along the chain (quantizer, back-end, reconstruction filter) to align with a synchronous clock. Depending on the part of the chain, jitter requirements can be different.

Fig. 2.22 Simple model for clock power estimation for a certain jitter



buffer thermal noise $\overline{V_{n,CK}^2}$, which, due to the unity gain, can be directly referred to the output. This noise can be calculated in a similar way as the quantizer noise (see Sect. 2.2.1, Eq. (2.16)). The buffer needs to charge C_{CK} ¹⁴ to V_{DD} , and we allocate a maximum of a quarter period of f_s to allow sufficient time for the actual sampling within half a period of f_s . The minimum required power consumed in the clock is then given as

$$\begin{aligned} P_{\text{jitter}} &= V_{DD} \cdot I_{CK} = 4V_{DD}^2 \cdot f_s \cdot C_{CK} \\ &= 4V_{DD}^2 \cdot f_s \cdot \frac{kT \cdot T_s^2}{16V_{DD}^2 \cdot t_{\text{jit}}^2}, \end{aligned} \quad [\text{V}^2] \quad (2.59)$$

where the Slew Rate (SR), which translates $\overline{V_{n,CK}^2}$ to t_{jit}^2 , has been written as voltage/time to provide V_{DD} within $0.25T_s$. By substituting t_{jitter}^2 from Eq. (2.56) for an input swing V_{FS} equal to V_{DD} and a Nyquist input frequency, the minimum power for a certain jitter is obtained

$$P_{\text{jitter}} = \frac{\pi^2}{2} kT f_s \cdot SNDR_{n,\text{jitter}}, \quad [\text{V}^2] \quad (2.60)$$

where $SNDR_{n,\text{jitter}} = \overline{V_{FS}^2} / \overline{V_{n,\text{jitter}}^2}$. Despite the several assumptions made to simplify the analysis, the above expression yields to a first-order a correct accuracy-power limit due to jitter, which is on par with the equivalent limits from the sampler and quantizer noise. The jitter-imposed limits of Eqs. (2.58) and (2.60) are plotted in Fig. 2.23 for several different parameters.

2.4.5 Mismatch Limit

At the beginning of this section, it was argued that errors associated with mismatch can be compensated with a small overhead, thus not introducing a fundamental trade-off between accuracy, speed, and power. Nevertheless, it is insightful to quantify the accuracy-speed and accuracy-power limits imposed by mismatch and

¹⁴ This capacitor includes the intrinsic buffer load and the sampling switch gate load.

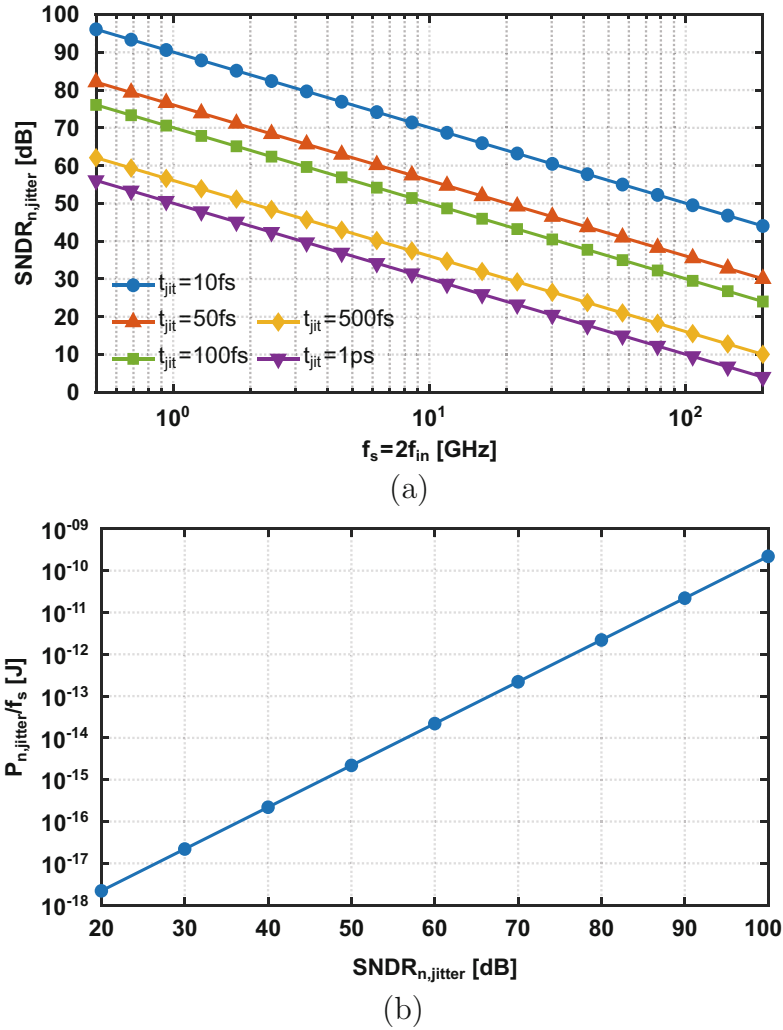


Fig. 2.23 Fundamental limits due to aperture jitter: (a) accuracy-speed and (b) accuracy-power

compare them to the derived ones imposed by noise, especially since the former are process dependent.

Similar to noise, mismatch is a random process as well, with a mean μ_M and a standard deviation σ_M (or variance σ_M^2). Assuming a differential pair with a mismatch dominated by the random variation in V_{TH} between the two devices, from Pelgrom’s law [28], we obtain the variance

$$\sigma_M^2 = \frac{A^2 V_{TH}^2}{WL}, \quad [V^2] \quad (2.61)$$

where $A_{V_{TH}}$ is a mismatch constant that depends on the process. σ_M^2 is inversely proportional to the area. Assuming also that the devices of the differential pair are biased in strong inversion, the input capacitance C_M is found

$$C_M = (2/3)WLC_{ox} = \frac{2A_{V_{TH}}^2 C_{ox}}{3\sigma_M^2}. \quad [F] \quad (2.62)$$

This capacitance together with the source and internal termination resistances creates an upper limit to the input bandwidth, as shown in Eq.(2.39) if C_S is replaced by C_M . If we then combine Eqs.(2.28), (2.39), and (2.62) and consider a $3\sigma_M$ confidence interval for the mismatch contribution, we finally reach to the accuracy-speed limit

$$SNDR_{\sigma,match} = 10 \log \left[\frac{V_{FS-diff}^2}{48\pi A_{V_{TH}}^2 C_{ox} (0.5R_{i,int} + R_S) f_s} \right]. \quad [dB] \quad (2.63)$$

The minimum power required to charge C_M can be derived similarly to the one for charging C_S in the sampler noise limit. We allocate half a period of f_s for the operation and assume an input swing V_{FS} equal to the supply voltage. If we also consider a $3\sigma_M$ mismatch confidence interval, re-employing Eq.(2.41) and keeping $SNDR_{\sigma,match}$ as a measure of accuracy, we end up with the accuracy-power limit due to mismatch

$$\begin{aligned} P_{\sigma,match} &= V_{DD} \cdot I_{match} = 2 \cdot 8 \cdot \overline{V_{FS}}^2 \cdot f_s \cdot C_M \\ &= 48A_{V_{TH}}^2 C_{ox} f_s \cdot SNDR_{\sigma,match}, \end{aligned} \quad [W] \quad (2.64)$$

Comparing the above two expressions with Eqs.(2.40) and (2.41) giving the equivalent limits due to noise, we see $A_{V_{TH}}^2 C_{ox}$ in the denominator instead of kT , plus an extra multiplication factor depending on the targeted σ_M confidence interval. Both $A_{V_{TH}}$ and C_{ox} are technology-dependent parameters, indicating the effect of the process on the matching limit, in contrast to the baseline noise limit. Table 2.2 shows typical values of these parameters for three different process nodes [12], while the derived mismatch limits are plotted in Fig. 2.24.

Table 2.2 Typical process parameters and comparison with kT

Process [nm]	$A_{V_{TH}}$ [mV- μ m]	C_{ox} [fF/ μ m ²]	$A_{V_{TH}}^2 C_{ox}/kT$
130	5	11	61.7
65	4	13	46.7
28	2	25	22.4

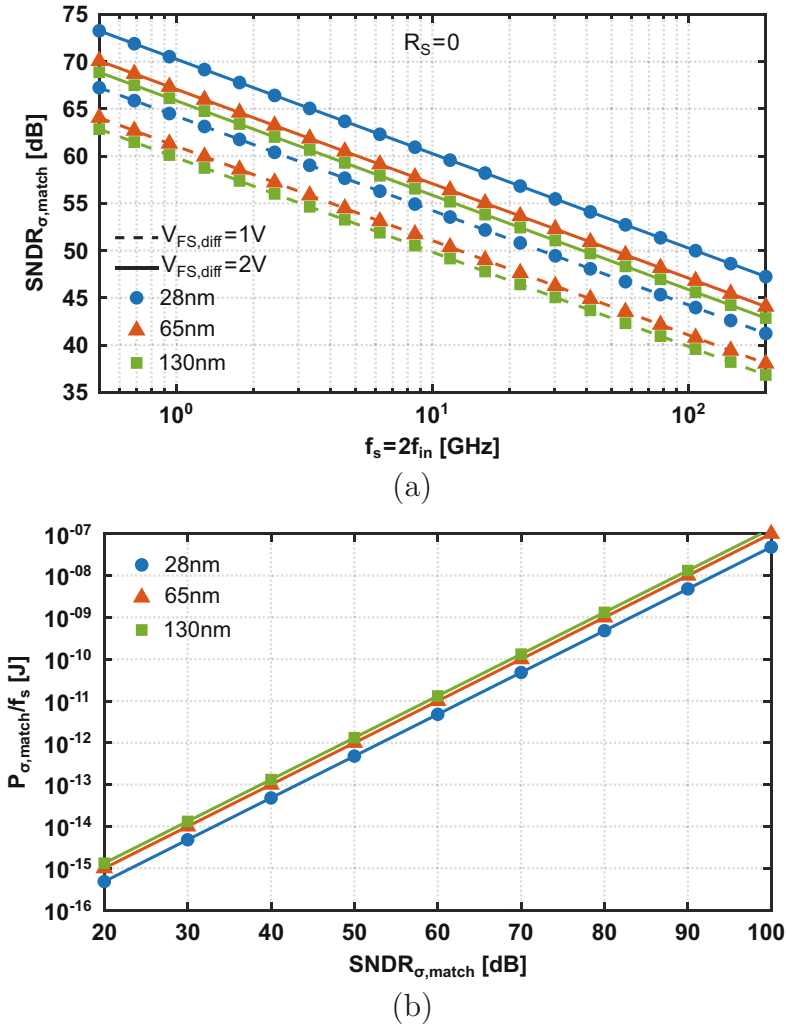


Fig. 2.24 Limits imposed by mismatch: (a) accuracy-speed and (b) accuracy-power

2.4.6 Heisenberg Uncertainty Principle

To complete our analysis, the Heisenberg uncertainty principle is also discussed based on [33], as the ultimate accuracy-speed limit in a converter’s performance, ultimately imposed by physics. The original principle [46] limiting what can be simultaneously known about the position and momentum of a quantum particle also applies to the energy-time complementary set stating




Fig. 2.25 Fundamental accuracy-speed limit due to Heisenberg

The more precisely the energy of a particle in a certain state is known, the greater the uncertainty in the interval of time, in which the particle possesses that particular energy.

The principle is described by the mathematical formula

$$\Delta E \cdot \Delta T \geq \frac{h}{4\pi}, \quad (2.65)$$

where ΔE may be interpreted as the required energy to be within $\pm \text{LSB}/2$ of a quantization level, ΔT is the time required to move from one level to another and assumed half a period of f_s , and $h = 6.62617 \cdot 10^{-34}$ J·s is the Planck constant. Under these assumptions and using $R_{i,\text{src}}$ from the model of Fig. 2.17, the above expression for a differential configuration can be written as

$$\frac{V_{\text{pp-diff}}^2}{2^{2ENOB_{\text{Heis}}} \cdot 8R_{i,\text{src}}} \cdot \frac{1}{(2f_s)^2} \geq \frac{h}{4\pi} \Rightarrow 2^{ENOB_{\text{Heis}}} \cdot f_s \leq \frac{V_{\text{pp-diff}}}{2\sqrt{2hR_{i,\text{src}}}}. \quad (2.66)$$

Finally, the maximum achievable SNDR dictated by the Heisenberg uncertainty principle can be obtained by utilizing Eq. (2.32) (Fig. 2.25)

$$\text{SNDR}_{\text{Heisenberg}} = 6.02 \log_2 \left(\frac{V_{\text{pp-diff}}}{2f_s \sqrt{2hR_{i,\text{src}}}} \right) + 1.76. \quad [\text{dB}] \quad (2.67)$$

2.4.7 Putting It All Together

To finalize our analysis, in Fig. 2.26, we plot all the previously derived accuracy-speed and accuracy-power limits for certain design choices and parameters. As seen in Fig. 2.26, the quantizer metastability for an a_{er} of $1e-5$ is the dominant accuracy limitation when increasing the sample rate above about 25 GS/s. Below this frequency, aperture jitter of 50 fs dominates the accuracy degradation down to about 4 GS/s. At lower sample rates, mismatch is the main limitation to the

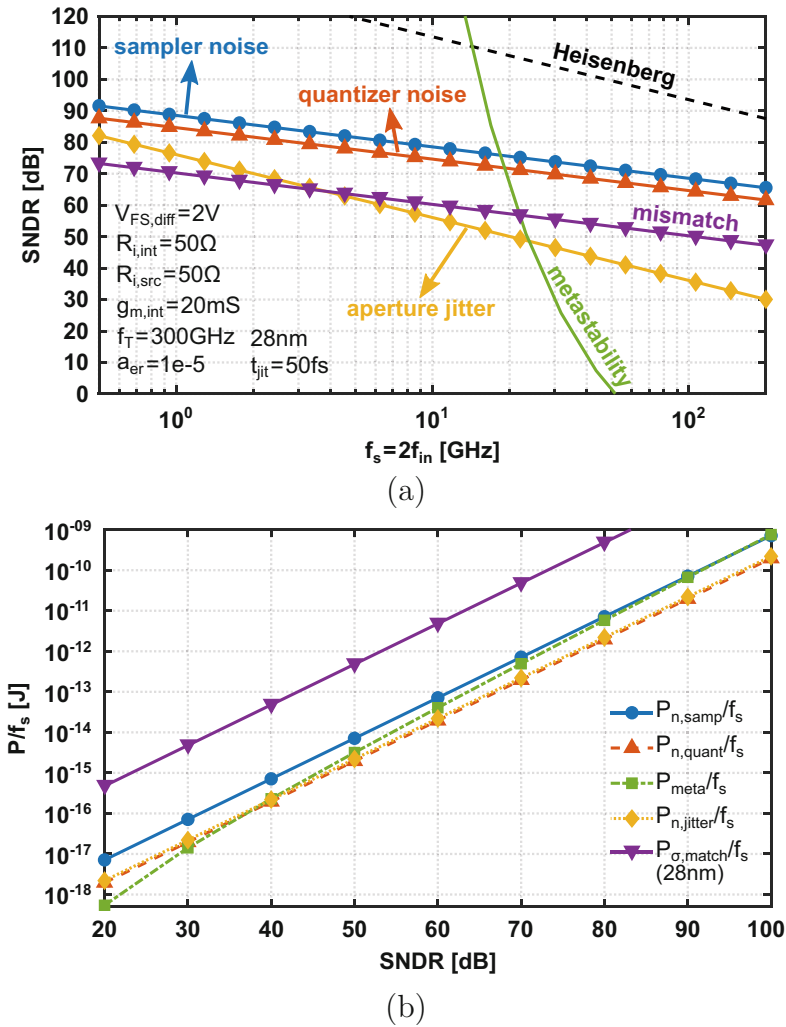


Fig. 2.26 Fundamental limit curves from all the error sources analyzed in this chapter: (a) accuracy-speed and (b) accuracy-power

achievable resolution. Assuming that mismatch is compensated, thermal noise starts limiting the achievable resolution for sample rates below 500 MS/s, with the quantizer as the dominant error source based on our derivations. This is expected at very low sample rates due to the steeper slope of the jitter-limited resolution. The physical Heisenberg uncertainty principle limitation is about 30 dB above the next limitation.

Regarding Fig. 2.26b, our simplified derivations indicate a maximum of about half an order of magnitude power consumption difference between the various noise and metastability limitations. For a 28 nm process, mismatch imposes a power consumption limit of about two orders of magnitude higher than the rest. In reality, the power of the sampler is expected to increase in the presence of an analog front-end with a certain settling requirement. Also, the power estimation for a certain jitter neglects multiple stages in the chain of Fig. 2.22 to realize a certain clock edge steepness, which will inevitably increase this power. Nevertheless, the important take from this first-order power analysis is that every contribution in a converter necessitates an equally careful optimization and/or compensation to yield the best overall results.

2.5 Conclusion

This chapter laid out the fundamental concepts of the A/D conversion process. Its two primary concepts of sampling (time discretization) and quantization (amplitude discretization) were thoroughly discussed. In order to prevent loss of information and yield the sampling process reversible, the Nyquist criterion dictates that the sample rate be at least twice the instantaneous bandwidth of the signal under sampling. The signal may be located in any of the Nyquist zones, and as long as it is band-limited within one, the Nyquist criterion is satisfied. Each sampled value is compared against 2^B discrete levels, and its amplitude is rounded to the nearest level by the quantizer. This rounding process introduces a deterministic quantization error ϵ_q , which under certain conditions can be approximated as white noise, and imposes the ideal single conversion error source. From this analysis, the maximum possible accuracy of a B -bit converter was derived in terms of its SQNR. The major error sources from the circuit blocks in a practical converter chain were identified to deteriorate the performance beyond the quantization error threshold. In the form of noise, these include the sampler thermal noise, the quantizer thermal noise, and the aperture jitter from the clock and input of the sampler. Simple models were introduced, and closed-form expressions were developed to quantify these errors in terms of design parameters. In the form of non-linearity, DNL and INL from the quantizer as well as INL and harmonic distortion from the other blocks in the chain (sampler and a potential front-end) were identified as the main contributors. Generally, any type of non-linearity originates from circuit imperfections and can be minimized either by proper design choices or by calibration, which was briefly overviewed as well. Further, several critical performance evaluation metrics,

including THD, SNR, SNDR, SFDR, as well as the two widely used figures of merit, FoM_W and FoM_S , were briefly discussed.

Equations serving as first-order guidelines were developed, which established the fundamental accuracy-speed-power limits imposed by (1) the sampler noise, (2) the quantizer noise, (3) the quantizer metastability, (4) the aperture jitter, and (5) ultimately physics under certain assumptions. The limits imposed by mismatch were also quantified and compared to the aforementioned ones. The derived equations provided an insight as to what may be ultimately achievable from the elementary building blocks in a converter and what has to be traded-off to maximize the ratio $accuracy \cdot speed \div power$. It was concluded that the contribution from every block needs to be equally carefully optimized and/or compensated to reach the best possible performance. More importantly, this insight allows a better circuit design optimization, avoiding excessive over-design or under-design that could potentially lead to poor power and/or speed performance for a certain accuracy.

Appendix A: Proper FFT Evaluation Setup

Assume we would like to sample at an f_s rate a one-tone sine wave with an input frequency f_{in} and evaluate its frequency spectrum by means of an FFT with N_{FFT} points. The total FFT evaluation time is found as

$$T_{FFT} = N_{FFT} \cdot 1/f_s. \quad (2.68)$$

The resolution bandwidth or FFT bin size is then given by

$$f_{bin} = 1/T_{FFT} = f_s/N_{FFT}. \quad (2.69)$$

For coherent sampling without using windowing, and to avoid spectral leakage, we must ensure an integer number of signal periods N_{PER} . The input frequency is therefore found as

$$f_{in} = N_{PER} \cdot f_{bin} = N_{PER} \cdot f_s/N_{FFT}. \quad (2.70)$$

The same setup can be followed for a two-tone sine wave¹⁵ with input frequencies f_{in1} and f_{in2} as well, provided that these frequencies fall exactly within FFT bins. One of many ways to guaranteeing this is the following:

$$\begin{aligned} f_{in1} &= N_{PER} \cdot f_{bin} - 2 f_{bin} = N_{PER} \cdot f_s/N_{FFT} - 2 f_s/N_{FFT} \\ f_{in2} &= N_{PER} \cdot f_{bin} + 2 f_{bin} = N_{PER} \cdot f_s/N_{FFT} + 2 f_s/N_{FFT}. \end{aligned} \quad (2.71)$$

¹⁵ It can be generalized to an m-tone sine wave with $f_{in1,2,\dots,m}$.

Finally, N_{PER} and N_{FFT} must be relatively prime, meaning that their only positive common divisor is 1. To give a numerical example, for an $f_s = 1$ GS/s and $N_{\text{FFT}} = 1024$, $N_{\text{PER}} = 79$ satisfies the above requirements, leading to an $f_{\text{in}} = 77.1484$ MHz for a one-tone and $f_{\text{in1}} = 75.1953$ MHz and $f_{\text{in2}} = 79.1016$ MHz for a two-tone sine wave, respectively.

Chapter 3

Architectural Considerations for High-Efficiency GHz-Range ADCs



One of the everlasting challenges in the design of data converters is how to get the maximum performance for the minimum amount of energy. This chapter aims to tackle this challenge by extending the block-level fundamental limit derivations of Chap. 2, to an architectural level, in the pursuit of determining the optimal architecture for maximizing $accuracy \cdot speed \div power$.


Section 3.1 begins by reviewing and interpreting the recent state-of-the-art standings, including ADCs presented in the foremost conferences, *ISSCC* and *VLSI*, in the last 10 years. Major architectures, such as flash, SAR, pipeline, and pipelined-SAR, are included. Sections 3.2–3.5 cover these architectures and their trade-offs in detail. After an operation description is given for each, models are derived to estimate and compare their accuracy-speed-power limits, offering a complete decomposition of the individual blocks' contributions. The power of these models is enhanced by including process effects, and the comparison is extended over four deep-scaled CMOS process nodes, building unique insight into both architectural and technological capabilities [47].

Section 3.7 discusses time-interleaving as a popular way of extending the speed of a standalone converter and focuses on key aspects such as interleaving errors and interleaver architectures. Finally, a model is developed to compare the different interleavers in terms of achievable bandwidth and sampling accuracy. The chapter closes with an overview and important conclusions in Sect. 3.8.

3.1 State of the Art


Our investigation on the optimal architectural choice and its associated trade-offs to efficiently achieve the best possible speed and resolution/accuracy begins by examining the State-of-the-Art (SotA) standings. Figure 3.1 illustrates the accuracy-speed (Fig. 3.1a) and accuracy-energy (Fig. 3.1b) performance of SotA

tech. \geq 32nm: \circ (TI-) SAR, \triangle (TI-) Pipe, \square (TI-) PipeSAR, \diamond (TI-) Flash, ∇ Others
 tech. \leq 32nm: \bullet (TI-) SAR, \blacktriangle (TI-) Pipe, \blacksquare (TI-) PipeSAR, \blacklozenge (TI-) Flash, \blacktriangledown Others



(a)

tech. \geq 32nm: \circ (TI-) SAR, \triangle (TI-) Pipe, \square (TI-) PipeSAR, \diamond (TI-) Flash, ∇ Others
 tech. \leq 32nm: \bullet (TI-) SAR, \blacktriangle (TI-) Pipe, \blacksquare (TI-) PipeSAR, \blacklozenge (TI-) Flash, \blacktriangledown Others



(b)

Fig. 3.1 State-of-the-art performance of various ADC architectures with data points taken from [36] (a) accuracy-speed and (b) accuracy-energy

CMOS ADCs published in the foremost conferences in the last 10 years with an SNDR \geq 30 dB and a Nyquist sample rate \geq 400 MS/s. These standings include several architectures, such as SAR, pipeline, pipelined-SAR, flash, and others (time domain and $\Sigma\Delta$), as well as their Time-Interleaved (TI) counterparts. To get a

clearer view on the evolution of the standings, all ADCs have been split into the ones using an older than 32 nm process node and those utilizing a node newer than 32 nm.

To a first extent, one can clearly notice that the most widely used architectures are the SAR and the pipeline, covering a wide range of specifications, while the pipelined-SAR has emerged more recently to combine the best of the two architectures it encompasses (see Sect. 3.5). The flash is mainly limited to lower-resolution levels due to its exponential growth in hardware per added bit (see Sect. 3.2). The remaining architectures (mainly time domain), despite showing impressive progress in recent years, have so far been lagging behind the pipeline/pipelined-SAR in the high-resolution regime and/or the SAR/flash in the low-to-medium-resolution regime. A possible explanation may be that voltage-domain ADCs require one operation (voltage \rightarrow bits), while time-domain converters necessitate two operations (voltage \rightarrow time \rightarrow bits), which unavoidably introduces extra accuracy degradation and/or complexity and overhead.

When a first-order curve fitting is applied on top of the data points of Fig. 3.1a, the corresponding curves for all architectures resemble the Bode diagram of an Operational Amplifier (OpAmp). This is explained by the fact that every ADC ends up with a comparator that generates noise. In order to reduce this noise and achieve a target SNDR, the comparator is preceded by gain stages either as pre-amplifiers (flash, SAR) or as residue amplifiers (pipeline, pipelined-SAR), whose noise is also budgeted in the overall contribution. The total amount of gain is proportional to the required SNDR. Comparing the different curves, we see that the pipeline is the winner in terms of absolute accuracy with very good speed, while the TI-SAR is the leader in absolute speed for medium resolution, with sample rates approaching 100 GS/s. One key note is that the SAR sample rate has benefited more than the pipeline from technology scaling. The pipelined-SAR is on par with the pipeline in terms of accuracy while also approaching similar speed levels. In fact, the flatness in the slope of its curve might predict that future TI-pipelined-SAR designs can perhaps surpass the traditional TI-pipeline and TI-SAR in both accuracy and speed.

First-order curve fitting is also applied to the data points of Fig. 3.1b to gain insight on the energy efficiency of each architecture. The resulting curves reveal that the SAR is leading in efficiency in the medium-to-high accuracy regime, while the pipeline can achieve higher absolute accuracy levels. Both architectures seem to have benefited from technology scaling in terms of efficiency, while the SAR shows an increasing trend of moving to lower accuracy levels (and higher speed) for an increased efficiency. The pipelined-SAR demonstrates the efficiency of the SAR while approaching the accuracy levels of the pipeline. Since almost all the pipelined-SAR ADCs included in this investigation have been implemented in advanced processes, a valid observation is that this hybrid clearly benefits from technology scaling as well. Finally, the flash shows a similar efficiency with the SAR in the low accuracy regime. To build an understanding on the above SotA standings, the key




Fig. 3.2 Block diagram of a B-bit flash ADC (the S/H is optional)

architectures involved (flash, SAR, pipeline, pipelined-SAR) are described in the following sections.¹

3.2 The Flash Architecture

3.2.1 Overview

The flash architecture, whose block diagram is depicted in Fig. 3.2, relies on a full parallelism of multiple comparators to quantize a signal in a single clock cycle [48–51]. For a B -bit resolution, it consists of $2^B - 1$ comparators to simultaneously compare the sampled signal with $2^B - 1$ equally spaced reference levels within FS provided by a resistive ladder. Upon clocking, each comparator evaluates the polarity of the difference between the sampled signal and its corresponding reference tap. The thermometer code output is a row of digital 0s for a negative polarity replaced by a row of digital 1s upon changing polarity. Therefore, a thermometer-to-binary encoder is necessary to translate this thermometer word into the final binary format at the ADC output.

Flash ADCs have the potential to achieve the highest speed among any single-channel architecture due to their full parallelism, necessitating only one clock cycle for conversion. Ideally, the speed of the converter is limited only by the delay of

¹ Although on its own the flash cannot achieve high accuracy levels efficiently, it is an essential part of the pipeline architecture; therefore, it is covered first.

the slowest comparator (on the transition between a 0 and a 1) and the delay of the encoder. From the metastability discussion in the previous chapter, we know that the comparator regenerates exponentially on a certain input. The time constant of this regeneration rate $\tau = C_L/g_{m,L}$ (see Chap. 2, Sect. 2.4.3) and knowledge of the input amplitude provide a good estimate of this architecture's speed. We also acknowledge that the time constant depends on the process unity gain frequency f_T , which typically increases as the technology scales to smaller process nodes (see Sect. 3.2.3). Thus, in theory, the flash architecture benefits from technology scaling. This is indeed the case in Fig. 3.1a, where we see the fastest flash is implemented in a node ≤ 32 nm. Further, due to its minimal latency, flash is a very good candidate for feedback systems.

On the contrary, since a B -bit flash requires $2^B - 1$ comparators, the exponential increase in comparator number sets an upper bound on resolution. For each additional bit, the comparator number increases by $2\times$, with each of them requiring twice the precision for the 6dB increase in SNR. For a noise-limited design, this translates to a $4\times$ power increase for each comparator and a total of $8\times$ increase in the converter power. This power is not utilized efficiently since in each conversion there is only one critical comparator resolving an input within $\pm\Delta/2$. Apart from the power, the capacitive load to the input grows at the same rate. This highly non-linear capacitive load, mainly contributed by the comparators' input devices, results in bandwidth loss and both amplitude- and frequency-dependent distortion, deteriorating the spectral purity. Up to date, flash converters are limited to resolutions of no more than 8 bits.

The resistor ladder implementation poses several challenges as well. In [12], the settling time constant of the ladder for a B -bit flash is found as

$$\tau_{\text{lad}} = \frac{R_{\text{tap}}C_{\text{tap}}}{\pi^2}2^{2B}, \quad [\text{s}] \quad (3.1)$$

where C_{tap} is the capacitance on each tap of the ladder, including the comparator input as well as parasitic contributions from the resistors and the wiring. R_{tap} is the ladder tap resistance given by

$$R_{\text{tap}} = \frac{\frac{i}{2^B} R_{\text{lad}} \cdot \frac{2^B-i}{2^B} R_{\text{lad}}}{\frac{i}{2^B} R_{\text{lad}} + \frac{2^B-i}{2^B} R_{\text{lad}}}, \quad \forall i = 0 \dots 2^B, \quad [\Omega] \quad (3.2)$$

with R_{lad} the total ladder resistance. To make the worst-case ladder settling ($i = 2^{B-1}$) smaller than the sampling time so as to not deteriorate the conversion speed, small resistor values are required. On top, minimizing the thermal noise of the ladder dictates small resistor values as well. However, smaller resistors increase the current required through the ladder for fixed reference voltage rails, resulting in a large power dissipation overhead. On a more practical front, realizing extremely small resistor values (in the range of $\sim 1 \Omega$) may impose a huge overhead in terms of parasitics or even be impossible depending on the technology at hand.

Another issue particularly exacerbated in flash ADCs relates to the kickback of their dynamic comparators [52]. Steep rail-to-rail transitions at the comparators' outputs parasitically couple to the input and reference taps, creating glitches, which need to sufficiently settle prior to the next comparator triggering. Because of the different impedance seen at every tap as well as the impedance difference between the reference taps and the input, this kickback must be minimized.² Finally, the difference in the input referred offset among the comparators results in non-linearity in the transfer characteristic (see Chap. 2, Sect. 2.2.2). Typical high-speed implementations size the comparator devices just enough so as to minimize the regeneration time constant for a given output load while avoiding self-loading and staying within a power budget. This comes at the expense of a large input referred offset for every comparator [28]; therefore, some form of cancellation is imperative to bring down this offset to the desired levels.

3.2.2 Flash Accuracy-Speed-Power Limits

The equations from our analysis in Chap. 2 (see Sect. 2.4) form the basis for the derivation of the fundamental accuracy-speed-power limits for all the ADC architectures under study. The jitter contribution is neglected from this analysis, assuming it imposes to a first-order equal limits to all architectures running at the same speed; therefore, it is considered an offset. For the flash converter, our derivation starts by estimating its total power consumption including contributions from the sampler, the comparator, the resistor ladder, as well as the digital logic (encoder) with the following expression:

$$P_{F,\text{tot}} = P_{F,\text{samp}} + (2^B - 1) \cdot P_{F,\text{comp}} + P_{F,\text{lad}} + P_{F,\text{dig}}. \quad [\text{W}] \quad (3.3)$$

$P_{F,\text{samp}}$ is given by Eq. (2.41) in a more generalized form

$$P_{F,\text{samp}} = V_{\text{DD}} \cdot NTC \cdot \theta_{F,\text{samp}} \cdot f_s \cdot C_S \cdot (\phi_{\text{sup}} V_{\text{DD}}), \quad [\text{W}] \quad (3.4)$$

where NTC corresponds to the number of settling time constants for a required precision ($NTC=1$ for 100% slewing), $\theta_{F,\text{samp}}$ captures the portion of the total converter period ($1/f_s$) allocated to the sampler, and ϕ_{sup} depicts the percentage of the supply utilized by the signal swing ($\phi_{\text{sup}} = V_{\text{FS}}/V_{\text{DD}}$). In the same way, $P_{F,\text{comp}}$, including the contribution of both input and latch stages, is obtained by the summation of Eqs. (2.47) and (2.55) in a more generic form

² Possible ways of achieving that are the addition of pre-amplification and/or cancellation circuitry. Minimizing the ladder resistance also helps reducing the kickback.

$$P_{F,\text{comp}} = V_{\text{DD}} \cdot \theta_{F,\text{comp}} \cdot f_s \cdot \left[\{C_I \cdot \Delta V_I\} + \{[(B - \log_2 A_I) \cdot \ln 2 + \ln(BER^{-1})] \cdot C_L \cdot V_{\text{GT}}/2\} \right]. \quad [\text{W}] \quad (3.5)$$

The first pair of the curly brackets contains the contribution of the noise critical input integrator, while the second pair depicts the bandwidth critical latch contribution. $\theta_{F,\text{comp}}$ captures the part of the total converter period occupied by the comparator, where it is assumed for simplicity that the input and latch share the same time. Allocating a portion of the total noise budget to the above expressions defines C_S , C_I , and C_L (see Chap. 2, Sect. 2.4), leading to the physical limits derived in the previous chapter for each block individually.

In reality, the utilized process puts a limit to the absolute minimum value these capacitances can take. Our analysis takes this effect into account by adopting as a minimum value the input capacitance of the practically smallest realizable gate at a given process, as will be shown in Sect. 3.2.3. Further, the capacitive loading to the sampler due to the comparator input stage imposes a considerable power overhead for the converter, making it essential to capture its contribution in our derivation. It is assumed for simplicity that this capacitance $C_{\text{in,I}}$ is on the same order as C_{GG} . Under this assumption, f_T can be expressed with the following equation:

$$f_T \approx \frac{g_{\text{m,I}}}{2\pi C_{\text{in,I}}}. \quad [\text{Hz}] \quad (3.6)$$

Equivalently, the comparator input stage integrator's bandwidth $f_{\text{bw,I}}$ can be approximated to a first-order by

$$f_{\text{I,bw}} \approx \frac{g_{\text{m,I}}}{2\pi A_I C_I}. \quad [\text{Hz}] \quad (3.7)$$

Furthermore, due to the integrator operation of the input stage (100% slewing), it is assumed that its allocated operation time equals its time constant; therefore,

$$f_{\text{I,bw}} \approx \frac{\theta_{F,\text{comp}} \cdot f_s}{2\pi}. \quad [\text{Hz}] \quad (3.8)$$

By incorporating Eq. (3.8) into (3.7) and dividing Eqs. (3.6) and (3.7), we end up with an estimation for $C_{\text{in,I}}$

$$C_{\text{in,I}} \approx A_I C_I \cdot \frac{2 \cdot \theta_{F,\text{comp}} \cdot f_s}{2\pi \cdot f_T}, \quad [\text{F}] \quad (3.9)$$

where the extra $2 \times$ in the numerator captures interconnect overhead and is included in all the subsequent parasitic capacitance derivations. The important take of the above expression is that in order to minimize $C_{\text{in,I}}$ for a fixed C_I , the ratio f_s/f_T must be minimized. For a fixed process f_T , this creates an upper bound

to the achievable converter speed for a tolerable accuracy degradation and power consumption overhead due to $C_{in,I}$.

To take our analysis one step further, we also include the parasitic loading at the output of each block, which will ultimately limit the achievable maximum speed depending on the f_T for a specific device biasing in the process at hand. Hence, C_I and C_L from Eq. (3.5) are revisited to take this loading into account

$$C_{I,tot} = C_I + C_{I,par} = C_I + \frac{g_{m,I} \cdot C_{I,par}}{g_{m,I}} \approx C_I + \frac{g_{m,I}}{\pi f_T}. \quad [F] \quad (3.10)$$

$$C_{L,tot} = C_L + C_{L,par} = C_L + \frac{g_{m,L} \cdot C_{L,par}}{g_{m,L}} \approx C_L + \frac{g_{m,L}}{\pi f_T}. \quad [F] \quad (3.11)$$

In the above expressions, it is clearly visible that in order to minimize the unwanted second term, g_m must be minimized for a fixed f_T or f_T must be maximized for a fixed g_m . Hence, the benefit of moving to process nodes with an increasing energy efficiency is indicated, for being able to get closer to the physical limits. This will become more evident in Sect. 3.2.3, where the impact of scaling will be discussed.

Now that we have considered all the important first-order effects to be captured in our derivation of the architectural limits, Eqs. (3.4) and (3.5) are revisited, and the contribution of these effects is included. Taking into account Eq. 3.9 as well as the process-limited minimum capacitance value, and incorporating them into Eq. (3.4), $P_{F,samp}$ takes its complete form

$$P_{F,samp} = V_{DD} \cdot NTC \cdot \theta_{F,samp} \cdot f_s \cdot \left[\max\{C_S, C_{min}\} + (2^B - 1) \cdot \max\left\{A_I C_I \cdot \frac{\theta_{F,comp} \cdot f_s}{\pi \cdot f_T}, C_{min}\right\} \right] \cdot (\phi_{sup} V_{DD}). \quad [W] \quad (3.12)$$

Furthermore, substituting the capacitor values from Eqs. (3.10) and (3.11) into Eq. (3.5), after some re-arranging and taking into account C_{min} and the basic MOS equation (2.45), we end up with the final form for $P_{F,comp}$ including all the aforementioned parameters

$$P_{F,comp} = V_{DD} \cdot \theta_{F,comp} \cdot f_s \cdot \left[\left\{ \frac{\max\{C_I, C_{min}\} \cdot \Delta V_I}{1 - \frac{2\Delta V_I}{\pi V_{GT}} \cdot \frac{\theta_{F,comp} \cdot f_s}{f_T}} \right\} + \left\{ \frac{[(B - \log_2 A_I) \cdot \ln 2 + \ln(BER^{-1})] \cdot \max\{C_L, C_{min}\} \cdot V_{GT}/2}{1 - \frac{[(B - \log_2 A_I) \cdot \ln 2 + \ln(BER^{-1})] \cdot \theta_{F,comp} \cdot f_s}{\pi \cdot f_T}} \right\} \right]. \quad [W] \quad (3.13)$$

To derive $P_{F,lad}$, we consider the static power dissipated by the resistor ladder for a resistance value corresponding to a certain portion of the total noise budget

Table 3.1 Core supply voltage and simulated minimum capacitance in four deep-scaled CMOS processes

Process	65 nm CMOS	40 nm CMOS	28 nm CMOS	16 nm CMOS
V_{DD}	1.2 V	0.9 V	0.9 V	0.8 V
C_{min}	960 aF	630 aF	570 aF	880 aF ^a

^a Only discrete transistor sizes are allowed with a minimum of two fins

$$\begin{aligned}
 P_{F,lad} &= V_{DD} \cdot \frac{(\phi_{sup} V_{DD})}{R_{lad}} \\
 &= V_{DD} \cdot \frac{4kT \cdot \theta_{F,samp} \cdot f_s}{V_{R_{lad}}^2} \cdot (\phi_{sup} V_{DD}), \quad [W] \quad (3.14)
 \end{aligned}$$

where it is assumed that the integration noise bandwidth is equal to the portion of f_s allocated to the sampler, since the ladder operation overlaps with sampling.

Finally, the digital power is computed to a first-order by estimating the number of gates necessary for the encoding operation. Assuming a Wallace encoder with minimum-sized gates as in [26] running at f_s , the following power is derived:

$$P_{F,dig} = V_{DD} \cdot f_s \cdot 5 \cdot (2^B - B) \cdot C_{min} \cdot V_{DD}. \quad [W] \quad (3.15)$$

3.2.3 Impact of Scaling

It was pointed out in the discussion of Sect. 3.2.2 that the utilized process sets a lower limit to the minimum capacitance, therefore to the minimum power required for charging/discharging that capacitance. The derived limits of Eqs. (3.12)–(3.15) capture to a great extent important process parameters, such as V_{DD} , C_{min} , f_T , and V_{GT} , shedding unique insight into how technology scaling may impact these limits. To get an idea in terms of absolute values, Table 3.1 reports the core supply voltage and simulated minimum capacitance in four deep-scaled CMOS processes, which will be henceforth used for the technological comparison of our derived limits. The minimum capacitance in each process is extracted as the gate capacitance of a 2×-minimum-sized³ inverter with a Fan Out (FO)-1 loading. The P-channel MOS (PMOS)/N-channel MOS (NMOS) ratio is set for the same drivability, and the gate voltage is set at mid-supply to cover usage of this cell as both a digital gate and an analog amplifier.

From a design perspective, a key identification for the capability of a process is how much speed (f_T) can be efficiently achieved by a transistor at a certain operating point (g_m/I_D through V_{GT}). To build a better understanding of the impact and importance of these parameters, Fig. 3.3 portrays a simplified small-signal model of an NMOS transistor. When increasing the transistor V_{GS} with $V_{DS} > 0$,

³ This is chosen to include somewhat the effect of layout interconnect parasitics.




Fig. 3.3 Simplified small-signal model of an NMOS transistor (bulk is omitted for simplicity)

three operating regions are distinguished [42]: 1) Weak-Inversion (W-I), 2) Strong-Inversion (S-I), and 3) Velocity-Saturation (V-S). g_m/I_D represents the relative biasing point and is summarized below for each region

$$\frac{g_m}{I_D} = \frac{2}{V_{GT}} = \begin{cases} \frac{K' \frac{W}{L} 4n U_T e^{V_{GS}/nU_T}}{K' \frac{W}{L} (2n U_T)^2 e^{V_{GS}/nU_T}}, & \text{W-I} \\ \frac{K' \frac{W}{L} 2(V_{GS} - V_{TH})}{K' \frac{W}{L} (V_{GS} - V_{TH})^2}, & \text{S-I} \\ \frac{W C_{ox} v_{sat}}{W C_{ox} v_{sat} (V_{GS} - V_{TH})}, & \text{V-S} \end{cases}, \quad [\text{S/A}] \quad (3.16)$$

where $K' = (\mu C_{ox})/(2n)$, $U_T = kT/q$, and $n = C_D/C_{ox} + 1 \approx 1.2 \dots 1.5$ with C_D and C_{ox} the depletion and oxide capacitances, respectively. If in the model of Fig. 3.3 we consider C_{GD} as open and approximate $C_{GS} \approx (2/3)WLC_{ox}$, f_T in the different operating regions may be expressed as follows:

$$f_T \approx \frac{g_m}{2\pi C_{GS}} = \begin{cases} \frac{3\mu}{2\pi L^2} e^{V_{GS}/nU_T}, & \text{W-I} \\ \frac{3\mu}{2\pi L^2} (V_{GS} - V_{TH}), & \text{S-I} \\ \frac{v_{sat}}{2\pi L}, & \text{V-S} \end{cases}. \quad [\text{Hz}] \quad (3.17)$$

The above equations indicate the theoretical speed improvement when scaling down the channel length L . In the weak-inversion and strong-inversion regions, this improvement is proportional to $1/L^2$, while in velocity-saturation, it only improves with $1/L$. Furthermore, the transistor's efficiency gradually degrades when going from weak-inversion to velocity-saturation since I_D increases at a higher rate compared to g_m , hence gradually reducing the achievable g_m/I_D [10]. The aforementioned arguments suggest that the velocity-saturation region should be avoided unless absolutely necessary for maximum speed.

The beauty of adopting g_m/I_D or V_{GT} as a single design parameter lies in the fact that it allows a fast and relatively accurate comparison of different process nodes, without requiring the absolute values of technology specifics (e.g., V_{TH} , μ , C_D , C_{ox} , t_{ox}). To provide more insight and serve as a guideline for the optimum process choice and device biasing, some important small-signal parameters are simulated for the four deep-scaled CMOS processes of Table 3.1, including both 2D-planar and 3D-FinFET technologies, and the results are plotted in Fig. 3.4. The simulations use




Fig. 3.4 f_T vs. g_m/I_D and $f_T \cdot g_m/I_D$ vs. g_m/I_D in four CMOS processes

representative low- V_{TH} RF devices provided by the available Process Design Kit (PDK)s. Interconnects up to Metal 3 are custom-added to these devices to partially capture the BEOL contribution. Furthermore, the dimensions are chosen such as to optimize the device f_T , while roughly the same W/L with the same number of fingers and minimum L are adopted in all the processes under comparison.

Plotting f_T vs. g_m/I_D (Fig. 3.4, left), the f_T benefit due to scaling from 65 nm to 40 nm and below is very evident across all operating regions. Comparing between 40 nm, 28 nm, and 16 nm, there is almost no more benefit when considering the peak- f_T . This might be partially attributed to the increased BEOL contribution, which becomes more dominant with scaling.⁴ However, the peak- f_T occurs at a slightly higher g_m/I_D as L scales down, indicating an increase in efficiency. More importantly, the f_T improvement due to scaling is still visible for $g_m/I_D \sim 10$ – 20 corresponding to Moderate-Inversion (M-I) (transition between W-I and S-I) and S-I regions. In 16 nm especially, the f_T benefit increases further when going deeper into W-I, attributed to the better control of the channel charge by the gate due to the FinFET structure [53].

$f_T \cdot g_m/I_D$ vs. g_m/I_D (Fig. 3.4, right) compares to a first-order the achievable energy efficiency for each process. A first observation is that all the optima occur for g_m/I_D values higher than the ones, where peak- f_T is reached. When considering only the planar processes, going to lower nodes results in a steeper optimum slowly shifting toward higher g_m/I_D . This is interpreted as a gradually disappearing S-I region ($f_T \cdot g_m/I_D \approx \text{const.}$), which matches the theoretical prediction in [10]. In the transition from planar to FinFET, the optimum keeps shifting toward W-I but with an improved flatness, bringing back part of the previously vanished S-I region. This

⁴ The lower metals in every process reduce in thickness and distance from the substrate, increasing both resistance and capacitance. The lower via resistance increases as well.




Fig. 3.5 Flash accuracy-speed-power limits: (a) for different f_s in 28 nm and (b) at $f_s = 4$ GHz in the processes under comparison

is equivalent to partially reviving the MOS “square-law” behavior, re-enhancing the applicability of Eqs. (3.16) and (3.17).

We now possess all the tools to predict an ADC’s accuracy-speed-power limits. Figure 3.5a depicts the flash limits derived in Eqs. (3.3) and (3.12)–(3.15), while Fig. 3.5b compares these limits at $f_s = 4$ GHz in 65 nm, 40 nm, 28 nm, and 16 nm for roughly the same f_T . The plots are taken for $\theta_{F,samp} = 2$, $\theta_{F,comp} = 2$, $\phi_{sup} = 1$, and $\Delta V_I = V_{DD}/2$, while V_{GT} is set through g_m/I_D based on Eq. (3.16) and Fig. 3.4. C_S , C_I , C_L , and R_{lad} are set by assigning $0.5 \cdot \bar{\epsilon}_q^2$ to the sampler, $0.35 \cdot \bar{\epsilon}_q^2$ to the

comparator, and $0.15 \cdot \bar{\epsilon}_q^2$ to the ladder. For low resolution, the process C_{\min} sets the baseline, while for higher resolution and relatively small f_s/f_T , noise determines these limits. When f_s/f_T increases, the parasitic capacitance from Eqs. (3.10) and (3.11) rises together with g_m , increasing the slope above the noise limit and eventually dictating the required power to achieve a certain resolution. When that power becomes infinite, certain resolutions cannot be achieved. The derived limits are very similar across the different processes in the noise-limited regime. This is explained by the fact that the scaling drawback of the reduced V_{DD} is to a first-order nullified by the increased g_m/I_D to achieve the same f_T .

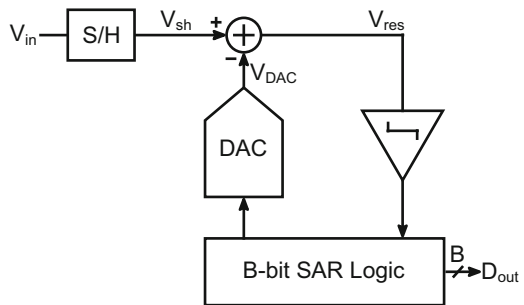
3.3 The SAR Architecture

3.3.1 Overview

The SAR ADC [54–57] is an algorithmic converter, which relies on a binary search Successive Approximation (SA) of the sampled input. Shown in Fig. 3.6, the basic components of a B -bit SAR ADC are the Sample-and-Hold (S/H) or Track-and-Hold (T/H) to sample the analog input, a DAC to provide binary-weighted references, and a comparator to evaluate the sign of the subtraction between the sampled value V_{sh} and the DAC output V_{DAC} . The SAR logic, apart from collecting the bits and interfacing with the outside world, is responsible for storing the comparator decision, based on which it dictates the next DAC reference to be generated such as to minimize the voltage difference at the summing node V_{res} .

During the conversion, the SAR logic divides the FS range by two in every cycle by assigning the binary weights to appropriately group the DAC elements. After the DAC has sufficiently settled to a new output, the comparator is strobed to evaluate the polarity between V_{sh} and V_{DAC} and give the first bit. If this polarity is positive, the group of elements with the next assigned weight is added on the existing elements to increase V_{DAC} , and the comparator is strobed again for a second evaluation. If the polarity is negative, the initial group of elements is removed, and only the new one remains. After the evaluation of all the bits with a sequence from

Fig. 3.6 Block diagram of a B -bit SAR ADC



the Most Significant Bit (MSB) to the least significant bit (LSB), the B -bit output is collected. Therefore, the SAR ADC resolves an input to B bits of accuracy in B clock cycles, but with minimum hardware. In this regard, SAR and flash can be seen as the “*yin and yang*” of data converters; the first is hardware efficient but time-consuming, whereas the second is hardware consuming but time efficient.

The operation of the binary SA algorithm can be illustratively described by means of a weighting scale, as shown in Fig. 3.7a. In the chosen example, we have four weighting trials to measure three gold bars of 10.7 kg with an accuracy within 1 kg. Having a FS from 0 to 15 kg, the available binary weights are [8 4 2 1] kg. After trying the SA algorithm four times by keeping or removing the appropriate weights to bring the scale to a balance, we end up measuring 11 kg, which approximates the original gold weight within the required accuracy. Figure 3.7b shows the voltage-domain equivalent waveforms of V_{sh} (weight of gold) and V_{DAC} (weights of the scale) and the residue voltage V_{res} revealing the sign of the evaluation at the comparator input and the resulting output bits.

It is worth mentioning that the SA algorithm does not necessarily have to use binary weights. The DAC elements can be grouped with different weights such that the algorithm still approximates the input within the required accuracy. In our weighting scale example, the weights [8 3 2 1 1] kg could have been used with the end result being equally accurate but requiring an extra weighting trial. The process of utilizing a non-binary algorithm with extra cycles in a SAR ADC is known as redundancy or OverRange (OR). Redundancy is a powerful tool, which allows the correction in the extra cycles of evaluation errors that might have occurred in the earlier cycles. A comprehensive analysis on the concept of redundancy and its trade-offs can be found in [58].

3.3.2 The DAC in a SAR

The most important block in a SAR converter is the DAC, effectively determining the accuracy, speed, power, and area of this architecture. Throughout the years, different DAC implementations have been adopted, including capacitive or charge redistribution [59, 60], resistive [61], hybrid capacitive-resistive [62], and current steering [63]. Due to the continuous accuracy improvement in the lithography when scaling to finer CMOS processes, the recently prevailing type is the Capacitive DAC (CDAC), showing superior linearity compared to its resistive or current steering counterparts. In a typical B -bit implementation, the total DAC capacitance is $2^B C_u$, grouped in a binary fashion ($2^{B-1} C_u, 2^{B-2} C_u, \dots, 2^0 C_u$). For fast settling reasons (see Chap. 5), it is highly beneficial to use a small unit capacitor C_u , whose minimum size is ultimately bound by matching and/or noise considerations. A major advantage of the capacitive topology is that it allows merging the S/H and DAC functions, hence reducing area and power. When the conversion finishes, the charge is re-distributed such as to make the comparator input a virtual ground within the given accuracy.




Fig. 3.7 (a) Scale equivalent of a binary SA algorithm and (b) waveform operation in the voltage vs. time domain

The dynamic operation of CDACs renders them highly efficient by nature. Nevertheless, various switching schemes have been proposed to further improve their efficiency. Although different schemes exist, the discussion below is limited to the operation and switching energy of a few noteworthy representatives, followed by their comparison in terms of energy efficiency. The discussion starts with an overview of the conventional switching scheme and its limitations. The interested reader is redirected to [64], for an analytical energy derivation in a switched-capacitor array.




Fig. 3.8 3-bit example of the conventional CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints

Conventional Switching Scheme

As mentioned above, the CDAC offers the possibility of merging the sampling and re-distribution or residue generation functions. Figure 3.8 illustrates a 3-bit example of the conventional CDAC switching scheme, annotating the energy drawn from the reference source V_{REF} for all possible switching transitions.⁵ During the sampling phase, the differential input is sampled on the CDAC array. After sampling, the MSB capacitor is charged to V_{REF} , while the rest of the capacitors are discharged to ground for the first bit cycle. It is easily noticed that this first transition dominates the energy consumption. Moreover, after the first transition, the “ascending” and “descending” transitions demonstrate a significant unbalance in terms of switching energy. This is due to the fact that depending on the particular transition, a different number of capacitors is charged/discharged between V_{REF} and ground. For the capacitors that are first charged and then discharged to ground, their charge is lost

⁵ Although the illustration describes a bottom-plate sampling operation, similar principles apply for top-plate sampling as well. Unless stated otherwise, throughout this manuscript, top plate refers to the CDAC side at the comparator input, while bottom plate refers to the side connected to the reference switches. Therefore, top-plate sampling refers to connecting the input during sampling to the comparator side, while bottom-plate sampling connects it to the references’ side with a fixed voltage on the top side.

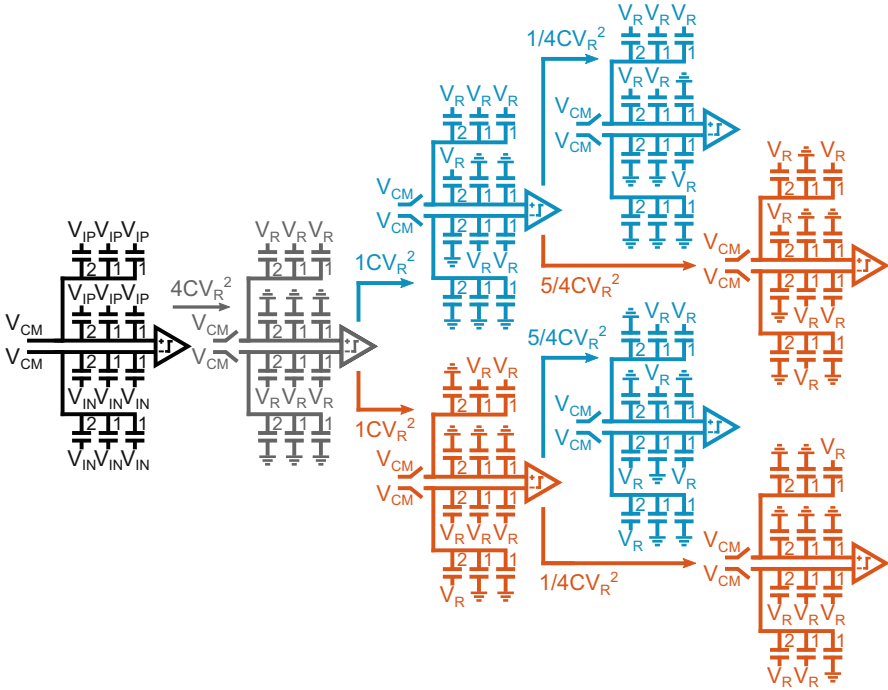


Fig. 3.9 3-bit example of the split-capacitor CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints

rather than recycled, which results in wasting energy. The average switching energy for a B -bit conventional CDAC switching scheme can be derived as

$$E_{conv} = \sum_{j=1}^B 2^{B+1-2j} (2^j - 1) \cdot CV_{REF}^2. \quad [J] \quad (3.18)$$

Split-Capacitor Switching Scheme

Identifying the aforementioned switching energy unbalance, in [64], a switching scheme was proposed to solve this problem. As it is depicted in Fig. 3.9, the MSB capacitor is split into an array identical to the remaining capacitors’ array. During sampling, the operation is exactly the same as the conventional scheme. After sampling, the MSB capacitor is charged to V_{REF} , while the rest of the capacitors are discharged to ground for the first bit cycle. If there is an “ascending” transition following, the MSB - 1 capacitor of the original array is charged to V_{REF} ; otherwise, the MSB - 1 capacitor of the split array is discharged to ground. This avoids having to discharge a previously charged capacitor and wasting the charge, resulting in a

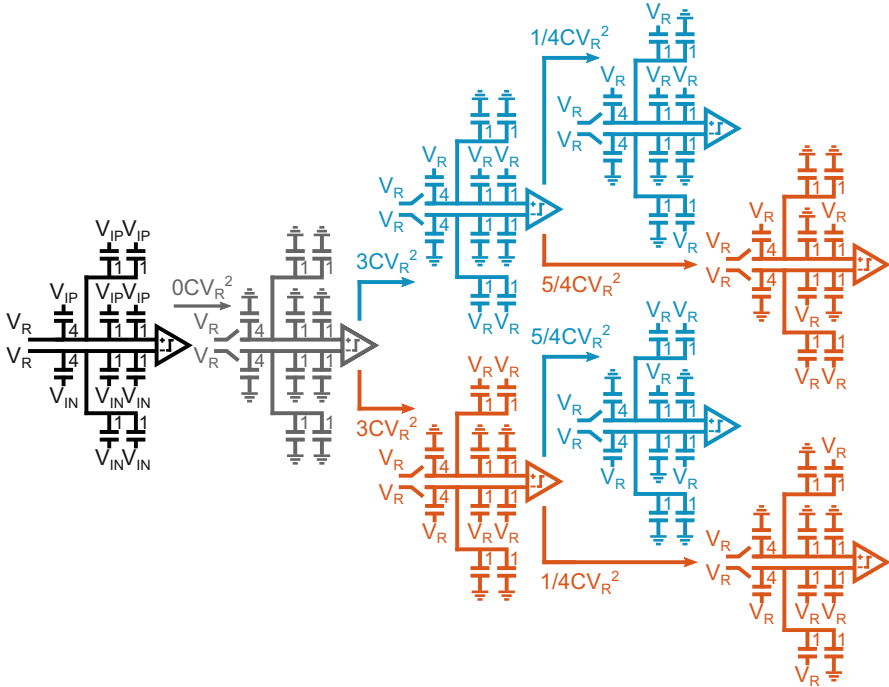


Fig. 3.10 3-bit example of the energy-saving CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints

better balance between “ascending” and “descending” transitions. The drawback is the increased number of switches as well as more complex logic to control the switching scheme. Nevertheless, there is a significant energy benefit, and the derived average switching energy for a B -bit split-capacitor CDAC switching scheme now becomes

$$E_{\text{split-cap}} = 2^{B-1} + \sum_{j=2}^B 2^{B+1-2j} (2^{j-1} - 1) \cdot CV_{REF}^2 \quad [\text{J}] \quad (3.19)$$

Energy-Saving Switching Scheme

To further reduce the switching energy, in [65], the split-capacitor technique was combined with a modification in the initial voltage that connects to the CDAC top plate, as shown in Fig. 3.10. During sampling, the differential input is sampled on the CDAC array identically to the conventional scheme with the only difference that V_{REF} instead of V_{CM} connects to the top plate. After sampling, all the capacitors are discharged to ground, resulting in zero energy consumption for the first bit cycle. The difference compared to the split-capacitor technique is that the MSB - 1 instead




Fig. 3.11 3-bit example of the monotonic CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints

of the MSB capacitor is split into an array, reducing the switching energy in the “descending” transitions. One drawback of this energy-saving scheme is that the top plate starts from V_{REF} but during conversion gradually approaches V_{CM} , which makes it susceptible to parasitic capacitance from the capacitors and the comparator. In a similar fashion as above, the obtained average switching energy for a B -bit energy-saving CDAC switching scheme can be calculated as follows:

$$E_{\text{energy-saving}} = 3 \cdot 2^{B-3} + \sum_{j=3}^B 2^{B+1-2j} (2^{j-1} - 1) \cdot C V_{REF}^2 \quad [\text{J}] \quad (3.20)$$

Monotonic Switching Scheme

To provide further energy savings, a monotonic switching scheme was introduced in [66]. Illustrated in Fig. 3.11, this scheme incorporates only discharges of capacitors, reducing the required switching energy in the transitions without requiring splitting or extra switches. At the same time, for a B -bit CDAC requires only 2^{B-1} units, rendering the MSB capacitor unnecessary. This is due to the fact that the first bit is evaluated by directly comparing the differential input without requiring any switching. This scheme is also susceptible to any parasitic capacitance, since the

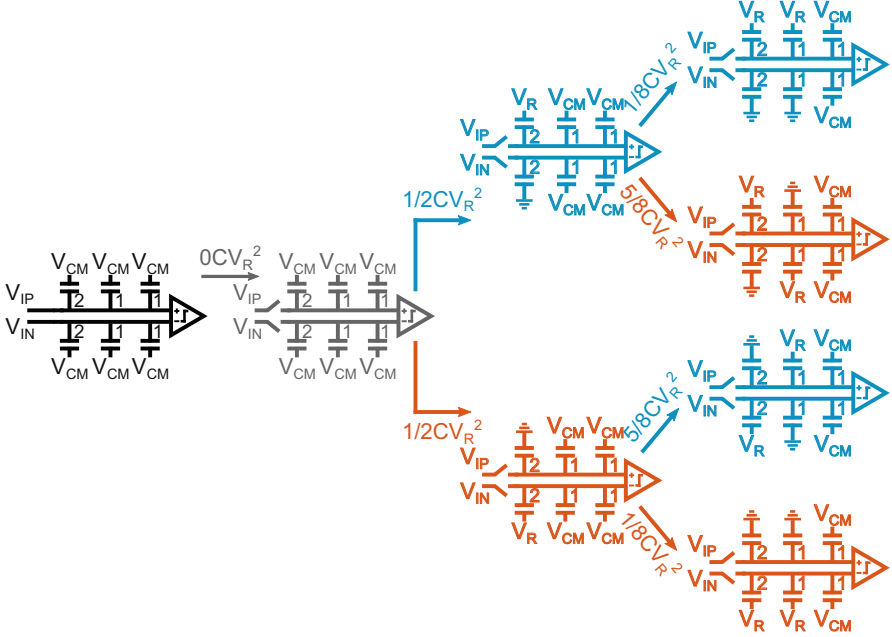


Fig. 3.12 3-bit example of the MCS CDAC switching scheme. V_{REF} is annotated as V_R to preserve clarity due to space constraints

top plate starts and ends the conversion with different voltages. Moreover, since top plate common mode is heavily changing, it introduces an extra burden on the comparator design. In [67], a solution to this problem was proposed by doubling the number of capacitors and switches and clocking each two capacitors with the same weight in a complementary fashion; hence, it is able to preserve a constant common mode. The average switching energy for a B -bit monotonic CDAC switching scheme can be obtained

$$E_{\text{monotonic}} = \sum_{j=1}^{B-1} (2^{B-2-j}) \cdot C V_{REF}^2. \quad [\text{J}] \quad (3.21)$$

Merged-Capacitor Switching Scheme

To enhance the energy savings even more, a Merged-Capacitor Switching (MCS) scheme was introduced in [68], depicted in Fig. 3.12 for top-plate sampling. During the sampling phase, the differential input is sampled at the top plate, while V_{CM} connected on the other plate of all the capacitors. After sampling, the first bit is evaluated immediately without any switching, identically to the monotonic switching. As a result, this scheme enjoys the same benefit of discarding the MSB




Fig. 3.13 Switching energy for the different CDAC switching schemes

capacitor, necessitating only 2^{B-1} units for a B -bit CDAC. In the second bit cycle, depending on the first bit evaluation, the largest capacitor is either charged to V_{REF} from V_{CM} or discharged from V_{CM} to ground. In contrast to the monotonic switching, where the switching energy is proportional to CV_{REF}^2 , in this scheme, it is proportional to $2 \cdot C(V_{REF}/2)^2$, leading to even higher energy savings. Additionally, this scheme preserves a constant common mode at the top plate, relaxing the comparator design in terms of common-mode rejection. When used in top-plate sampling, this scheme also shows a sensitivity to parasitic capacitance. This can be easily overcome by adopting bottom-plate sampling while seamlessly preserving all the energy benefits and not increasing complexity. The calculated average switching energy for a B -bit MCS CDAC switching scheme is given as

$$E_{MCS} = \sum_{j=1}^{B-1} 2^{B-3-2j} (2^j - 1) \cdot CV_{REF}^2. \quad [J] \quad (3.22)$$

For a better understanding, Fig. 3.13 plots the average switching energy for each of the aforementioned switching schemes while sweeping the number of bits. It can be seen that the MCS scheme is able to achieve the lowest switching energy. Due to its pronounced benefits and low complexity, the MCS scheme will be utilized in our prototype ADCs, both with top-plate sampling (see Chap. 5) and with bottom-plate sampling (see Chap. 6).

Despite the remarkable energy efficiency of the SAR ADC in the low-to-medium-resolution regime and sample rates of several MS/s (Fig. 3.1), the bit-at-a-time nature remains the main bottleneck in significantly extending the sample rate of the single channel while preserving the efficiency levels. Several speed-boosting techniques have been invented to tackle this bottleneck. These techniques will be

reviewed in Chap. 5 and assessed together with the introduced techniques in our prototype implementation. Furthermore, when increasing the resolution to levels larger than ~ 10 bits, driving the noise-limited input capacitance at sample rates approaching GS/s with low enough noise and distortion and high energy efficiency becomes extremely challenging.

3.3.3 SAR Accuracy-Speed-Power Limits

Similar to the flash ADC, our derivation of the SAR limits starts by estimating the total power consumption of a binary-weighted SAR including contributions from the sampler, the comparator, the CDAC, and the digital SA logic as follows:

$$P_{SA,tot} = P_{SA,samp} + B \cdot P_{SA,comp} + P_{SA,DAC} + P_{SA,dig} \cdot [W] \quad (3.23)$$

When incorporating the capacitive loading due to the comparator input and the process-limited C_{min} into $P_{SA,samp}$, its final form becomes

$$P_{SA,samp} = V_{DD} \cdot NTC \cdot \theta_{SA,samp} \cdot f_s \cdot \left[\max\{C_S, C_{min}\} + (\max\{A_I C_I \cdot \frac{\theta_{SA,comp} \cdot f_s}{\pi \cdot f_T}, C_{min}\}) \right] \cdot (\phi_{sup} V_{DD}), \quad [W] \quad (3.24)$$

where $\theta_{SA,samp} = B + 1$ and $\theta_{SA,comp} = (B + 1) \cdot 1/\beta$, $0.9 \geq \beta \geq 0.5$ capture the portions of the total converter period allocated to the sampler and the comparator, respectively. A synchronous B -bit converter is assumed with one out of $B + 1$ cycles allocated to the sampler, while the comparator and CDAC share each of the other cycles through smart timing (see Chap. 5). Similarly, $P_{SA,comp}$, including the contribution of both input and latch as well as all the parasitic and process effects introduced in the previous section, is written as

$$P_{SA,comp} = V_{DD} \cdot \theta_{SA,comp} \cdot f_s \cdot \left[\left\{ \frac{\max\{C_I, C_{min}\} \cdot \Delta V_I}{1 - \frac{2\Delta V_I}{\pi V_{GT}} \cdot \frac{\theta_{SA,comp} \cdot f_s}{f_T}} \right\} + \left\{ \frac{[(B - \log_2 A_I) \cdot \ln 2 + \ln(BER^{-1})] \cdot \max\{C_L, C_{min}\} \cdot V_{GT}/2}{1 - \frac{[(B - \log_2 A_I) \cdot \ln 2 + \ln(BER^{-1})] \cdot \theta_{SA,comp} \cdot f_s}{\pi}} \right\} \right] \cdot [W] \quad (3.25)$$

A distinct difference between the above two expressions and the equivalent ones for the flash regards the portion of the allocated period to the sampler and the comparator with the resolution; for the flash, it remains constant, while for the SAR, it increases with resolution posing stricter timing and power requirements.

For the power consumption of the CDAC, we consider Eq. (3.22) and the MCS scheme for a reduced power consumption as shown in Fig. 3.13. To be more precise, V_{REF} for the CDAC is provided by a regulating reference circuit drawing current

from V_{DD} . To capture somewhat the contribution of this circuit, V_{REF}^2 is swapped with $V_{DD} \cdot V_{REF}$, and an efficiency factor λ_{ref} is included. Therefore, $P_{SA,DAC}$ takes its final form taking into account also the process C_{min}

$$P_{SA,DAC} = \frac{V_{DD}}{\lambda_{ref}} \cdot \theta_{SA,DAC} \cdot f_s \cdot \left[\sum_{j=1}^{B-1} 2^{B-3-2j} (2^j - 1) \right] [W] \quad (3.26)$$

$$\times \max \left\{ \frac{C_S}{2^B}, C_{min} \right\} \cdot V_{REF}.$$


In the above expression, the term in the square brackets represents the average switching activity over one full converter period. It was also mentioned above that the CDAC shares each bit cycle with the comparator. Therefore, the timing portion of the CDAC over one full period is $\theta_{SA,DAC} = [(B+1)/B] \cdot 1/(1-\beta)$. Further, in Eq. (3.26), the corresponding unit capacitors could theoretically take values smaller than C_{min} , realized by a custom placement of two parallel metal layers available at any process, as shown in [69, 70]. Nevertheless, in our derivation, we preserve C_{min} as the smallest achievable value. Finally, to account for a reasonable efficiency for the regulating reference circuit, we will assume hereafter $\eta_{ref} = 60\%$, which is an achievable value with a class-AB circuit.

To conclude our derivation of the SAR limits, for the estimation of the SA logic power, we will assume for each bit three minimum-sized gates for the control of the MCS CDAC (one per switch), one gate for clocking each bit cycle, and one gate for storing the bits. This is the minimum number of gates and might be somewhat optimistic, but provides a good baseline considering that the logic is typically not dominating the power consumption. $P_{SA,dig}$ is then expressed as


$$P_{SA,dig} = V_{DD} \cdot f_s \cdot 5 \cdot B \cdot C_{min} \cdot V_{DD}. \quad [W] \quad (3.27)$$

The accuracy-speed-power limits of a SAR, derived in Eqs. (3.23)–(3.27) with the assumptions established above and making use of Eq. (3.16) and Fig. 3.4, can be numerically computed. These limits are portrayed in Fig. 3.14a for three different f_s values in 28 nm⁶ and compared in 65 nm, 40 nm, 28 nm, and 16 nm for roughly the same f_T in Fig. 3.14b. The plots are taken for allocating in each cycle 70% of the time to the comparator and 30% to the CDAC ($\beta = 0.7$). The rest of the fixed parameters common to both flash and SAR (NTC , ϕ_{sup} , A_1 , ΔV_1 , V_{GT}) retain the same values. Furthermore, the noise partitioning assigns a common $0.5 \cdot \bar{\epsilon}_q^2$ between the sampler and the CDAC, and the rest $0.5 \cdot \bar{\epsilon}_q^2$ is allocated to the comparator. The curves in Fig. 3.14a follow the same trend as for the flash but with better efficiency and flatter slope due the linear increase in power with resolution ($\propto B$) compared to the exponential of the flash ($\propto 2^B$). In Sect. 3.6, all architectures will be compared against each other in terms of their accuracy-speed-power limits.

⁶ The 28 nm CMOS process is referred frequently throughout this book. The reason is that three out of the four prototypes in Chaps. 4–7 are implemented in this node.



(a)



(b)

Fig. 3.14 SAR accuracy-speed-power limits: (a) for different f_s in 28 nm and (b) for $f_s = 500$ MHz in the processes under comparison

When comparing the different processes in Fig. 3.14b, we again see the derived limits being very similar in the noise-limited regime. In fact, between 65 nm and 40 nm, there is no power benefit at all due to the significant supply drop, but there is a slight improvement going to 28 nm and 16 nm. This is attributed to the mere existence of any supply-limited contributions,⁷ while on the other hand, the digital contribution only benefits from lowering the supply. It is worth mentioning that the

⁷ The comparator latch (second term in Eq. 3.25) is the only such contribution.

same trend is seen in the flash limits (Fig. 3.5b). In the process-limited regime, the C_{min} increase in 16 nm is compensated by the lower V_{DD} , therefore places it roughly together with the 40 nm and the 28 nm.

3.4 The Pipeline Architecture

3.4.1 Overview

One of the most popular architectures that has been (and still being) used in industry to simultaneously realize high resolution and high speed is the pipeline ADC [71–74]. The block diagram of a B -bit pipeline ADC is shown in Fig. 3.15. It incorporates a cascade of m -stages each one resolving B_s bits, where $s = 1, \dots, m, < B$. Consecutive stages operate in opposite phases of the sampling clock, and the entire pipeline at the end of the conversion outputs $B_1 + B_2 + \dots + B_m$ bits. The digital logic combines the stage bits in a weighted sum fashion and generates a new output at the sample rate, but with a latency of m clock periods. Thanks to the pipeline operation, the throughput of this topology is bound by the speed of one stage, enabling very-high-speed conversion. Unlike flash ADCs, the number of components increases linearly with resolution rather than exponentially, allowing a higher resolution while retaining the speed.

In a typical implementation, each stage contains a S/H, a B_s -bit flash sub-ADC, a B_s -bit sub-DAC, and a Residue Amplifier (RA) block with a gain A_s to amplify




Fig. 3.15 Block diagram of a B -bit m -stages pipeline ADC

the difference between the sampled value and the output of the sub-DAC.⁸ The sub-DAC, the subtraction, and the residue amplification are combined into a single block known as the Multiplying DAC (MDAC). The amplified difference by the MDAC serves as the input to the next stage, which has the same FS range with the current stage given that its gain matches its resolution as follows:

$$A_s = 2^{B_s} \iff B_s = \log_2(A_s). \quad (3.28)$$

With a stage gain of exactly 2^{B_s} , there is no margin to absorb any sub-ADC error (e.g., comparator offset) or amplifier offset/gain error in that stage, which can shift the amplified residue $A_s V_{\text{res}}$ out of its ideal range. The erroneous residue will be processed by the back-end down the pipe, therefore deteriorating the overall converter transfer characteristic. Figure 3.16a shows the residue plot of stage- s for an ideal case without errors, where the residue occupies the entire range without any

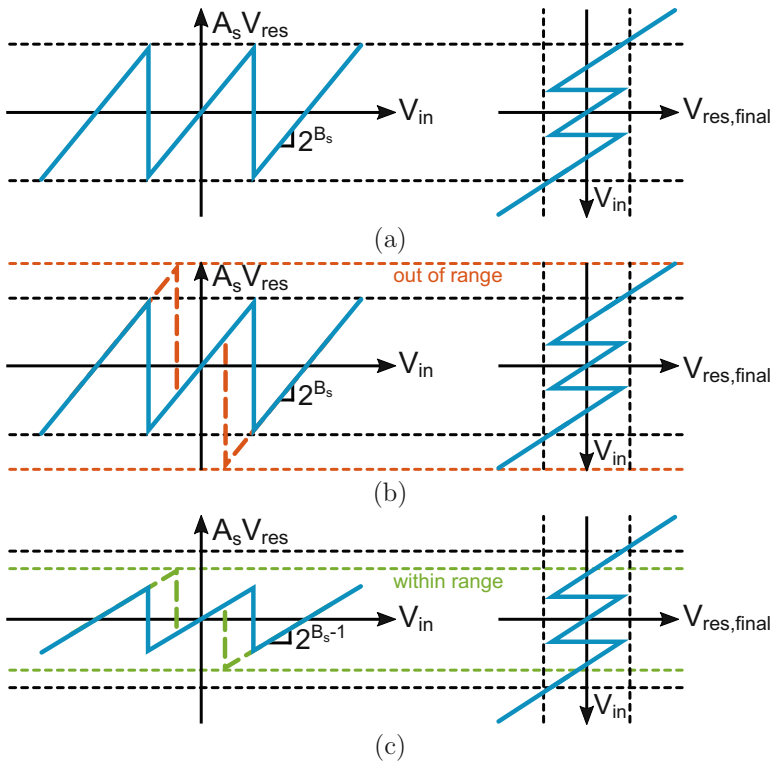


Fig. 3.16 Residue plot of stage- s : (a) ideal case with $A_s = 2^{B_s}$, (b) $A_s = 2^{B_s}$ with error and no OR, and (c) $A_s = 2^{B_s-1}$ with error and $2\times$ -OR

⁸ This is effectively the quantization error of the stage's sub-ADC.

margin available. Figure 3.16b shows a case with errors shifting the residue out of its ideal range without the possibility of recovering. The concept of redundancy or OR is often incorporated in the form of designing a gain less than 2^{B_s} and allocating the rest of the range for the absorption of errors. This is depicted in Fig. 3.16c, where a 2^{B_s-1} half gain is adopted, making the other half of the range available for error absorption. Any gain value can be chosen arbitrarily, but using a power of 2 minimizes the complexity of the digital logic, which implements that gain digitally to combine the bits accordingly (Fig. 3.15).


Two critical pipeline design considerations are the stage scaling and the stage resolution. From a noise perspective, each stage contributes a noise proportional to kT/C , and the noise from any later stage is referred to the input divided by the total gain squared up to that stage. Therefore, the front-end stages impose the most stringent noise and power requirements. One extreme would be to size all the stages equally, leading to a poor efficiency due to the excess power consumed by the back-end. Another extreme would be to size the stages for an equal noise contribution, resulting in a poor total noise and requiring additional power to bring that noise down to tolerable levels. A shallow optimum between these extremes exists in using a scaling factor of roughly the gain stage [75]. Regarding the stage resolution, from a pure speed standpoint, the minimum number of bits is desirable [76], allowing for a faster MDAC settling and a lower flash sub-ADC capacitance (see Sect. 3.2). This comes at the expense of more pipeline stages for a certain aggregate resolution. On the other hand, allocating more bits per stage [77] reduces the stage count and relaxes the precision requirements on each stage residue [26]. However, the increased loading from the flash sub-ADCs can induce a power burden when pushing the speed.

Although the pipeline is more hardware efficient than the flash and faster than the SAR, its main limitation is the requirement for accurate amplification with stringent bandwidth requirements, increasing the power consumption significantly. More and more designs have adopted open-loop amplifiers [78] or integrators [79], which result in power savings, provided that the necessary calibration to correct their errors does not introduce a significant overhead. Amplifier sharing [80] has been also adopted to reduce their number in exchange for faster remaining amplifiers. Yet, the need for good analog components might explain the inferior scalability in speed of the pipeline compared to the SAR (Fig. 3.1), while its latency can be an issue in feedback systems.

3.4.2 Pipeline Accuracy-Speed-Power Limits

To start the derivation of the pipeline ADC accuracy-speed-power limits, we consider a $(B_s - 1)$ -bit/stage effective resolution pipeline, where each stage is implemented with B_s bits to allow a 2 \times -OR between stages, and we study the effective resolution cases of 1,2,3,4-bit/stage. The stage scaling assumes the optimum value of $2^{-(B_s-1)}$, and the number of m -stages is determined by adding the appropriate

Fig. 3.17 Basic $g_m - C$ amplifier for modeling the RA gain stage



$\leq B_s$ bits to the last stage for realizing the necessary aggregate resolution B . It is also assumed that there is no dedicated input S/H and the sampling is performed in the stage-1 MDAC (S/H-less), which saves power [26]. Taking into account the contributions from the sampler, the RA, the comparator, the resistor ladder, and the digital logic, the total power consumption of an m -stage B_s -bit/stage pipeline can be generally expressed as

$$P_{P,\text{tot}} = P_{P,\text{samp}} + P_{P,\text{RA,tot}} + (2^{B_s} - 1) \cdot P_{P,\text{comp,tot}} + m \cdot (P_{P,\text{lad}} + P_{P,\text{dig}}), \quad [\text{W}] \quad (3.29)$$

where $P_{P,\text{RA,tot}}$, $(2^{B_s} - 1) \cdot P_{P,\text{comp,tot}}$ encompass the contribution of all the RAs and the comparators along the pipeline, as will be derived. We start by deriving the expressions for the RA, since this is the newly introduced block with respect to the flash and the SAR, and it is typically the dominant contributor to the pipeline power consumption. We study the basic open-loop $g_m - C$ amplifier of Fig. 3.17 as the single RA with a gain expression

$$A_s = 2^{B_s-1} = g_{m,\text{RA}} \cdot r_{o,\text{RA}} \cdot (1 - e^{-\frac{T_{\text{RA}}}{\tau_{\text{RA}}}}). \quad (3.30)$$

Furthermore, we consider only linear settling to $1/4$ LSB accuracy of the back-end at a percentage ζ_{set} of the allocated time to the RA. When including the parasitic loading at the RA output C_{RA} (see Eqs. (3.10) and (3.11)), $g_{m,\text{RA}}$ with the process limits is found from the exponential settling (see Appendix B)

$$g_{m,\text{RA}} = A_s \cdot \frac{\theta_{\text{P,RA}}}{\zeta_{\text{set}}} \cdot f_s \cdot \frac{(B - B_s + 3) \cdot \ln 2 \cdot \max\{C_{\text{RA}}, C_{\text{min}}\}}{1 - A_s \cdot \frac{(B - B_s + 3) \cdot \ln 2}{\pi} \cdot \frac{\theta_{\text{P,RA}} \cdot f_s}{\zeta_{\text{set}} \cdot f_T}}. \quad [\text{S}] \quad (3.31)$$

In the above expression, $\theta_{\text{P,RA}} = 2/\psi$ captures the portion of the conversion allocated to the RA. In a pipeline with flash sub-ADCs, each interstage RA shares half of the total converter period ($2/f_s$) with the sub-flash of the previous stage. Commonly used values for ψ and ζ_{set} are 0.7 [72] and 0.5 [81], respectively, which are also adopted in this analysis. The factor $(B - (B_s - 1) + 2)$ depicts the settling accuracy to

1/4 LSB of the back-end. C_{RA} can be substituted from the RA input-referred noise expression

$$\overline{V_{n,RA}^2} = \frac{kT}{A_s C_{RA}}. \quad [V^2] \quad (3.32)$$

By assigning a portion of the total budget to this noise contribution, the value of C_{RA} can be determined. One more contribution that should be added to Eq. (3.31) is the input capacitive loading of the next stage comparator $C_{in,I}$, as calculated in Eq. (3.9)

$$C_{in,I} \approx A_I C_I \cdot \frac{2 \cdot \theta_{P,comp} \cdot f_s}{2\pi \cdot f_T}, \quad [F] \quad (3.33)$$

which in the final expression should be multiplied by the sub-flash comparator number to capture this loading correctly. The quantity $\theta_{P,comp} = 2/(1 - \psi)$ captures the rest of the $2/f_s$ period occupied by the sub-flash comparators. The next addition is to include in the $g_{m,RA}$ expression the contribution from all the RAs along the pipeline. This is done by considering the asymptotic expansion of the C_{RA} , $C_{in,I}$ sum while incorporating the stage scaling factor as follows:

$$\begin{aligned} g_{m,RA,tot} &= A_s \cdot \frac{\theta_{P,RA}}{\zeta_{set}} \cdot f_s \cdot \frac{(B-B_s+3) \cdot \ln 2}{1 - A_s \cdot \frac{(B-B_s+3) \cdot \ln 2}{\pi} \cdot \frac{\theta_{P,RA} \cdot f_s}{\zeta_{set} \cdot f_T}} \\ &\times \left[\sum_{i=0}^{m-2} \max\{C_{RA} \cdot 2^{-(B_s-1) \cdot i}, C_{min}\} \right. \\ &\left. + (2^{B_s} - 1) \cdot \sum_{i=0}^{m-2} \max\{C_{in,I} \cdot 2^{-(B_s-1) \cdot i}, C_{min}\} \right]. \end{aligned} \quad [S] \quad (3.34)$$

A final addition to the open-loop RA power includes a linearity factor η_{lin} [26], capturing the overhead for a precision greater than the back-end by one bit

$$\eta_{lin} \approx \sqrt{\frac{32}{3}} \cdot 10^{\frac{HD3}{10}}, \quad HD3 \approx 6.02 \cdot (B - (B_s - 1) + 1) + 1.76. \quad (3.35)$$

With all the above considerations, the total pipeline RA power consumption can now be derived as

$$P_{P,RA,tot} = V_{DD} \cdot I_{RA,tot} = \frac{V_{DD}}{\eta_{lin}} \cdot g_{m,RA,tot} \cdot V_{GT}/2. \quad [W] \quad (3.36)$$

It is worth mentioning that the above expression might be overestimating somewhat the power, since it assumes the same linearity requirement for all the RAs. On the other hand, our analysis does not capture any calibration overhead that might be necessary to compensate the gain mismatch between the different RAs in the pipeline. Therefore, the slight overestimation of the linearity overhead can be

considered offset by the underestimation of the calibration overhead, preserving the good relative accuracy of our analysis.

Similar to the comparator input loading, we need to introduce the amplifier capacitive input loading $C_{in,RA}$ in our analysis. To do this, we can rewrite Eqs. (3.6) and (3.7) and substitute the subscripts for the RA. If we also consider its settling requirement to 1/4 LSB accuracy of the back-end at its allocated portion of the total conversion period, $C_{in,RA}$ is estimated as

$$C_{in,RA} \approx A_s C_{RA} \cdot \frac{2 \cdot \theta_{P,RA} \cdot (B - B_s + 3) \cdot \ln 2 \cdot f_s}{\zeta_{set} \cdot 2\pi \cdot f_T}. \quad [F] \quad (3.37)$$

Proceeding with the sampler power, when taking into account the process limitations as well as contributions from the comparators' and amplifier loading, $P_{P,samp}$ can be expressed with its final form

$$\begin{aligned} P_{P,samp} = & V_{DD} \cdot NTC \cdot \theta_{P,samp} \cdot f_s \cdot \left[\max\{C_S, C_{min}\} \right. \\ & + (2^{B_s} - 1) \cdot \max\left\{A_I C_I \cdot \frac{\theta_{P,comp} \cdot f_s}{\pi \cdot f_T}, C_{min}\right\} \\ & \left. + \max\left\{A_s C_{RA} \cdot \frac{\theta_{P,RA} \cdot (B - B_s + 3) \cdot \ln 2 \cdot f_s}{\zeta_{set} \cdot \pi \cdot f_T}, C_{min}\right\} \right] \cdot (\phi_{sup} V_{DD}). \end{aligned} \quad [W] \quad (3.38)$$

The above expression depicts to a first-order correctly the major parameters influencing the pipeline sampler power. Similar to the flash, half of the total period of the pipeline is allocated for the sampling function; therefore, $\theta_{P,samp} = 2$.

To derive the total comparator power along the pipeline, the implemented OR should be considered, as it can absorb comparator decision errors due to noise and offset. Therefore, the comparator noise and power can be relaxed by scaling down both C_I (input integrator) and C_L (output latch) accordingly. To capture this effect in our analysis, we define the Overrange Relaxation Factor (ORF) and relate it to the effective stage and aggregate resolutions as follows:

$$ORF = 2^{B - (B_s - 1)}. \quad (3.39)$$

This factor is used to scale down the aforementioned capacitors, hence to relax the noise-limited sub-flash comparator power. The power of the resistor ladder is also relaxed by the same amount, as will be shown. The take of Eq. (3.39) is that for a fixed stage resolution and OR, the sub-ADC noise-limited power savings increase with the aggregate resolution. Following the same procedure that led to Eqs. (3.13) and (3.25) for the flash and the SAR comparator powers, we can estimate $P_{P,comp}$ including both physical and process bounds. Further, just like with the RA, the contribution from all the pipeline stages is taken into account through the asymptotic expansion of the C_I, C_L sum and the stage scaling factor. The total power for a single comparator then becomes

$$\begin{aligned}
P_{P,\text{comp,tot}} = V_{\text{DD}} \cdot \theta_{P,\text{comp}} \cdot f_s \cdot & \left[\left\{ \frac{\sum_{j=0}^{m-1} \max\{ \frac{C_I}{ORF} \cdot 2^{-(B_s-1) \cdot j}, C_{\min} \}}{1 - \frac{2\Delta V_I}{\pi V_{\text{GT}}} \cdot \frac{\theta_{P,\text{comp}} \cdot f_s}{f_T} \cdot (1/\Delta V_I)} \right\} \right. \\
& \left. + \left\{ \frac{[B_s \cdot \ln 2 + \ln(BER^{-1})] \cdot \sum_{j=0}^{m-1} \max\{ \frac{C_I}{ORF} \cdot 2^{-(B_s-1) \cdot j}, C_{\min} \} \cdot V_{\text{GT}}/2}{1 - \frac{[B_s \cdot \ln 2 + \ln(BER^{-1})]}{\pi} \cdot \frac{\theta_{P,\text{comp}} \cdot f_s}{f_T}} \right\} \right]. \quad [\text{W}] \quad (3.40)
\end{aligned}$$

Compared to Eqs. (3.13) and (3.25), the term $(B - \log_2 A_1)$ has been replaced with the stage resolution B_s to keep it fixed regardless the aggregate resolution. It is worth mentioning that our analysis assumes complete allocation of the OR to the comparator and resistor ladder noise, neglecting other practical noise sources (e.g., supply/ground noise), which would normally be given part of the OR. This is affecting the accuracy of our results minorly, since the RA power can be considerably larger than the power of the sub-ADC blocks.

To reach to the resistor ladder power, the same procedure and assumptions as for the equivalent Eq. (3.14) in the flash are followed. One extra factor to include here, as mentioned above, is the ORF, which can be effectively seen as lowering the power of the ladder by the same amount. This power is then expressed as

$$P_{P,\text{lad}} = V_{\text{DD}} \cdot \frac{(\phi_{\text{sup}} V_{\text{DD}})}{ORF \cdot R_{\text{lad}}}, \quad [\text{W}] \quad (3.41)$$

Finally, the same procedure as for the flash is also adopted to estimate the digital power in each sub-flash of the pipeline. On top of that, the contribution of the “align and combine” logic is accounted for by adding the power of two extra gates per bit in each pipeline stage, leading to the final estimation

$$P_{P,\text{dig}} = V_{\text{DD}} \cdot f_s \cdot [5 \cdot (2^{B_s} - B_s) + 2 \cdot B_s] \cdot C_{\min} \cdot V_{\text{DD}}. \quad [\text{W}] \quad (3.42)$$

Equations (3.29), (3.36), (3.38), and (3.40)–(3.42) allow us to predict and compare the pipeline accuracy-speed-power limits by capturing to a great extent all the major contributions. These limits are portrayed in Fig. 3.18 for the effective stage resolutions under study and cover three different f_s values in 28 nm. The assumptions and the values of $\theta_{P,\text{samp}}$, $\theta_{P,\text{RA}}$, $\theta_{P,\text{comp}}$, and ζ_{set} leading to the shown plots are covered in their corresponding text. Further, the noise partitioning to set C_{RA} , C_{S} , C_{I} , C_{L} , and R_{lad} is done by assigning $0.4 \cdot \bar{\epsilon}_{\text{q}}^2$ to the RA, $0.4 \cdot \bar{\epsilon}_{\text{q}}^2$ for the sampling, $0.2 \cdot 0.7 \cdot \bar{\epsilon}_{\text{q}}^2$ to the comparator, and $0.2 \cdot 0.3 \cdot \bar{\epsilon}_{\text{q}}^2$ to the ladder. For a low enough f_s/f_T (Fig. 3.18a), the noise-limited region (60 dB and above) favors the larger per stage effective resolution with reducing benefits as the stage resolution gets higher than 2-bit/stage. This is because every additional per stage bit reduces the linearity and settling requirements on the amplifier residue. With the $2 \times C_{\text{RA}}$ reduction compensating the $2 \times$ higher gain, there is an overall reduction in the dominant RA power. However, this benefit gets shallower for more than 2-bit/stage, since the exponentially increasing sub-flash power starts counteracting the RA power reduction. Also, C_{RA} cannot scale indefinitely and is ultimately limited by C_{\min} . When f_s/f_T increases, the higher RA gain in Eq. (3.34) and the

increased parasitic contribution start bridging the gap (Fig. 3.18b) or make the multi-bit approaches less efficient, even not achieving certain resolutions compared to single-bit (Fig. 3.18c). Therefore, 1-bit/stage is the optimum for achieving a high resolution while preserving f_s in the GHz range. For moderate f_s values, 2-bit/stage and 3-bit/stage are comparable, with lower resolutions favoring the former due to reduced sub-flash and parasitic overhead.

Figure 3.19 compares the limits at $f_s = 500$ MHz for the 1,2,3-bit/stage cases in the four different processes under study. The relative relationship and crossover regions between the different per stage resolution cases are very similar across these processes. Moreover, the results are following similar trends as for the flash and the SAR showing that the reduced supply is approximately offset by the higher g_m/I_D to realize a certain f_T . However, one distinct difference in the pipeline is that the process-limited regions are extended to higher resolutions, which is more profound for increasing the stage resolution. This is due to the fact that by increasing the bits/stage, scaling down the capacitance by $2^{-(B_s-1)}$ leads to a faster saturation toward C_{\min} , ending up utilizing the process-limited rather than the noise-limited value for an extended range of resolutions.

3.5 The Pipelined-SAR: A Powerful Hybrid

3.5.1 Overview

The pipelined-SAR [82–87] has been a blooming hybrid ADC architecture that combines appealingly the SAR and pipeline concepts. The block diagram of a typical B -bit two-stage pipelined-SAR is shown in Fig. 3.20. It comprises a coarse B_1 -bit SAR_1 , a fine B_2 -bit SAR_2 , and an interstage RA. The “align and combine” logic collects the bits and provides a new output with a latency of 2 clock periods at the sample rate of a single stage. Similar to the pipeline, in order to reduce power, the RA can be implemented as an open-loop dynamic amplifier, while the S/H_2 can be integrated with the RA into a single block. The concept of OR is incorporated in this hybrid architecture as well, serving as the coarse SAR_1 error absorption mechanism. To understand the circumstances, under which the combination of SAR and pipeline improves on its composing counterparts, the benefits of this hybrid on each of the two previous architectures are laid out next.

In Sect. 3.3, the noteworthy low-to-medium-resolution energy efficiency and inherent scalability of the SAR were discussed. It was also argued that when the resolution increases beyond a certain threshold (e.g., 10 bits), driving fast and efficiently enough the noise-limited DAC capacitance can be non-trivial. Besides, although a single comparator is employed, every added bit results in a $4\times$ increase

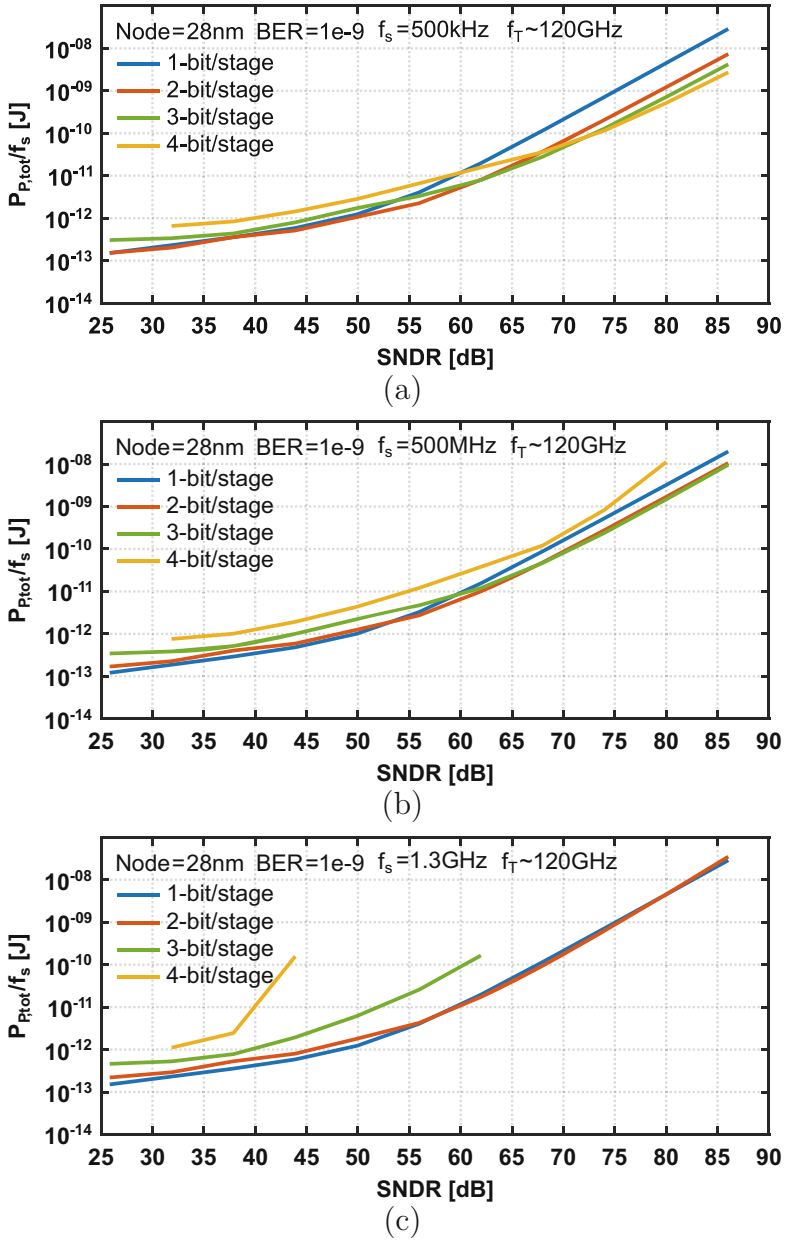


Fig. 3.18 Pipeline with 1,2,3,4-bit/stage effective resolution accuracy-speed-power limits in 28 nm: (a) $f_s = 500$ kHz, (b) $f_s = 500$ MHz, and (c) $f_s = 1.3$ GHz

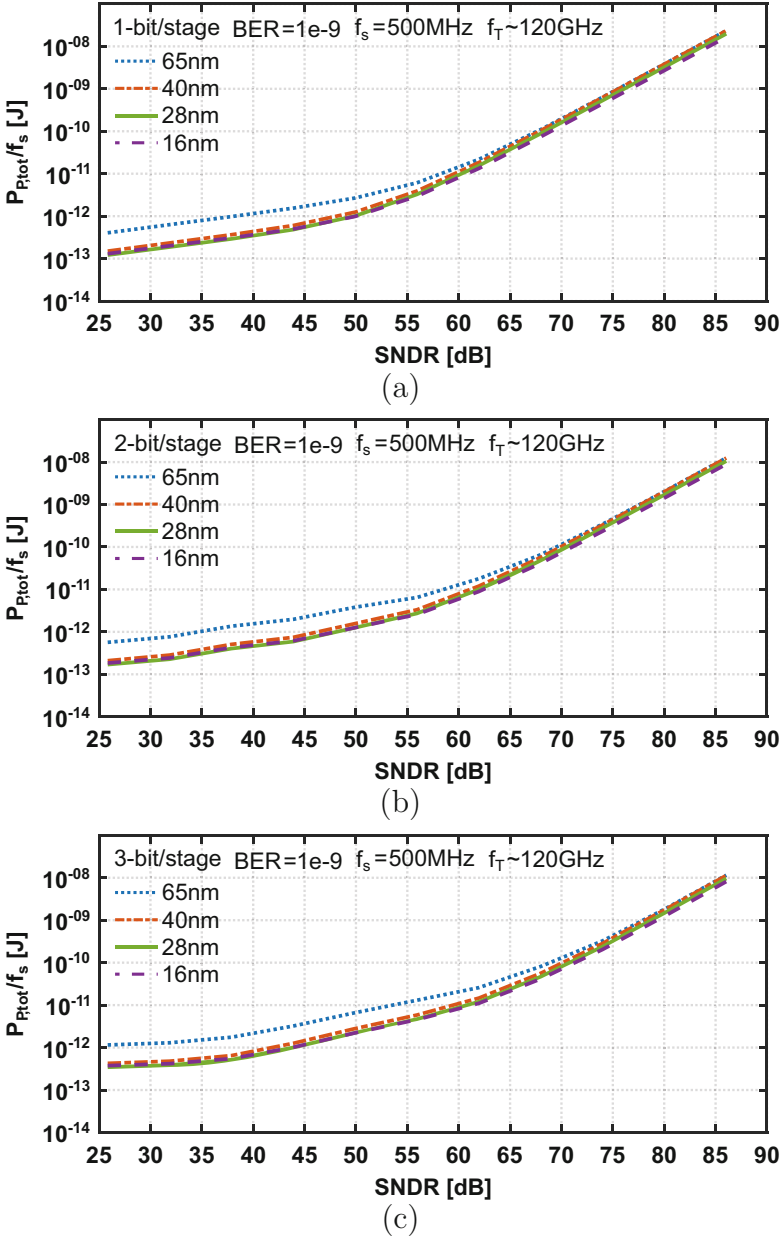


Fig. 3.19 Pipeline accuracy-speed-power limits across different processes at $f_s = 500\text{MHz}$: (a) 1-bit/stage, (b) 2-bit/stage, and (c) 3-bit/stage




Fig. 3.20 Block diagram of a B -bit two-stage pipelined-SAR ADC

in power⁹ to meet the low noise specification. As illustrated in Fig. 3.21a, going from MSB to LSB, the conversion transitions from high-noise (low-energy) to low-noise (high-energy) events, with the highest energy levels required only in the last couple of cycles. Using the same comparator results in energy waste, since it has to be designed for these last cycles consuming unnecessary energy in the rest of the conversion. Figure 3.21b illustrates the energy profile of a two-stage pipelined-SAR with a 2x-OR (1-bit extra) and an interstage RA. The sub-SAR bit allocation is done such that combined with the chosen OR, it enables a high-noise (low-energy) SAR₁ comparator. A high-noise (low-energy) comparator is also made possible in SAR₂, relaxed by the gain of the RA. This leaves the RA the sole low-noise (high-energy) contributor, whose efficiency can be on the same order as that of a low-noise comparator in a SAR (see Chap. 6). On top, pipelining boosts the converter sample rate by roughly the amount of sequential cycles saved from each SAR

$$f_s = \begin{cases} \frac{1}{B+1} \cdot f_{\text{cycle}}, & \text{SAR} \\ \frac{1}{B_1+2} \cdot f_{\text{cycle}}, & \text{pipe-SAR} \end{cases} \quad [\text{Hz}] \quad (3.43)$$

where it is assumed that the sampling, the conversion, and the amplification are all allocated one bit cycle.

In Sect. 3.4, it was mentioned that increasing the pipeline stage resolution can result in power savings by relaxing the settling and linearity requirements on each

⁹ In reality, even more power is needed to compensate for the extra parasitics if the same speed is to be preserved.




Fig. 3.21 Illustration of conversion energy requirement in (a) a binary SAR ADC and (b) a two-stage pipelined-SAR ADC

residue. Further, the number of stages is reduced, and so does the calibration power overhead to compensate for any mismatches between different RAs. However, there exists a limit to this benefit due to the exponential power increase of the sub-flash comparators and the imperfect stage scaling, which reaches C_{\min} at a faster rate. This benefit even vanishes when f_s reaches values such that the sub-flash parasitic loading to the sampler and the RA (Eqs. (3.34) and (3.38)) starts dominating. Using a sub-SAR instead, the comparator power increases only linearly instead of exponentially, resulting in an increased energy efficiency. Further, for the same f_s/f_T between flash and SAR, the comparator parasitic loading to the sampler and the RA also increases linearly ($\propto \theta \propto B$) rather than exponentially ($\propto 2^B$). Therefore, the range of resolutions for which a multi-bit/stage pipeline is superior to its single-bit/stage counterpart can be extended up to values that allow a low energy comparator for the sub-SAR.

The vast majority of the up-to-date pipelined-SAR ADCs found in literature are two-stage. There barely exist a handful of examples with a higher pipelining

order [88–91].¹⁰ From a theoretical standpoint, analogous to the regular pipeline with flash sub-ADCs, there is no fundamental reason preventing the pipelining of more than two SARs. In fact, intuition and Eq. (3.43) predict that higher-order pipelined-SARs should be able to achieve a higher absolute sample rate due to further reduction of the sequential cycles in each sub-SAR. The energy efficiency of > 2-stage pipelined-SARs should also increase when opting for a high resolution and an f_s in the GHz range due to reduced parasitic loading. This efficiency might potentially surpass one of the regular pipelines provided that the sub-SAR is more efficient than the sub-flash for the chosen per stage sample rate and resolution. The following mathematical derivation tries to shed some light on whether this is indeed the case.

3.5.2 Pipelined-SAR Accuracy-Speed-Power Limits

The derivation of the pipelined-SAR ADC accuracy-speed-power limits begins by constructing a B -bit converter with m -stages, both of which determine the B_s bits per stage, including a 2 \times -OR between stages ($B_s - 1$ effective bits). The number of stages under study is $m = 2, 3, 4, 5$, and the bit partitioning is done as shown in Table 3.2. The total converter power, including the sampler, the RA, the comparator, the CDAC, and the digital logic, can be generally expressed as

Table 3.2 Bit partitioning for different aggregate resolutions in a 2,3,4,5-stage pipelined-SAR including 2 \times -OR between stages

Total bits [B]	Bits/stage 2-stage	Bits/stage 3-stage	Bits/stage 4-stage	Bits/stage 5-stage
4	2–3	2-2-2	2-2-2-1	2-2-2-1-1
5	3–3	2-2-3	2-2-2-2	2-2-2-2-1
6	3–4	2-3-3	2-2-2-3	2-2-2-2-2
7	4–4	3-3-3	2-2-3-3	2-2-2-2-3
8	4–5	3-3-4	2-3-3-3	2-2-2-3-3
9	5–5	3-4-4	3-3-3-3	2-2-3-3-3
10	5–6	4-4-4	3-3-3-4	2-3-3-3-3
11	6–6	4-4-5	3-3-4-4	3-3-3-3-3
12	6–7	4-5-5	3-4-4-4	3-3-3-3-4
13	7–7	5-5-5	4-4-4-4	3-3-3-4-4
14	7–8	5-5-6	4-4-4-5	3-3-4-4-4

¹⁰ Chapter 6 details the implementation and experimental verification of [91], being one of the proposed prototypes in this book.

$$P_{\text{PS,tot}} = P_{\text{PS,samp}} + P_{\text{PS,RA,tot}} + B_s \cdot P_{\text{PS,comp,tot}} + P_{\text{PS,DAC,tot}} + m \cdot P_{\text{PS,dig}}, \quad [\text{W}] \quad (3.44)$$

where $P_{\text{PS,RA,tot}}$, $B_s \cdot P_{\text{PS,comp,tot}}$, and $P_{\text{PS,DAC,tot}}$ contain the total RA, comparator, and CDAC contribution along the pipeline. Using our previous analyses, we possess all the necessary tools to estimate the sampler power

$$P_{\text{PS,samp}} = V_{\text{DD}} \cdot NTC \cdot \theta_{\text{PS,samp}} \cdot f_s \cdot \left[\max\{C_s, C_{\min}\} + \max\{A_I C_I \cdot \frac{\theta_{\text{PS,comp}} \cdot f_s}{\pi \cdot f_T}, C_{\min}\} + \max\{A_s C_{\text{RA}} \cdot \frac{\theta_{\text{PS,RA}} \cdot (B - B_s + 3) \cdot \ln 2 \cdot f_s}{\zeta_{\text{set}} \cdot \pi \cdot f_T}, C_{\min}\} \right] \cdot (\phi_{\text{sup}} V_{\text{DD}}). \quad [\text{W}] \quad (3.45)$$

In our basic SAR model (see Sect. 3.3), we allocated one cycle to the sampler. However, now the RA also needs to fit within one sub-ADC conversion period. Allocating one cycle to the sampler and one to the RA would be overly pessimistic and not entirely realistic. In our pipelined-SAR model, we allocate one and a half cycles to both the sampler and the RA. In terms of power, this is effectively seen as adding one extra conversion cycle to the sub-SAR. Thus, in Eq. (3.45), $\theta_{\text{PS,samp}} = \theta_{\text{PS,RA}} = (B_s + 1 + 2)/1.5$ and $\theta_{\text{PS,comp}} = (B_s + 1 + 2) \cdot 1/\beta$, with $\beta = 0.7$ as in the SAR. Also, $\zeta_{\text{set}} = 0.5$ as for the regular pipeline.

For estimating the power of the RA, we again adopt the open-loop $g_m - C$ model of Fig. 3.17 and retain all the aforementioned assumptions regarding linearity and settling accuracy from Sect. 3.4.2. The power consumption of the RA in the pipelined-SAR, including the parasitic loading and the process contribution, can be then expressed as

$$P_{\text{PS,RA,tot}} = A_s(B_s) \cdot \frac{\theta_{\text{PS,RA}}}{\zeta_{\text{set}}} \cdot f_s \cdot \frac{(B - B_s + 3) \cdot \ln 2}{1 - A_s(B_s) \cdot \frac{(B - B_s + 3) \cdot \ln 2}{\pi} \cdot \frac{\theta_{\text{PS,RA}} \cdot f_s}{\zeta_{\text{set}} \cdot f_T}} \times \left[\sum_{i=0}^{m-2} \max\{C_{\text{RA}} \cdot 2^{-(B_s-1) \cdot i}, C_{\min}\} + \sum_{i=0}^{m-2} \max\{C_{\text{in,I}} \cdot 2^{-(B_s-1) \cdot i}, C_{\min}\} \right] \cdot \frac{V_{\text{GT}} - V_{\text{DD}}}{2 \cdot \eta_{\text{lin}}}. \quad [\text{W}] \quad (3.46)$$

It is important to explain one difference regarding the RA gain between the regular pipeline and the pipelined-SAR and how it is reflected in our derivation. In the former, A_s remains constant when increasing B , since B_s stays the same and m increases. However, in the latter, for a fixed m , when increasing B , the bit partitioning is not constant, making A_s a function of B_s . For example, in the case of a 6-bit, 3-stage pipelined-SAR in Table 3.2, $A_1 = 2$ and $A_2 = 4$. These values are both captured in Eq. (3.46), each one multiplied either with its output C_{RA} according to noise and stage scaling considerations or with C_{min} .

The total power of a single comparator along the pipeline can be derived by employing Eq. (3.40), including ORF and both physical and process bounds

$$P_{\text{PS,comp,tot}} = V_{\text{DD}} \cdot \theta_{\text{PS,comp}} \cdot f_s \cdot \left[\left\{ \frac{\sum_{j=0}^{m-1} \max\left\{ \frac{C_{\text{I}}}{\text{ORF}} \cdot 2^{-(B_s-1) \cdot j}, C_{\text{min}} \right\}}{1 - \frac{2\Delta V_{\text{I}}}{\pi V_{\text{GT}}} \cdot \frac{\theta_{\text{PS,comp}} \cdot f_s}{f_{\text{T}}} \cdot (1/\Delta V_{\text{I}})} \right\} \right. \\ \left. + \left[\frac{[B_s \cdot \ln 2 + \ln(BER^{-1})] \cdot \sum_{j=0}^{m-1} \max\left\{ \frac{C_{\text{I}}}{\text{ORF}} \cdot 2^{-(B_s-1) \cdot j}, C_{\text{min}} \right\} \cdot V_{\text{GT}}/2}{1 - \frac{[B_s \cdot \ln 2 + \ln(BER^{-1})]}{\pi} \cdot \frac{\theta_{\text{PS,comp}} \cdot f_s}{f_{\text{T}}}} \right] \right]. \quad [\text{W}] \quad (3.47)$$

For the power consumption of all CDACs, the MCS scheme is considered with the appropriate switching activity for each sub-SAR, depending on the bit partitioning. In addition, the stage scaling is taken into account by dividing each back-end CDAC capacitance with the gain product up to that stage. In the case of a pipelined-SAR, this capacitance is the C_{RA} itself. The expression covering all stages under study and bits/stage according to Table 3.2 is given as

$$P_{\text{PS,DAC,tot}} = \theta_{\text{PS,DAC}} \cdot f_s \cdot \left[\sum_{i=2}^5 \left(\sum_{j=1}^{B_s(i)-1} 2^{B_s(i)-3-2j} (2^j - 1) \right) \right. \\ \left. \times \max \left\{ \frac{C_{\text{S}}}{2^{B_s(i)}} \prod_{l=0}^4 \frac{1}{A_{\text{sl}}(B_s(i))}, C_{\text{min}} \right\} \right] \cdot \frac{V_{\text{REF}} \cdot V_{\text{DD}}}{\lambda_{\text{ref}}}. \quad [\text{W}] \quad (3.48)$$

In the above expression, $l=0$ corresponds to the front-end CDAC ($A_{\text{sl}} = 1$). Further, $\theta_{\text{PS,DAC}} = [(B_s + 1 + 2)/B_s] \cdot 1/(1 - \beta)$ represents the timing portion allocated to the CDAC over one full converter period.

To conclude our architectural limits' derivation, the power of the digital logic is estimated assuming the same number of gates for the sub-SAR as in the regular SAR and adding two extra gates per bit in each stage as in the regular pipeline

$$P_{\text{PS,dig}} = V_{\text{DD}} \cdot f_s \cdot [5 \cdot B_s + 2 \cdot B_s] \cdot C_{\text{min}} \cdot V_{\text{DD}}. \quad [\text{W}] \quad (3.49)$$

The limits derived by Eqs. (3.44)–(3.49) are portrayed in Fig. 3.22 for different stages and three different f_s values in 28 nm, for the aforementioned values of $\theta_{\text{PS,samp}}$, $\theta_{\text{PS,RA}}$, $\theta_{\text{PS,comp}}$, and ζ_{set} . Further, C_{RA} , C_{S} , C_{I} , and C_{L} are set by assigning $0.4 \cdot \bar{\epsilon}_{\text{q}}^2$ to the RA, $0.4 \cdot \bar{\epsilon}_{\text{q}}^2$ to the sampler, and $0.2 \cdot \bar{\epsilon}_{\text{q}}^2$ to the comparator, while B_s in Eq. (3.44) is taken as $\max\{B_s\}$ according to Table 3.2. For low f_s , reducing the stages improves efficiency in the noise-limited region, but the benefits are less profound than increasing the bits/stage in the regular pipeline. When increasing the aggregate resolution, the RA has less time available due to the extra sequential cycles in the sub-SAR. This is partially cancelling the advantage of the relaxed linearity and settling. Increasing f_s/f_{T} , similar trends as in the regular pipeline are seen, with the higher RA gain and increased parasitics making the higher-stage approaches more suitable in the GHz range.

The derived limits are once more compared in the four different processes under study at $f_s = 500$ MHz for the 3,4,5-stage cases (Fig. 3.23). The crossover regions between the different stage cases are very similar across these processes. Increasing the number of stages keeps the process-limited regions roughly in the same range of resolutions. In contrast to the regular pipeline, the sub-SAR resolution in these regions is more similar across the different stage cases.

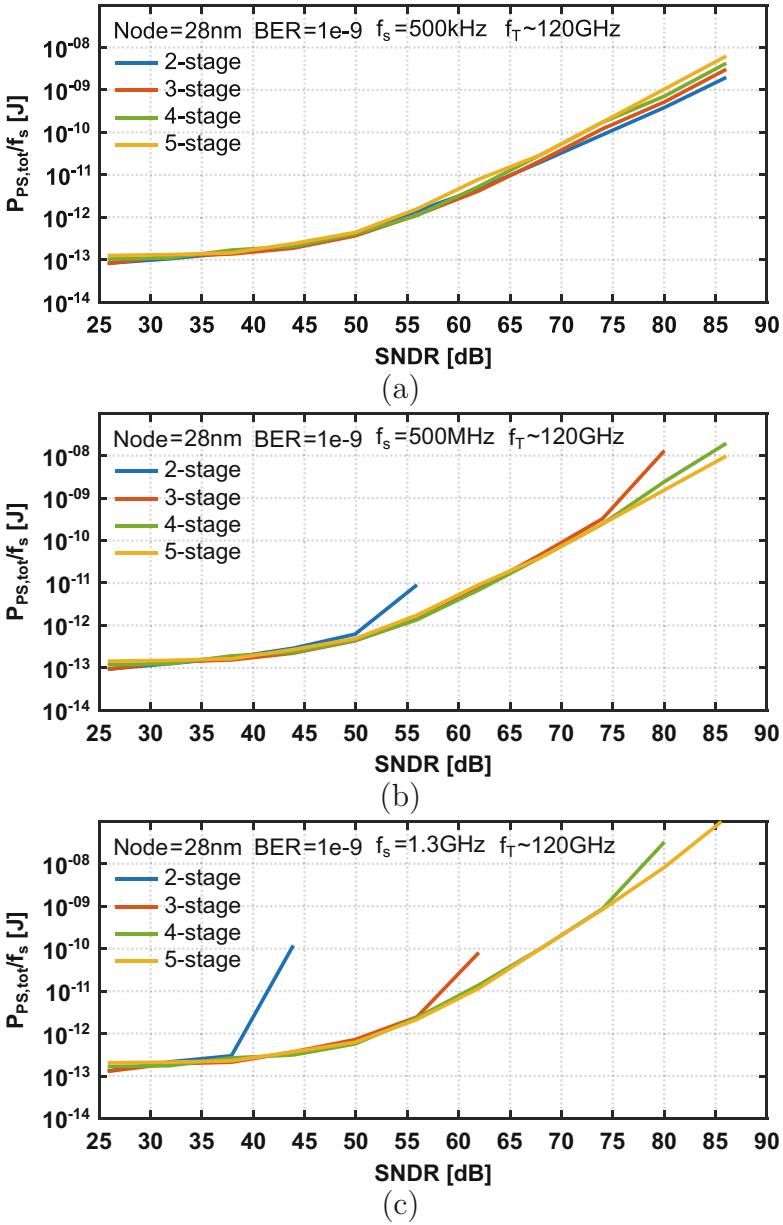


Fig. 3.22 2,3,4,5-stage pipelined-SAR with accuracy-speed-power limits in 28 nm: (a) $f_s = 500$ kHz, (b) $f_s = 500$ MHz, and (c) $f_s = 1.3$ GHz




Fig. 3.23 Pipelined-SAR accuracy-speed-power limits across different processes at $f_s = 500$ MHz: (a) three-stage, (b) four-stage, and (c) five-stage

3.6 Architectural Limits' Comparison

After having derived the fundamental plus process accuracy-speed-power limits for the most important ADC architectures (flash, SAR, pipeline, pipelined-SAR), it only makes sense to compare them against each other and build an insight on which architecture performs best in which speed and/or resolution region. This is shown in Figs. 3.24, 3.25, and 3.26. All the plots are taken employing the assumptions laid out in the corresponding sections of the architectures under comparison. The SAR derived limits consider a binary converter, without involving the concept of OR or other comparator noise relaxation technique [92]. It is non-trivial to develop a generic estimation of their benefits and overhead, since their optimum utilization is tailored to the needs of a specific design.

For $f_s = 500$ kHz in Fig. 3.24, the parasitic contribution from the RA and the comparators is negligible; therefore, the slopes of the different curves are first process-limited and then noise-limited. As expected, above about 40 dB SNDR, the flash is the most energy-inefficient with a slope $\propto 2^B$. The pipelines show a similar slope in their noise-limited regions but with a better overall efficiency compared to the flash due to sub-ranging and OR. The SAR is very hard to beat up to about 45 dB, after which its energy increases with a slope $\propto B$. The pipelined-SARs follow a similar noise-limited slope and show the best efficiency from 45 dB to 65 dB. In the range of 70 dB and above, they are slightly more efficient than the multi-bit/stage pipelines. When $f_s = 500$ MHz (Fig. 3.25), the flash remains the most energy-inefficient above 40 dB. However, the slopes of the other architectures have increased, with the parasitic contribution deteriorating their efficiency more than the flash, due to more stringent internal timings and/or high RA gain. The SAR is still superior to all the other architectures up to 40 dB while closely competing with the pipelines up to 50 dB. The two-stage pipelined-SAR stops at 56 dB, while the 4-bit/stage pipeline and the three-stage pipelined-SAR stop at 80 dB. The 1,2-bit/stage pipelines and the 4,5-stage pipelined-SARs can achieve the highest resolution at a good efficiency, with the 2-bit/stage pipeline and the five-stage pipelined-SAR leading this race above 75 dB. At $f_s = 1.3$ GHz (Fig. 3.26), the flash remains almost intact, while the exacerbated parasitic contribution in the other architectures further deteriorates their efficiency, with more stopping at certain SNDR values. The SAR retains the highest efficiency up to 40 dB, while the 1,2-bit/stage pipelines are the most efficient above 75 dB. The 4,5-stage pipelined-SARs are winning in the range 40–65 dB and competing with the pipelines up to 75 dB, with the three-stage following up to 56 dB. Extending the pipelined-SAR to more than five stages can catch or even surpass the 1-bit/stage pipeline in high-resolution efficiency. We may conclude that for the same stage count, the pipelined-SAR is more efficient and potentially faster than the pipeline for an extended range of resolutions.

It is worth noting that there are techniques to enhance the efficiency of some of the studied ADC architectures beyond the analytically predicted in the previous sections. For example, a fully settled RA was assumed for the pipeline and pipelined-SAR, and its transconductance was derived based on the number of time




Fig. 3.24 Accuracy-speed-power limits for the different ADC architectures studied at $f_s = 500\text{ kHz}$

constants for a given settling accuracy. Instead, using an unsettled integrator that minimizes the number of time constants to one is becoming a popular low-power choice [84, 93]. Appendix C derives the transconductance of an RA operating in the integrator mode and compares it with the derived one for the fully settled mode. A similar argument can be made about the efficiency of the SAR and pipelined-SAR,

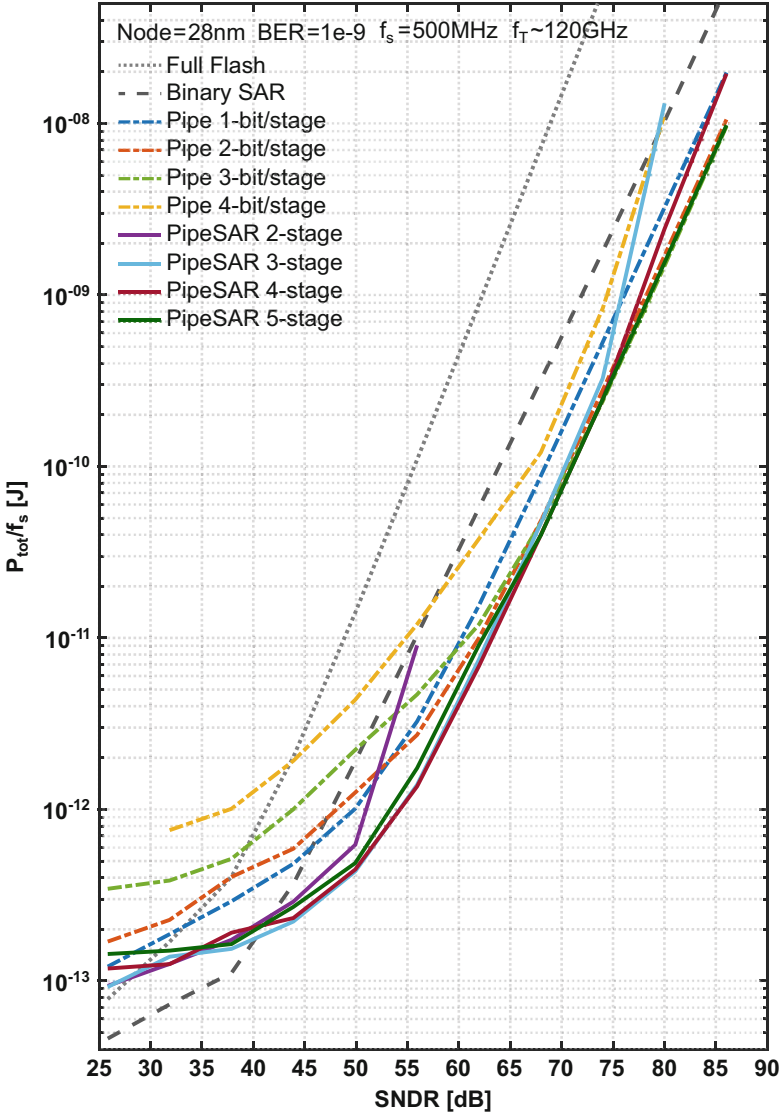


Fig. 3.25 Accuracy-speed-power limits for the different ADC architectures studied at $f_s = 500$ MHz

which could be further enhanced by employing the concept of OR or relaxing the BER through smart asynchronous timing. Both these techniques would relax the comparator power for a certain noise/speed contribution. However, the overhead of all the aforementioned techniques due to potential extra circuitry (e.g., more




Fig. 3.26 Accuracy-speed-power limits for the different ADC architectures studied at $f_s = 1.3\text{GHz}$

logic and/or calibration) should be carefully assessed to determine the net efficiency benefits.

Additionally, there are aspects omitted by our analysis that could potentially favor the multi-stage pipelined-SAR even more compared to the traditional pipeline for a similar stage count. In the S/H-less topology assumed for the pipeline for power

reasons, there exists a mismatch between the first stage MDAC and flash paths, resulting in large skew. One way to mitigate the effect of this mismatch is to allocate a big portion of the first stage OR for its absorption. However, this would take away an equivalent amount of the ORF for scaling down the capacitors of the first stage comparators. Increasing these capacitors to keep the same noise contribution would in turn increase the comparators' power as well as the power of the sampler driving their increased input capacitance. Further, the resistor ladder power was estimated for resistors sized for a certain portion of the total quantization noise, leading to rather large values. In reality, the resistors in a practical design would have to be sized smaller, due to kickback and settling considerations at the input of the flash comparators, leading to a higher power consumed by the ladder.

Finally, it is worth clarifying the reason for the SAR scalability in lower process nodes not being entirely obvious in the plots of Sect. 3.3. This is partly due to the fact that, as discussed in Chap. 2, noise imposes a fundamental limitation; hence, the noise-limited performance of any architecture does not improve much with scaling. Further, the comparison between the different processes is made for the same f_T , resulting in a larger g_m/I_D going to lower nodes (Fig. 3.4). Performing the comparison for the same g_m/I_D leads to a smaller f_s/f_T going to lower nodes. This favors SAR-based more than flash-based converters in achieving a higher f_s and/or a reduced power, due to accumulated benefits over multiple internal cycles. Such a trend is also depicted in the SotA standings (Fig. 3.1), which include several practical error sources on top of noise.

3.7 Time-Interleaving

3.7.1 Overview

Up to this point, several standalone converter architectures and their trade-offs have been analyzed. Depending on their hardware and internal timings, these architectures are capable of achieving a different *accuracy · speed ÷ power* product. Pipelining is one way to improve this product for both SAR and flash by extending their speed and/or accuracy. The f_T of a certain process is what ultimately determines what speed can be achieved for a certain accuracy and how much power needs to be spent. One very popular way to boost the sample rate beyond the capabilities of a standalone converter, extending to a first-order the process f_T limit, is to run N identical sub-ADCs (or channels) with clocks shifted in time in a TI configuration [94–97]. Figure 3.27 shows a high-level diagram of an N -channel TI-ADC (Fig. 3.27a) alongside its time-domain sampling with an N -interleaved Dirac sequence (Fig. 3.27b). Each sub-ADC converts every N^{th} sample operating at an f_s sample rate and a $2\pi/N$ -rad phase shift with respect to its preceding and succeeding channel, resulting in a total converter sample rate of $N \cdot f_s$. The digital output




Fig. 3.27 (a) High-level block diagram of an N -channel TI-ADC and (b) sampling of a signal using an N -interleaved Dirac pulse sequence




Fig. 3.28 Power vs. frequency illustration of a non-TI- and a TI-ADC

streams of all the N channels are combined into a single output stream by a digital $N:1$ Multiplexer (MUX) running at the total sample rate.

Any ADC type, including the architectures described in the previous sections, can be turned into a TI-ADC. The potential interleaving advantages as well as the optimum interleaving factor N to maximize the overall *accuracy · speed ÷ power* depend on the required specifications and the capabilities of the chosen sub-ADC architecture. For a total sample rate $N \cdot f_s$, the lower N is, the higher the individual sample rate f_s of each non-interleaved converter must be. When increasing the sample rate of an ADC, its power consumption first increases linearly with frequency, and beyond a certain threshold, this increase is super-linear, as illustrated in Fig. 3.28. As it was proven in the derived expressions of the previous sections, when the f_s/f_T rises, the exacerbated parasitic loading increases the power vs. frequency slope, eventually dictating the required power. From our previous

analysis, it can be easily deduced that the threshold between the linear and super-linear regions is architecture and process node dependent.

Ideal interleaving would pick a sub-ADC operating at its optimal f_s just before its super-linear region and extend its f_s by the interleaving factor, provided that a total input bandwidth $\geq N \cdot f_s/2$ is guaranteed. The power vs. frequency of the TI-ADC initially preserves the linear behavior of the sub-ADC. By increasing N , the increasing capacitive loading at the output of the front-end buffer/amplifier and its reduced available settling time eventually lead to a super-linear increase in power. If Schreier's FoM is adopted from the previous chapter as a measure of the power efficiency, the sub-ADC and its TI counterpart should theoretically achieve the same FoM

$$\begin{aligned} FoM_{\text{sub-ADC}} &= SNDR + 10 \log \left[\frac{f_s}{2P} \right] \\ FoM_{\text{TI-ADC}} &= SNDR + 10 \log \left[\frac{N \cdot f_s}{2(N \cdot P)} \right] = FoM_{\text{sub-ADC}}. \end{aligned} \quad [\text{dB}] \quad (3.50)$$

In reality, the TI-ADC is always less efficient than its non-TI constituent (Fig. 3.28) due to the Interleaving Overhead (ILO) associated with practical interleaving. This overhead comprises the additional power needed for buffering and distribution of the high-quality analog input and reference signals to all sub-ADCs, the generation and distribution of multiple accurately controlled clocked phases, and the digital back-end MUX. Additionally, the calibration circuitry necessitated to compensate for errors associated with interleaving (see Sect. 3.7.2) further increases this overhead. Taking the aforementioned into account, the FoM of a practical TI-ADC becomes

$$\begin{aligned} FoM_{\text{TI-ADC,pr}} &= SNDR + 10 \log \left[\frac{N \cdot f_s}{2(N \cdot ILO \cdot P)} \right] \\ &< FoM_{\text{sub-ADC}}, \quad ILO > 1, \end{aligned} \quad [\text{dB}] \quad (3.51)$$

where in both Eqs. (3.50) and (3.51) the ADC power consumption is denoted by P .

3.7.2 Interleaving Errors

The concept of interleaving, although very powerful and occasionally the only option to achieve the highest possible speed, does not come free of challenges. Besides the already discussed error sources of the individual sub-ADCs, TI-ADCs suffer from additional errors stemming from the nature of interleaving to process different portions of the signal by different circuits [98, 99]. Imbalances between these circuits lead to four types of mismatch errors: (1) offset OS ; (2) gain G ; (3) timing δT ; and (4) bandwidth BW . These errors are depicted in Fig. 3.29 for a four-channel TI-ADC, and the characteristics of each are discussed next.

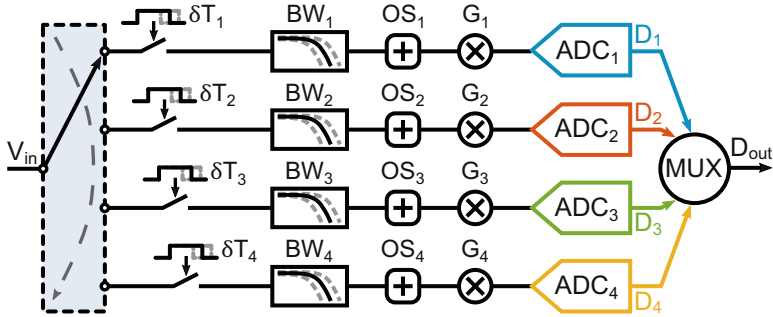


Fig. 3.29 Illustration of mismatch errors in a four-channel TI-ADC example




Fig. 3.30 Graphical illustration of sub-ADC offset mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum

Offset Mismatch Errors

The origin of offset mismatch errors is typically the offset variation between the comparators, the DACs, and/or the RAs of the different sub-ADCs. The offset of a non-interleaved ADC would create a DC term that can be either removed or ignored, therefore not deteriorating the converter’s SNDR. However, the offset mismatch between the sub-ADCs of a TI converter results in a periodic error signal with a period equal to N/F_s , where the total converter sample rate $N \cdot f_s$ is denoted by F_s for brevity. The different samplers (Fig. 3.29) introduce their own offsets, and mismatch between them further enhances this error signal. As a consequence, the output spectrum contains the fundamental tone at f_{in} as well as spurious tones due to offset mismatch at frequencies

$$f_{\text{tone, OS}} = k \cdot F_s / N, \quad k = 1, 2, \dots, N. \quad [\text{Hz}] \quad (3.52)$$

The magnitude of these tones depends on the number of sub-ADC channels and the amplitude of the error signal between them but is independent of the amplitude and frequency of the input signal. To gain more insight, Fig. 3.30a shows the outputs of a four-channel TI-ADC with different offsets for a sinusoidal input. The difference




Fig. 3.31 Graphical illustration of sub-ADC gain mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum

of each output from an ideal output is each channel's offset sampled at an $F_s/4$ frequency and a $2\pi/4$ -rad phase shift with respect to the other channels [100]. The spectrum up to $F_s/2$ is shown in Fig. 3.30b highlighting the resulting tones due to offset mismatch. The DC term is due to the average offset of the sub-ADCs, similar to a non-interleaved ADC.

Offset mismatch errors can be brought to desired levels by employing proper design techniques such as device scaling [28]. Alternatively, auxiliary calibration circuits may be employed to correct these errors either at start-up [91] or in the background [94]. Any of these approaches requires extra hardware, which unavoidably adds somewhat to the area/power overhead of the converter.

Gain Mismatch Errors

Errors due to gain mismatch originate from the variation in the unit elements of the DAC and the RA gains as well as the difference in the reference voltages between the sub-ADCs. Variations in the gain of the samplers through their on-resistance or charge injection further increase these errors. Similar to the offset mismatch case, gain mismatch results in a core error signal with a period of N/F_s , but the multiplicative nature of this error modulates its amplitude by the input frequency. For a sinusoidal input with a frequency of f_{in} , the multiplication of the input signal with the periodic error signal results in spurious tones at the output spectrum with frequency locations

$$f_{\text{tone},G} = k \cdot F_s/N \pm f_{in}, \quad k = 1, 2, \dots, N. \quad [\text{Hz}] \quad (3.53)$$

On par with the offset mismatch case, the magnitude of these tones depends on the number of sub-ADCs and the error signal amplitude between them and is independent of the input frequency. However, now there is a dependency on the input signal amplitude, which is not the case for the offset mismatch. The outputs of a four-channel TI-ADC with different gains are shown in Fig. 3.31a for a sinusoidal input. In this case, subtracting each output from an ideal output gives the channel's gain error, which is a scaled sinusoid at the input frequency sampled at $F_s/4$ with a




Fig. 3.32 Graphical illustration of sub-ADC timing mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum

$2\pi/4$ -rad phase shift relative to the other channels. This results in scaled versions of the input around $F_s/4$, $F_s/2$, as highlighted in the Nyquist spectrum of Fig. 3.31b.

Gain mismatch errors can also be either minimized by employing proper design techniques or corrected by foreground [91] and/or background calibration [94]. Due to their static nature, offset/gain mismatch errors are the easiest ones to tackle. Typically, a single point calibration at a low input frequency suffices to bring them to the desired levels in the entire band of interest.

Timing Mismatch Errors

Timing mismatch errors attribute their origin to the variation in the sampling instants between the samplers preceding the sub-ADCs (Fig. 3.29). Random variations in the clock generation and distribution path, including both circuits (switches, buffers) and routing, convert the ideal phase difference of $2\pi/N$ between adjacent channels to $2\pi/N + 2\pi f_{in} \delta T$, with δT the timing deviation. Timing mismatch generates a core error signal with a period of N/F_s and its amplitude modulated by the input frequency but with a $\pi/2$ -rad phase shift compared to the gain mismatch case. When subtracting two sinusoids with different gains and equal phase, their maxima occur at their peaks, while in the case of equal gain and different phases, the maxima occur at the zero-crossings. The spectrum contains tones at the same frequencies as the gain mismatch case

$$f_{\text{tone},\delta T} = k \cdot F_s/N \pm f_{in} = f_{\text{tone},G}, \quad k = 1, 2, \dots, N. [\text{Hz}] \quad (3.54)$$

Unlike offset/gain mismatch, the magnitude of these tones depends on the input frequency. A fixed δT leads to a larger phase deviation as f_{in} increases. The main difference between timing mismatch and jitter is that the former only affects TI-ADCs and is deterministic, while the latter exists in any ADC and is random. The outputs of a four-channel TI-ADC with different phases are shown in Fig. 3.32a for a sinusoidal input. Their difference gives once more scaled versions of the input around $F_s/4$, $F_s/2$, as shown in the spectrum of Fig. 3.32b.




Fig. 3.33 Graphical illustration of sub-ADC bandwidth mismatch errors in a four-channel TI-ADC: (a) time waveform and (b) frequency spectrum

Timing mismatch errors can be eliminated to a first-order by properly choosing the interleaver architecture (see Sect. 3.7.3). Calibration may be employed to minimize the errors. Estimation is mainly done in the digital domain and correction in the analog, for a mixed-signal solution [101], although fully digital schemes with Finite-Impulse-Response (FIR) filters exist. Due to the errors' dynamic nature, the highest frequency should be chosen for calibration.

Bandwidth Mismatch Errors

Bandwidth mismatch errors emerge from variations in the on-resistance and sampling capacitance as well as the interconnect resistance and capacitance between the samplers prior to the sub-ADCs. This generates gain and timing mismatch between the sub-ADCs, both of which are frequency dependent. The resulting spurious tones are added to the existing ones from the gain/timing mismatch cases at the same frequencies

$$f_{\text{tone}, BW} = k \cdot F_s / N \pm f_{\text{in}} = f_{\text{tone}, G}, \quad k = 1, 2, \dots, N [\text{Hz}] \quad (3.55)$$

The outputs of a four-channel TI-ADC with different bandwidths are shown in Fig. 3.33a, while the corresponding spectrum with the tones as scaled versions of the input around $F_s/4$, $F_s/2$ is depicted in Fig. 3.33b. The significance of the gain and timing mismatch on the magnitude of these tones depends on the nominal channel bandwidth. From the transfer function of our simple RC model (see Chap. 2, Sect. 2.2), a gain and a phase component can be extracted [100]

$$G_{BW}(f) = 1/\sqrt{1 + (f_{\text{in}}/BW)^2} \quad \theta_{BW}(f) = -\arctan(f_{\text{in}}/BW), \quad (3.56)$$

The input frequency vs. the channel bandwidth must be minimized to reduce the tones to a desired level. Increasing the bandwidth unavoidably increases the ADC power overhead. Calibration can be applied, but due to the dynamic nature of the bandwidth-induced gain/timing errors, any type of static calibration can only correct them at the applied frequency f_{cal} . However, if f_{cal} is chosen high enough [102],

their effect becomes negligible at lower frequencies, while above f_{cal} , the bandwidth is relaxed due to a change in the errors' dependency on f_{in}/BW . The dynamic gain/timing errors might not be able to be corrected to a satisfactory level across the entire band of interest with a static calibration. A dynamic calibration, however, can alter the magnitude response of these errors over the band of interest and, in principle, completely remove them. Typically, hardware-intensive digital filters are employed, with adaptive coefficients [103]. Apart from their complexity and power overhead, such filters may require part of the band for correction, thus band-limiting the signal. The optimum choice lies with assessing the overhead and effectiveness of dynamic calibration vs. increasing the bandwidth to minimize the errors within the band of interest.

The effect of the aforementioned interleaving errors on the SNDR of a four-channel TI-ADC has been modeled, and the simulation results are shown in Fig. 3.34. The SNDR degradation due to timing (Fig. 3.34c) and bandwidth mismatch (Fig. 3.34d) increases for an increased input frequency and reduced channel bandwidth, respectively. When the bandwidth-induced gain/phase mismatch effects are separated (Fig. 3.34e), the phase mismatch dominates SNDR for small f_{in}/BW , with the two effects getting closer as the channel bandwidth decreases. Figure 3.34f captures the combined mismatch effect on the SNDR.

3.7.3 Interleaver Architectures

One of the most important design considerations in a TI-ADC is the interleaver architecture. It determines to a great extent the converter's analog bandwidth and sampling accuracy, as well as presents a considerable power/area overhead. Additionally, the calibration complexity is largely dependent on the interleaver, since it contributes a portion of the offset/gain mismatch spurious tones and it is the exclusive source of the timing/bandwidth mismatch tones. Therefore, the interleaver choice is a key point to be decided early on during the design procedure as it affects the overall performance of the TI-ADC.

Interleaver architectures found in literature can be classified into two primary categories: direct interleavers and hierarchical interleavers (Fig. 3.35). The fundamental distinction between these two categories is the number of samplers prior to the sub-ADCs. Direct interleavers sample the input signal in a single stage of N samplers, while hierarchical interleavers cascade the N samplers into multiple stages. Hierarchical interleavers can be further divided into two sub-categories: de-multiplexing interleavers and re-sampling or sub-sampling interleavers. These interleaver architectures are discussed in detail with their benefits and drawbacks. In addition, a model is introduced that compares them in terms of achievable bandwidth and sampling accuracy, and this comparison is extended to the different process nodes of Fig. 3.4. This analysis aims to serve as a foundation in determining the optimum interleaver for a given set of specifications.




Fig. 3.34 Simulated SNDR vs. (a) σ_{OS}/V_{DD} , (b) σ_G/G , (c) $\sigma_{\Delta T}/T_{s, TI}$, (d) σ_{BW}/BW , and (e) σ_{BW}/BW with separated gain/phase and (f) combined errors




Fig. 3.35 Classification tree for different interleaver architectures




Fig. 3.36 (a) Direct interleaver architecture and timing diagram for $N = 8$ (b) with 50% duty-cycle clocks and (c) with (1/8)-100% duty-cycle clocks

Direct Interleaver

Direct interleaving offers the simplest and most commonly adopted interleaving configuration [88, 91, 96, 104]. The analog input is connected to all the channels as shown in Fig. 3.36a and directly sampled on the sub-ADC sampling capacitors through a single stage of samplers. Each channel comprises its own sampler, clock, and sub-ADC and operates at a sample rate of f_s for a total sample rate of $N \cdot f_s$. One or more channels may be tracking the input at a time with a maximum of $N/2$. One option is to have $N/2$ channels tracking the input with a 50% duty-cycle clock (Fig. 3.36b, $N = 8$). This results in a reduced bandwidth due to multiple sampling capacitors loading the input. This load can be relaxed by having only one channel tracking the input with a $100/N\%$ duty-cycle clock (Fig. 3.36c, $N = 8$) in exchange for a shorter tracking time.

This interleaver necessitates the minimum amount of hardware, which results in a high energy efficiency. However, it may present significant challenges at very high sample rates. For a large channel count, the increased input capacitance of the samplers and the interconnect result in a bandwidth degradation. The calibration complexity is also maximum with both timing and bandwidth mismatches being present for all channels on top of the existing offset/gain mismatches that are exacerbated by the contribution of the parallel samplers. This interleaver is particularly attractive for a limited channel count ($N \leq 8$).

Hierarchical Interleaver: De-multiplexing

De-multiplexing (demux) hierarchy [105, 106] cascades several stages to sample the analog input on the sub-ADC sampling capacitors through series samplers. A




Fig. 3.37 (a) Interleaver architecture with a hierarchical $N = L \times K$ de-multiplexing and (b) timing diagram for $N = 8$ with $L \times K = 2 \times 4$

two-stage demux interleaver with L front-end samplers each branching out to K back-end channels for a total channel count of $N = L \times K$ is shown in Fig. 3.37a. The front-end samplers operate at a rate of $N \cdot f_s / L$ with the back-end channels running at f_s but with a N/K longer tracking time compared to the direct interleaver. This is illustrated in the $8 = 2 \times 4$ timing diagram of Fig. 3.37b.

A noteworthy benefit of this interleaver compared to its direct counterpart is the reduction of the timing/bandwidth mismatch critical samplers from N to L , which relaxes the clock generation and/or calibration overhead. In the special case of a single front-end master sampler ($L = 1$), the timing/bandwidth mismatch is theoretically eliminated, together with its calibration, if the output of the master sampler can settle sufficiently. The downside is that the master sampler has only half a period of the Master Clock (MC) for tracking ($1/2Nf_s$). Regarding the bandwidth, the resistance goes up by the number of cascaded stages, while the input capacitance is split in an $L + K$ fashion compared to the full capacitance of N samplers in the direct interleaver. For a large $N \geq 16$ and a proper $L \times K$ allocation, there can be a reduction in the capacitance to overcome the increase in resistance and yield an overall bandwidth improvement.

Hierarchical Interleaver: Re-sampling

Re-sampling (resamp) hierarchy [74, 107] demonstrates a method of partly hiding the total capacitance of the N channels from the input without increasing the resistance in the signal path. A two-stage resamp interleaver with L samplers in the front-end and K channels in each back-end branch is shown in Fig. 3.38a, while the timing diagram of a 2×4 example is shown in Fig. 3.38b. The analog input is




Fig. 3.38 (a) Interleaver architecture with $N = L \times K$ re-sampling hierarchy and (b) timing diagram for $N = 8$ with $L \times K = 2 \times 4$

connected to the L samplers, one of which samples it on its output capacitor. The sampled input is buffered and re-sampled on a sub-ADC sampling capacitor. When a front-end sampler is in track mode, its back-end branch is in hold mode, thus not increasing the critical resistance. The corresponding buffer tracks the input signal loaded only by the back-end samplers and not the sub-ADCs.

This interleaver merges the benefits of the minimum series resistance of direct interleavers with the reduced timing/bandwidth mismatch critical samplers of the demux. Hence, both the bandwidth is enhanced and the clock generation and calibration overhead are relaxed. Furthermore, the bandwidth is determined completely by the front-end stage and not the sub-ADC, making it possible to increase N (through K) without degrading the bandwidth or requiring more timing/bandwidth calibration. On the downside, the re-sampling process increases the noise due to the additional capacitor. To bring down the noise, the intermediate and sampling capacitors have to be upscaled, which can hinder the bandwidth benefit of this interleaver. Additionally, the use of buffers increases the total power, noise, and non-linearity of the converter. A way of alleviating this drawback is to remove the buffers altogether, resulting in the so-called passive re-sampling interleaver [107]. However, this approach suffers from signal attenuation due to charge redistribution between the front-end and back-end. Further, the back-end load is not hidden anymore from the input.

To compare the described interleavers in terms of input bandwidth and sampling accuracy, the equivalent RC model of Fig. 3.39a is constructed. It encompasses the sampling capacitor C_S and equivalent input resistance due to termination $R_{i,eq}$ as




Fig. 3.39 Equivalent RC model for (a) interleaver and (b) simple switch

well as the on-resistance and capacitance of the sampling switches. For a fixed C_S targeting a certain noise and a fixed $R_{i,eq}$ due to the termination network, they ultimately limit the bandwidth to its maximum theoretical value. The switches further reduce the bandwidth, and a simple model is shown in Fig. 3.39b. The on-resistance is denoted by R_{on} , and the on-capacitance on each side C_{on} is assumed equal to $C_{GG}/2$, where C_{GG} represents the total capacitance at the switch gate. The off-capacitance C_{off} , typically about three times smaller than C_{on} [21], is assumed $C_{on}/2$ to capture some interconnect overhead. The components of our RC model can be then expressed in terms of the different interleavers' parameters with the following expressions:

$$R_1 = R_{i,eq} = 0.5R_{i,int}, \quad C_1 = N_{on}C_{on1} + N_{off}C_{off1} \quad (3.57a)$$

$$R_2 = R_{on1}/N_{on}, \quad C_2 = N_{on}C_S + N_{on}C_{on1}.$$

$$R_1 = R_{i,eq} = 0.5R_{i,int}, \quad C_1 = L_{on}C_{on1} + L_{off}C_{off1}$$

$$R_2 = R_{on1}/L_{on}, \quad C_2 = L_{on}C_{on1} + L_{on}K_{on}C_{on2} + L_{on}K_{off}C_{off2} \quad (3.57b)$$

$$R_3 = R_{on2}/L_{on}, \quad C_3 = L_{on}C_S + L_{on}K_{on}C_{on2}.$$

$$R_1 = R_{i,eq} = 0.5R_{i,int}, \quad C_1 = L_{on}C_{on1} + L_{off}C_{off1} \quad (3.57c)$$

$$R_2 = R_{on1}/L_{on}, \quad C_2 = L_{on}\alpha C_S + L_{on}C_{on1}.$$

Equations (3.57a), (3.57b), and (3.57c) give the RC components for the direct, demux, and resamp interleavers, respectively. The upscaling factor α captures the necessary increase of the intermediate and the sub-ADC sampling capacitors to preserve the same SNR in the resamp interleaver. Also, for a device in triode region, R_{on} and C_{on} can be linked through the process f_T [21]

$$f_T = \frac{g_m}{2\pi C_{GG}} = \frac{1/R_{on}}{2\pi C_{GG}}. \quad (3.58)$$

Including all the above and incorporating the peak- f_T from Fig. 3.4 for the four different processes, the simulated bandwidth vs. channel count is shown in Fig. 3.40. A typical $50\ \Omega$ input termination is assumed, while C_S is sized at $200\ \text{fF}$ for an above 10-bit SNR under typical signal swings. Also, it is assumed that $\alpha = 2$ in Eq. (3.57c).

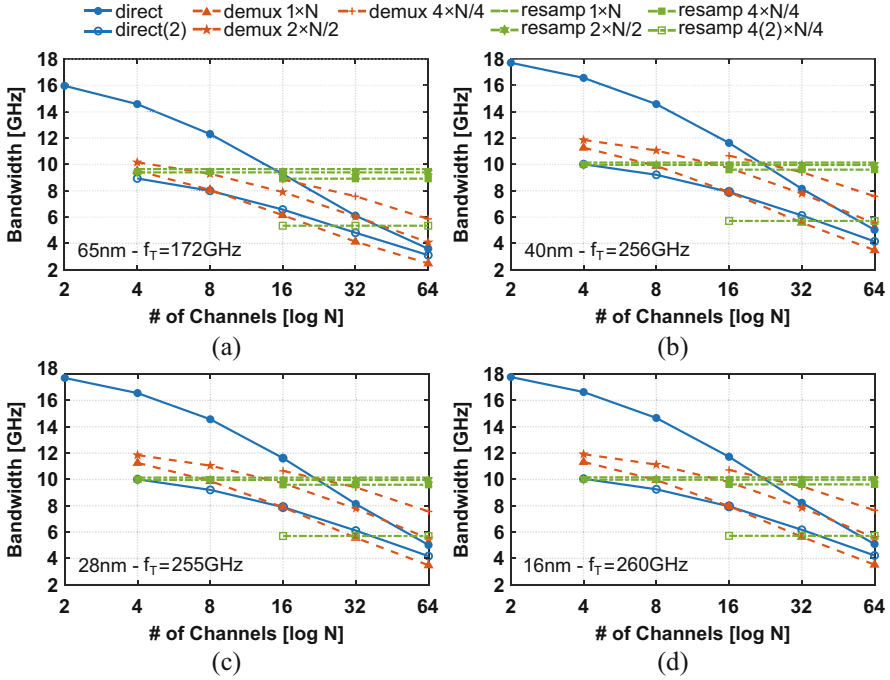


Fig. 3.40 Bandwidth vs. channel count for different interleavers in (a) 65 nm, (b) 40 nm, (c) 28 nm, and (d) 16 nm CMOS

An R_{on} of $10\ \Omega$ is adopted for all switches regardless of their stage, and C_{on} and C_{off} are determined as explained above for each process node. The interleavers under comparison include direct with one or two channels simultaneously tracking the input (denoted inside brackets); two-stage demux with one, two, or four front-end samplers; and resamp with the same settings and one or two front-end samplers on simultaneously. Architectures with more than two channels tracking the input as well as higher than two-stage hierarchy are not depicted since they did not show benefits up to $N=64$. It can be observed that due to their simplicity, direct interleavers with one channel tracking at a time achieve the highest bandwidth for $N \leq 8$, but their bandwidth rapidly drops for $N \geq 16$. With two channels tracking, the absolute bandwidth is smaller due to excess parasitic capacitance and the fact that in contrast to R_{on} , $R_{i,eq}$ is not halved to compensate for the double C_s . Demux interleavers cannot compete with the direct ones for $N \leq 8$ due to the increased resistance of the series switches, but their bandwidth degrades with a smoother slope for increasing N due to the hierarchical nature. From those architectures, the least beneficial are the ones with a front-end master sampler ($L=1$), which can be effectively seen as direct interleavers with an extra series resistance. As already discussed, the potential benefits of this architecture start showing up for $L \geq 4$ and $N \geq 16$, since the hierarchical capacitance splitting starts compensating

the increase in the resistance. The resamp interleavers for the assumed doubling in C_s to preserve the SNR ($\alpha = 2$)¹¹ start from a lower bandwidth compared to direct interleavers with one channel tracking, but they preserve the same bandwidth due to their architecturally scalable nature when increasing N by increasing only the back-end channels. As such, for $N \geq 16$, they start becoming competitive with the other architectures, and for $N \geq 32$, they are the most beneficial architecture in terms of achievable bandwidth. A similar relative bandwidth drop as for the direct interleaver is noticed for the case of two channels tracking.

Regarding the technology effect, Eq.(3.58) reveals that the $R_{on}C_{GG}$ product reduces for increasing f_T resulting in a bandwidth increase. Since the peak- f_T for 40 nm and below does not change much, the achievable bandwidth in these processes is comparable. However, when these three processes are compared with the 65 nm, there is a noticeable improvement due to larger peak- f_T . Something noteworthy is that the relative process benefit is larger for the direct and demux architectures due to a greater effect of the parasitics on their achievable bandwidth compared to the resamp architecture, as was hinted by Eqs. (3.57a)–(3.57c).

The other important specification determining an optimum interleaver is the achievable sampling accuracy in terms of bits for a given tracking time. The bandwidth of each interleaver from Fig. 3.40 can be used to estimate this accuracy. Assuming a first-order exponential settling, the achievable accuracy can be computed at a certain sample rate as follows:

$$B = 2\pi\theta_{\text{track}} \frac{T_s \cdot BW}{\ln 2} - 1, \quad \theta_{\text{track}} = 0.5, 1, 2, \quad (3.59)$$

where $T_s = 1/Nf_s$ and the factor θ_{track} captures the front-end tracking time depending on the channels tracking the input at the same time. For the direct interleaver with one channel tracking at a time and the hierarchical ones with $L = 1$, $\theta_{\text{track}} = 1$. When two channels are tracking at a time, $\theta_{\text{track}} = 2$, while for the front-end master sampler, $\theta_{\text{track}} = 0.5$. Including all the above, Fig. 3.41 plots the achievable sampling accuracy vs. channel count for the interleavers at four different sample rates. Since the different processes affect their relative sampling accuracy in the same way as the bandwidth, only 28 nm is shown.

Direct interleavers achieve the highest sampling accuracy for $N \leq 8$. Between the two versions considered, the one with two channels tracking shows a superior performance, despite its lower bandwidth. This is explained by Eq.(3.59), since between one and two channels tracking, θ_{track} increases by $2\times$, while the bandwidth drops by a smaller amount yielding an overall better *accuracy · speed*. Similar conclusions can be drawn for the resamp interleavers using one or two channels tracking. For $N = 16$, the direct interleavers with one channel tracking become

¹¹ Different upscaling ratios between the intermediate and the sub-ADC sampling capacitors can achieve higher bandwidths and preserve the SNR (e.g., 1.5–3). However, they might increase the buffer power for the back-end to sample with the same accuracy. The optimum ratio is subject to the specifics of each design; therefore, the generic 2 - 2 ratio is assumed here.

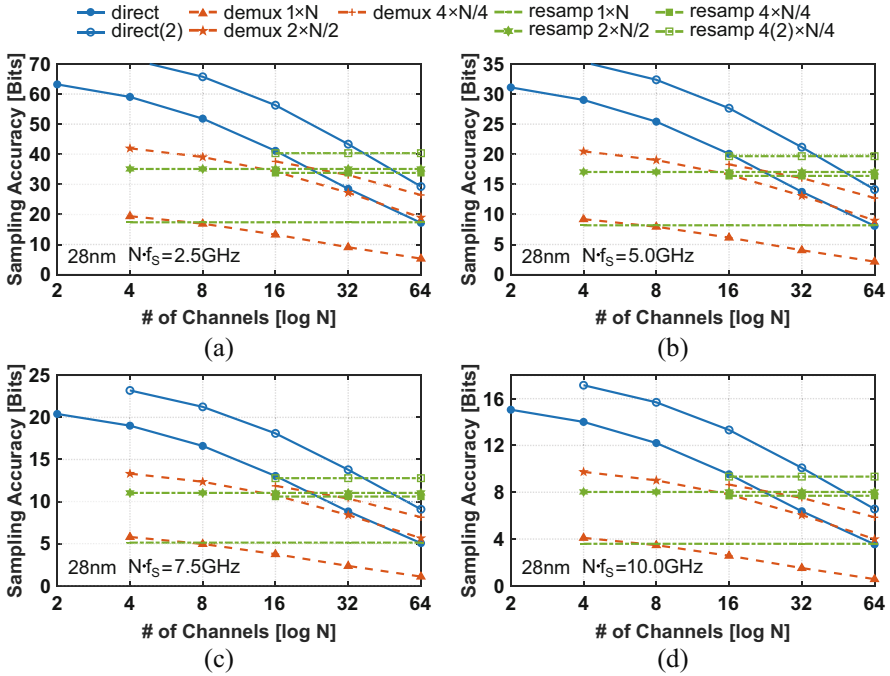


Fig. 3.41 Sampling accuracy vs. channel count for the interleavers in 28 nm at (a) $Nf_s = 2.5 \text{ GHz}$, (b) $Nf_s = 5.0 \text{ GHz}$, (c) $Nf_s = 7.5 \text{ GHz}$, and (d) $Nf_s = 10 \text{ GHz}$

comparable and eventually worse than the hierarchical ones, while this transition for two channels tracking happens for $N = 32$. From the hierarchical interleavers, the demux with $L \geq 2$ compete with the equivalent resamp ones up to $N = 16$. Beyond that point, the resamp with two channels tracking achieve the best accuracy. For both hierarchical architectures, the lowest accuracy is achieved when employing the master front-end sampler.

Based on the introduced interleaver model and analysis, some general guidelines regarding the optimum interleaver choice may be derived. For $N \leq 8$ and if the dynamic interleaving errors' calibration complexity is not prohibitive, direct interleavers should be preferred due to their simplicity, energy efficiency, as well as high bandwidth and *accuracy · speed*. Depending on the specifics of each design, the version with one or more channels tracking the input at the same time can be adopted. For $N \geq 16$ and in case a hierarchical architecture cannot be avoided, if the absolute performance is the priority, resamp interleavers should be the preferred choice due to their superior performance stemming from an architectural advantage. Demux interleavers could offer a more efficient hierarchical solution without the buffers, if the drop in bandwidth and sampling accuracy can be tolerated. For even larger N ($N > 64$), it might be even possible that a combination between the two can yield optimum results [105]. It is worth noting that in order to convincingly

reach to a final decision, the non-linearity due to the switch resistance and parasitic capacitance as well as that from the buffers should be carefully evaluated and depends on the process at hand.

3.8 Conclusion

This chapter built on the block-level investigation of Chap. 2 and conducted an investigation and comparison among different architectures, aiming to determine the optimal ADC architecture for maximizing $accuracy \cdot speed \div power$. This investigation started by examining the recent SotA standings, focusing on the highest-performance architectures, such as flash, SAR, pipeline, and pipelined-SAR, as well as their TI counterparts. These standings revealed that currently published pipeline/TI-pipeline ADCs are leading in accuracy with very good speed but not the best efficiency. On the other hand, current TI-SAR ADCs, highly benefiting from technology scaling, achieve the highest speed with very good efficiency, but their accuracy is inferior to the pipeline especially when pushing the speed. The pipelined-SAR hybrid ADCs, preserving the scaling benefits of SARs, are catching up to the pipelines in terms of accuracy and speed. Flash ADCs are surpassed in accuracy and efficiency by SARs, but they are still the fastest existing standalone converter architecture and an essential part of traditional pipelines.

To better understand the mechanics of the aforementioned ADC architectures, they were covered in detail with their trade-offs highlighted. This understanding introduced new mathematical models to quantitatively estimate and compare their accuracy-speed-power limits, offering a complete decomposition of the individual blocks' contributions. The architectures under comparison involved flash, binary SAR, 1,2,3,4-bit/stage pipeline, and 2,3,4,5-stage pipelined-SAR. These models' power was further enhanced by including technology parasitics and the process V_{DD} , f_T , C_{min} , and g_m/I_D . Four deep-scaled mainstream CMOS processes, namely, 65 nm, 40 nm, 28 nm, and 16 nm, were included.

In a $power/f_s$ vs. accuracy plot, at low sample rates, the slopes of the different ADC curves are first process-limited and then noise-limited. Above 40 dB SNDR, the flash has the worst efficiency. The pipelines show a better overall efficiency in their noise-limited regions due to sub-ranging. The SAR is very hard to beat up to about 45 dB, after which it starts becoming inefficient. The pipelined-SARs show the best efficiency from 45 dB to 65 dB. In the range of 70 dB and above, they are slightly more efficient than the multi-bit/stage pipelines. At intermediate sample rates, the SAR is still the most efficient up to 40 dB and close to the pipelines up to 50 dB. The two-stage pipelined-SAR stops 56 dB, while the 4-bit/stage pipeline and the three-stage pipelined-SAR stop at 80 dB. The 1,2-bit/stage pipelines as well as the 4,5-stage pipelined-SARs can achieve the highest resolution. At GHz-range sample rates, the SAR remains the most efficient architecture up to 40 dB, while the 1,2-bit/stage pipelines are the most efficient above 75 dB. The 4,5-stage pipelined-SARs are winning in the range 40 dB–65 dB and competing with the pipelines up to

75 dB, with the three-stage following this race up to 56 dB. This analysis revealed that increasing the stage count of the pipelined-SAR can potentially surpass the traditional pipeline in speed and efficiency for an extended range of resolutions, indicating a promising future research direction.

Time-interleaving was also covered as a popular way to extend a converter's sample rate beyond what is allowed by the process f_T . Interleaving errors due to mismatch, namely, offset, gain, timing, and bandwidth, were discussed, and their impact on accuracy was quantified through simulations. The main interleaver architectures, namely, direct, demux, and resamp, were dealt with as being one of the most important design considerations affecting the overall performance of TI-ADCs. In addition, a model was introduced to compare them in terms of achievable bandwidth and sampling accuracy providing insight in determining the optimum interleaver depending on the design. Direct interleavers achieve the best results for $N \leq 8$, while for $N \geq 16$ resamp are superior to demux in terms of accuracy and bandwidth but with a higher power due to the buffers.

Appendix B: Transconductance—Settled RA

Assume the settled RA in Fig. 3.17. Its time constant including the output parasitic loading is given by

$$\begin{aligned} \tau_{\text{RA,set}} &= r_{\text{oRA,set}}(C_{\text{RA}} + C_{\text{RA,par}}) \\ &= \frac{A_s}{g_{\text{mRA,set}}} C_{\text{RA}} + \frac{A_s}{\pi f_T}, \end{aligned} \quad [\text{s}] \quad (3.60)$$

where the second part is derived in a similar way as in Eq. (3.10). From the input-output characteristic of the RA, for a settling accuracy of $B+M$ at a percentage $\zeta_{\text{set}} \leq 1$ of the allocated operation time, and using the above expression for the time constant, we have

$$\zeta_{\text{set}} T_{\text{RA}} = (B + M) \cdot \ln 2 \cdot \left[\frac{A_s}{g_{\text{mRA,set}}} C_{\text{RA}} + \frac{A_s}{\pi f_T} \right]. \quad [\text{s}] \quad (3.61)$$

Multiplying both parts of the above expression and solving for $g_{\text{mRA,set}}$, we end up with the final expression

$$g_{\text{mRA,set}} = A_s \cdot \frac{f_{\text{RA}}}{\zeta_{\text{set}}} \cdot \frac{(B + M) \cdot \ln 2 \cdot C_{\text{RA}}}{1 - A_s \cdot \frac{(B+M) \cdot \ln 2}{\pi} \cdot \frac{f_{\text{RA}}}{\zeta_{\text{set}} \cdot f_T}}. \quad [\text{S}] \quad (3.62)$$

This expression matches the derived Eq. (3.31) for $f_{\text{RA}} = \theta_{\text{RA}} \cdot f_s$, where θ_{RA} captures the portion of the conversion allocated to the RA.

Appendix C: Transconductance—Integrator RA

Assume now that the RA in Fig. 3.17 has to operate in the unsettled $(g_m/C) \cdot T$ integrator mode. The transition from settled to integrator mode is indicated by the condition where the time constant matches the integration time

$$\zeta_{\text{int}} T_{\text{RA}} \leq \tau_{\text{RA,int}}, \quad [\text{s}] \quad (3.63)$$

where $\zeta_{\text{int}} \leq 1$ accounts for the integration occurring at a percentage of the total allocated RA time, while the remaining time is allocated to reset. Since the integrator is a dynamic circuit that needs reset, this is a reasonable assumption. The time constant is now expressed as follows:

$$\begin{aligned} \tau_{\text{RA,int}} &= r_{\text{oRA,int}}(C_{\text{RA}} + C_{\text{RA,par}}) \\ &= \frac{A_s}{g_{\text{mRA,int}}} C_{\text{RA}} + \frac{A_s}{\pi f_T}. \end{aligned} \quad [\text{s}] \quad (3.64)$$

To meet the integration mode condition, $r_{\text{oRA,int}}$ needs to be larger than $r_{\text{oRA,set}}$ by a factor N . In order to achieve the same gain as the settled mode for the same output load, $g_{\text{mRA,int}}$ should be reduced by the same factor N with respect to $g_{\text{mRA,set}}$. Combining the two above expressions, we get

$$\zeta_{\text{int}} T_{\text{RA}} = \left[\frac{A_s}{g_{\text{mRA,int}}} C_{\text{RA}} + \frac{A_s}{\pi f_T} \right]. \quad [\text{s}] \quad (3.65)$$

Similar to the settled mode, multiplying and solving the above expression for $g_{\text{mRA,int}}$, we end up with the final expression

$$g_{\text{mRA,int}} = A_s \cdot \frac{f_{\text{RA}}}{\zeta_{\text{int}}} \cdot \frac{C_{\text{RA}}}{1 - A_s \cdot \frac{1}{\pi} \cdot \frac{f_{\text{RA}}}{\zeta_{\text{int}} f_T}}. \quad [\text{S}] \quad (3.66)$$

Comparing Eqs. 3.62 and 3.66 and assuming a small f_{RA}/f_T for simplicity, the relation between the two transconductances becomes

$$\frac{g_{\text{mRA,set}}}{g_{\text{mRA,int}}} \approx (B + M) \cdot \frac{\zeta_{\text{int}}}{\zeta_{\text{set}}} = N \quad (3.67)$$

The basic MOS Eq. (2.45) leads to an equivalent relation between the currents. Note that in the integrator mode, the output is a function of time. To ensure the gain accuracy, the integration time must be precisely controlled by additional circuitry, which adds overhead and makes the benefit smaller than the predicted.

Chapter 4

Ultrahigh-Speed High-Sensitivity Dynamic Comparator



The significance of the comparator, both as a standalone block and within an ADC, was already pointed out in our analyses of Chaps. 2–3. Its speed determines the ADC sample rate either fully (flash) or to a great extent (SAR, pipeline, pipelined-SAR). For flash and SAR ADCs especially, its power consumption to achieve a certain noise and metastability significantly impacts the overall converter efficiency.

This chapter focuses on the design of dynamic regenerative comparators, which is currently the prevailing high-speed low-power choice. Section 4.1 highlights the importance of comparators in modern mixed-signal systems and reviews two widely used circuits. In Sect. 4.2, an ultrahigh-speed three-stage fully dynamic comparator is presented. The design and operation of the novel comparator are described, and its simulated performance is compared with the two widely used circuits. Section 4.3 discusses the experimental verification and provides a standing against state-of-the-art comparators. The two widely used circuits are included in the prototype IC as well, enabling a fair comparison. Finally, Sect. 4.4 draws the conclusion of this chapter.

Parts of this chapter were presented at the 2019 *European Solid-State Circuits Conference (ESSCIRC'19)* in Kraków, Poland, and concurrently published in the *Solid-State Circuits Letters (SSC-L'19)* in September 2019 [108]. An extended analysis and measurements on the prototype comparator and the two prior art circuits was also published in the *Transactions on Circuits and Systems I: Regular Papers (TCAS-I'22)* in September 2022 [109].

4.1 Dynamic Regenerative Comparators

Comparators are indispensable blocks in any mixed-signal system, from ADCs and wireline data links to memories and clock generation circuits. Their role is particularly critical in high-speed ADCs, where they have to extract the digital




Fig. 4.1 Single-stage strong-ARM comparator and its signal waveforms

representation of analog input signals with the utmost accuracy and speed while dissipating the minimum amount of power. Depending on the architecture, one or multiple comparators may be employed in an ADC, imposing a significant contributor on the total converter power. Dynamic regenerative comparators are established as the topology of choice. Their clocked nature with positive feedback offers a high-speed operation for a zero static power. Although many comparator variants exist, two (or their modifications) are widely adopted in ADCs: (1) the single-stage strong-ARM [110] and (2) the two-stage double-tail [111]. These two are analyzed next, prior to the proposed improved comparator.

4.1.1 Single-Stage Latch-Based Strong-ARM Comparator

The classic single-stage latch-based SAC [110] is shown in Fig. 4.1. This circuit has been largely used in virtue of its high-speed operation, rail-to-rail output, and high power efficiency. It includes a clocked differential pair M_{1P}/M_{1N} with a tail current source T_1 , a cross-coupled inverter pair M_{2P}/M_{2N} , M_{3P}/M_{3N} , and four reset devices M_{4P}/M_{4N} , M_{5P}/M_{5N} . The operation of this comparator can be described in four primary phases.

In the first phase of reset ($CLK = V_{SS} = 0\text{ V}$), T_1 is off, and M_{4P}/M_{4N} , M_{5P}/M_{5N} pull nodes $X+/X-$ and $O+/O-$, respectively, V_{DD} . The second phase, that of the primary integration, starts with $CLK = V_{DD}$. T_1 turns on, and the input pair devices M_{1P}/M_{1N} are in saturation. Nodes $X+/X-$ discharge through the input pair with slopes proportional to the differential current, which in turn depends on the differential input voltage ΔV_1 . The gain generated during this phase is given as

$$\frac{|\Delta V_X|}{|\Delta V_I|} = \frac{g_{m1} t_{int1}}{C_X}, \quad (4.1)$$

where g_{m1} is the transconductance of the input pair, C_X is the capacitance on each of nodes $X+/X-$, and t_{int1} is the primary integration time

$$t_{int1} = 2 \frac{C_X V_{TH}}{I_{T1}}. \quad [s] \quad (4.2)$$

In the above, the current on each side $I_{T1}/2 \pm g_{m1} \Delta V_I$ is approximated with $I_{T1}/2$, which is fair for small ΔV_I . The expression above indicates that the second phase lasts until the intermediate nodes $X+/X-$ drop one NMOS threshold voltage¹ below V_{DD} , initiating the third phase, namely, the secondary integration. M_{2P}/M_{2N} are on, and the output nodes $O+/O-$ discharge with a differential voltage starting from ΔV_X and providing further amplification. The built-up gain at the output during both these phases is then found as

$$\frac{|\Delta V_O|}{|\Delta V_I|} = \frac{g_{m1} (t_{int1} + t_{int2})}{C_O}, \quad (4.3)$$

where C_O is the capacitance on each of $O+/O-$ and t_{int2} is given by

$$t_{int2} = 2 \frac{C_O |V_{TH}|}{I_{T1}}. \quad [s] \quad (4.4)$$

Again, the above expression reveals that this third phase lasts until the output nodes $O+/O-$ drop one PMOS threshold voltage below V_{DD} , prior to initiating the fourth and final phase of the comparator. By combining Eqs. (4.2)–(4.4), the total integration time and the gain, respectively, can be approximated by the following two expressions:

$$t_{SAC,int} = 2 \frac{|V_{TH}|}{I_{T1}} \cdot (C_X + C_O) \approx 2 \frac{|V_{TH}|}{I_{T1}} C_O, \quad C_O \gg C_X. \quad [s] \quad (4.5)$$

$$\frac{|\Delta V_O|}{|\Delta V_I|} = 2g_{m1} \frac{|V_{TH}|}{I_{T1}} \cdot \left(\frac{C_X + C_O}{C_O} \right) \approx 2g_{m1} \frac{|V_{TH}|}{I_{T1}}, \quad C_O \gg C_X. \quad (4.6)$$

The simplifications above generally hold true since C_O includes both the input capacitance of any following circuit and the comparator self-load; hence, it can be significantly larger than C_X .

¹ Assuming $V_{TH,NMOS} = V_{TH} \approx |V_{TH,PMOS}| = |V_{TH}|$. This is fair for deep-scaled processes due to adjustments in the dopant implantation and mechanical stress bringing the NMOS and PMOS devices closer in terms of threshold voltage and mobility.

The final phase, when latching occurs, starts with turning on M_{3P}/M_{3N} . Due to the cross-coupled inverter pair, the output of the PMOS, of which the gate drops at a faster rate, is pulled to V_{DD} again, while the other output is pulled to V_{SS} . The delay of the circuit during this phase to deliver a clear logical level differential output (typically $V_{DD}/2$) arises from the exponential input-output latch characteristic (see Chap. 2, Sect. 2.4.3)

$$t_{\text{SAC,latch}} = \frac{C_O}{g_{\text{m,eff}}} \cdot \ln \left[\frac{V_{DD}/2}{|\Delta V_O|} \right]. \quad [\text{s}] \quad (4.7)$$

$g_{\text{m,eff}}$ is the effective transconductance of the latch devices. This primarily comprises g_{m3} since M_{3P}/M_{3N} are mainly responsible for starting the latching. Combining Eqs. (4.5)–(4.7), the total SAC delay with respect to ΔV_I is obtained

$$\begin{aligned} t_{\text{SAC,tot}} &= t_{\text{SAC,int}} + t_{\text{SAC,latch}} \\ &= 2 \frac{|V_{\text{TH}}|}{I_{\text{T1}}} C_O + \frac{C_O}{g_{\text{m,eff}}} \cdot \ln \left[\frac{V_{DD}(I_{\text{T1}}/g_{\text{m1}})}{4|V_{\text{TH}}||\Delta V_I|} \right]. \end{aligned} \quad [\text{s}] \quad (4.8)$$

The above expression reveals the dependency of the comparator delay on several parameters. The output capacitance C_O is present in both terms of Eq. (4.8), therefore directly influencing the total delay. The differential input also influences the delay. A smaller ΔV_I leads to a larger total delay, but depending on the input amplitude, one of the two terms may dominate over the other. The supply voltage also affects the delay. A larger V_{DD} increases I_{T1} , which reduces $t_{\text{SAC,int}}$. This creates a smaller initial difference for the latch to regenerate on. On the other hand, $g_{\text{m,eff}}$ increases, whose effect on $t_{\text{SAC,latch}}$ is stronger than ΔV_O (linear vs. logarithmic). Therefore, $t_{\text{SAC,tot}}$ overall reduces by increasing V_{DD} . Finally, the dependency of the delay on the input common-mode V_{CM} is implied. Increasing V_{CM} reduces $t_{\text{SAC,int}}$ by increasing the branch current ($\approx I_{\text{T1}}/2$). However, this also reduces ΔV_O on which the latch regenerates, hence somewhat increasing $t_{\text{SAC,latch}}$. Overall, $t_{\text{SAC,tot}}$ reduces by increasing V_{CM} up to a certain voltage, after which the two aforementioned opposite effects cancel each other and there is no more delay benefit in increasing V_{CM} further [58].

Besides the delay, which is a key specification for enabling higher-speed systems, other metrics, such as noise, offset, and kickback, are also important. Considering noise, in-depth and excellent analyses were carried out in [22, 112]. Here we limit the discussion to the contribution of the input pair M_{1P}/M_{1N} , whose noise typically dominates provided there is gain during the primary and secondary integration phases to suppress the noise from the latch devices. The output noise variance is expressed as the convolution of the white noise power spectral density (PSD), S_i , and the magnitude squared impulse response from the input noise source to the output. At the end of t_{intl} , the noise power at $X+/X-$ is

$$\overline{u_{\text{SAC,opre}}^2} = \frac{S_i}{2} \int_0^{t_{\text{intl}}} \frac{1}{C_X^2} dt, \quad S_i = 8kT\gamma g_{\text{m1}}. \quad [\text{V}^2] \quad (4.9)$$

Substituting t_{int1} from Eq. (4.2) and dividing $\overline{v_{\text{SAC,opre}}^2}$ by the gain squared from Eq. (4.1), the input-referred noise can be expressed as

$$\overline{v_{\text{SAC,inpre}}^2} = 2 \frac{kT\gamma(I_{T1}/g_{m1})}{C_X V_{\text{TH}}} \quad [\text{V}^2] \quad (4.10)$$

Besides the inverse proportionality between noise and C_X , Eq. (4.10) reveals the dependency of V_{CM} (through g_{m1}/I_{T1}) on the total comparator noise. Increasing V_{CM} increases the bandwidth; hence, noise is integrated over a larger bandwidth. Further, reducing $t_{\text{SAC,int}}$ reduces the gain during this phase, increasing the contribution of the latch when referred to the input. The same can be said about the input-referred offset. The choice of V_{CM} indicates a trade-off between delay and noise/offset, a direct limitation of the single current path of the SAC, dictating both the integration and latching conditions. A small V_{CM} (small I_{T1}) is desirable to generate a high gain during integration, while a large V_{CM} (large I_{T1}) is needed to enable a fast latching. This is non-ideal in systems with a wide V_{CM} range. Further, the SAC stacks several devices requiring at least two threshold voltages plus two overdrive voltages for the stacked devices to overlap in saturation. The constant supply reduction in finer CMOS processes makes it challenging to provide a high integration gain, to minimize the latch delay, and to suppress its noise/offset while also maximizing $g_{m,\text{eff}}$. Finally, stacking the latch and the input pair makes this topology susceptible to latch kickback [52]. Differential variations on $X+/X-$ and $O+/O-$ couple to the inputs through C_{GS} and C_{GD} of $M_{3\text{P}}/M_{3\text{N}}$ and $M_{1\text{P}}/M_{1\text{N}}$ and may lead to evaluation errors.

4.1.2 Two-Stage Double-Tail Latched Comparator

To overcome the limitations of the single current controlling both the integration and the latching as well as reduce device stacking, the split input-latch DTC of Fig. 4.2 was proposed in [111]. The reduced stacking allows this topology to operate properly at decreased supply voltages. The pre-amplifier stage includes a clocked differential pair $M_{1\text{P}}/M_{1\text{N}}$ with reset devices $M_{2\text{P}}/M_{2\text{N}}$ and a tail current source T_1 controlling the integration. The latch contains a cross-coupled inverter pair $M_{4\text{P}}/M_{4\text{N}}$, $M_{5\text{P}}/M_{5\text{N}}$, and a tail current source T_2 controlling the latching. The intermediate $M_{3\text{P}}/M_{3\text{N}}$ amplify ΔV_X while transferring it to the latching stage.

The operation of this comparator can be described in the following phases. During the reset phase ($\text{CLK} = V_{\text{SS}}$, $\overline{\text{CLK}} = V_{\text{DD}}$), T_1 , T_2 are off and $M_{2\text{P}}/M_{2\text{N}}$, $M_{3\text{P}}/M_{3\text{N}}$ pull nodes $X+/X-$ and $O+/O-$ to V_{DD} and V_{SS} , respectively. After reset, when $\text{CLK} = V_{\text{DD}}$, $\overline{\text{CLK}} = V_{\text{SS}}$, nodes $X+/X-$ discharge through the input pair with slopes are proportional to the differential current in the pre-amplifier. Almost simultaneously, the output nodes charge with proportional slopes to the differential latch current, and $M_{3\text{P}}/M_{3\text{N}}$ transfer ΔV_X to the latch while providing




Fig. 4.2 Double-tail comparator and its signal waveforms

further amplification. Assuming M_{1P}/M_{1N} and M_{3P}/M_{3N} are largely overlapping cascaded integrators, the amplification toward the latch is

$$\frac{|\Delta V_O|}{|\Delta V_I|} = \int_0^{t_{\text{DTC,int}}} \frac{g_{m1}}{C_X} t \cdot \frac{g_{m3}}{C_O} dt \approx \frac{g_{m1}g_{m3}}{C_X C_O} \cdot t_{\text{DTC,int}}^2, \quad (4.11)$$

where $t_{\text{DTC,int}}$ is the time, during which $O+/O-$ are integrated on C_O until one of them rises one NMOS threshold voltage above V_{SS} , for M_{4P}/M_{4N} to start the exponential latching

$$t_{\text{DTC,int}} = 2 \frac{V_{\text{TH}}}{I_{T2}} C_O \approx 2 \frac{V_{X,\text{DC}}}{I_{T1}} C_X. \quad [\text{s}] \quad (4.12)$$

The term $V_{X,\text{DC}}$ represents the voltage drop on $X+/X-$ such that the intermediate devices are no longer able to clamp the outputs to V_{SS} . In both stages, the current on each side is approximated as $I_{T1}/2$ and $I_{T2}/2$, which is fair for small enough ΔV_I and ΔV_X , respectively. The above expressions reveal the influence of some parameters on the pre-amplifier and intermediate devices' gain, prior to the latch regeneration. To maximize the total gain, g_{m1} and g_{m3} should be maximized. This also reduces the noise/offset contribution of the latch devices referred to the input. Finally, a higher gain improves the latch delay as well, which regenerates on a larger initial ΔV_O . The delay during the latching phase is given by Eq. (4.7) as for the SAC. Combining Eqs. (4.7), (4.8), (4.11), and (4.12), the total DTC delay with respect to its differential input voltage ($t_{\text{DTC,int}} + t_{\text{DTC,latch}}$) is obtained

$$\begin{aligned}
t_{\text{DTC,tot}} &= t_{\text{DTC,int}} + t_{\text{DTC,latch}} \\
&= 2 \frac{V_{\text{TH}}}{I_{\text{T2}}} C_{\text{O}} + \frac{C_{\text{O}}}{g_{\text{m,eff}}} \ln \left[\frac{V_{\text{DD}}(I_{\text{T2}}^2/g_{\text{m1}}g_{\text{m3}})}{8V_{\text{TH}}^2|\Delta V_{\text{I}}|(C_{\text{O}}/C_{\text{X}})} \right]. \quad [\text{s}] \quad (4.13)
\end{aligned}$$

To estimate the DTC noise, we consider the noise developed at the pre-amplifier outputs during an initial time $\alpha t_{\text{DTC,int}}$ ($\alpha < 1$), where the input pair $M_{1\text{P}}/M_{1\text{N}}$ only is in saturation with $M_{3\text{P}}/M_{3\text{N}}$ still in triode. Following the same approach as for the SAC, the noise power on nodes $X+/X-$ is

$$\overline{v_{\text{DTC,opre}}^2} = \frac{S_{\text{i}}}{2} \int_0^{\alpha t_{\text{DTC,int}}} \frac{1}{C_{\text{X}}^2} dt, \quad S_{\text{i}} = 8kT\gamma g_{\text{m1}}. \quad [\text{V}^2] \quad (4.14)$$

The gain during $\alpha t_{\text{DTC,int}}$ is found similarly to Eq. (4.1) for the SAC

$$\frac{|\Delta V_{\text{X}}|}{|\Delta V_{\text{I}}|} = \frac{g_{\text{m1}}\alpha t_{\text{DTC,int}}}{C_{\text{X}}}, \quad (4.15)$$

where $\alpha t_{\text{DTC,int}}$ can be written as

$$\alpha t_{\text{DTC,int}} = 2 \frac{C_{\text{X}} V_{\text{X,DCpre}}}{I_{\text{T1}}}, \quad [\text{s}] \quad (4.16)$$

and $V_{\text{X,DCpre}}$ is the voltage drop on $X+/X-$ during this initial $\alpha t_{\text{DTC,int}}$. Following the same approach as for the SAC, and dividing $\overline{v_{\text{DTC,opre}}^2}$ by the above gain squared, leads to the input-referred noise

$$\overline{v_{\text{DTC,inpre}}^2} = 2 \frac{kT\gamma(I_{\text{T1}}/g_{\text{m1}})}{C_{\text{X}} V_{\text{X,DCpre}}}, \quad [\text{V}^2] \quad (4.17)$$

where $V_{\text{X,DCpre}}$ is also assumed equal to V_{TH} , which marks the $M_{3\text{P}}/M_{3\text{N}}$ transition from triode to saturation region.

As evident by the above-derived expressions, the DTC exhibits an extra design degree of freedom by virtue of the two separate tail currents, as seen in the second term of Eq. (4.13). I_{T1} can be optimized for a certain pre-amplifier noise/offset and a gain such that the latch contributions are sufficiently suppressed. On the other hand, I_{T2} can be optimized for a fast latch regeneration, showing less dependency on the V_{CM} and allowing a correct operation for a wider common-mode range. A necessary condition to guarantee noise/offset suppression is to ensure that the input pair $M_{1\text{P}}/M_{1\text{N}}$ does not enter triode too early prior to latching. The splitting of the input and latch stages also reduces the required headroom to two threshold voltages and one overdrive voltage (one less than the SAC), making this topology more suitable to lower supply voltages.

A potential drawback of the DTC versus the SAC is the lack of cascode devices in the pre-amplifier. Given the cascode devices remain in saturation for a meaningful period of time, they provide further amplification prior to the latching phase, hence

further reducing the output noise when referred to the input. The cross-coupled cascodes, apart from the extra isolation, discharge $X+/X-$ to about V_{TH} rather than a full discharge to V_{SS} . This reduces somewhat the power consumption, considering that these nodes need to fully charge back to V_{DD} during reset. Further, by keeping these nodes to about one V_{TH} would allow the intermediate M_{3P}/M_{3N} to slightly enhance the $g_{m,eff}$, which now contains primarily g_{m4} . Based on Eqs. (4.8) and (4.13), it would be highly desirable to minimize the comparator delay by enhancing both $g_{m,eff}$ and ΔV_O prior to latching. It would be ideal to achieve these two without an equivalent increase in C_O or power consumption.

4.2 Prototype IC: A 28 nm CMOS Three-Stage Triple-Latch Feed-Forward Comparator

This section presents a three-stage Triple-Latch Feed-Forward (TLFF) fully dynamic comparator, with an achievable data rate of 13.5 Gb/s and a BER of less than 10^{-12} for differential inputs as small as 5 mV [108, 109]. The combination of a high-gain three-stage configuration and an extra parallel feed-forward path results in a maximum CLK-OUT delay of about 27 ps ($V_{CM} = 0.5$ V, $V_{DD} = 1$ V) and a delay vs. $\log(\Delta V_I)$ slope of -6.4 ps/decade within an input range of 5 and 50 mV. Additionally, the cascaded triple-latch arrangement with reduced device stacking enables a delay of less than 70 ps down to 0.6 V V_{DD} and across a wide range of V_{CM} . The prototype comparator is fabricated in 28 nm bulk CMOS. It occupies a core area of $78 \mu\text{m}^2$ while dissipating 2.2 mW (including its output inverters) from a 1 V supply at a full speed (13.5 Gb/s) and a minimum differential input (5 mV). The two already analyzed SAC and DTC comparators are fabricated on the same die and compared against the proposed TLFF with delay, input-referred noise, energy/comparison, and area measurements, highlighting the benefits and trade-offs of the proposed solution.

4.2.1 Circuit Operation and Analysis

The circuit schematic of the proposed TLFF fully dynamic comparator is shown in Fig. 4.3. It constitutes a stage-1 amplifier M_{1P}/M_{1N} with a cascoded NMOS half-latch M_{2P}/M_{2N} , followed by a stage-2 amplifier M_{12P}/M_{12N} with a PMOS half-latch M_{4P}/M_{4N} . Both these stages drive the final stage NMOS latch M_{5P}/M_{5N} through a parallel direct path M_{23P}/M_{23N} and a feed-forward path M_{13P}/M_{13N} . The multi-stage nature with cascaded latches and the parallel direct/feed-forward paths enable a very high total gain. Further, they allow for a separate optimization of each stage, significantly reducing the total delay, the noise/offset, and sensitivity to V_{CM} altogether. Stage-1 is designed to provide a certain noise/offset and a relatively high

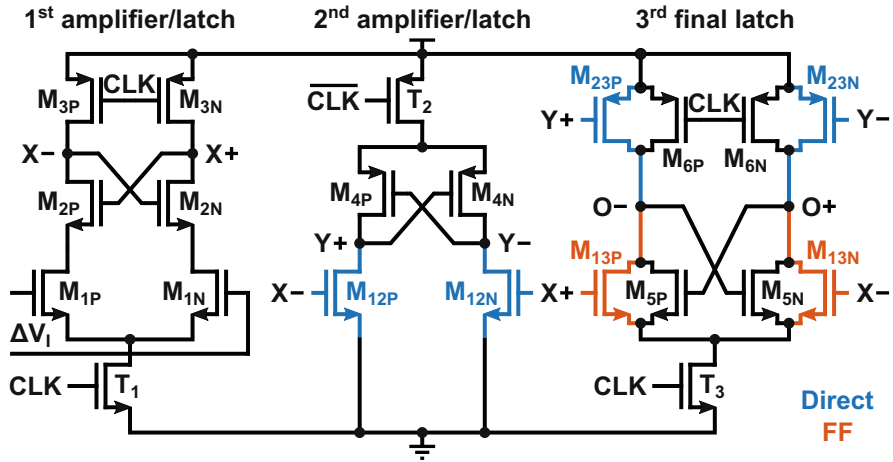


Fig. 4.3 Proposed three-stage TLFF dynamic comparator

gain to sufficiently attenuate the noise/offset and kickback of the following stages. Stage-3 is optimized for a minimum regeneration time given its total load (following circuit and self-load). Stage-2 serves as an extra gain stage increasing its differential output in an exponential fashion prior to the final latch. By horizontally cascading instead of vertically stacking the cross-coupled pairs, the required headroom of this topology reduces by at least one threshold voltage compared to the SAC and the DTC, reducing its minimum V_{DD} for a proper operation by the same amount.

The operation of the proposed TLFF comparator involves several phases. To properly analyze its behavior, the Linear Time-Variant (LTV) model of Fig. 4.4 is constructed. During the reset phase ($CLK = V_{SS}$, $\overline{CLK} = V_{DD}$), all stages are off, and nodes $X+/X-$ and $O+/O-$ are pulled to V_{DD} , while intermediate nodes $Y+/Y-$ are discharged to V_{SS} . The intermediate M_{12P}/M_{12N} , M_{23P}/M_{23N} perform a dual role, namely, both as reset devices and as transconductors. This alleviates the need for additional explicit reset devices, which would increase the capacitive load on these critical nodes. However, due to the very high data rate, the near minimum-sized M_{6P}/M_{6N} are still added, gated by the CLK (Fig. 4.3), to speed up reset and minimize any memory effect. When employed in a SAR ADC, the comparator reset is equally important to its decision, as either of them can dominate the bit cycle [113].

The integration phase starts by turning on the three tails, T_1 , T_2 , and T_3 , simultaneously through the symmetrical CLK/\overline{CLK} ($CLK = V_{DD}$, $\overline{CLK} = V_{SS}$). Concurrently, the cascaded stage-2/stage-3 PMOS/NMOS cross-coupled pairs also turn on together with a large initial overdrive voltage of about $V_{DD} - |V_{TH}|$. This is a key speed-boosting feature of the proposed TLFF compared to both the SAC and the DTC, where their cross-coupled pairs turn on sequentially rather than concurrently. Nodes $X+/X-$ discharge with slopes proportional to the differential current, and a differential voltage starts building up at the stage-1 outputs, similar to Eq. 4.11

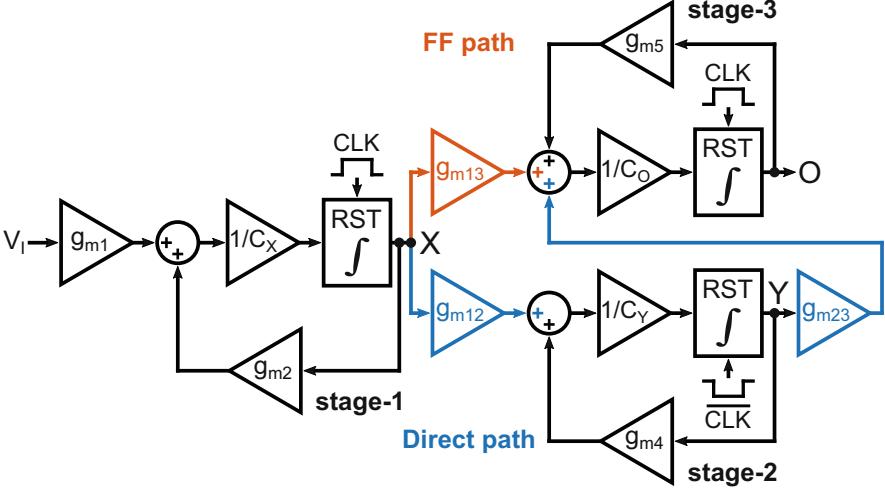


Fig. 4.4 LTV representation of the proposed TLFF comparator

from the DTC. The initial delay t_{int1} is the time it takes for $X+/X-$ to obtain a sufficient voltage drop $V_{X,DC}$, such that devices M_{12P}/M_{12N} and feed-forward M_{13P}/M_{13N} are no longer able to keep the intermediate $Y+/Y-$ and the outputs $O+/O-$, respectively, at identical voltage levels. To simplify the derived equations, this time is assumed to largely overlap with the time t_{int2} needed for $Y+/Y-$ to sufficiently rise by $V_{Y,DC}$, for devices M_{23P}/M_{23N} to not be able to clamp $O+/O-$ to V_{DD} , and for the time t_{int3} , during which a voltage drop of $V_{O,DC}$ is built on $O+/O-$

$$t_{int1} = 2 \frac{V_{X,DC}}{I_{T1}} C_X, \quad t_{int2} = 2 \frac{V_{Y,DC}}{I_{T2}} C_Y, \quad t_{int3} = 2 \frac{V_{O,DC}}{I_{T3}} C_O. \quad [s] \quad (4.18)$$

The initial ΔV_X , undergoing further amplification by M_{12P}/M_{12N} , produces a differential voltage ΔV_Y at the stage-2 outputs, on which the PMOS cross-coupled pair M_{4P}/M_{4N} starts regenerating according to the following:

$$\frac{|\Delta V_Y|}{|\Delta V_X|} = 2 \frac{V_{Y,DC}}{I_{T2}} g_{m12} \cdot e^{\frac{g_{m4}}{C_Y} \cdot t_{int2}}. \quad (4.19)$$

The multi-stage nature of the TLFF and the simultaneous turning on of all stages with large overdrive on the devices make t_{int1} and t_{int2} as well as the latching to partially overlap. Hence, the above expression could be already considered the onset of the latching phase. Furthermore, the parallel direct/feed-forward paths make the integration time of this comparator more dependent on the ΔV_I compared to the prior art

$$\{t_{\text{int}2}, t_{\text{int}3}\} \leq t_{\text{TLFF,int}} \leq \{t_{\text{int}2} + t_{\text{int}3}\}. \quad [\text{s}] \quad (4.20)$$

ΔV_Y is transferred to the final stage outputs by the direct M_{23P}/M_{23N} , where it is combined with the transferred ΔV_X by the feed-forward M_{13P}/M_{13N} , producing ΔV_O , prior to the final latching by the NMOS cross-coupled pair M_{5P}/M_{5N} . The operation is a cascaded M_{1P}/M_{1N} – M_{12P}/M_{12N} integration prior to M_{4P}/M_{4N} starting latching combined with a cascaded M_{1P}/M_{1N} – M_{13P}/M_{13N} integration prior to M_{5P}/M_{5N} completing the latching. Combining Eqs. (4.1) and (4.18)–(4.20), the partially regenerated ΔV_O is obtained as follows:

$$|\Delta V_O| = \left[8 \frac{V_{Y,\text{DC}}^2 \cdot V_{O,\text{DC}}}{I_{T2}^2 I_{T3} / (g_{m1} g_{m12} g_{m23})} \cdot \frac{C_Y}{C_X} \cdot e^{\frac{g_{m4}}{C_Y} \cdot t_{\text{int}2}} + 4 \frac{V_{O,\text{DC}}^2}{I_{T3}^2 / (g_{m1} g_{m13})} \cdot \frac{C_O}{C_X} \right] \cdot |\Delta V_I|. \quad (4.21)$$

In the first term of the above expression, $t_{\text{int}2}^2$ and $t_{\text{int}3}$ are included for the integration, while in the second term, $t_{\text{int}3}^2$ is included. Keeping Eqs. (4.20) and (4.21) in mind, the total TLFF delay ($t_{\text{TLFF,int}} + t_{\text{TLFF,latch}}$) is

$$\begin{aligned} t_{\text{TLFF,tot}} &= t_{\text{TLFF,int}} + t_{\text{TLFF,latch}} \\ &= t_{\text{TLFF,int}} + \frac{C_O}{g_{m13} + g_{m23} + g_{m5}} \cdot \ln \left[\frac{V_{\text{DD}}/2}{|\Delta V_O|} \right]. \end{aligned} \quad [\text{s}] \quad (4.22)$$

The above expressions capture to a first-order and highlight the major speed enhancements of the TLFF comparator over its aforementioned prior art circuits. The larger ΔV_O due to the larger gain in the signal path significantly reduces $t_{\text{TLFF,latch}}$. This gain stems from a longer $t_{\text{TLFF,int}}$ (see the right side of Eq. (4.20)), having the devices in all stages offering their maximum transconductances. Since $t_{\text{TLFF,latch}}$ dominates the total delay for smaller ΔV_I [58], this proves extremely beneficial. For larger ΔV_I , the longer dominant integration time of a three-stage comparator would render its total delay worse compared to a single-stage or a two-stage equivalent. This is because the devices quickly enter the triode region; thus, they behave more as digital delay gates rather than analog transconductors. The feed-forward path combined with the stage-1 cross-coupled cascodes (Fig. 4.3) mitigates this problem. One of M_{13P}/M_{13N} stays on since one of $X+/X-$ stays high enough, thanks to cross-coupling. Hence, latching already starts with minimum $t_{\text{TLFF,int}}$ (see the left side of Eq. (4.20)), without waiting for the intermediate devices' delay. The importance of the simultaneously on cascaded triple-latch configuration is also seen in the effective latch transconductance increase. It now comprises g_{m5} , g_{m23} , and g_{m13} , effectively minimizing the latch time constant. In both the SAC and the DTC, this time constant was primarily determined by the g_m of a single device. Although more stages draw current from the supply, a lower supply can be tolerated. Also, the stage-1 outputs stay at about one V_{TH} instead of a full V_{SS} discharge. Finally, it is worth noting that to get the above gain and delay

benefits, it is essential to guarantee the proper timings of each stage's operation. Given the freedom of the three stages, this is not hard, by properly dimensioning the stages and choosing the right V_{TH} transistor flavors.

In terms of noise, although the TLFF includes more devices than both the SAC and the DTC, the dominant contribution in a proper design is still the input pair M_{1P}/M_{1N} noise integrated on $X+/X-$ before transferring ΔV_X to the direct/feed-forward paths. The noise power at the stage-1 outputs during an initial time $\beta t_{TLFF,int}$ ($\beta < 1$), where only the input pair M_{1P}/M_{1N} develops gain, is given by

$$\overline{v_{TLFF,opre}^2} = \frac{S_i}{2} \int_0^{\beta t_{TLFF,int}} \frac{1}{C_X^2} dt, \quad S_i = 8kT\gamma g_{m1}. \quad [V^2] \quad (4.23)$$

With $\beta t_{TLFF,int}$ and the gain during this time found as in Eqs. (4.15) and (4.16), dividing by the gain squared, the input-referred noise is found as

$$\overline{v_{TLFF,inpre}^2} = 2 \frac{kT\gamma (I_{T1}/g_{m1})}{C_X V_{X,DCpre}}, \quad [V^2] \quad (4.24)$$

where $V_{X,DCpre}$ is again assumed equal to V_{TH} , as for the DTC. Comparing Eqs. (4.9), (4.10), (4.14), (4.17), (4.23), and (4.24), for the same I_{T1}/g_{m1} and C_X , the dominant noise contribution for all the comparators is roughly the same. In reality, the TLFF and the DTC may exhibit slightly higher noise since, prior to regeneration, this noise propagates to the latching nodes through cascaded integrators without charge transfer. In the SAC, the noise propagation resembles a single integrator. However, both the TLFF and the DTC offer the freedom to overcome this by properly designing the g_m/I and capacitances of their other stages with minimum to zero speed loss, a key limitation of the SAC.

A fundamental speed limitation of latched comparators is metastability. This phenomenon refers to the situation where the input difference is so small that for its allowed evaluation time, the latch cannot produce a sufficiently large differential output for the following circuit to unambiguously perceive as a clear logical level. In the absence of noise, metastability is a deterministic phenomenon that occurs with a probability P_{meta} equal to 1 when the input difference of the latch falls within a certain interval ΔV_{meta} and 0 outside this interval

$$P_{meta}(\Delta V_O, t_{latch}) = \begin{cases} 1, & |\Delta V_O| \leq |\Delta V_{meta}| \\ 0, & |\Delta V_O| > |\Delta V_{meta}| \end{cases}, \quad (4.25)$$

where ΔV_O is the initial latch input difference prior to regeneration and how much ΔV_{meta} is considered small enough depends on the latch available time t_{latch} . Writing the second term of any of Eqs. (4.8), (4.13), and (4.22) in its most generic form, the exponential input-output latch characteristic is given as

$$|\Delta V_{OUT}| = |\Delta V_O| \cdot e^{\frac{g_{m,eff} \cdot t_{latch}}{C_O}} = A |\Delta V_I| \cdot e^{\frac{g_{m,eff} \cdot t_{latch}}{C_O}}. \quad (4.26)$$

A is the gain due to any pre-amplification prior to the latch, and ΔV_{OUT} is the $V_{DD}/2$ acceptable output difference. Assuming a uniformly distributed comparator input difference across V_{DD} , the error probability or the bit error rate (BER) due to metastability can be computed by employing Eq. (4.26)

$$BER = \frac{|\Delta V_I|}{V_{DD}} = \frac{1}{2A} \cdot e^{-\frac{g_{m,eff} \cdot t_{latch}}{C_O}}. \quad (4.27)$$

Assuming a fixed t_{latch} , C_O , and ΔV_I , two important design parameters in minimizing BER are $g_{m,eff}$ and A , the first with an exponential dependency and the second in a linear manner. Comparing the above generic expression with the second term of each of Eqs. (4.8), (4.13), and (4.22) showcases the benefits of the proposed TLFF in achieving a superior BER by means of its larger signal gain and latch $g_{m,eff}$. It is worth mentioning that t_{latch} is not always identical for the three comparators, as this is shared with t_{int} . Also, the different phases in the operation of each comparator are self-timed and dependent on the input common-mode and the capacitances of the intermediate nodes. However, for small input differences (in the mV range and below), t_{latch} mainly dominates t_{tot} , hence assumed to a first order fixed and equal for all comparators.

In the presence of thermal noise, the input difference of the latch becomes a random Gaussian distributed variable with $\mu_{\Delta V_O}$ mean and $v_{\Delta V_O}^2$ variance. Noise turns metastability into a statistical (rather than a deterministic) phenomenon. It reduces the probability of occurrence for input differences within the initial ΔV_{meta} interval and increases the probability of occurrence for input differences outside this interval by roughly the same amount [114]. Overall, noise does not alter the total area of the metastability probability distribution function but rather shapes it from a uniform to a normal distribution. Additionally, as was already discussed and depicted in Eqs. (4.10), (4.17), and (4.24), the input-referred noise for all three comparators is roughly the same and mainly comprises the noise of the differential input pair. Although the DTC and the proposed TLFF include more devices and stages than the SAC, the noise of these extra stages is progressively attenuated by the total gain squared along the chain, when referred to the input. Therefore, the validity of Eqs. (4.26) and (4.27) in comparing the BER for a given speed (or the speed for a given BER) between the three presented comparators is not impacted by the presence of thermal noise. All the aforementioned speed benefits of the proposed TLFF are still preserved. This is also verified in Sect. 4.3.2 later in this chapter.

4.2.2 Simulation and Comparison with Prior Art

The developed theory on the proposed TLFF is verified with post-extracted simulations. In the timing waveforms of Fig. 4.5, the comparator delay is evaluated together with the different stages' outputs, in consecutive cycles of a 13.5 GHz clock. The overdrive recovery test settings are applied [115]. Within four consec-




Fig. 4.5 Simulated timing waveforms of the proposed TLFF comparator

utive cycles, the differential input switches between a rail-to-rail and a very small voltage with opposite polarity and then between a rail-to-rail and a very small voltage with the same polarity. Such a combination of differential inputs evaluates both the delay and the memory effect of the comparator. As depicted in Fig. 4.5, at a 13.5 GHz clock, the TLFF is able to overcome these extreme switching settings and deliver memory-free digital outputs for a ΔV_I as low as 5 mV. For a small ΔV_I , all stages contribute their maximum gain toward the final latch, with the stage-2 having already started the exponential regeneration. For a large ΔV_I , stage-1 has regenerated significantly, and thanks to the feed-forward path, the final outputs reach the rails almost at the same time as the stage-2 outputs. Finally, it is seen that, thanks to the triple-latch topology, the common-mode voltage of stage-1/stage-2 outputs does not go through a full V_{SS}/V_{DD} discharge/charge but stops at about one threshold voltage. Since these nodes need to go to the opposite rail during reset, the energy of these stages is reduced from $2 \cdot C \cdot V_{DD}^2$ to about $2 \cdot C \cdot V_{DD} \cdot (V_{DD} - V_{TH})$.

To support the theory on the delay benefits of the proposed TLFF comparator, its delay versus ΔV_I is simulated and, together with its outputs, compared with the comparators discussed in the previous section, as shown in Fig. 4.6. For a fair comparison, all the comparators are dimensioned for a similar input-referred noise ($\leq 1 \text{ mV}_{\text{rms}}$) and offset ($\leq 10 \text{ mV}_{\text{rms}}$) while driving the same output load. Furthermore, the critical paths for $g_{m,\text{eff}}$, t_{int} , and t_{latch} utilize devices with ultra-low V_{TH} for maximum speed. For the delay comparison, a three-stage comparator without a feed-forward path is also included, to highlight the benefits across a large ΔV_I range.

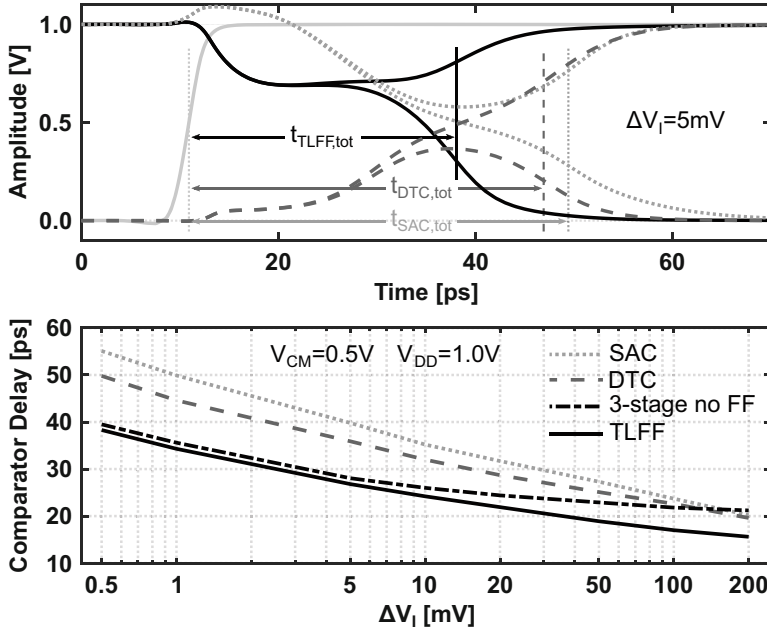


Fig. 4.6 Simulated outputs and delay versus ΔV_I for different comparators

In the upper half, the TLFF outputs exhibit a superior total delay compared to the SAC and DTC, thanks to the increase in the signal gain and the effective latch transconductance. The output common-mode goes down by about a threshold voltage due to the turn on of the output latch and the feed-forward devices with a large initial overdrive voltage, quickly regenerating on the signals from the parallel direct/feed-forward paths. A similar common-mode effect but to a lesser extent is noticed in the outputs of the DTC, due to the turn on of the PMOS cross-coupled pair. For the smallest simulated ΔV_I of 0.5 mV, the enhanced gain TLFF demonstrates a 1.4 \times and 1.3 \times shorter delay compared to the SAC and the DTC, respectively. The three-stage comparator without feed-forward [60], for the same constraints, achieves almost the same delay as the TLFF due to the similar total signal gain.² For a ΔV_I of 200 mV, the TLFF still offers a 1.25 \times delay advantage over the SAC and the DTC counterparts. It also shows about 1.34 \times shorter delay than the three-stage comparator without feed-forward, whose relative delay has increased, as explained (see Sect. 4.2.1). Although the absolute delay for such large ΔV_I is inherently short, minimizing it still yields benefits when accumulated over several cycles (e.g., within an asynchronous SAR).

² The small difference is attributed to the extra g_{m13} of the TLFF.

The time constant τ_{comp} of the comparators can be estimated from Fig. 4.6 and the delay expressions from the difference in delay Δt_{comp} when varying ΔV_I by 10^x , according to the following formula [116]:

$$\Delta t_{\text{comp}} = x \ln 10 \cdot \tau_{\text{comp}}. \quad [\text{s}] \quad (4.28)$$

This is a first-order approximation since τ_{comp} is assumed constant across the entire ΔV_I range, which might not necessarily be the case. Further, t_{comp} contains both integration and latching times, and the former is also not necessarily constant across the input range. However, the approximation is accurate for smaller ΔV_I , where latching dominates t_{comp} over integration. For the TLFF comparator, τ_{TLFF} is estimated to 4.4 ps in the ΔV_I decade between 0.5 and 5 mV. For the DTC and the SAC, τ_{DTC} and τ_{SAC} are estimated to 5.8 ps and 6.4 ps, respectively. As explained in the metastability discussion, a small τ_{comp} (by means of maximizing $g_{\text{m,eff}}$) is key to minimizing errors due to metastability.

Regarding the energy consumption of the TLFF, the simulated energy per comparison for the smallest differential input of 0.5 mV (Fig. 4.6) is about 150 fJ/comparison, at 1 V V_{DD} and 5 GHz clock. This energy is higher compared to the one of the SAC and the DTC by about 1.5 \times and 1.25 \times , respectively. This is a reasonable trade-off in ultrahigh-speed mixed-signal systems (e.g., multi-GHz ADCs for multi-Gb/s wireline data links). In such systems, the comparator typically consumes only a fraction of the total energy, while its speed may limit the overall system speed. Considering that the TLFF achieves a shorter delay than both the SAC and the DTC, thus enabling a higher absolute system speed, this benefit potentially overcomes the higher energy consumption. As discussed in Sect. 4.3.2, lowering V_{DD} reduces/nullifies this energy overhead while increasing the relative TLFF delay benefit due to its reduced stacking.

4.2.3 On-Chip Delay Evaluation Setup

The implemented test chip with the on-chip delay evaluation is shown in Fig. 4.7. The test chip includes the proposed TLFF as well as the SAC and DTC comparators in their dedicated channels. Each channel contains the comparator core as well as its local output, clock, and control circuitry. The differential clock is generated globally from an on-chip terminated sinusoidal signal and routed to all the channels symmetrically. Within each channel, a local unit provides the clock to each comparator with sharp edges. Further, an implemented controller turns the channels on/off, for individual evaluation and minimum inter-channel interference. The inputs are on-chip terminated, with separate V_{CM} voltages, for offset cancellation prior to performance characterization. They are routed to the channels in a matched fashion, similar to the clock.

The comparator outputs, after two cascaded inverters/buffers, are fed to a capture latch after two cascaded inverters and then further buffered by an inverter chain

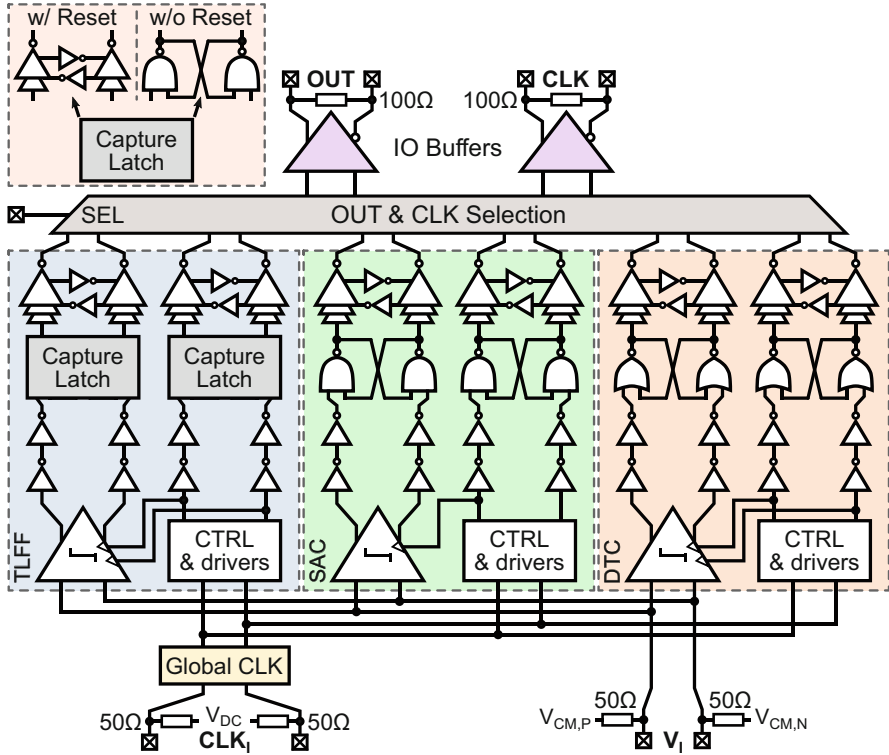


Fig. 4.7 Top-level diagram of the multiple comparators test chip

with cross-coupling. For the TLFF, two versions are included, one with a proper latch to hold the signal during reset and one with simple inverters to also capture the reset. Each comparator and its adjacent inverters include their own supply for a proper energy evaluation and better switching noise isolation. To accurately characterize the absolute CLK-OUT delay, all the circuits following the comparators (including interconnect and surroundings) are replicated using the clock as their input (Fig. 4.7). The OUT and CLK of each channel pass through a matched selection block, after which they are brought to the pads by two identical on-chip terminated IO buffers. Overall, circuits and paths across the OUT/CLK chain (chip, board, cables) are kept as much as possible identical for each and all comparators.

4.3 Experimental Verification

The implemented multiple comparators test chip with the prototype TLFF is fabricated in a single-poly ten-metal (1P10M) 28 nm bulk CMOS process. A die

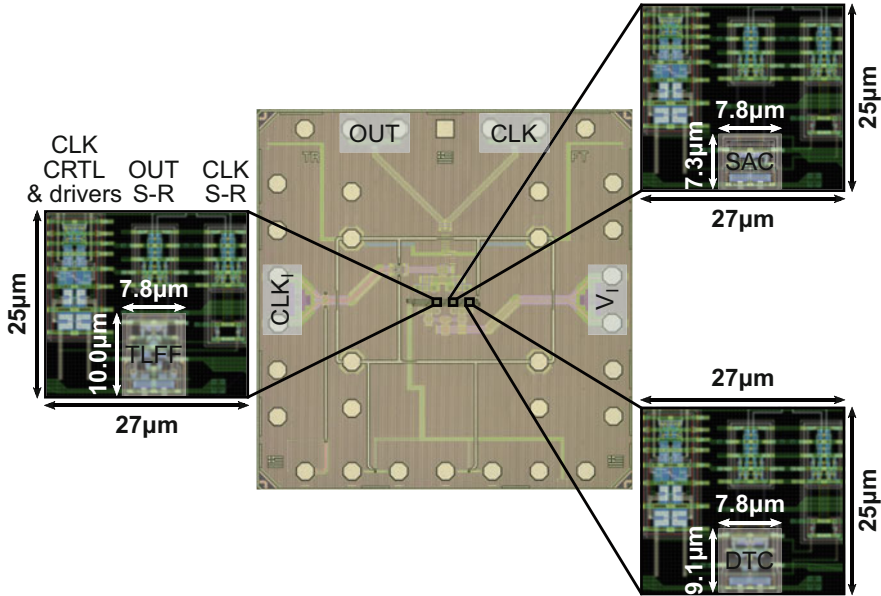


Fig. 4.8 Die photo of the 28 nm IC with zoomed-in comparator layout views

photo is shown in Fig. 4.8, where the zoomed-in layout views of the comparator channels are included. The TLFF occupies a compact area of $7.8 \times 10 \mu\text{m}$. This is about $1.36\times$ and $1.1\times$ larger than the area of the SAC and the DTC, respectively. In each channel, the OUT propagation path is placed right above the comparator core, while the identical CLK path is located to its right. The local clock unit, including its on/off controller and drivers, is placed at the left side of each channel. In each block, differential symmetry is maintained as much as possible. This compact block arrangement results in a small area for each channel of about $27 \times 25 \mu\text{m}$. The OUT and CLK are collected at the top side of the chip. The master clock and input are coming from opposite sides and distributed equidistantly to all the channels. The control signals for the local clock units and the channel selection block are coming from the bottom.

The test chip utilizes three 1 V supply domains, each domain having its dedicated ground by using guard rings. The comparator cores and their output inverters are placed in one domain, for isolation and characterization purposes. The clock circuitry occupies its own domain, while the S-R latches and the inverter chain with cross-coupling share the ground with the comparator cores. The IO circuits with part of the selection unit and the buffers occupy the remaining domain. All cross-domain transitions are differential to preserve signal integrity.

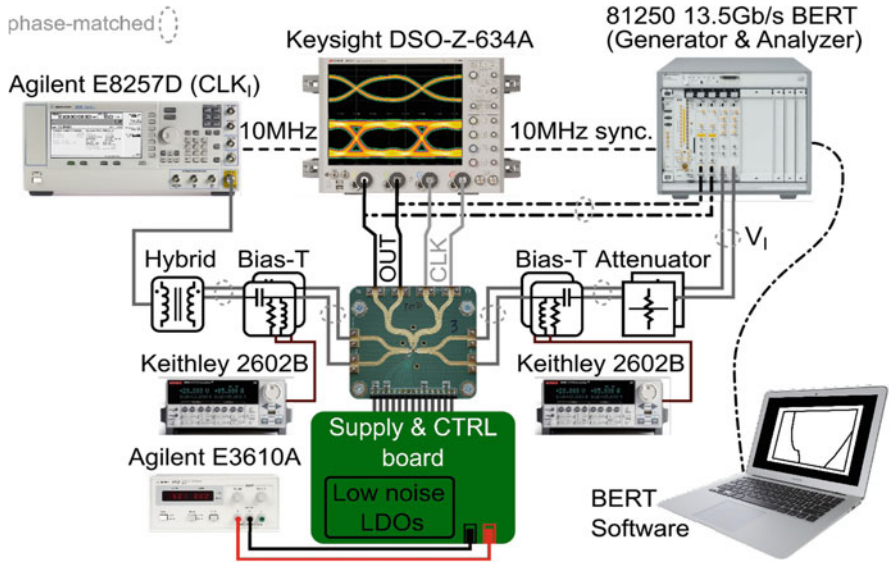


Fig. 4.9 Measurement setup of the multiple comparators test chip with the prototype 13.5 Gb/s TLFF comparator

4.3.1 Measurement Setup

A detailed view of the measurement setup used to evaluate the comparators’ performance is given in Fig. 4.9. An Agilent E8257D signal generator provides the sinusoidal master clock of up to 13.5 GHz to the chip. This clock signal, after converted into differential by a wideband 180° hybrid, is AC-coupled to the chip through wideband bias-Ts and phase-matched cables. The differential input to the chip is provided by the generator module of an Agilent 13.5 Gb/s parallel bit error rate tester (ParBERT). The inputs are also AC-coupled to the chip, and the common-mode voltage to each side is set independently by a dual-channel Keithley sourcemeter. This approach is adopted in order to compensate the offset of each comparator prior to the rest of the characterization. Matched attenuators are used selectively to get input differences as small as 5 mV, while for 100 mV and above, the ParBERT alone suffices. Phase-matched cables are used in all the sections of the input chain.

The differential OUT and CLK are connected through quad phase-matched cables to a Keysight DSO-Z-634A 63 GHz scope. The scope serves as a waveform analyzer, through which the offset is compensated and the delay and noise are characterized. First, the raw offset from each comparator is measured by applying a zero differential input and the same common-mode voltage from its corresponding sourcemeter. This voltage is adjusted differentially in the two channels until the OUT shows an equal probability of logical 0s and 1s, dictated by noise. After that,

the CLK-OUT delay and the noise are characterized by sweeping the differential input. It is worth noting that the offset compensation procedure is repeated for every different common-mode and supply voltage that the delay and/or the noise needs to be characterized. After delay and noise characterization, the OUT is captured by the analyzer module of the ParBERT for BER analysis in a PC using the BERT software. All the devices are synced in a master-slave fashion through their 10 MHz sources.

The custom chip board employs carefully optimized ultra-wideband controlled impedance traces on a R4000 Series high-frequency low-loss laminate material. In the same way as the on-chip paths, the board traces, especially the CLK and OUT, are kept as much as possible alike. The necessary supply and bias voltages are generated with dedicated low-noise Low-DropOut (LDO)s on the custom motherboard and provided to the chip after further low-pass filtering with discrete components as well as additional on-chip filtering. The control signals to select each comparator channel for individual characterization are also coming from the motherboard.

4.3.2 Measurement Results

The key characterization goal of this test chip involves the delay (thus speed) of the proposed TLFF comparator in order to verify its benefit over the two widely used SAC and DTC, as analyzed in Sect. 4.2.1. However, the energy overhead due to more stages is also important to assess and compare against the delay benefit. Figure 4.10 plots the measured CLK-OUT delays under various conditions. The delay is measured for a '1010' input pattern from the ParBERT. When sweeping the differential input after having compensated the offset (Fig. 4.10a), the TLFF demonstrates a delay of 26.8 ps for a ΔV_I of 5 mV, which drops to 17.6 ps at 200 mV ΔV_I . For the same ΔV_I range, these delays are 1.44 \times –1.19 \times /1.31 \times –1.14 \times shorter compared to the measured ones for the SAC and the DTC, respectively. The delays for the opposite sign inputs differ by less than 5% for all the comparators. The delay versus $\log(\Delta V_I)$ slope of the TLFF is -6.4 ps/decade, superior to both the DTC (-9.5 ps/decade) and the SAC (-11.3 ps/decade).

The measured delays for a V_{CM} between 0.3 V and 0.8 V and a 5 mV ΔV_I are shown in Fig. 4.10b. The pattern is similar for all the comparators, with a decreasing delay for an increasing V_{CM} up to 0.7 V, after which further V_{CM} adversely affects the delay. The TLFF exhibits the shortest absolute delay across the entire common-mode range, more than 1.34 \times shorter than the other comparators. It also shows the smallest delay variation within 20 ps across the entire V_{CM} range, attributed to its reduced common-mode sensitivity. Figure 4.10c shows the measured comparator delays for a 5 mV ΔV_I when sweeping V_{DD} between 0.6 and 1.1 V. The reduced stacking of the TLFF due to the concurrently on horizontally cascaded latches allows for the shortest absolute delay among all the comparators, with increasing benefits as the supply scales down. At a V_{DD} of 0.6 V, the below 70 ps delay of the

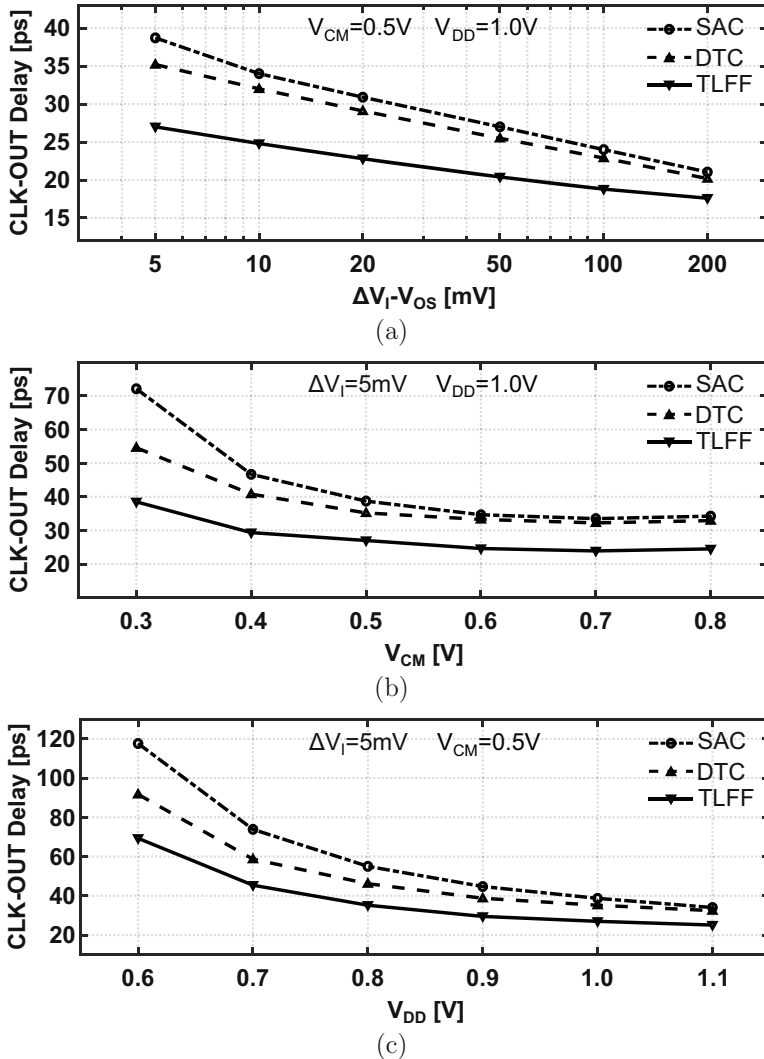


Fig. 4.10 Measured CLK-OUT delays for the TLFF, SAC, and DTC comparators versus (a) ΔV_I , (b) V_{CM} , and (c) V_{DD}

TLFF is about $1.7\times$ and $1.3\times$ shorter than that of the SAC and the DTC. The TLFF also achieves the smallest delay variation within 50 ps across the swept V_{DD} range, making it a favorable ultrahigh-speed high-sensitivity candidate for a wide range of common-mode and supply voltages. It is worth noting that the offset of each comparator is compensated for every V_{CM} as well as V_{DD} setting, prior to sweeping ΔV_I for characterizing the delay.

The input-referred noise of all the comparators is measured by sweeping a DC ΔV_I , and fitting the occurring ratio of logical 0s and 1s to a Gaussian Cumulative Distribution Function (CDF), as shown in Fig. 4.11. At a V_{CM} of 0.5 V and a V_{DD} of 1 V, the input-referred noise is similar for all the comparators, with a $1\text{-}\sigma$ value within 0.82–0.89 mV_{rms}. The slightly higher noise of the TLFF is attributed to the extra contribution of the feed-forward M_{13P}/M_{13N} devices. This noise could be reduced by adding reset devices at the sources of the M_{2P}/M_{2N} cascodes for extra gain, at the cost of a slightly larger delay. Sweeping V_{CM} between 0.3 and 0.8 V, the noise for all comparators increases with V_{CM} (Fig. 4.11d) due to the higher integration current, but the TLFF shows less dependency due to its reduced stacking cascaded latch arrangement. Finally, all comparators exhibit a comparable input-referred noise across V_{CM} .

The measured energy per comparison for the comparators (including their adjacent inverters) is shown in Fig. 4.12. The clock frequency for this measurement is 4 GHz to cover the necessary time for a correct comparison of the slowest comparator under the hardest setting ($\Delta V_I = 5$ mV, $V_{DD} = 0.6$ V). Sweeping V_{DD} for $\Delta V_I = 5$ mV and $V_{CM} = 0.5$ V (Fig. 4.12a), the TLFF consumes about 1.46× and 1.22× more energy than the SAC and the DTC, respectively, at 1 V. This is expected given the extra stages of the TLFF. As this is almost entirely offset by the TLFF delay benefit over the SAC and the DTC at 1 V, one might assume that increasing their devices' sizes to match the TLFF energy would yield an equivalent delay reduction. This is not true in an optimized multi-GHz comparator, whose latch devices are upscaled until its time constant $\tau_{comp} = (C_L + C_p)/g_m$ is just dominated by the load C_L and not the parasitic capacitance C_p . Therefore, further device upscaling yields only minor delay benefits. The TLFF energy overhead over the SAC and the DTC quickly reduces with reducing V_{DD} and is completely eliminated at 0.6 V (Fig. 4.12b). The reason for this is that the latter two reside longer in their metastable state due to a longer latching time (see Eqs. (4.8), (4.13), and (4.22)), thus preserving a direct path from V_{DD} to V_{SS} for a bigger portion of a clocking period. From Figs. 4.10c and 4.12a, the shorter delay and delay variation vs. V_{DD} of the TLFF allow for a lower supply operation with a shorter delay and a lower energy altogether than both the SAC and the DTC. For example, the TLFF at 0.8 V yields a 1.13× and 1.05× shorter delay with 1.15× and 1.38× less energy than the SAC and the DTC, respectively, at 1 V.

The efficiency comparison becomes clearer and more interesting when plotting the energy-delay product vs. V_{DD} (Fig. 4.12b). For V_{DD} values below 1 V, the TLFF demonstrates a considerably lower energy-delay product compared to both the SAC and the DTC, thanks to its increasing delay benefit and diminishing energy overhead. With the constant supply reduction due to scaling (0.8 V in 16 nm, ≤ 0.65 V in 7 nm, 5 nm), the architecturally superior TLFF presents an increasingly attractive ultrahigh-speed solution with a combined shorter delay, enhanced V_{CM} and V_{DD} robustness, greater design flexibility, and competitive energy.

The data rate of the TLFF for a given input sensitivity is characterized by performing BER measurements. A $2^{31} - 1$ Pseudo Random Binary Sequence (PRBS) differential pattern is applied to the comparator input, and errors at the captured




Fig. 4.11 Measured noise cumulative distribution and Gaussian distribution fitting curve for (a) the SAC, (b) the DTC, (c) the TLFF, and (d) measured input-referred noise vs. V_{CM}




Fig. 4.12 Measured (a) energy consumption of the TLFF, SAC, and DTC versus V_{DD} and (b) energy-delay product versus V_{DD}

differential output are recorded. Figure 4.13a shows a generated 5 mV ΔV_i eye at 13.5 Gb/s. The TLFF output errors are recorded when sampling this input pattern at different time instants, and the bathtub curve of Fig. 4.13b is constructed for two different input amplitudes. The lowest measured BER is found at a 0.6 unit-interval (UI) sampling instant and remains below 10^{-12} for a ± 1.8 ps (± 0.025 UI) sampling offset. The curve follows the shape of the input pattern, limited by the eye opening of the generator. When ΔV_i increases to 10 mV, a BER lower than 10^{-12} is maintained for a sampling offset of ± 5.6 ps (± 0.075 UI). For completeness, the measured CLK (Fig. 4.13c) and the $2^{31} - 1$ OUT (Fig. 4.13d) eyes are also shown. The 13.5 GHz (27 Gb/s) CLK eye height is 260 mV with an eye width of about 30 ps (0.81 UI). The 13.5 Gb/s OUT eye height is 310 mV with an eye width of 67 ps (0.9 UI). Although the delay measurements (Fig. 4.10) suggest that for a 5 mV ΔV_i the TLFF should be able to properly work up to 18 Gb/s, the generator/analyzer ParBERT modules are limiting the data rate measurements to 13.5 Gb/s.




Fig. 4.13 (a) Measured ΔV_1 eye, (b) measured bathtub curve of the TLFF, (c) measured CLK eye, and (d) measured OUT eye

4.3.3 State-of-the-Art Comparison

Table 4.1 summarizes the measured TLFF performance along with that of the SAC and the DTC on the same test chip. Table 4.2 provides a state-of-the-art comparison of the proposed TLFF with the fastest deep-scaled CMOS comparators in literature, including a two-stage comparator [117], a 2 \times -interleaved single-stage SAC-like [118], a three-stage comparator with two pre-amplifiers [107], and a comparator with dynamic-bias pre-amplifier [119]. The introduced TLFF achieves the highest reported data rate of 13.5 Gb/s with mV-range input sensitivity while maintaining a BER below 10^{-12} . Further, it exhibits the smallest delay vs. $\log(\Delta V_1)$ (-6.4 ps/decade) and delay variations vs. V_{CM} ($<\pm 3.5$ ps) and V_{DD} ($<\pm 2.7$ ps). The input-referred noise of 0.89 mV_{rms} and the energy/comparison of 163 fJ are on par with the state of the art. The works in [107, 118, 119] employ a very similar process feature size while having the SOI advantage of reduced parasitics and better channel current control. Porting them to the 28 nm bulk CMOS process at hand, and assuming a similar input-referred noise target, [107] is expected to achieve a similar performance as the simulated three-stage comparator without feed-forward (Fig. 4.6), while [118] is expected to perform similarly to the designed

Table 4.1 Summary of the TLFF, SAC, and DTC on the same test chip

	Proposed TLFF	Single-stage SAC	Two-stage DTC
V_{DD} range [volts]	0.6–1.1	0.6–1.1	0.6–1.1
V_{CM} range [volts]	0.3–0.8	0.3–0.8	0.3–0.8
ΔV_I range [mV]	5.0–200.0	5.0–200.0	5.0–200.0
Delay slope [ps/dec]	–6.4	–11.3	–9.5
Delay vs. V_{DD} [ps] ^a	69.3–25.0	118.0–34.1	91.8–32.5
Delay vs. V_{CM} [ps] ^a	38.4–24.4	72.2–34.3	54.6–32.9
Data rate [Gb/s] ^{a,b}	13.5	9.0	10.0
Core area [μm^2]	78.0	57.0	71.0
1- σ noise [mV] ^c	0.89	0.82	0.85
Energy vs. V_{DD} [fJ] ^d	45.2–226.2	43.7–146.1	44.9–178.2

^a@ min ΔV_I ^bBER < 10^{-12} ^c@ $V_{CM} = 0.5$ V^d@ 4 GHz clock and including output inverters**Table 4.2** TLFF comparison with state-of-the-art comparators

	This work ESSCIRC'19	Goll [117] ISSCC'09	Kull [118] ^a ISSCC'13	LeTual [107] ^a ISSCC'14	Bindra [119] ISSCC'22
Technology [nm]	28 CMOS	65 CMOS	32 SOI	28 FDSOI	22 FDX
Topology	TLFF	Two-Stage	SAC-Like	Tri-Stage	Dyn-Bias
V_{DD} [volts]	1.0	1.2	1.0	1.0	0.8
Min ΔV_I [mV]	5.0	280.0	2.0	2.4	1.0
Slope [ps/dec]	–6.4	–20.0	–10.5	–11.9	–20.0
Del vs. V_{DD} [ps] ^b	< ± 2.7	N.A.	N.A.	N.A.	< ± 50.0
Del vs. V_{CM} [ps] ^b	< ± 3.5	< ± 7.0	< ± 8.0	N.A.	< ± 25.0
D-rate [Gb/s] ^c	13.5	7.0	$\sim 12.0^d$	10.0	< 1.0^f
Area [μm^2]	78.0	319.5	100.0	N.A.	57.0
1- σ noise [mV]	0.89	N.A.	1.35	1.0	0.2
Power [mW] ^c	2.2 ^e	1.3	0.8	1.8	N.A.
Energy [fJ] ^c	163.0 ^e	185.0	67.0	180.0	75.0 ^f

^aEmbedded in an ADC^bFor a ± 0.1 V variation^c@ min ΔV_I and BER < 10^{-12} ^d2 \times -interleaved, within SAR^eWith output inverters^fBER not reported

SAC (Table 4.1). Regarding [119], it is questionable whether it can achieve data rates higher than a few GHz, or small enough delay variations, due to its dynamic-bias pre-amplifier topology. Finally, [117] is expected to perform similarly to the designed DTC (Table 4.1). The architectural enhancements of the proposed TLFF generally hold true against these comparators as well, making TLFF a

favorable candidate for ultrahigh-speed mixed-signal systems, with increasingly more profound advantages in lower supply deep-scaled CMOS nodes.

4.4 Conclusion

This chapter covered the analysis and design of ultrahigh-speed latch-based dynamic comparators, as they are vital blocks in numerous high-performance mixed-signal systems. Their role is of significant value in high-speed ADCs, where they have to efficiently extract the digital representation of analog input signals with maximum speed, sufficient accuracy, and good power and robustness.

First, the two widely adopted topologies, the single-stage SAC and the two-stage DTC, were reviewed, and the various parameters affecting their delay were analyzed. The single current path of the SAC presents the main limitation in optimizing its delay as well as its variability to parameters such as V_{CM}/V_{DD} . On top, its increased device stacking makes scaling of its proper operation to lower supply deep-submicron processes particularly challenging. The DTC splits the input and latch stages, providing a better control over the parameters that influence the delay and reducing the V_{CM}/V_{DD} variability. The split input-latch also reduces the device stacking, allowing a proper operation at reduced supply voltages. However, both these topologies stack their latch devices vertically, still occupying significant headroom. The series turn on of these devices reduces the effective regeneration rate. Further, a higher signal gain prior to latching would be desirable, especially for small differential inputs, to reduce the total delay by creating a larger initial difference for the latch to regenerate on.

To improve on the above issues, this work introduced a three-stage triple-latch comparator topology with a reduced stacking and parallel direct/feed-forward paths to improve the delay across a large input range. The multi-stage nature with cascaded latches enabled a very high total gain in an exponential fashion prior to the final latching. The concurrent turn on of the latch devices with a large overdrive voltage rather than in series increased the effective regeneration rate by increasing the effective device transconductance. Further, the horizontal cascading instead of the vertical stacking of the latch devices further reduced the required headroom of this topology compared to the SAC and the DTC, allowing a favorable operation at lower supply voltages.

Fabricated in a 28 nm bulk CMOS process along with the other comparators, the proposed TLFF achieves the smallest absolute delay, delay slope, and variation across a wide V_{CM}/V_{DD} range, with a similar input-referred noise. The energy overhead, gradually disappearing at lower supply voltages, trades off against the reduced delay and increased robustness. Finally, a BER below 10^{-12} is achieved for mV-range differential inputs at a data rate of 13.5 Gb/s, the highest reported among SotA non-interleaved non-pipelined comparators. All these features make TLFF a favorable candidate for ultrahigh-speed operation in deep-scaled CMOS, either standalone or within a mixed-signal system.

Chapter 5

High-Speed Wide-Bandwidth Single-Channel SAR ADC



The traditionally slow SAR ADC has been the center of attention in various high-speed applications for about 10 years now. The rapid down-scaling to finer CMOS processes has rendered the efficient design of high-speed high-accuracy analog amplifiers particularly challenging due to the supply voltage drop. On the other hand, the highly digital nature of the SAR having the comparator as the single supply-limited block makes it more immune to such challenges. Additionally, a number of techniques have emerged to significantly enhance its speed.

Section 5.1 of this chapter reviews the conventional SAR clocking scheme and discusses some noteworthy speed-boosting techniques that have enabled the SAR to be a high-speed protagonist. Section 5.2 introduces the proposed prototype SAR ADC and elaborates on the speed-boosting architectural and circuit principles employed. The experimental verification, including the measurement setup, measured results, and a state-of-the-art comparison, are the focus of Sect. 5.3. Finally, the conclusion is drawn in Sect. 5.4.

Parts of this chapter were previously presented at the 2015 *Conference on Ph.D. Research in Microelectronics and Electronics (PRIME'15)* in Glasgow, Scotland [120], and the 2017 *European Solid-State Circuits Conference (ESSCIRC'17)* in Leuven, Belgium [121], and published in the *Journal of Solid-State Circuits (JSSC'18)* and in the *Electronics Journal (MDPI Electronics'19)* in July 2018 [60] and January 2019 [122], respectively.

5.1 Pushing the SAR Conversion Speed

The admirable energy efficiency of SAR ADCs for low-to-medium resolutions, as seen in [36] and our analysis in Chap. 3, has made them the dominant architecture for high-speed wireline (>112–224 Gb/s PAM-4) and short-range wireless (Fifth-Generation (5G), Sixth-Generation (6G)) systems. However, their bit-at-a-time




Fig. 5.1 Timing sequence illustration of a B -bit SAR with a conventional synchronous clocking scheme

operation necessitates a large channel count interleaving ($N \geq 36$) to reach the targeted multi-GHz sample rates [8]. Interleaving is the most attractive option to achieve such high sample rates, but its introduced errors (see Chap. 3, Sect. 3.7) often require complex and power-hungry calibration. Further, increasing the channel count results in a bandwidth degradation due to a larger parasitic input load. Therefore, improving the single-channel SAR speed is highly desirable to reduce the channel count and the associated interleaving overhead. While doing so, it is crucial to maximize the *accuracy · speed ÷ power* while minimizing its area, such that it can be smoothly integrated into the TI system without exacerbating the overhead and complexity.

5.1.1 Conventional Synchronous Clocking Scheme

The conventional synchronous clocking scheme for a B -bit SAR is shown in Fig. 5.1. An internal high-speed clock divides the total conversion period into equally spaced cycles to accommodate one sampling/tracking cycle T_s and each of the bit cycles [107, 123]. Every bit cycle comprises three critical sequential timings: (1) the comparator evaluation/resolving time t_{comp} , (2) the DAC settling time t_{DAC} , and (3) the delay of the digital SAR logic $t_{\text{logic,sync}}$. The allocated time for the comparator is typically designed to meet the requirements of the worst-case (slowest) scenario $t_{\text{comp,max}}$. From the exponential input-output characteristic of a regenerative latch (see Chap. 2, Sect. 2.4.3), this scenario occurs for a comparator input within $\pm \text{LSB}/2$. The same holds for the allocated DAC time, whose worst-case scenario $t_{\text{DAC,max}}$ occurs during the MSB cycle since it has to cover the largest range for a certain settling accuracy. Depending on which of the two dominates, the internal clock duty cycle can be designed accordingly. The logic delay is more or less fixed across the cycles. It involves the state memory part responsible for storing the comparator output and initiating the next DAC settling, while its value depends




Fig. 5.2 Timing sequence illustration of a B -bit SAR with an internally asynchronous clocking scheme

on the implementation. The critical timing for each bit cycle can be expressed as

$$t_{crit,sync} = t_{comp,max} + t_{DAC,max} + t_{logic,sync}, \tag{5.1}$$

while the total conversion period (excluding sampling) for a B -bit synchronous SAR converter can be written as

$$T_{sync} = B \cdot (t_{comp,max} + t_{DAC,max} + t_{logic,sync}). \tag{5.2}$$

5.1.2 Speed-Boosting Techniques

Asynchronous Processing

In a binary SAR, the worst-case comparison cycle occurs only once, while in the other cycles, the comparator evaluation is faster. Synchronous processing must satisfy that worst-case scenario in every cycle, resulting in time wasting in the other faster cycles and thus limiting the converter speed. Asynchronous processing, whose timing illustration is shown in Fig. 5.2, saves time from the faster comparison cycles and distributes it where necessary [124–126] resulting in a shorter total comparison time. The high-speed clock is eliminated, and the SAR logic asynchronously controls the bit cycles by locally generated signals. When the differential comparator output crosses a certain threshold, a comparison detection block generates a “ready” signal to initiate the comparator reset, and at the same time, the output is passed to the state memory logic to initiate the next DAC settling. This results in a variable time for the comparator $t_{comp,var}$, and the critical timing for each bit cycle as well as the total conversion period for a B -bit asynchronous SAR converter can be, respectively, expressed as

$$t_{\text{crit,async}} = t_{\text{comp,var}} + t_{\text{DAC,max}} + t_{\text{logic,async}}, \quad t_{\text{comp,var}} \leq t_{\text{comp,max}}, \quad (5.3)$$

$$T_{\text{async}} = \sum_{i=1}^B t_{\text{comp,var}} + B \cdot (t_{\text{DAC,max}} + t_{\text{logic,async}}). \quad (5.4)$$

Comparing Eqs. (5.2) and (5.4), it is evident that the time savings provided by asynchronous processing are dependent on the number of bits. The total time savings are particularly effective as a buffer margin to cope with the metastability of the ADC (see Chap. 2, Sect. 2.4.3) [127, 128]. A synchronous SAR has to deal with metastability within each cycle as an equal probability event. However, metastability occurs with an increasing probability going from MSB to LSB and typically only once. Offering a single accumulated buffer margin at the end of the conversion for each bit cycle to utilize as necessary, asynchronous processing can reduce the metastability error probability.

The reduction in comparison time of the asynchronous processing comes with an increased logic complexity, which now has to cover more functionality. Apart from the existing state memory part, the logic has to internally provide the signal to switch the comparator between evaluation and reset modes quickly enough to meet the target speeds. This requires a number of extra gates in the critical path making $t_{\text{logic,async}} > t_{\text{logic,sync}}$. At tens of GHz internal speeds, the logic delay can be a significant part of the bit cycle, while its power is not negligible either compared to the comparator and DAC. The comparator load is increased, deteriorating its time constant, while its reset time may be prolonged to dominate the DAC settling. Thus, the extra logic overhead and its impact partially negate the comparator time savings and metastability margin.

Multi-bit per Cycle

A B -bit SAR conversion period is a cascade of B single-bit internal cycles. There is nothing that limits each cycle to a single bit, and the speed can be enhanced by evaluating more than one bit per cycle, as has been widely demonstrated [95, 129, 130]. This can be seen as embedding a small flash ADC inside the SAR loop. The number of cycles reduces by a factor equal to the number of bits per cycle, ideally increasing the conversion speed by the same factor. Without the loss of generality, Fig. 5.3 shows the timing illustration of a 2-bit per cycle example with three comparators. For an even B , the cycles are halved, while for an odd B , there is one extra single bit cycle. Either synchronous or asynchronous processing can be employed internally, which defines $t_{\text{crit,mult}}$.

Multi-bit per cycle schemes provide theoretically the highest speed improvement among existing speed-boosting techniques. However, they require multiple comparators and DACs, exponentially increasing with every added bit per cycle. Further, one or multiple calibration schemes are necessary to compensate the difference in offset and gain between these blocks and reduce the non-linearity stemming




Fig. 5.3 Timing sequence illustration of a B -bit SAR with a multi-bit per cycle resolving scheme (2-bit per cycle shown in the example)




Fig. 5.4 Timing sequence illustration of a B -bit SAR with multiple comparators loop unrolling clocking scheme

from this difference. This calibration can often be complex and power/area hungry, degrading the converter efficiency. Finally, the increased layout interconnect results in extra parasitic capacitance and power consumption, diminishing the theoretical speed and efficiency benefits of such schemes.

Loop Unrolling

Among the three critical timings inside a bit cycle, t_{comp} and t_{DAC} must occur sequentially since the latter is switching in the direction dictated by the sign of the former. However, t_{logic} could be partially or fully eliminated from the critical path by overlapping with one of the other two. The loop unrolling scheme [113, 131] uses this concept to increase the conversion speed. Its timing illustration is provided in Fig. 5.4. Instead of the same single comparator being switched between evaluation

and reset modes in every cycle, this scheme uses B comparators to evaluate B bits. One comparator is triggered in each cycle and stays latched until the end of the conversion. The DAC switches to the comparator output without the intermediate t_{logic} , which now overlaps with t_{DAC} and controls the start of the next cycle. A single cycle resets all B comparators prior to next sampling. This scheme is typically combined with asynchronous processing, and the critical timing for each cycle is given as

$$t_{\text{crit,unrol}} = t_{\text{comp,var}} + \max\{t_{\text{DAC,max}}, t_{\text{logic,unrol}}\}, \quad t_{\text{comp,var}} \leq t_{\text{comp,max}}. \quad (5.5)$$

Similar to the multi-bit per cycle, this scheme also requires calibration to compensate the different offsets of B comparators. The calibration increases the ADC complexity and can also reduce the comparator speed or increase its noise by adding a tunable output element bank (capacitive or current) [113] or extra input devices [67, 132]. Additionally, the increased input capacitance from the different comparators can lead to a non-linear gain error at the DAC output. This non-linearity is further enhanced by the multi-comparator kickback, which requires extra power to suppress sufficiently (e.g., through pre-amplification).

Redundancy

The thus far mentioned techniques focus on saving comparator or logic time from the critical path but requiring full accuracy on the DAC settling. One popular and powerful approach to relax the settling accuracy of the DAC is to use redundancy or OR. This is done either in the form of sizing each higher-rank DAC unit element (in this case capacitor) smaller than the sum of all the lower-rank elements (sub-radix-2) [133] or by keeping the binary ratios and using occasionally repetitive compensation steps [134]. In this way, shorter, incomplete settlings can be tolerated, and any potential errors can be absorbed in the extra cycles used at the end of the conversion. Figure 5.5 shows a timing illustration of a B -bit $B+1$ -cycle redundant SAR. Redundancy is particularly beneficial when applied to the longer MSB settlings, and the total amount of redundancy dictates the extra cycles necessary to achieve an aggregate quantization [58].

The speed benefit of redundancy relies on the fact that the exponential gain in speed, thanks to incomplete settling, outmatches the linear loss from the extra cycles. If the comparator or the logic timings dominate the SAR cycle delay, redundancy may not necessarily result in speed improvement. This is typically the case for low-to-medium-resolution designs since the DAC capacitance is already minimized to the noise or matching limit to guarantee a large input bandwidth. Further, redundancy in its most common sub-radix-2 form can potentially increase the metastability error probability of the ADC to more than once per conversion. Finally, extra digital correction or arithmetic circuits may be required, adding power, area, and latency [133].




Fig. 5.5 Timing sequence illustration of a B -bit SAR with extra cycles and redundancy implemented (one extra cycle shown in the example)

Pipelining

As it was extensively analyzed in Chap. 3, the combination of pipelining and SAR concepts can be highly advantageous under the right circumstances, as it boosts the converter sample rate by roughly the amount of sequential cycles saved from each SAR. The speed improvement can be on the same order as the multi-bit per cycle, however avoiding the multiple comparators and DACs per stage, and their associated calibration overhead. Pipelining can be combined with one or more of the aforementioned techniques, by properly assessing the benefits and trade-offs of each combination for a required set of specifications.

On the downside, the necessary fast residue amplifiers with accurate gain, low noise, wide bandwidth, and good linearity may significantly increase the complexity and design effort while deteriorating the efficiency of the converter. Hence, these architectures have been proven most beneficial beyond a certain resolution (> 8 bits), when the efficiency of the noise-limited single SAR comparator starts degrading. This was also demonstrated by our derivation in Chap. 3 leading to Figs. 3.24, 3.25, and 3.26.

DAC Switching Schemes

In Chap. 3, several DAC switching schemes were analyzed from an energy efficiency perspective. However, some of them can also increase the conversion speed. The monotonic and the MCS schemes specifically, combined with top-plate sampling, do not require any settling prior to the MSB evaluation, which can immediately occur after sampling. Eliminating the longest DAC settling from the critical path can result in considerable time savings. It is worth mentioning that moving from bottom-plate to top-plate sampling alone without any special DAC switching scheme can hide the initial settling cycle prior to the MSB evaluation by

performing it simultaneously with the sampling. This is what has been implicitly assumed in all the timing diagrams of this section and is typically the norm in high-speed low-to-medium-resolution SAR ADCs. However, top-plate sampling comes with extra non-linearity due to input-dependent switch charge injection, preventing its use at high resolutions.

5.2 Prototype IC: A 1.25 GS/s 7-bit SAR ADC in 28 nm CMOS

This section presents a 1.25 GS/s 7-bit single-channel SAR ADC that achieves a Nyquist SNDR/SFDR of 40.1/52.0 dB and consumes 3.56 mW from a 1 V supply in 28 nm CMOS. The ADC maximizes the *accuracy · speed ÷ power* with a “minimalist” approach [60, 121], reckoning that at very high sample rates, any unnecessary hardware reduces speed and bandwidth, increases power and complexity, and eventually degrades SNDR and robustness. A single-bit per cycle, single-comparator topology is chosen, where a semi-asynchronous processing is utilized that eliminates the logic delay from the critical path. Additional features include an improved bootstrapped input switch, a triple-tail dynamic comparator, and a Unit-Switch-Plus-Cap (USPC) DAC. These features enable a sample rate and a bandwidth in excess of 1 GS/s and 5 GHz, respectively, with little performance degradation across the entire band without any calibration, allowing for a smooth integration of this ADC into a TI system.

5.2.1 High-Level Design

The top-level ADC architecture along with its timing diagram is shown in Fig. 5.6. The 1.25 GHz generated sampling clock SAM with a 12.5% duty cycle drives the T/H and initiates the SAR conversion on its falling edge. Top-plate sampling is adopted to accommodate the stringent speed requirements. The input signal is sampled onto the DAC, directly at the comparator input, through a bootstrapped NMOS switch to ensure good sampling linearity and resilience to the input common mode. The input/DAC capacitance is sized to achieve a low enough thermal noise and dynamic power with a better than 7-bit matching. Combined with the switch on-resistance, the termination equivalent resistance, and the ElectroStatic Discharge (ESD) capacitance, an input bandwidth of about 6 GHz is attained by this ADC’s input network.

The SAR logic performs several operations. It is responsible for generating the clock that switches the comparator between evaluation and reset modes CLK_{COMP} . It also generates the bit phases $PHASE_6 - PHASE_1$, aligned with the comparator evaluation mode. In each of these phases, the comparator output is stored in its




Fig. 5.6 Top-level architecture of the proposed ADC and its timing diagram

corresponding memory element. In parallel, the corresponding DAC capacitors are switched based on the evaluation sign to close the SAR loop. The stored outputs are serially collected off-chip at the full data rate of 10 Gb/s for performance evaluation.

5.2.2 Semi-asynchronous Processing w/o Logic Delay

The timing sequence employed in this work to boost the speed of this ADC is detailed in Fig. 5.7. It combines the merits of simple logic and cycle control from synchronous processing with the dynamically allocated internal timing of asynchronous processing. The input signal is sampled periodically with a 12.5% duty cycle, and the same fixed time is allocated in every bit cycle, governed by the full rate clock. This makes off-chip capturing of the serial data readily available without any re-timing compared to a fully asynchronous approach. However, within each bit cycle, the time is asynchronously shared, resulting in a variable time for both the comparator $t_{comp,var}$ and the DAC $t_{DAC,var}$. The SAR logic is triggered in parallel to the comparator by the synchronously generated bit phases (Fig. 5.6), enabling an immediate DAC settling (see Sect. 5.2.6). Considering that t_{logic} can occupy as much as 30–40% of a bit cycle as short as 100 ps, eliminating it from the critical path brings significant speed benefits. The critical timing of this scheme is

$$t_{crit,semi-async} = t_{comp,var} + t_{DAC,var}, \quad t_{comp,var} \leq t_{comp,max}. \quad (5.6)$$

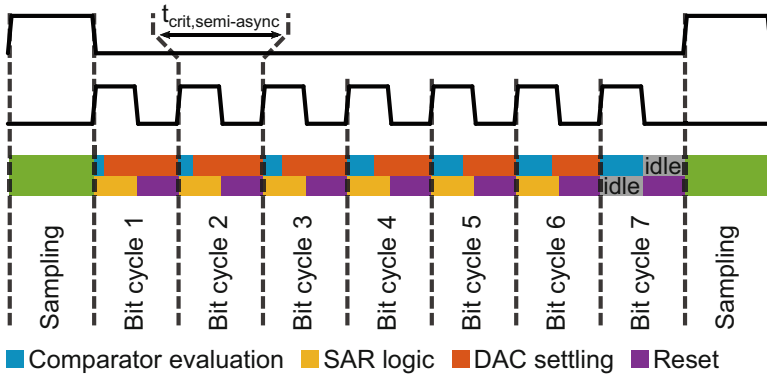


Fig. 5.7 Implemented semi-asynchronous scheme with the logic delay eliminated from the critical path

Regarding the comparator-DAC internal time sharing, the implemented scheme fully utilizes the SAR nature of an unlikely occurrence of the slowest comparator evaluation and DAC settling in the same cycle. The slowest evaluation occurs with the highest probability in the LSB cycle, where the DAC settling is not important anymore since there are no more evaluations required. By designing the comparator to resolve the smallest required voltage in that cycle, given the allocated duty cycle and tolerable power budget, the ADC sample rate can be directly determined. In the other cycles, the fast evaluations are exploited to improve the DAC accuracy by significantly extending its settling time to more than half a cycle. In combination with the implemented MCS switching scheme to eliminate the slowest MSB DAC settling (see Sect. 5.2.4), redundancy does not bring additional speed benefits, therefore is not employed in this ADC.

It is worth noting that the proposed semi-asynchronous processing is inferior to fully asynchronous schemes regarding metastability. However, its simplicity and combined merits can still make it an attractive candidate for GHz-range operation. The metastability probability increases from MSB to LSB, but its significance increases vice versa. The error from a metastable LSB is negligible if the comparator thermal noise is in the same range. Further, the memory element following the comparator (see Sect. 5.2.6) continues to regenerate on comparator reset, borrowing cycle time from the DAC. If metastability occurs in a cycle, time borrowing is possible since the fast next evaluation imposed by the SAR nature does not require the full DAC settling in that cycle. The DAC speed determines how much time may be borrowed to not incur a wrong next evaluation and that time roughly sets the ADC metastability error probability. If a certain probability is not possible by efficient design, correction schemes as in [96, 135] may be applied, with similar overhead as asynchronous processing.

5.2.3 Dual-Loop Bootstrapped Input Switch

In every ADC, the linearity of the input switch directly impacts and might even dominate the total converter spectral purity, with increasing gravity at GHz sample rates and bandwidths. The non-linear behavior of the switch is mainly attributed to its input signal-dependent on-resistance and parasitic capacitance, both of which generate harmonic distortion that is exacerbated when sampling high-frequency signals. Fast switch bootstrapping to ensure a constant voltage across its gate and source is imperative to achieve a larger than 50 dB sampling linearity at GHz-input frequencies and reduce the amplitude- and frequency-dependent impedance modulation at the T/H input.

The typical bootstrap circuit proposed in [136] is shown in Fig. 5.8. The speed critical loop $M_1 - C_B - M_4$ comprises the series on-resistors R_{M1} and R_{M4} and the combination of C_B and $C_{VG} + C_{par}$, with C_{VG} the gate capacitance of M_S and C_{par} the lumped capacitance of all remaining parasitic contributions at node VG. Assuming this a single-pole loop, its time constant can be approximated as

$$\tau_{BS} = (R_{M1} + R_{M4}) \cdot \frac{C_B(C_{VG} + C_{par})}{C_B + C_{VG} + C_{par}}. \quad (5.7)$$

During the HOLD phase ($SAM = V_{SS}$, $\overline{SAM} = V_{DD}$), M_S is off with its gate pulled to V_{SS} through M_7 and M_8 , and the conversion is occurring on the DAC output based on the sampled value of the previous sampling instant. In parallel, VP is pulled to V_{DD} by M_6 , and V_{DD} is applied across C_B through M_2 and M_3 . In the original circuit, M_3 is implemented as an NMOS device requiring a charge pump to provide a boosted clock between V_{DD} and $2V_{DD}$. Alternatively, M_3 can be replaced by a




Fig. 5.8 Typical bootstrap circuit with its speed critical loop highlighted




Fig. 5.9 Improved dual-loop bootstrap circuit proposed in this work

PMOS device [66], discarding the area consuming charge pump. At the beginning of the TRACK phase ($SAM = V_{DD}$, $\overline{SAM} = V_{SS}$), M_4 turns on, but M_1 does not fully turn on, until VG has reached a sufficiently large value; therefore, its large R_{M1} increases τ_{BS} considerably. This mechanism, in combination with the large parasitics at VG and a large C_B to avoid loss of overdrive due to charge-sharing, limits the bootstrap bandwidth, causing a significant on-resistance modulation of M_S for a substantial portion of the TRACK phase, and results in a loss of sampling linearity at high frequencies.

This work proposes the circuit shown in Fig. 5.9 to alleviate these limitations and minimize τ_{BS} . A PMOS M_3 is adopted in the branch charging C_B with its gate bootstrapped by VG . M_9 in the typical circuit can also be removed if the highest input is one threshold below V_{DD} to keep M_5 on. The key difference with existing works lies with the control of M_1 . It is disconnected from the speed critical loop, relieving VG of its load. Instead, a separate loop utilizing devices $M_{10} - M_{12}$ is added for its control, operating in parallel to the main bootstrap loop. During the HOLD phase, the operation is identical to the typical circuit with the exception of the M_1 gate being pulled to V_{SS} through M_{10} and M_{11} . When the TRACK phase starts, M_4 turns on, but almost simultaneously, M_1 also turns fully on through M_{12} , completely decoupled from node VG . Therefore, both M_1 and M_4 track the input signal together fully bootstrapped with maximum gate-source voltage, thus minimum and constant R_{M1} and R_{M4} from the start of the TRACK phase. The benefit of the separate loop in Eq. (5.7) is twofold, reducing both R_{M1} and C_{par} . This enables VG to track the input signal faster, reducing significantly the impedance modulation of M_S . The reduction in C_{par} improves the VG fall transient as well, leading to a steeper falling edge and a better controlled sampling instant. This is illustrated in Fig. 5.10a.




Fig. 5.10 (a) Timing illustration and (b) simulated M_S on-resistance for the typical and the proposed bootstrap circuit

To further improve τ_{BS} , the bulks of speed critical devices are tied to their sources for minimum on-resistance. In particular, the bulk of M_S is not connected directly to its source but to the bottom plate of C_B (Fig. 5.9). During TRACK phase, the situation is identical to the case of tying the bulk directly to the source since the bottom plate of C_B is shorted with the input through M_1 . During HOLD phase, the bulk of C_B is connected to V_{SS} rather than the input, which reduces the necessity of cross-coupled devices to compensate for signal feed-through. From a layout perspective, the arrangement of the grouped wells minimizes the parasitics, yielding the bulk connections most effective.

The effectiveness of the aforementioned techniques in realizing a low and constant M_S on-resistance is verified and compared to the typical single-loop bootstrap with extracted simulations, the results shown in Fig. 5.10. The sizes of the critical transistors and C_B are kept the same. The TRACK period of 100 ps and the total input/DAC capacitance C_{IN} (see Sect. 5.2.4) necessitate an on-resistance below 60 Ω to ensure a sufficient settling accuracy well before the end of the TRACK period

$$t_{\text{TRACK}} = (B + 1) \cdot \ln 2 \cdot (R_{MS} + 25 \Omega) \cdot C_{IN}, \quad (5.8)$$

where the 25 Ω is the input equivalent resistance from the 50 Ω source resistance and the internal termination. R_{MS} of the typical circuit experiences a significant modulation for inputs above 300 mV since the VG transient is not fast enough to provide a maximum gate-source overdrive. The proposed circuit preserves R_{MS} around the designed 40 Ω across the entire input range. These simulations further demonstrate a linearity boost of 7 dB when sampling Nyquist inputs at 1.25 GS/s. It should be noted that there have been improvements to the typical circuit of Fig. 5.8,

with some representatives given in [63, 137, 138]. However, most of them, except for [63] that requires one extra phase for successful operation, retain the problematic $M_1 - C_B - M_4$ loop unaltered. Therefore, the comparisons in Fig. 5.10 hold broadly true for such improved circuits as well.

5.2.4 Unit-Switch-Plus-Cap DAC

The CDAC in this work is depicted in Fig. 5.11. Top-plate sampling is employed, and the MCS scheme [68] is adopted due to its reduced energy and symmetry, allowing for a constant, signal-independent current to be drawn from the references. During sampling, the bottom plates of all the capacitors are tied to V_{CM} , while the input signal is sampled on their shared top plate. When sampling is completed, the bottom plates of the capacitors are consecutively switched to either V_{RP} or V_{RN} depending on the comparator evaluation sign. The maximum digital levels of V_{SS} (0 V) and V_{DD} (1 V) are used as reference voltages to provide maximum switch overdrive, and V_{CM} is set to 500 mV to facilitate the comparator trade-offs (see Sect. 5.2.5).

A scheme with an explicit V_{CM} is preferred over splitting each capacitor in two halves to generate it internally. This avoids any matching degradation due to half-sized units. It also saves one wire per bit coming from the logic since four instead of three wires would be required if splitting would be applied. As a compromise, three reference voltages are required in the CDAC instead of two. These references are sufficiently decoupled on-chip and not shared with any other ADC parts to minimize performance degradation in the CDAC. Since this design does not utilize the full signal swing, a fixed capacitance C_H is tied to V_{CM} , which reduces the CDAC signal range to 400 mV on each side. To avoid a possible direct current path between V_{CM} and one of V_{RP}/V_{RN} , each capacitor is first disconnected from V_{CM} (*break*) before it is connected to one of V_{RP}/V_{RN} (*make*). The speed and power benefits of this




Fig. 5.11 DAC topology with a constant V_{CM} and C_H to set the signal range

switching scheme over the conventional [59] or split-capacitor [64, 65] schemes are based on the elimination of the longest and most power-consuming MSB capacitive settling prior to the first comparator evaluation. This eventually removes the MSB capacitor itself, thus requiring only 2^{B-1} unit cells for a B -bit quantization. Further, the common-mode voltage is kept constant during conversion, unlike [66] and [139] where it drops after every CDAC switching, affecting comparator accuracy and compromising the achievable linearity of the ADC.

It was highlighted in Sect. 5.1 that the CDAC settling imposes one of the major delays in every SAR ADC; therefore, minimizing it while still meeting the required accuracy is of great importance. Typically, a settling accuracy better than $\text{LSB}/2$ is required to prevent dynamic errors. The settling time constant τ_{CDAC} is determined by the unit capacitance C_U and its corresponding reference switch on-resistance R_{ON} . On top, the wiring parasitic resistance R_W and capacitance C_W can add a significant contribution. Again, approximating the switch-CDAC path by a single-pole network, τ_{CDAC} is given as

$$\tau_{\text{CDAC}} = (R_{\text{ON}} + R_W) \cdot (C_U + C_W). \quad (5.9)$$

The CDAC settling time can then be expressed as

$$t_{\text{CDAC, set}} = (B + 1) \cdot \ln 2 \cdot (R_{\text{ON}} + R_W) \cdot (C_U + C_W). \quad (5.10)$$

In conventional Unit-Cap (UC) SAR CDACs, the reference switches are placed somewhere along the path between the SAR logic and the CDAC [107]. This can lead to a large R_W and/or C_W , which can significantly increase τ_{CDAC} , hence prolonging $t_{\text{CDAC, set}}$. From these parasitics, C_W could even dominate the total capacitance in case very small unit capacitors ($C_U \sim 1$ fF) are employed. The schematic of a conventional UC CDAC highlighting the aforesaid is shown in Fig. 5.12a, together with its post-extracted simulated settling. The settling time to reach $\text{LSB}/2$ accuracy is within the 50 ps half cycle. This does not guarantee that the timing can be easily met across all corners. Further, it minimizes the time borrowing margin for the memory element to keep regenerating in case of a metastable comparator evaluation.

In [140], an attempt is made to reduce the settling time by placing the local decoders under the CDAC. Although the distance between the unit capacitors and their switches is reduced, metal shielding is necessary to prevent some of the unwanted digital activity coupling to the CDAC. This shielding can create unnecessary increase in C_W , yielding the settling time reduction sub-optimal. Furthermore, the used shielding metals are right below the capacitors, which can result in losing significant signal range owing to large capacitive division.

This work introduces a USPC CDAC topology, which simultaneously minimizes the C_W and R_W contribution to the CDAC settling by merging the reference switches with C_U into a single cell, making them part of the CDAC. The schematic of the proposed topology is shown in Fig. 5.12b. Both C_W and R_W are massively reduced in the critical path, while the switches are kept small and still easy to drive without

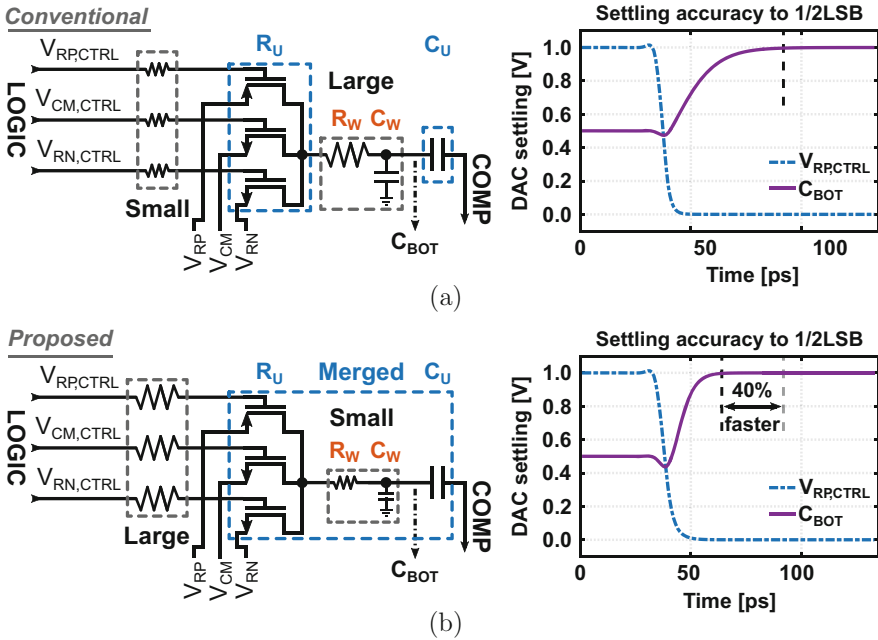


Fig. 5.12 Schematic and simulated settling time of (a) a conventional UC CDAC and (b) the proposed USPC CDAC

too much excess delay from the logic, despite the increase in their gate resistance. This increase is attributed to the longer interconnect of the logic control signals to the reference switches' gates. Routing these signals on higher metals makes their parasitic effect on CDAC settling much more benign than the critical C_W and R_W . Since the critical parasitics are minimized, the settling not only becomes faster, but it is also more uniform across the cells, determined by the “clean” resistance and capacitance in the path. The post-extracted simulated settling in Fig. 5.12b shows about 40% shorter settling per cycle compared to the typical topology. This allows an equivalent time to be borrowed for extra regeneration or an aggregate 14% ADC sample rate boost accumulated over the seven cycles.

A partial 4-bit layout of the single-sided USPC CDAC is shown in Fig. 5.13. The CDAC implementation is done in two columns: the left column contains the switchable unit cells, while the right column incorporates the unit elements for C_H . Common-centroid arrangement is followed, and the reference switches are connected to the SAR logic through vertical wires. Dummies are placed on both sides of the CDAC (not shown) to guarantee identical environment for all the units. Also, the area below the units is kept empty to prevent accuracy degradation of the capacitors. An aspect ratio of roughly 1:2.5 is used in the complete CDAC to avoid too long wires coming from the logic and excessive parasitic capacitance at the comparator input. The latter has been taken into account when designing

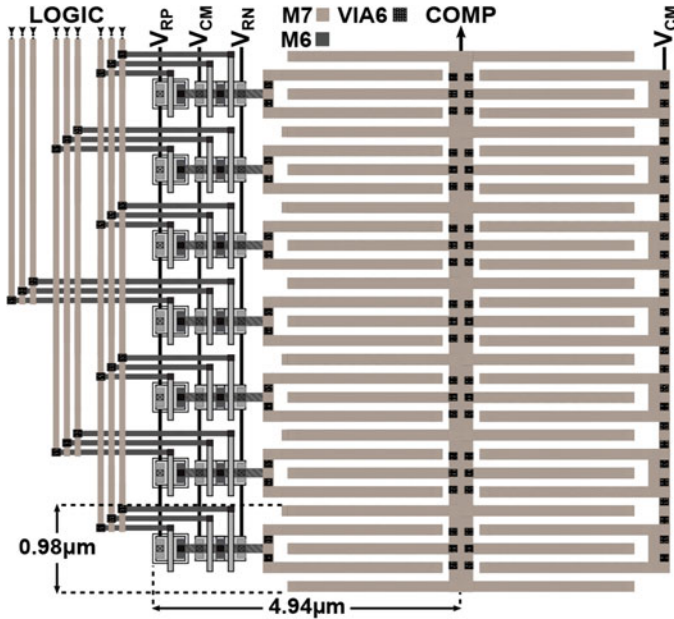


Fig. 5.13 Single-ended partial layout of the USPC CDAC (the actual implementation is differential)

C_H in order to compensate for signal range loss due to capacitive division. Such an implementation can be easily adopted in designs with different resolutions. Depending on the intended design requirements, a proper aspect ratio can be realized to balance the various trade-offs.

Each switchable unit cell comprises a custom-designed Metal-Oxide-Metal (MOM) capacitor and its corresponding reference switches. The distance between capacitor and switches is kept minimum to eliminate simultaneously R_W and C_W while still ensuring small parasitic coupling between the top plate and each bottom plate as well as between the different bottom plates. Metals 6 and 7 are used due to their distance from the substrate, to realize a unit capacitance of 1.25 fF. The single-ended CDAC capacitance of 200 fF is larger than the noise-limited value to ensure a raw matching above 7 bits. This was verified by mismatch simulations of standard library plate capacitors with roughly the same value and area. After determining the unit capacitance, the switch R_{ON} is designed to minimize the CDAC time constant and yield the settling in Fig. 5.12b. NMOS devices are used for V_{RN} and PMOS devices for V_{RP} , sized for matched impedances. NMOS devices are also used for V_{CM} , equal in size to the ones for V_{RN} , whose R_{ON} suffices for the allocated sampling time.

5.2.5 Triple-Tail Dynamic Comparator

The comparator is an integral block in every high-speed ADC [12, 92, 110, 111], with increasing importance in a SAR due to the bit-at-a-time operation. Its noise, kickback, and common-mode sensitivity determine to a large extent the total ADC accuracy. Its resolving ability has a major influence on the speed and metastability, while its overall design renders it a significant contributor to the total power budget. The two main design parameters, speed and noise, impose two fundamental limits and were analyzed in Chap. 2. These two parameters adversely affect each other; therefore, significant effort should be allocated to properly address this trade-off.

The schematic of the implemented comparator in this work is drawn in Fig. 5.14. It comprises a cascoded integrator as the first pre-amplifier followed by a second pre-amplifier, which acts as both integrator and half-latch, driving finally the output latching stage in a triple-tail fully dynamic arrangement. The multi-stage configuration softens the trade-off between the different design parameters, providing a more orthogonal optimization for each parameter. This allows the comparator to achieve both high speed and low noise/offset. The first pre-amplifier defines the noise/offset; therefore, it is optimized for that. Its gain also helps attenuating the noise/offset of the following stages. The NMOS cascode devices on top of M_{1P}/M_{1N} help isolate nodes XP/XN from the parasitic capacitance of the input pair during integration. The cascodes also isolate the input pair from the kickback generated on those nodes upon reset. The latching stage sets the speed; therefore, it is optimized to have a very low time constant τ_{comp} . The second pre-amplifier suppresses the output noise and provides further signal gain, enhanced by the cross-coupling, prior to the latch, thus minimizing its regeneration time. The intermediate devices M_{2P}/M_{2N} and M_{3P}/M_{3N} act both as gain devices and as reset devices to provide further shielding from the

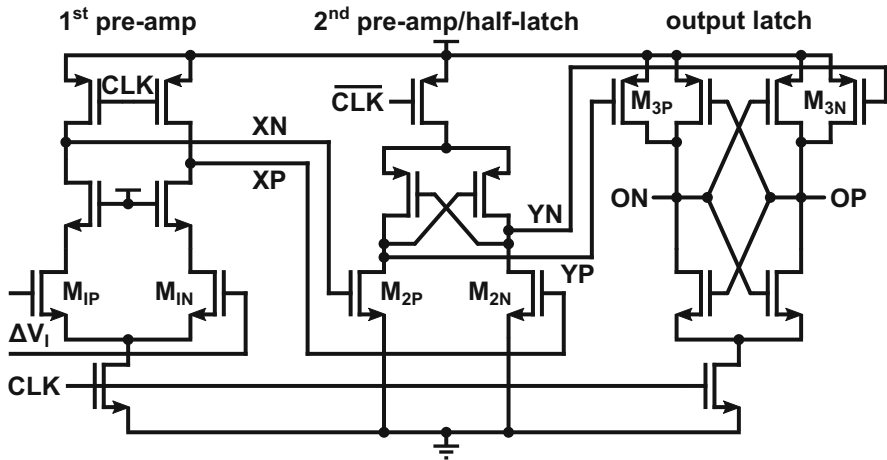


Fig. 5.14 Schematic of the implemented triple-tail dynamic comparator




Fig. 5.15 Simulated performance of the triple-tail comparator

latch output noise, as well as reset devices for nodes YP/YN and OP/ON. Hence, the need for additional reset transistors at these nodes is obviated, which reduces their capacitance.

When CLK is low ($\overline{\text{CLK}}$ is high), nodes XP/XN and OP/ON are reset to the supply voltage, while YP/YN are pulled to ground. When CLK goes high (CLK goes low), the drain currents of the input pair discharge nodes XP/XN toward ground with different slopes depending on the input signal, while nodes YP/YN are charged toward the supply with an increased slope difference due to the extra gain. At the same time, the NMOS pair of the output latch is activated, pulling down OP/ON whose slope difference is further increased due to the intermediate transconductors M_{2P}/M_{2N} and M_{3P}/M_{3N} . When one of them reaches one PMOS threshold below the supply voltage (about 270 mV for ultra-low threshold devices), latching takes place. For large differential inputs, the second stage suffices as a latch due to its positive feedback, and the final stage can be seen as an extra digital buffer toward the SAR logic.

The aforementioned are depicted in Fig. 5.15, where the post-extracted performance of this comparator has been characterized under the overdrive recovery test, the most stressful performance assessment. In two consecutive cycles of 10 GHz, the differential input ΔV_1 toggles between the supply rails and a very small signal with opposite polarity. For the large input, the differential pair steers all the current to one side, producing a large difference for the following stages to resolve. When the input changes polarity, the two amplification stages have to recover and change the polarity of nodes YP/YN in time for the latching to yield a correct sign. This comparator is able to evaluate to a sufficiently large differential level, free of




Fig. 5.16 Simulated outputs of the triple-tail comparator (top) and one of the logic memory latches (bottom)

any memory effect, input differences smaller than $\text{LSB}/10$ within the maximum allocated 50 ps half cycle.

It is important for the comparator to be able to evaluate as small as possible input differences within the targeted time, in order to reduce the overall metastability error probability. As discussed in Sect. 5.2.2, the comparator is followed by one memory element (latch) for each bit as part of the SAR logic (see Sect. 5.2.6). Each latch operates in parallel to the comparator, triggered by one of the generated bit phases, and continues regenerating during comparator reset. The regeneration ability of the two combined with the settling speed of the CDAC defines to a first extent the probability of certain code errors for the targeted ADC sample rate.

To investigate this effect, the simulated outputs of both the comparator and the memory latch for input differences down to 1 nV ¹ are shown in Fig. 5.16. The comparator is loaded by all memory latches of the logic, and one of them is triggered to capture its output. For input voltages of about $10 \mu\text{V}$ and below, the comparator cannot provide a differential output greater than $V_{\text{DD}}/2$ within its allocated 50 ps. However, for some of these voltages, the memory latch can still provide valid levels and switch the CDAC, allowing some time for partial settling within the current bit cycle. For input differences between $1 \mu\text{V}$ and 1 nV , the differential output of the

¹ It is possible to simulate smaller input differences (e.g., 1 fV). However, systematic offset stemming from the circuits as well as simulator tolerances requires accurate compensation and/or settings to yield correct results.

memory latch (buffered or inverted) is still large enough to leave about 15 ps for the CDAC to settle. In our design with the USPC topology, this would result in a partial settling to 4 LSB accuracy. This can still yield correct bits since a full CDAC settling is not a must prior to a fast next evaluation, imposed by the SAR nature.

The estimated τ_{comp} from the waveforms in Fig. 5.16 is about 6 ps. Taking into account that an input change of 10^x changes the evaluation time by $x \ln 10 \cdot \tau_{\text{comp}}$, the metastability error probability to allow a maximum of 85 ps for the combined comparator and memory latch is about $1e^{-6}$. Possible ways of partially improving this value without compromising the ADC sample rate include creating more gain in the comparator and/or reducing τ_{comp} , which increase the power. Alternatively, detection and correction schemes as in [96, 135] may be applied to reduce metastability to tolerable levels.

Input common-mode voltage $V_{\text{CM,I}}$ is another important aspect of the comparator design, which affects both its resolving time and noise. A higher $V_{\text{CM,I}}$ increases the current through the input pair, leading to a shorter integration time. However, this increases the noise integration bandwidth as well; therefore, its value is of significant importance to achieve the optimum between them. The effect of changing $V_{\text{CM,I}}$ on the resolving time and input-referred noise is shown in Figs. 5.17a and b, respectively. For comparison, the single-stage strong-ARM [110] and the two-stage double-tail [111] counterparts are simulated and plotted as well. All the comparators are sized for similar input-referred noise/offset and latching strength. The triple-tail comparator shows a faster resolving time with a lower common-mode dependence for a wide range of voltages compared to the strong-ARM and double-tail circuits due to the extra stage, which allows for more design flexibility. An optimum exists between 500 mV and 600 mV. This is explained by the fact that too high a $V_{\text{CM,I}}$ reduces the amplified voltage difference seen by the latching stage due to a shorter integration time, slowing down the latch.

Input-referred noise increases almost linearly with $V_{\text{CM,I}}$ and is very similar for all the comparators. As a result, in this design, a $V_{\text{CM,I}}$ of 500 mV was chosen for a near-optimum resolving time and a small enough input-referred noise with respect to the LSB size. This voltage comes directly from the input and preserves the MCS switching symmetry in the CDAC. The input-referred offset is also simulated with a $1 - \sigma$ value of 9–10 mV_{rms} for all comparators. This offset is typically not a problem in a single-comparator SAR ADC since it results in a global offset.

The resolving time and energy/comparison versus the input difference for the three different comparators are also shown in Figs. 5.17c and d, respectively. For very small input differences around the LSB range, where the speed of the comparator determines to a great extent the maximum ADC sample rate, the proposed comparator offers more than 20% resolving time improvement for the aforementioned sizing conditions. When the input difference increases to levels around 50 mV and above, the proposed comparator shows a slightly larger resolving time. This is attributed to the three stages adding more gate delay compared to the single-stage strong-ARM and the two-stage double-tail. As the comparator resolving time is sufficiently short for such inputs, this larger value is not limiting the ADC sample rate.

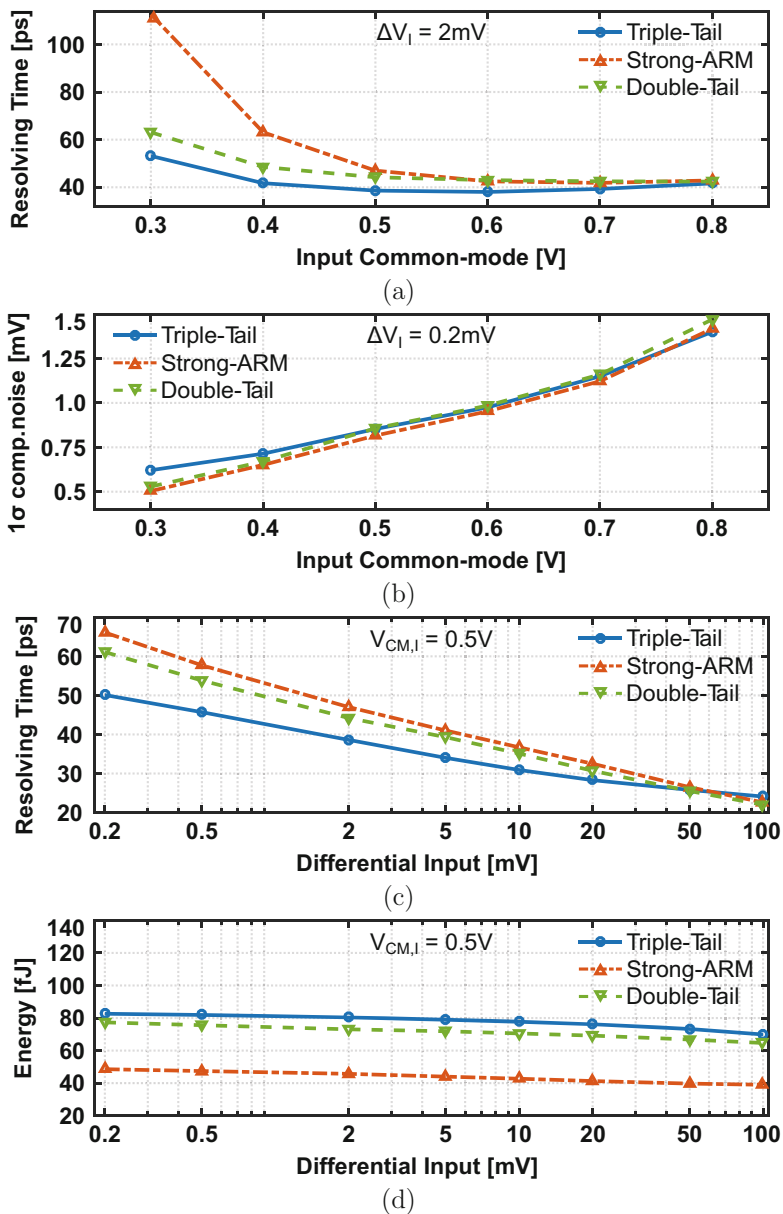


Fig. 5.17 (a) Simulated comparator resolving time and (b) input-referred noise versus $V_{CM,I}$ and (c) resolving time and (d) energy versus ΔV_1

Energy/comparison is computed by dividing the simulated comparator power with the maximum frequency it can resolve the smallest shown input (0.2 mV), while clocked at that frequency. Under this setup, the triple-tail comparator achieves

an energy/comparison on the same order as the double-tail and about 35% higher than the strong-ARM latch. The comparator contributes about 29% of the total ADC power (see Sect. 5.3) while dominating the total ADC sample rate. Therefore, the speed benefit overcomes the higher energy/comparison and, accumulated over the SAR cycles, improves the overall ADC *accuracy · speed ÷ power*.

5.2.6 Custom SAR Logic

The custom SAR logic in this work comprises two core parts: (1) the clock generation, responsible for providing the clock for the comparator and the bit phases, and (2) the state memory, responsible for storing the comparator output and switching the CDAC based on the evaluation sign in each of the provided bit phases. The top-level logic diagram is depicted in Fig. 5.18, where these parts are highlighted. The inverted comparator outputs are connected to a pair of memory cells that switch the CDAC from MSB-1 to LSB. The MCS scheme, together with one CDAC settling, avoids its corresponding memory cells as well.

The clock generation combines the 10 GHz full rate clock and the 1.25 GHz sampling pulse SAM to generate the comparator clock with simple combinational logic (NAND, NOR). At the same time, $B - 1$ master latches ML_i , and slave latches SL_i are employed that are controlled by the full rate clock. Their outputs are combined by simple gates to provide the bit phases $PHASE_i$ sequentially. To attain a maximum sample rate, matched critical paths are ensured between the full rate 10 GHz clock, the 12.5% duty-cycle sampling pulse SAM, and the outputs of

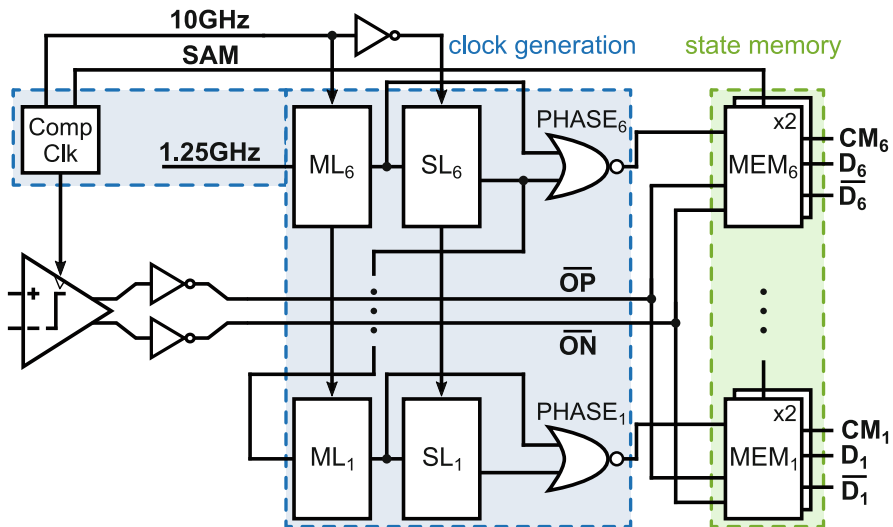


Fig. 5.18 Custom SAR logic including the comparator clock, bit phases, and memory elements




Fig. 5.19 Schematic of one memory cell with optimized critical path toward the CDAC reference switches (top). Timing diagram and truth table of the memory cell (bottom)

the comparator clock and phase generators $PHASE_i$. These phases are aligned with the comparator’s evaluation time for its output to propagate immediately, allowing the maximum remaining cycle time for CDAC settling.

One cell of the state memory part containing the memory latch covered in the comparator discussion is shown in Fig. 5.19. Each cell connects to the inverted differential output of the comparator through pass transistors and provides the control signals CM_i , D_i , and \bar{D}_i for the V_{CM} , V_{RP} , and V_{RN} switches, respectively. All cells are simultaneously reset during sampling, and one cell is activated during every bit cycle by its corresponding phase $PHASE_i$.

When the sampling starts, CM_i of all the cells are high, passing V_{CM} to the bottom plates of the CDAC capacitors, while D_i and \bar{D}_i are such that both V_{RP} and V_{RN} switches are turned off. At the same time, the outputs of all memory latches LP_i and LN_i are reset to ground. These latches are implemented as cross-coupled tri-state pairs with a similar latching strength as the comparator final latch. After the sampling is finished and one of the bit phases $PHASE_i$ is generated, CM_i goes low, turning off its corresponding V_{CM} switch. At the same time, the critical path with the memory latch is transparent to the output of the comparator. Depending on the evaluation sign, D_i and \bar{D}_i go either both low or both high to turn on one of the reference switches in the CDAC. In this path, the number of gates has been optimized for a minimum rise/fall time and propagation delay product, given the loading of the CDAC and interconnect. During this mode, D_i and \bar{D}_i of a memory cell pair controlling the differential CDAC capacitor of the same bit have opposite sign.

5.3 Experimental Verification

The prototype ADC is fabricated in a single-poly ten-metal (1P10M) 28 nm bulk CMOS process. A die micrograph is shown in Fig. 5.20, measuring a total area of $790\ \mu\text{m} \times 1380\ \mu\text{m}$. The SAR core is also shown, occupying an area of $49\ \mu\text{m} \times 145\ \mu\text{m}$. The CDAC takes up about half of the area, with the capacitors of the bootstrapped input switch and the complete SAR logic following. The three wires of the USPC compared to the one of a conventional implementation are offset by the bootstrap capacitors, thus not imposing any area overhead.

The placement of the core ADC blocks is carefully optimized to minimize long interconnect in the critical path in order to enhance speed and minimize power. The differential input signal is applied at the bottom of the chip, whereas the clock is coming from the left, and the outputs are collected serially from the top. The input is sampled onto the CDAC, right above the bootstrapped switch. The comparator interacts with the CDAC as well as with the SAR logic. Therefore, it is located in between these two blocks to minimize the wiring. The logic block is placed on top of the comparator and connects to the CDAC reference switches, closing the SAR loop. Differential symmetry is kept as much as possible in the ADC, adding dummies to create the same environment around critical blocks.

The nominal ADC input swing is $800\ \text{mV}_{\text{pp,diff}}$ centered around a $500\ \text{mV}$ common mode. The ADC utilizes multiple core $1\ \text{V}$ supply domains and the measured core power consumption of $3.56\ \text{mW}$ (excluding clock generation and output buffers) at $1.25\ \text{GS/s}$ partitions into $0.47\ \text{mW}$ for the bootstrapped input




Fig. 5.20 Die micrograph of the 28 nm IC with a zoomed-in view of the SAR core occupying an active area of $0.0071\ \text{mm}^2$

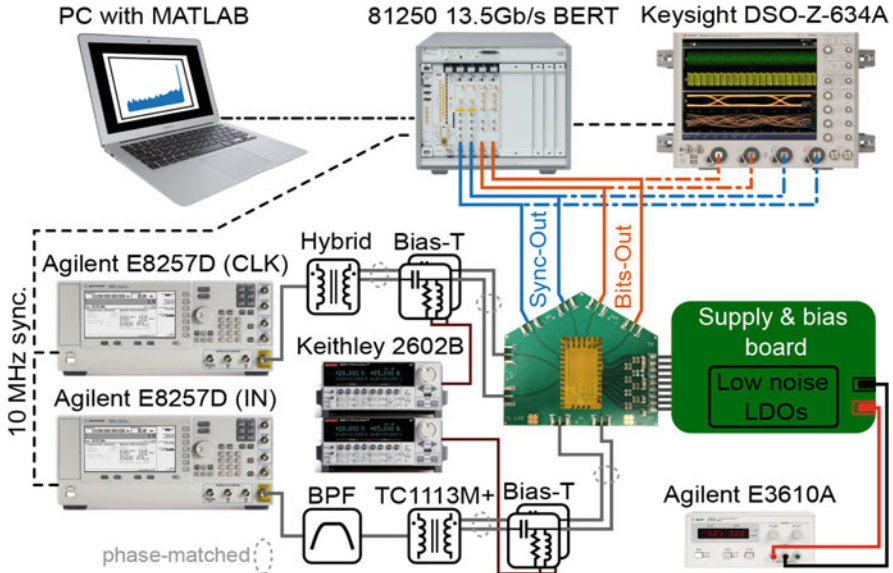


Fig. 5.21 Measurement setup of the 1.25 GS/s 7-bit SAR prototype

switch, 0.6 mW for the USPC CDAC, 1.06 mW for the triple-tail comparator, and 1.43 mW for the phase and state SAR logic, based on measurement results.

5.3.1 Measurement Setup

The complete measurement setup used to evaluate the ADC performance is depicted in Fig. 5.21. An Agilent E8257D signal source is used to generate the input signal. Its spectral purity is improved by adding a ninth-order band-pass filter for every frequency under test. An identical signal source is employed to generate the 10 GHz sinusoidal ADC clock. The combined integrated jitter of both the signal and the clock sources is below $100 f_{\text{srms}}$ in a bandwidth from 1 kHz to 5 GHz. Both input and clock signals are converted into differential signals by wideband baluns and AC-coupled to the chip through custom-designed bias-Ts and phase-matched cables.

The signal generators are synchronized with each other as well as with a Keysight DSO-Z-634A scope and an Agilent BERT. There are two differential outputs coming from the chip serially at the full 10 Gb/s data rate, Bits-Out and Sync-Out. These are first connected to the scope for waveform inspection and then captured by the BERT for logic analysis. The captured data are finally processed to a PC in MATLAB. The last memory cell output of the SAR logic is buffered and serves as the Sync-Out output. This signal is reset during sampling and activated only at the LSB+1 cycle (see Sect. 5.2.6). In every conversion period, if the LSB+1 is a digital

“1”, Sync-Out is a digital “0” and remains in this state until the next sampling. In all other cases, this signal is a digital “1”, with this function being incorporated into the output buffers. The position that such a transition occurs for the first time is identified, and from there, the MSB of the next digital word is located after two cycles. To ensure a proper alignment, Bits-Out and Sync-Out are buffered and routed identically on both chip and board level. Further, phase-matched cables are used to connect them to the BERT for performance evaluation.

The required supply and bias voltages for the different domains in the chip are generated with dedicated low-noise LDOs on the custom bias board and provided to the chip board after sufficient low-pass filtering. Two dual-channel Keithley sourcemeters are used to provide the input and clock common-mode voltages, respectively. The option of the two channels is particularly useful for the input, in order to compensate the comparator offset. The offset is characterized by using the same 500 mV common-mode voltage at both inputs and monitoring the digital word. The common-mode voltages are then adjusted to the values, where the ADC produces the mid-code word.

5.3.2 Measurement Results

Several measurements are performed in order to characterize this prototype. These include both static and one-tone spectral measurements. First, the DNL and INL are characterized by means of the histogram (code density) test [27]. The measured characteristics are plotted in Fig. 5.22 for a sinusoidal input of 160 kHz at a sample rate of 1.25 GS/s. Both DNL and INL lie within $+0.14/-0.46$ LSB (Fig. 5.22a) and $+0.37/-0.41$ LSB (Fig. 5.22b), respectively, verifying the possibility of achieving above 7-bit matching accuracy with the implemented CDAC topology and layout. Still, there exist systematic DNL and INL jumps around the quarter, half and three quarter of the full-scale input, especially profound around the half. Although two dummy USPC cells are placed at the top and bottom sides of the CDAC, they are not switched to provide exactly the same environment for all units. Also, the interconnect from the logic was not replicated in these dummies since it would require more space between the comparator, input switch, and CDAC, increasing the speed critical interconnect. Since units of the partitioned MSB USPC cells exist at the sides of the CDAC (Fig. 5.13), they are affected the most. Missing these jumps during simulation could be partially attributed to the inaccuracy of the extraction tool, given the custom-designed MOM capacitors.

Spectral measurements are performed for different input frequencies and different sample rates. At 1.25 GS/s, the measured output spectra for input frequencies of 623 MHz (1st Nyquist zone) and 4.93 GHz (8th Nyquist zone, folded), respectively, are depicted in Fig. 5.23. The most important metrics are annotated in the plots. At




Fig. 5.22 Measured static performance with the histogram (code density) test at 1.25 GS/s for a sinusoidal input of 160 kHz: (a) DNL and (b) INL

a Nyquist input frequency,² the achieved SNDR is 40.1 dB (Fig. 5.23a), limited by thermal noise, whose main contribution is the comparator, followed by the quantization noise and the CDAC thermal noise. The SNR includes both thermal noise and clock jitter, with the latter being negligible at this frequency. At an $8\times$ Nyquist input frequency (Fig. 5.23b) the achieved SNDR of 36.4 dB is limited largely by the loss of signal gain, while the accumulated jitter from both on-chip and off-chip sources also comes into play. In both spectra, the SNDR is SNR-dominated, while the tones that stand out the most are the harmonic tones originating from the input switch.

Something worth pointing out is that at the Nyquist input frequency (Fig. 5.23a), the 3rd harmonic is suppressed and the 5th harmonic is the one that dominates the SFDR. This is verified in multiple samples, but there is no clear explanation of its root cause. One speculation is that somehow another non-linear effect in the signal chain counteracts the 3rd harmonic for a fundamental around the Nyquist frequency.

² The input frequency is chosen slightly smaller than the actual Nyquist and mutually prime with the sample rate. This is applied to all input frequencies henceforth.




Fig. 5.23 Measured output spectra at 1.25 GS/s for (a) a Nyquist input frequency and (b) an $8\times$ Nyquist input frequency

A similar effect is noticed in the 7th and 11th harmonics, but these are inherently low to begin with.

Figure 5.24 plots the measured SFDR/SNDR versus the input frequency at 1.25 GS/s (Fig. 5.24a) and versus the sample rate at a 76 MHz input (Fig. 5.24b) for four different samples. The high linearity and internal bandwidth of the proposed bootstrapped switch allow for a flat SFDR in the range of 50 dB and above up to 5 GHz, making this ADC a good candidate for larger system integration. The SNDR is relatively flat and above 41 dB up to around 300 MHz and stays above 40 dB at Nyquist and above 36 dB up to 5 GHz. When sweeping the sample rate, the SFDR remains above 50 dB up to about 1.125 GS/s, while also the SNDR is relatively flat and larger than 41 dB for the same sample rate range. At sample rates higher than 1.125 GS/s, they both start degrading gradually as the cycle time becomes too short for the tracking and conversion to complete properly, given that the corresponding circuits are not optimized for these sample rates. The different samples achieve very similar performance within $\pm 1\text{--}2$ dB spread, indicating the robustness of the proposed techniques.

Finally, Fig. 5.25 plots the measured FoM versus the input frequency at 1.25 GS/s (Fig. 5.24a) and versus the sample rate at a 76 MHz input (Fig. 5.24b) for the aforementioned four samples. A nearly constant FoM of less than 36 fJ/conv-step

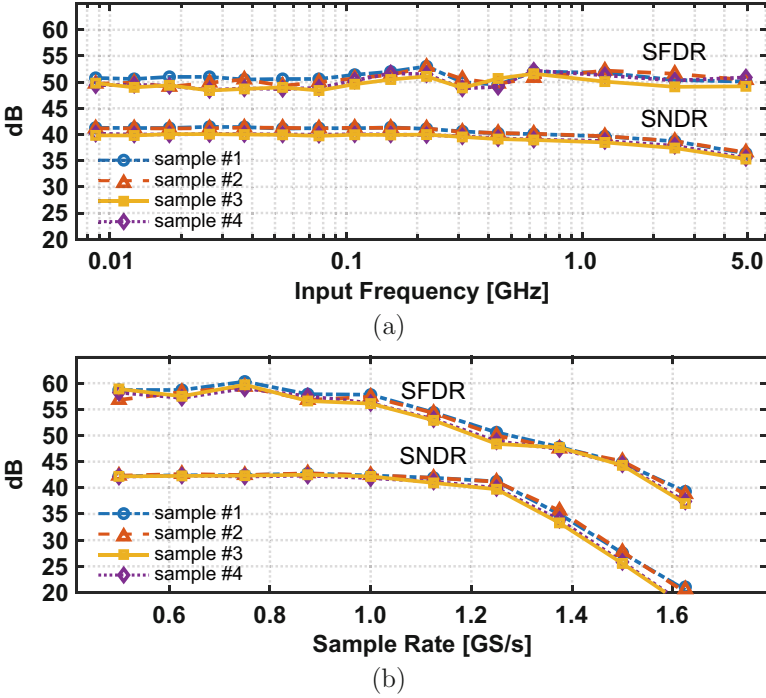


Fig. 5.24 Measured SFDR/SNDR versus (a) input frequency at 1.25 GS/s and (b) sample rate for a 76 MHz input

up to Nyquist is preserved for all samples, which then deteriorates smoothly. Across the various sample rates, there is a shallow optimum FoM between 1.125 GS/s and 1.25 GS/s of about 30 fJ/conv-step, which degrades above 1.25 GS/s following the SNDR trend.

5.3.3 State-of-the-Art Comparison

The measurement results of this work are summarized in Table 5.1, together with a recent SotA comparison among SAR ADCs of similar performance in deep-scaled CMOS processes [36]. The comparison includes metrics from single-channel SAR ADCs or from one channel of a TI system. The presented speed-boosting techniques allow this ADC to achieve among the highest sample rates of 1.25 GS/s. The optimized minimalist approach shows the lowest SNDR drop from the designed aggregate quantization level at both low and Nyquist input frequencies of 2.6 dB and 3.9 dB, respectively. This ADC also attains an input sampling ability of $8\times$ its Nyquist frequency with an overall power dissipation, area, and FoM that is on




Fig. 5.25 Measured FoM versus (a) input frequency at 1.25 GS/s and (b) sample rate for a 76 MHz input

par with the SotA, without employing any calibration. Compared to works with the same sample rate [95, 107], it achieves a larger SNDR/SFDR and a lower FoM, while compared to works with a similar SNDR/SFDR (within ± 1 dB), it demonstrates a higher sample rate.

5.4 Conclusion

This chapter discussed techniques and challenges for extending the sample rate of low-to-medium-resolution single-channel SAR ADCs in the GHz range while not compromising their highly digital nature, excellent efficiency, and simplicity. Such ADCs are of particular importance in a variety of applications, both as standalone blocks and integrated into larger systems.

First, the conventional SAR algorithm was studied from a timing perspective, in order to understand its fundamental speed-limiting factors. It was identified that in a conventional synchronous clocking scheme, three sequential critical timings occur within each bit cycle: (1) the comparator evaluation time t_{comp} , (2) the DAC settling

Table 5.1 Performance summary and comparison with state-of-the-art SAR ADCs

	This work	Wei [129]	Kull [118]	Le Tual [107]	Chan [95]	Choo [139]	Chan [130]
	ESSCIRC'17	ISSCC'11	ISSCC'13	ISSCC'14	ISSCC'15	ISSCC'16	ISSCC'17
Technology	28 nm CMOS	65 nm CMOS	32 nm SOI	28 nm FDSOI	65 nm CMOS	40 nm CMOS	28 nm CMOS
Architecture	1b/cycle	2b/cycle	2 comp.	1b/cycle	3b/cycle	1b/cycle ci	1+2b/cycle
Interleaving factor	1x	1x	1x	8x	4x	1x	2x
Resolution [bits]	7	8	8	6	6	6	7
Supply [volts]	1.0	1.2	1.0	1.0	1.0	1.0	0.9
Sample rate [GS/s]	1.25	0.4	1.2	1.25 ^a	1.25 ^a	1.0	1.2 ^a
Max. f_{in} [GHz]	5.0	0.2	0.6	20.0	2.5	0.5	1.2
SFDR @ low f_{in} [dB]	50.8	55.0	N.A.	N.A.	44.1	N.A.	54.1
SNDR @ low f_{in} [dB]	41.4	44.5	39.6	33.9	32.0	35.1	40.0
SFDR @ Nyquist [dB]	52.0	53.0	49.8	41.1	43.1	49.7	54.3
SNDR @ Nyquist [dB]	40.1	40.4	39.3	33.8	30.8	34.6	40.0
Power consum. [mW]	3.56	4.0	3.1	4.0 ^b	1.375 ^a	1.26	2.5 ^{a,c}
Active area [mm ²]	0.0071	0.024	0.0031	0.00072 ^a	0.0225 ^a	0.00058	0.0022 ^a
Calibration	NO	YES	YES	YES	YES	N.A.	YES
FoM ^b [f]/conv-step]	34.4	116.9	34.0	80.4	39.0	28.7	25.3 ^c

^a Metric per single-channel^b @ Nyquist f_{in} ^c Power of multiple DACs not reported

time t_{DAC} , and (3) the digital logic delay $t_{\text{logic, sync}}$. The main speed limitation stems from the need of each cycle to accommodate the longest t_{comp} (LSB) and t_{DAC} (MSB). However, these do not appear in every cycle, resulting in time wasting.

Next, some noteworthy prior art techniques that have successfully tackled the above speed limitation were discussed, highlighting their advantages and drawbacks. These techniques included (1) asynchronous processing, (2) multi-bit per cycle, (3) multi-comparator loop unrolling, (4) redundancy, (5) pipelining, and (6) various DAC switching schemes. All of them dealt with reducing or removing one or more of the three critical timings mentioned above to yield an overall faster SAR conversion with affordable trade-offs.

Finally, a single-channel SAR ADC was presented, with several proposed techniques to maximize $\text{accuracy} \cdot \text{speed} \div \text{power}$. On the architectural level, a single-bit per cycle, single-comparator topology was chosen for its minimal complexity. A semi-asynchronous processing was introduced to combine the merits of simple logic and cycle control from synchronous schemes with the dynamically allocated internal timing of asynchronous schemes. Further, the logic delay was eliminated from the critical path by overlapping with the comparator evaluation. On the circuit level, a dual-loop bootstrapped input switch was proposed to improve the input bandwidth and high-frequency linearity. A USPC CDAC topology and a triple-tail dynamic comparator were also proposed, to reduce the settling and evaluation times, respectively.

Fabricated in a 28 nm bulk CMOS process, the prototype SAR ADC employing the proposed techniques achieves a sample rate of 1.25 GS/s with a Nyquist SNDR/SFDR of 40.1/52.0 dB, which remain still 36.4/50.1 dB at a 5 GHz input frequency without any calibration. The FoM of 34.4 fJ/conv-step is achieved while consuming 3.56 mW from a 1 V supply. With a core area of 0.0071 mm², this ADC can be smoothly integrated into a larger system.

Chapter 6

High-Resolution Wide-Bandwidth Time-Interleaved RF ADC



The previous chapter emphasized techniques to efficiently boost the low-to-medium-resolution (6–8-bit) single-channel SAR speed as a favorable base candidate for integrating into a larger system. This chapter delves deeper into architectural and circuit capabilities to enable a higher ADC resolution (>10 bits) while preserving the multi-GHz sample rate and bandwidth and maximizing the efficiency. Such high-resolution, multi-GHz sample rate and bandwidth, low-power RF ADCs are of interest in next-generation wideband communication, data acquisition, and instrumentation applications.

Section 6.1 of this chapter overviews the needs and the challenges for efficiently realizing such RF ADCs by means of their role in the receiver chain. Common ADC architectural choices are briefly reviewed and their trade-offs discussed. Section 6.2 presents the prototype TI hybrid RF ADC and details its performance-enabling principles. The sub-ADC architecture and interleaving choices are greatly motivated by the analyses developed in Chap. 3. The experimental verification, including the detailed measurement setup, the measured results, and a comparison with recent state of the art, are treated thoroughly in Sect. 6.3. Finally, the conclusion is drawn in Sect. 6.4.

Parts of this chapter were previously presented at the 2019 *International Solid-State Circuits Conference (ISSCC'19)* in San Francisco, CA, USA [91], and published in the *Journal of Solid-State Circuits (JSSC'20)* in June 2020 [93].

6.1 RF Sampling ADCs: Needs and Challenges

The constant demand for higher throughput and bandwidth in next-generation wireless and wireline communications, such as 5G massive Multiple-Input Multiple-Output (MIMO) and Data Over Cable Service Interface Specification (DOCSIS) 4.0 over Hybrid Fiber Coax (HFC) networks, has triggered the need for multi-GS/s

ADCs to digitize several GHz of bandwidth with high spectral purity and power efficiency. Such wideband/multi-band applications occupy multiple channels up to 6 GHz (5G sub-6 GHz FR1) [141] or extend the signal bandwidth up to 3 GHz (extended spectrum DOCSIS 4.0) [142]. Low-power ADCs in deep-scaled CMOS are desirable to enable a direct sampling at these frequencies, offering an increased flexibility and integration capability with lower cost and footprint compared to traditional approaches.

6.1.1 The ADC Role in the Receiver

To better understand the increasing importance and challenges of wideband RF ADCs to enable direct sampling in the aforementioned applications, the direct RF sampling receiver architecture is discussed, highlighting the ADC role in the chain. Such a receiver, in its most ideal form, contains only an ADC, while the rest of the functionality and signal processing occurs in the digital domain [7]. However, such an ideal architecture has not yet been implemented in hardware. A more realistic direct RF sampling receiver implementation with increasing popularity and research interest is portrayed in Fig. 6.1 [143]. The wideband RF ADC is led by an anti-aliasing BPF and an RF Low-Noise Amplifier (LNA)/Variable Gain Amplifier (VGA). The latter serves to drive the ADC input load while providing the expected signal range to prevent the ADC from saturating and improve the dynamic range of the total receiver. The VGA may be integrated as a front-end within the converter, allowing an even higher integration and a better controlled interface.

This receiver architecture aims to take full advantage of the constant performance advancements and the improved flexibility of DSP with scaling into finer CMOS processes. In contrast to the more commonly used heterodyne [144], homodyne [145], or low-IF [146] receivers, this architecture removes the majority of the analog conditioning/translation components, such as one or more mixers, BP/LP filters, amplifiers, and the demodulator. Additionally, a single fixed frequency Local Oscillator (LO) provides the clock to the ADC, reducing the complexity




Fig. 6.1 Generic block diagram of a direct RF sampling receiver


regarding the Phase-Locked Loop (PLL) design. Placing the converter closer to the receiver input enables the direct digitization of multi-GHz wideband or multi-band RF signals (L-, S-, and C-bands as per the applications' needs) concurrently, minimizing the receiver chain count to only a few or even one. The remainder of the functionality, such as band/channel selection, mixing, and downconversion, are handled by the inherently flexible and integration-friendly digital back-end. Several standards, including the SDR, can be supported with an increased flexibility. Finally, providing an ADC sample rate of at least twice the total bandwidth of interest can significantly simplify the frequency planning and relax the anti-aliasing filter's roll-off requirements. In this topology, the ADC emerges as the utmost critical block, but also the only block that can turn it from dream to reality. It is worth noting that the RF sampling receiver has been long in demand, with the ADC sample rate and bandwidth being its showstopper due to process limitations. However, its viability has been growing rapidly, driven by the continuous circuit, architectural, and technology advancements.

Despite its integration advantages, area/cost savings, and promise for wideband or multi-band communications, the direct RF sampling architecture entails considerable challenges, tightly coupled to the required ADC performance. By removing several signal conditioning blocks from the receiver chain, the ADC needs to correctly digitize signals over several GHz of bandwidth with spectral purity levels (both SFDR and IMn) of about 65–70 dB across the entire band of interest [101]. If time-interleaving is used, the spurs due to interleaving errors need to be suppressed even further (see Sects. 6.1.2 and 6.2.4), either by design or by calibration. Regarding noise, guaranteeing a better than 60 dB jitter-limited SNR at the highest frequency of interest (6 GHz) results in jitter values below 50 fs. The sensitivity performance also becomes important, where NSD values better than -153 dBFS/Hz are required [101]. Achieving these specifications can considerably increase the ADC power consumption. Hence, investigating power-efficient architectural and circuit techniques, and their scalability into finer processes, is key to improving the viability of this architecture.

6.1.2 ADC Architectural Trade-Offs

To enable the required sample rates in excess of 3-4-5 GS/s with a resolution higher than 10 bits, recent SotA RF ADCs have been extensively employing time-interleaving to boost the overall converter speed through high-efficiency slower sub-ADCs [87, 89, 94, 101, 147–149]. Nevertheless, TI-ADCs come with interleaving errors due to offset, gain, timing, and bandwidth mismatches (see Chap. 3, Sect. 3.7.2), the latter ultimately limiting the achievable performance at high input frequencies. Furthermore, the input front-end loading, routing, clock generation/distribution, and calibration circuitry to compensate for sub-ADC and interleaving errors impose an additional design overhead. Hence, the interleaving factor N and the sub-ADC architecture become highly critical choices in realizing

Fig. 6.2 Sub-ADC and interleaving overhead vs. channel count illustration



Strategy #1	Strategy #2
Minimum interleaving high-speed pipeline sub-ADCs ($N \leq 8$)	Massive interleaving low-speed SAR sub-ADCs ($N \geq 32$)
<div style="border: 1px solid black; border-radius: 10px; padding: 5px; text-align: center;"> small front-end loading less interleaving errors max. f_s up to 18GS/s high sub-ADC power less scaling benefit </div>	<div style="border: 1px solid black; border-radius: 10px; padding: 5px; text-align: center;"> large front-end loading more interleaving errors max. f_s up to 6.4GS/s low sub-ADC power more scaling benefit </div>

Fig. 6.3 Major design strategies regarding the choice of the sub-ADC and the interleaving factor

an efficient sub-ADC and minimizing interleaving overhead to achieve optimal performance. This trade-off is illustrated in Fig. 6.2. The slope of each curve and the optimum point to yield the minimum combined overhead depend on several factors, such as the targeted sample rate, the sub-ADC and the interleaver architecture, the nature and severity of errors to be calibrated, and the available process.

Currently, two major design strategies are prevailing in choosing the sub-ADC architecture and the interleaving factor. These two design strategies are depicted in Fig. 6.3 with some trade-offs highlighted. The first incorporates the faster but less efficient pipeline sub-ADC in order to interleave as few channels as possible [74, 101, 148, 150]. This approach results in a relatively easy to drive TI-ADC with the calibration complexity kept under control. The clock and reference distribution are also easier to handle. However, the MDACs and the sub-flash stages inside the pipeline do not put this approach at the top of the scalability chain. More recently, the pipelined-SAR sub-ADC with open-loop RAs has been interleaved under this design strategy [88, 89], benefiting from the advantages of the hybrid and demonstrating an improved converter efficiency.

The second strategy takes advantage of the superior efficiency and inherent scalability of a slow SAR sub-ADC and massively interleaves a large number of channels to reach the desired multi-GS/s sample rate [63, 69, 94, 149]. To relax the front-end loading, this approach usually adopts a hierarchical interleaver architecture, which also relaxes some of the calibration complexity. Nevertheless,




Fig. 6.4 Accuracy-speed standings of the ADCs adopting the two design strategies. Points taken from [36]

routing all the necessary signals in and out of the channels with appropriate synchronization, sufficient isolation, and small enough capacitance is not a trivial task. Thus far, this second design strategy has not been able to surpass the 12-bit quantization barrier with about 9.4 ENOB and a maximum sample rate of 6.4 GS/s [149]. The first one has demonstrated quantization levels up to 14 bits with 9.3 ENOB at 5 GS/s [72], while sample rates of 10 GS/s with 8.8 ENOB [101] and 18 GS/s with 7.7 ENOB [151] have been made possible (Fig. 6.4).

To ensure a sufficiently large input bandwidth, the highest-performing converters from both strategies adopt a static front-end unity gain buffer (source follower). This reduces the loading and impedance variations seen by the ADC preceding circuitry and actively drives the interleaved array of large sampling capacitors ($C_S > 500$ fF). In order to achieve a low enough output impedance ($Z_{out} \approx 1/g_m$), this buffer usually dissipates a higher power compared to the rest of the ADC blocks combined. Furthermore, the buffer significantly reduces the available swing, while its additive noise and non-linearity deteriorate the ADC spectral purity. The large devices used reduce the isolation at the highest frequencies, eventually degrading the converter performance due to the buffer dynamic non-linearity and large input/output capacitive loading. Hence, the front-end buffer, although simplifying the interfacing, remains among the primary performance and efficiency bottlenecks at the highest frequencies. A novel front-end solution with reduced power and increased functionality is discussed in Chap. 7.

The efficiency benefits of achieving a multi-GS/s operation without the use of a front-end buffer can be observed in [149]. In that work, a buffer-less front-end $32 \times$ TI-SAR was able to achieve a sample rate of up to 6.4 GS/s dissipating only 225 mW. Nevertheless, the large interleaving factor necessitated to achieve the aggregate sample rate, in combination with the chosen hierarchical scheme and distribution network, introduced a significant loading at the ADC input. This, in turn, limited the achievable bandwidth and the frequency of acceptable spectral

purity levels to about 1 GHz ($< \text{Nyquist}/4$). The next section presents a different solution extending the bandwidth beyond 6 GHz.

6.2 Prototype IC: A 5 GS/s 12-bit Hybrid TI-ADC in 28 nm CMOS

This section presents a 5 GS/s 12-bit passive-sampling 8 \times -interleaved hybrid RF ADC, achieving a low-frequency SFDR/SNDR of 75.2/62.4 dB and a Nyquist SFDR/SNDR of 65.4/58.5 dB [91, 93]. A significant reduction in power is enabled by an on-chip terminated very fast settling buffer-less input front-end. Several circuit and layout techniques are introduced that minimize the total resistance/capacitance in the signal path, to attain an input bandwidth in excess of 6 GHz. A three-stage pipelined-SAR sub-ADC with a single comparator per stage and a Dynamic Residue Amplifier (DRA) is employed that maximizes the efficiency for the given resolution, speed, and technology.¹ Moreover, an on-chip clock conditioning/distribution chain with an additive jitter as low as 11 fs is designed that improves the SNR at high frequencies, critical for sub-sampling operation at a higher Nyquist zone. Finally, on-chip co-designed analog-digital calibrations deal with sub-ADC and interleaving errors to improve the spectral performance across the entire band of interest. The 28 nm bulk CMOS prototype consumes a total power of 158.6 mW from a 1 V supply.

6.2.1 High-Level Design

The complete TI-ADC architecture is illustrated in Fig. 6.5. Eight sub-ADCs, each running at 625 MS/s, are directly interleaved to achieve the aggregate sample rate of 5 GS/s. The AC-coupled input and clock signals are protected by minimized custom laid-out ESD diodes and are 50 Ω on-chip terminated (100 Ω differentially). The termination components are realized with parallel high-R polysilicon resistors, showing superior voltage and temperature characteristics compared to other resistor types available by the process at hand. Each sub-ADC samples the input signal by means of bottom-plate sampling through bootstrapped switches, to ensure a high sampling linearity [26].

The raw digital outputs from each sub-ADC are collected by the synthesized digital calibration engine, which corrects the sub-ADC errors as well as the interleaving errors because of offset and gain mismatches. Timing mismatches

¹ The architectural conception of the proposed three-stage pipelined-SAR sub-ADC took place in the late summer of 2016. This occurred prior to and independent of the first available in open literature work of [88], which was in February 2017.

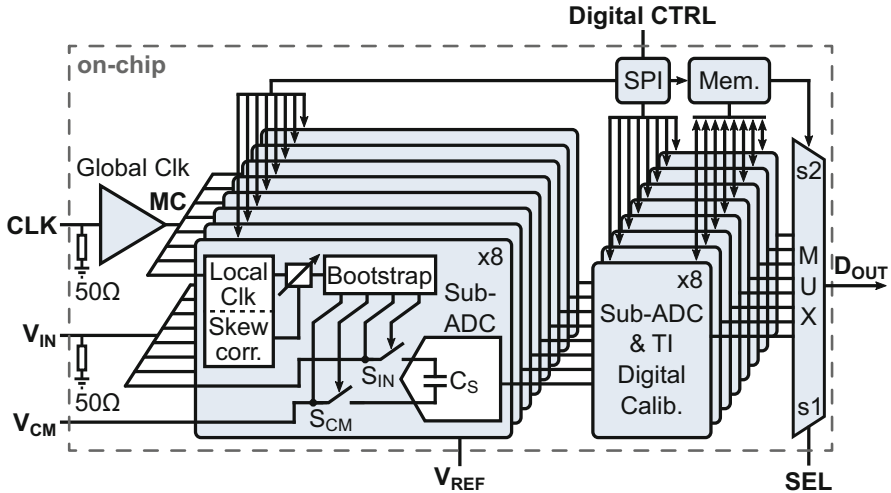


Fig. 6.5 Top-level diagram of the complete 5 GS/s 12-bit TI-ADC architecture (single-ended shown for simplicity)

between the sub-ADCs are corrected by tuning their sampling edges locally by a very fine step analog skew correction block. This is preferred over a digital FIR filter approach [103], for its higher accuracy and lower power. All the errors are calibrated in the foreground, and the calibration engine is controlled through a Serial Peripheral Interface (SPI) protocol. The calibrated outputs are stored in a high-speed memory block capturing 8192 samples for each sub-ADC (65,536 in total) and brought off-chip to a logic analyzer for performance evaluation.

6.2.2 Interleaving Factor and Sub-ADC Architecture

The interleaving factor and the sub-ADC architecture are critical choices in order to achieve a 5 GS/s sample rate and an SNDR of about 60 dB with maximum efficiency at the given 28 nm. To have the SNDR thermal and not quantization noise limited, a 12-bit aggregate quantization level is decided. First, it quickly becomes clear that interleaving must be employed since each of the potential sub-ADC candidates either cannot achieve the required sample rate and accuracy or operates at an extremely inefficient power vs. frequency point, based on our analysis from Chap. 3. Since interleaving is unavoidable, design strategy #1 from Fig. 6.3 is preferred to relax the front-end loading, signal distribution, and calibration overhead. Further, a power of two interleaving is chosen, offering the advantages of a balanced and symmetrical layout. Finally, a $2\times$ interleaving is avoided because going from a standalone converter to an interleaved one usually entails an equivalent overhead to running the standalone at double the speed on top of the interleaving errors [99].

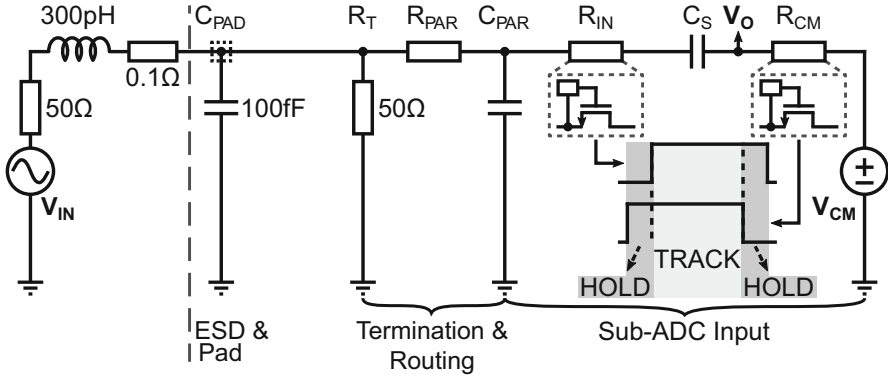


Fig. 6.6 Passive front-end model of this ADC (single-ended shown)

This overhead is relaxed going to $4\times$ and $8\times$ due to a reduced additional interleaving calibration overhead and an improved efficiency of the slower sub-ADC.

The sub-ADC architecture and input bandwidth considerations dictate the final decision between $4\times$ and $8\times$ to achieve maximum performance and efficiency. From Chap. 3, the optimal candidates to achieve a sub-ADC sample rate between 625 MS/s ($8\times$) and 1.25 GS/s ($4\times$) and an SNDR between 60 dB and 70 dB are the 3,4,5-stage pipelined-SARs and 1,2,3-bit/stage pipelines. Although a $4\times$ interleaving could theoretically work, $8\times$ is chosen to include some margin for performance degradation due to increased BEOL contribution of the 28 nm process. For the required sample rate and SNDR, the 3,4,5-stage pipelined-SARs show very similar efficiency, outperforming the pipelines. Out of the three pipelined-SAR options, the three-stage pipelined-SAR is finally preferred, due to its reduced number of RAs and smaller footprint compared to the other two. Finally, a direct interleaver architecture is adopted for its clock generation simplicity and energy efficiency, with minimum extra calibration complexity. With one sub-ADC tracking at any time, this interleaver achieves the largest bandwidth for $N \leq 8$ (see Chap. 3, Sect. 3.7.3). In this design, it is able to attain an input bandwidth well above the 2nd Nyquist zone, as shown next.

6.2.3 Passive Input Front-End

The inefficiency and additive noise and non-linearity of the front-end buffer typically used in high-performance RF ADCs were already stressed. This work aims to explore the input bandwidth, drivability, and spectral purity limits of bufferless front-end multi-GS/s RF ADCs while maximizing their efficiency. A simplified equivalent model of this front-end is shown in Fig. 6.6. The sampling capacitor is given by C_S , while R_{IN} and R_{CM} capture the on-resistances of the switches S_{IN}

and S_{CM} that perform the bottom-plate sampling. R_{PAR} is the resistance from the input routing toward the sub-ADCs, while C_{PAR} includes the contributions from both the routing and the switches' C_{on}/C_{off} . The equivalent input resistance of $25\ \Omega$ results from the $50\ \Omega$ source resistance and internal termination. For completeness, the extracted ESD and pad capacitance of $100\ \text{fF}$ as well as a bondwire inductance of $300\ \text{pH}$ with a series $0.1\ \Omega$ are included.

To achieve the highest possible performance and efficiency from such a front-end, both resistance and capacitance in the signal path must be minimized. The sampling capacitor C_S is among the dominant factors of the front-end loading. To minimize its thermal noise contribution, a relatively large input swing is chosen, leading to a C_S of $256\ \text{fF}$ (single-ended) for an SNR of about 11 bit. This capacitance guarantees that the impedance of each sub-ADC within the $2.5\ \text{GHz}$ Nyquist band is significantly higher than the on-chip $50\ \Omega$ termination during both TRACK and HOLD phases between sub-ADCs. Therefore, the dynamic impedance variations are kept low. In addition, the bottom-plate sampling is implemented by simply delaying the V_{IN} clock with respect to the V_{CM} , as shown in Fig. 6.6, which simplifies the clock generation, while preserving the linearity benefits of the traditional bottom-plate sampling [26]. In the TRACK and HOLD transitions between sub-ADCs, the V_{CM} side of each sub-ADC disconnects before the V_{IN} side of the next sub-ADC connects, leaving C_S shortly floating. This guarantees that the load seen at the input alternates between one and zero C_S , keeping the variations low without compromising the input bandwidth. The amount of the delay between V_{CM} and V_{IN} is a trade-off between realizing the bottom-plate sampling and allowing a sufficient tracking time for the targeted bandwidth and sampling linearity. In this design, a delay of about $15\ \text{ps}$, implemented with inverters, guarantees both the aforementioned.

With an allocated tracking time of $200\ \text{ps}$ for each sub-ADC, the sampling switches also play a major role in achieving a high input bandwidth and sampling linearity. Therefore, both R_{IN} and R_{CM} contributions are minimized by employing boosting circuits with very steep rise/fall edges. The bootstrap circuit for minimizing R_{IN} ($\sim 10\ \Omega$) is shown in Fig. 6.7a, with its important node waveforms illustrated in Fig. 6.7b. It is an optimized version of the one in the previous chapter, to achieve a higher linearity at larger input swings. The separate loop introduced outputting VN, in addition to the existing loop providing VP, allows the critical S_1 and S_2 to turn on with maximum overdrive almost simultaneously while significantly reducing the parasitic capacitance on VG. Combined with the grouped bulk connections of the critical devices, this enables very steep VG transitions, significantly reducing the impedance modulation of S_{IN} and resulting in an improved high-frequency sampling linearity. To ensure VP properly follows V_{IN} during TRACK time, S_3 is added with its gate tied to VG. To minimize R_{CM} ($\sim 10\ \Omega$), a constant voltage clock booster V_{SS} to $V_{DD} + V_{CM}$ is implemented as in [101], which allows S_{CM} to also operate with maximum overdrive.

To minimize the routing parasitic contribution, the symmetrical differential intertwined input/clock Y-tree structure shown in Fig. 6.8 is introduced, offering several parasitic reductions. First, the small tree area of $930\ \mu\text{m} \times 260\ \mu\text{m}$ in




Fig. 6.7 (a) Bootstrap circuit employed for S_{IN} and (b) timing waveforms of the important nodes

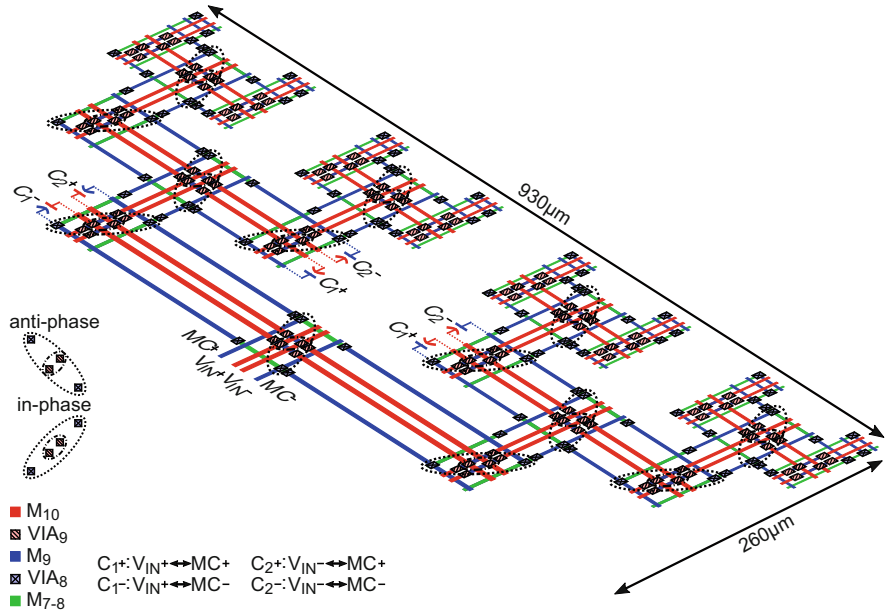


Fig. 6.8 Proposed intertwined input/clock Y-tree structure to minimize the front-end loading

combination with the utilization of thick top metals results in a R_{PAR} of only 4Ω on each input. Furthermore, the input and clock are routed side by side to minimize the systematic timing skew between the sub-ADCs. Since the input pair is more critical in terms of bandwidth, it is routed on the inside and on the highest metal to minimize its parasitic contribution to neighboring nodes. The trace spacing is

designed to achieve a flat transfer characteristic over the band of interest. The clock is routed on the metal right below the input and is intertwined around the input by arranging their vias either in-phase (short/long input and clock turns on the same side) or anti-phase (short/long input turns and long/short clock turns on the same side) in the different turns of the Y-tree (Fig. 6.8). The spacing between the input and clock is designed for a common-mode input-clock coupling smaller than the mutual differential input coupling. After the number of in-phase and anti-phase turns is fixed, this spacing is optimized in every turn such that their overall mutual differential coupling cancels out. That is, overall, $MC+$ couples with V_{IN+} by the same amount as with V_{IN-} , and the same holds for $MC-$. In this design, the simulated differential input coupling is about 50 fF with a common-mode input-clock coupling of about 20 fF due to the small tree area and the routing of the input and clock on different metals. The proposed intertwisting eliminates the need for the typical ground shielding when the input and clock are routed alongside [149], tremendously reducing the interconnect capacitance ($>2\times$ compared to the ground shielding approach), resulting in a total C_{PAR} of about 160 fF (including the switches' C_{on}/C_{off}).

The benefits of the techniques described above, in passively achieving a wide input bandwidth and small in-band impedance variations to simplify interfacing with an off-chip source, are validated via extracted front-end simulations, as depicted in Fig. 6.9. In Fig. 6.9a, the S-parameters in both the TRACK and HOLD modes are shown. $S_{11,TRACK}$ and $S_{11,HOLD}$ stay below -10 dB and -15 dB, respectively, up to 5 GHz, while both $S_{21,TRACK}$ and $S_{21,HOLD}$ show a very wide bandwidth well above 5 GHz. Figure 6.9b plots the simulated input impedance Z_{IN} looking at the pad. $Z_{IN,HOLD}$ stays relatively flat up to 5 GHz, while $Z_{IN,TRACK}$ is about $47\ \Omega$ at Nyquist and about $39\ \Omega$ at 5 GHz. The input current profiles are also plotted in Fig. 6.10 for a near sub-ADC Nyquist frequency (300 MHz, Fig. 6.10a) and a near-total Nyquist frequency (2.4 GHz, Fig. 6.10b). The current glitches due to switching transitions between the sub-ADCs are below 5–10% of the total current with a recovery well within the 200 ps TRACK interval.

One potential issue that might affect the performance of this front-end with the proposed routing approach is the unbalancing between the input and clock. This results in a non-zero differential coupling, causing even-order harmonics to rise and deteriorating the spectral purity at high input frequencies. Figure 6.11 plots the HD2 at a 5 GS/s sample rate and a near-2.5 GHz input frequency versus the percentage of the differential input-clock coupling. The compact area of the Y-tree together with the routing of the input and clock on different metals results in a small input-clock coupling, allowing for an HD2 better than -80 dB even in the unlikely scenario of a complete unbalancing.




Fig. 6.9 Simulated (a) S-parameters and (b) input impedance of this front-end

6.2.4 Clock Generation and Distribution

The detailed clocking diagram of this ADC is drawn in Fig. 6.12. The converter is clocked from a single clock, and after conditioning, the eight-phase sampling pulses are synchronously generated and distributed to all the sub-ADCs. An external 5 GHz differential sinusoid with a jitter of about 40 fs_{rms} is terminated on-chip, and a 50% duty-cycle MC is generated, able to drive the Y-tree and sub-ADCs with very sharp edges. The generated MC is divided by eight, and the 625 MHz 50% duty-cycle output shifts along the sub-ADCs (SH_{7-0}), where it is re-timed locally by eight consecutive phases of the 5 GHz MC. Subsequently, combinational logic receives locally SH_{7-0} in order to create the 12.5% duty-cycle differential sampling pulses. These pulses are re-timed once more to preserve the synchronous operation and sharp edges and finally buffered and delayed to create $SAM_{7-0,CM}$ and $SAM_{7-0,IN}$, which realize the bottom-plate sampling.

In our bottom-plate sampling scheme, the sampling instant is primarily defined by the S_{CM} switches (Fig. 2.17). Therefore, the high quality especially of the $SAM_{7-0,CM}$ pulses with respect to jitter and timing skew is key to achieving the desired SNR and non-harmonic SFDR at high input frequencies. Figure 6.13a plots the SNR at Nyquist versus σ_{jitter} , annotating also the 12-bit quantization level, while Fig. 6.13b plots the Nyquist non-harmonic SFDR versus σ_{skew} (average of




Fig. 6.10 Input current profile of this front-end for (a) 300 MHz and (b) 2.4 GHz input frequencies




Fig. 6.11 Simulated HD2 vs. differential input-clock coupling (unbalancing) for a near-2.5 GHz input

100 Monte-Carlo iterations). An SNR above 62 dB necessitates a total jitter of less than $50 \text{ fs}_{\text{rms}}$, including both the external clock source and the on-chip contribution. Furthermore, a timing skew of less than $10 \text{ fs}_{\text{rms}}$ is required to guarantee a non-harmonic (interleaving tone dominated) SFDR of better than 80 dB [147].

To minimize the on-chip jitter contribution, the clock conditioning chain shown in Fig. 6.14 is proposed. The global clock unit consists of a Current-Mode Logic




Fig. 6.12 Timing diagram with the generated clocks of the TI-ADC

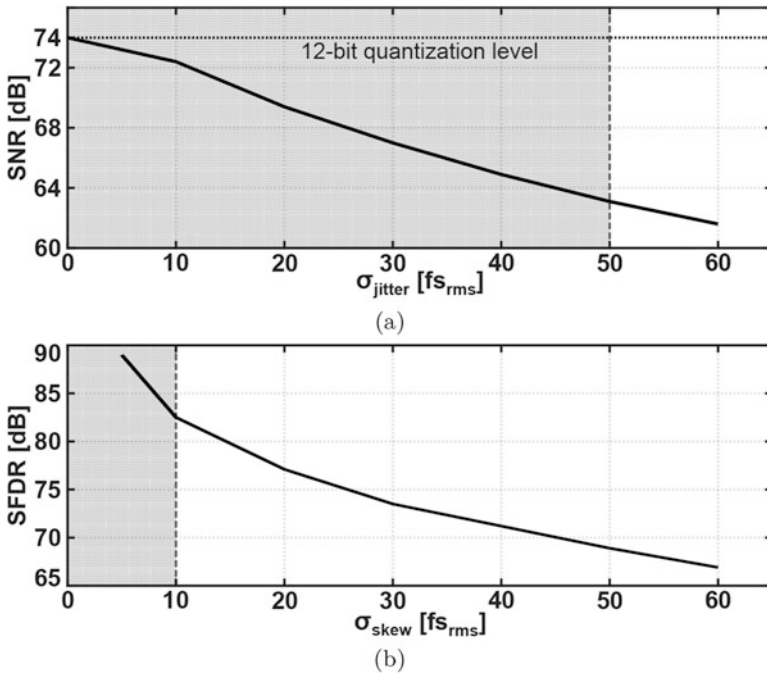


Fig. 6.13 (a) Simulated SNR vs. σ_{jitter} and (b) SFDR vs. σ_{skew} at Nyquist

(CML) stage, a CML-to-CMOS converter, and a cascaded FO-2 CMOS duty-cycle correction block with a 4:1 ratio between its driver inverters and its cross-coupled




Fig. 6.14 Block diagram of the proposed clock conditioning chain for this ADC

pair. This global unit generates the differential MC and distributes it to the clock divider and the sub-ADCs via the Y-tree. The CML-CMOS combination is chosen over a CMOS-only solution to enhance robustness against supply/substrate noise and Process-Voltage-Temperature (PVT) variations, in exchange for a higher current due to the partial class-A operation. The MC is divided by eight ($\div 8$) through a cascade of three custom divide-by-two master-slave flip-flops (FFs) and timed once with the MC. The divided and timed clock propagates through the local sub-ADC clock units, each containing a custom re-time/shift FF, a NAND/NOR gate for the differential sampling pulses, and a final re-time FF with a latch and a driver that offers the shortest path (one transistor and one inverter) between the triggering MC pulse and the critical SAM_{CM} .

To correct the timing skew between the sub-ADCs, a single-stage Digital-to-Delay Tuner (DDT) with MOS devices is introduced, shown in Fig. 6.15a. This circuit exploits the properties of the MOS device with shorted drain and source, and fixed bulk, to act as a variable capacitor upon changing its gate voltage. To match the capacitance of each unit cell for both rise and fall edges of the clock, complementary PMOS/NMOS devices are adopted. The segmented DDT comprises a coarse 6-bit binary part determining the tuning range and a fine 3-bit unary part dictating the tuning step. This segmentation makes it possible to achieve both a fine step and a sufficient tuning range in a single stage, with a smaller total area and step uniformity compared to a coarse unary fine binary approach. Our compact single-stage implementation minimizes the total area and jitter contribution, compared to the typical multi-stage approaches [96, 147, 149, 152], for the same load, rise/fall edges, and power in the clock buffers. The latter was already hinted by our jitter

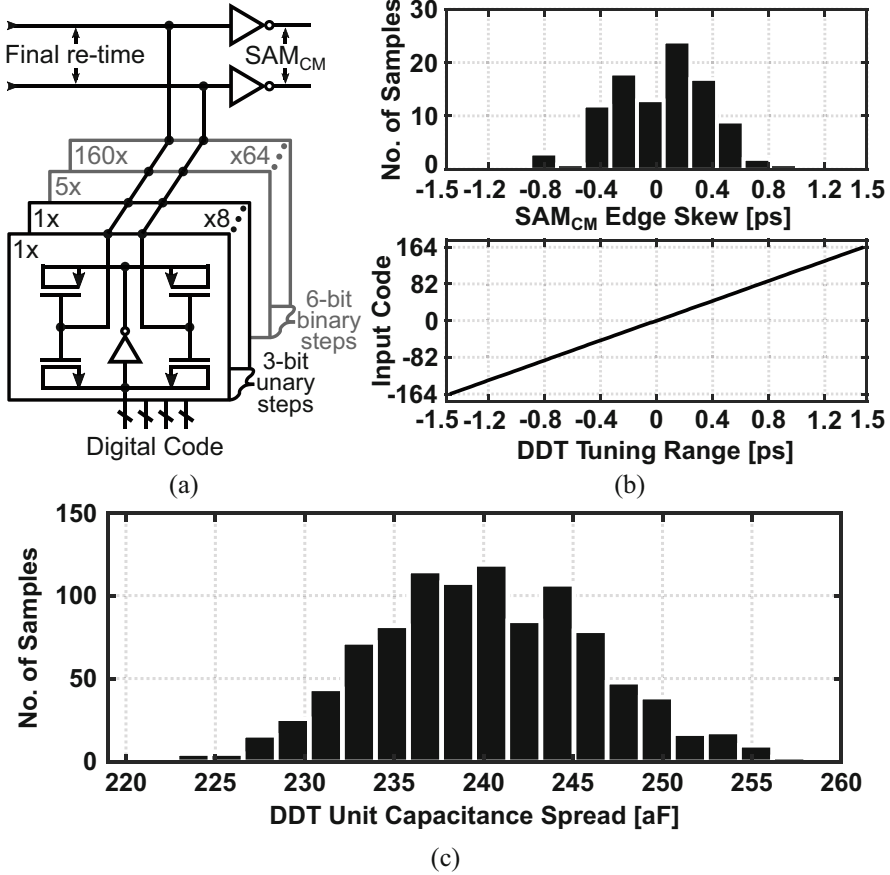


Fig. 6.15 (a) DDT circuit with simulated (b) sampling edge skew and tuning range and (c) capacitance spread of one DDT unit cell

limits’ analysis of Chap. 2 (see Sect. 2.4.4). Isolating the clock buffer current I_{CK} from our derived Eq. (2.60), the relation between this current and the thermal noise dominated jitter at the clock chain output t_{jit} can be approximated as

$$t_{jit} \approx \sqrt{kT \frac{C_{CK}}{I_{CK}^2}}, \quad [s] \quad (6.1)$$

where C_{CK} is the total load at the buffer output. The above expression reveals that if a single buffer with a current I_{CK} charging a capacitor C_{CK} is replaced by n cascaded buffers each with I_{CK}/n charging C_{CK}/n , the jitter of the latter case is \sqrt{n} × that of the former case. This assumes the same edge steepness and no jitter correlation between the n buffers. In reality, the jitter benefits between a single-stage and a multi-stage

approach are somewhat less profound than the above estimation. One reason is that C_{CK} does not comprise solely the DDT capacitance, which is only a fraction of the total clock chain load. Further, the edge steepness at each stage is to some extent a function of the number of stages and the properties of the externally applied signal at the clock chain input (e.g., square wave vs. sinusoid), therefore not constant.

The tuning range of ± 1.5 ps corrects the simulated SAM_{CM} skew to better than $3 - \sigma$ level (Fig. 6.15b), while the ~ 9 fs tuning step aims for the desired non-harmonic SFDR. The simulated $3 - \sigma$ capacitance spread of a DDT unit cell (Fig. 6.15c) in combination with the output resistance of the circuit driving the DDT alters the tuning step by less than 1 fs. Finally, sufficient overrange is allocated between the binary and unary parts to account for this spread and to prevent missing codes. The total clock chain is co-optimized together with the DDT, and the extracted additive jitter on SAM_{CM} ranges from a minimum of 11 fs_{rms} when loaded by a fully off DDT to a maximum of 14 fs_{rms} with a fully on DDT loading it. This is about 3–4 \times smaller than the external source contribution, therefore negligible in their squared sum.

6.2.5 Hybrid Sub-ADC Design

Motivated by our developed architectural analysis, the sub-ADC topology of choice includes a hybrid three-stage pipelined-SAR. For the targeted sample rate and resolution, it yields an improved energy efficiency compared to a two-stage equivalent, while being roughly on par with the 4,5-stage counterparts, considering also practical implications. The implemented architecture is detailed in Fig. 6.16. It comprises three single-comparator SAR stages with two open-loop DRAs in between. The sampling capacitor of 256 fF is scaled down to 160 fF in the following stages to reduce power and area. The resolution partitioning in each stage involves several considerations. Allocating more bits in SAR_1 relaxes the power, noise, and linearity requirements on the amplified residue and the back-end stages. However, a higher SAR_1 resolution directly reduces the sub-ADC sample rate, determined by the SAR_1 conversion time and the RA residue transfer time. Another limiting factor to raising the SAR_1 resolution is its comparator noise and offset, governing an exponentially larger OR between the stages or a larger power to tackle.

In this work, the priority is set in achieving the targeted 625 MS/s sample rate with an increased efficiency and a low design complexity. Therefore, the bit partitioning is chosen as 4-4-6-bits/stage with 1-1-bit interstage OR, totaling a quantization level of 12 bits. A 4-bit resolution in SAR_1 with the utilized signal swing ($1.6\text{--}1.8 V_{pp,diff}$) offers a sufficiently large OR, which relaxes the relative offset between the comparator and the RA while efficiently achieving the targeted sample rate. To minimize the complexity, 4 bits are allocated in SAR_2 as well, allowing for the same RA design between all stages with negligible power overhead. The remaining 6 bits for SAR_3 are chosen to accommodate the aggregate quantization level while not compromising the sample rate. The RA is designed




Fig. 6.16 Detailed block diagram of the implemented 12-bit three-stage pipelined-SAR sub-ADC

with a half gain of $4\times$, instead of the typical $8\times (2^{4-1})$ for a 4-bit stage, which reduces the swing at the output of RA_1 and RA_2 to half and quarter of the input, respectively. This choice is necessary to satisfy the linearity of especially RA_1 [83], whose back-end resolves 8 bits. Although this increases the input-referred noise of SAR_2 and SAR_3 , it has a minor impact on the overall noise budget. In addition, the LSB voltages of SAR_2 and SAR_3 remain sufficiently large to allow a high-speed, low-power comparator design. In SAR_1 and SAR_2 , a single-stage comparator is used [153], to minimize power, while in SAR_3 , a three-stage comparator based on [60] is adopted, for its increased gain and speed in the presence of smaller input differences. All the comparators are dimensioned for a similar input-referred noise and offset with $1 - \sigma$ values of about $700 \mu V_{rms}$ and $8 mV_{rms}$, respectively.

The details of the capacitive DACs in all SAR stages are also given in Fig. 6.16. Bottom-plate sampling is preferred for DAC_1 for its linearity merits, while top-plate sampling is adopted in DAC_2 and DAC_3 to maximize the conversion speed. A tri-reference (V_{REFP} , V_{CM} , V_{REFN}) switching [60] is employed in all stages, for its very low and symmetrical energy. To accommodate the RA_1 and RA_2 half and quarter output swings, a fraction of the DAC_2 (C_{H2}) and DAC_3 (C_{H3}) capacitance is used to attenuate their reference range. The reference voltages are provided externally (see Sect. 6.3), and sufficient custom on-chip decoupling is employed to supply the dynamic currents, so as to minimize any transient glitch to less than $1/4$ LSB. For measurement flexibility and monitoring, separate references are used in DAC_1 and $DAC_2 - DAC_3$. Regarding the design of the switches, NMOS devices are used




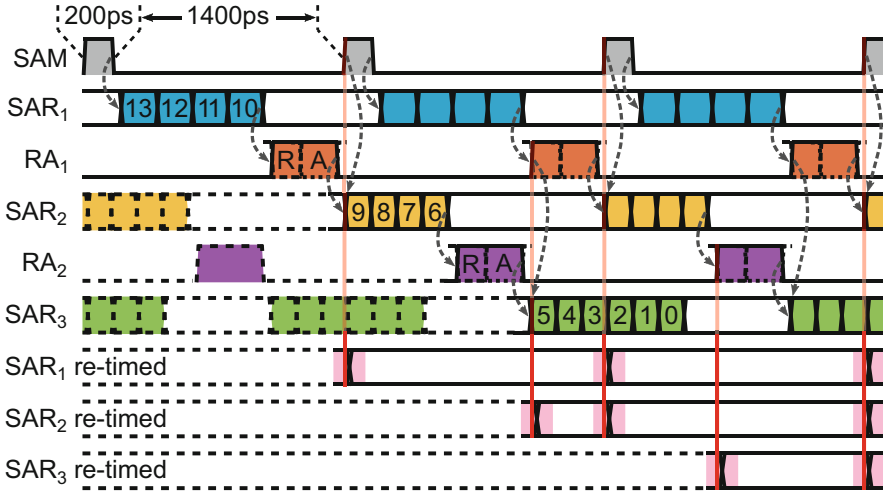
Fig. 6.17 Dynamic integrator RA with simulated SAR₁ - SAR₂ residue

for V_{REFN} and PMOS devices for V_{REFP} , optimized for matched impedances. The V_{CM} switches are also implemented as NMOS devices, equal in size to those for V_{REFN} . All the switches in DAC₁ are sized to meet the longest V_{CM} settling, while in DAC₂ – DAC₃, their size to meet any of the V_{REFP} and V_{REFN} settlings suffices since V_{CM} in those is applied simultaneously to the about 4–5× slower residue transfer.

The circuit of the single RA design is shown in Fig. 6.17a, and its simulated waveforms are plotted in Fig. 6.17b. An open-loop cascoded integrator-based clocked amplifier is adopted [84, 154], which efficiently offers a high-speed and a low-noise operation, provided that its gain and timing can be precisely controlled. Such a structure enhances the ADC efficiency beyond the analytically predicted in Chap. 3, due to the unsettled operation. For a given output common-mode drop, integration time, and overdrive voltage of the input pair, the gain is fixed. The cascode devices are responsible for isolating the outputs from the drains of the input pair, avoiding the need for explicit series switches. The cascode devices are initially switched off, and a brief initial integration on their source nodes V_{YP}/V_{YN} occurs, followed by the main integration on their drains V_{OP}/V_{ON} to provide the final amplified output. The integration time of this RA is controlled by a CM-detect circuit that combines the concepts from [154] and [155]. When V_X crosses a threshold voltage around 0.5 V, the CM-detect turns off the cascodes and the tail current, providing a steady differential output for the following stage. Combined with the utilization of ultra-low threshold devices biased at strong inversion, a proper output common mode to ensure that the devices are always in saturation, and the reduced gain, this structure satisfies the linearity to drive its 8-bit resolution back-end. For linearity and process variation considerations, the gain and integration

Table 6.1 RA gain variation with temperature (typical-typical corner)

Temperature [°C]	27	−40	125
Gain [V/V]	4	4.96	3.12

**Fig. 6.18** Sub-ADC internal asynchronous timing sequence with re-timing

time are further controlled by adjusting the initial condition of V_X through V_{CTRL} . The simulated $1 - \sigma$ values for the RA input-referred noise and offset are about $150 \mu V_{rms}$ and $6 mV_{rms}$, respectively. Table 6.1 shows the simulated raw RA gain variation (i.e., no V_{CTRL} adjustment) for three different temperatures in a typical-typical corner. Even at the extremes of -40 and 125 °C, the RA output still occupies less than a quarter of the allocated OR, which was budgeted during design time.

Figure 6.18 illustrates the sub-ADC internal timing sequence. Asynchronous stage clocking is utilized for optimal timing allocation and easier RA timing generation. Delay overlapping among the DAC, comparator, and logic is applied here as well [60], to further increase the conversion speed. The falling edge of SAM initiates the 4-MSB SAR_1 conversion. Upon its completion, the RA_1 is asynchronously triggered to dynamically amplify its input residue. The output of RA_1 is sampled by SAR_2 on the synchronous SAM – RA_1 combination to ensure the integrity of the conversion. The intermediate 4-ISB conversion triggers asynchronously RA_2 , which amplifies its own residue. The output of RA_2 is sampled by SAR_3 for the final 6-LSB conversion. The bits from each stage are re-timed by the “align and combine” logic, with the sequence shown in Fig. 6.18. After a latency of three sampling periods, all the data becomes available in parallel for propagating to the calibration. For full flexibility in this prototype, the timings of all the DACs and RAs can be externally controlled.




Fig. 6.19 One slice top-level diagram of the 8× synthesized correction block

6.2.6 Digital Calibration

As already mentioned, this ADC employs both analog and digital calibration to correct for several sub-ADC and interleaving errors and to accomplish the desired spectral performance levels. Figure 6.19 depicts the top-level diagram of one out of the 8×-interleaved synthesized digital correction blocks. Calibration is performed in the foreground by first applying a low-frequency test sinusoid and capturing the raw data. The correction coefficients to minimize the different errors are then estimated using the reconstruction formula

$$\begin{aligned}
 D_{\text{OUT}} = & \sum_{i=10}^{13} W_i \cdot b_i + \frac{1}{G_{\text{RA1}}} \cdot \sum_{j=6}^9 W_j \cdot b_j \\
 & + \frac{1}{G_{\text{RA1}} \cdot G_{\text{RA2}}} \cdot \sum_{k=0}^5 W_k \cdot b_k + D_{\text{OS}}, \quad (6.2)
 \end{aligned}$$

where b_i , b_j , and b_k are the bits of the SAR₁, SAR₂, and SAR₃ stages; W_i , W_j , and W_k are each stage's DAC reconstruction weights; G_{RA1} and G_{RA2} are the RAs' gain weights; and D_{OS} is the digital weight for the sub-ADC offset.

Mismatches on the DAC capacitors of SAR₁ and SAR₂ are corrected by adjusting the fixed-point 25-bit programmable W_i and W_j . Since these mismatches are supply voltage and temperature independent, this correction need not be executed continuously. Gain errors in RA₁ and RA₂ are corrected by adjusting the 27-bit G_{RA1} and 23-bit G_{RA2} , respectively, to minimize the RMS errors between the sub-

ADCs. These two corrections scale the ranges of the different stages appropriately to effectively correct the gain error of each sub-ADC. After that, the 25-bit D_{OS} are adjusted, such that the offset errors of the sub-ADCs are equalized. The comparator and RA offsets are lumped into the above corrections, absorbed by the OR. The correction coefficients' lengths are determined after extensive simulations and optimized such that the calibrated errors are not limiting the final TI-ADC performance. The output stream of the correction block is truncated to the aggregate ADC resolution prior to being collected. The digital coefficients are represented in a binary weighted format, and the truncation discards the LSBs, therefore resulting in a minimal accuracy degradation. The timing errors are corrected by applying a Nyquist sinusoid and comparing the relative phases between adjacent channels. The DDT (see Sect. 6.2.4) is then configured to minimize the difference between the actual relative phases and the ideal $2\pi/8$, by adjusting the sampling edges of the sub-ADCs.

During reconstruction, $SAR_1 W_i$ are masked with b_i , and $SAR_2 W_j$ are masked with b_j and accumulated. $SAR_3 W_k$ are multiplied by G_{RA2} , and the output is shifted and added to the weight-corrected SAR_2 output. This, in turn, is multiplied by G_{RA2} , and the result is shifted and added to the weight-corrected SAR_1 for the final output prior to offset correction. For the synthesized calibration block to support the 625 MS/s sub-ADC sample rate, pipelining is used across the correction stages, as indicated in Fig. 6.19.

6.3 Experimental Verification

The prototype TI ADC is fabricated in a baseline single-poly ten-metal (1P10M) 28 nm bulk CMOS process. A complete die micrograph is shown in Fig. 6.20, measuring a pad-limited total area of $1900\ \mu\text{m} \times 2400\ \mu\text{m}$. The active part of the die, including the sub-ADCs, the clock generation and distribution, the calibration engine, and all the control and combine circuitry, occupies an area of $990\ \mu\text{m} \times 1180\ \mu\text{m}$. One hybrid three-stage pipelined-SAR sub-ADC is also shown, with a compact area of $75\ \mu\text{m} \times 200\ \mu\text{m}$, including the “align and combine” logic.

The differential input and clock are applied at the bottom of the chip. The input together with the generated global MC is distributed to all the sub-ADCs through the Y-tree. Each sub-ADC channel receives the $\div 8$ MC and generates locally the sampling pulses as described in Sect. 6.2.4. The aligned outputs from the three stages of each sub-ADC are sent to the on-chip calibration and memory storage block prior to being buffered at the top of the chip, where they are collected for performance evaluation. Extra circuitry is foreseen across the chip to provide the flexibility of distinguishing between each sub-ADC output as well as between calibrated and raw data, for debugging purposes. The control signals and reference voltages to the sub-ADCs are coming from the sides. Dense custom-made MOS + MOM grids are placed inside, between, and around the sub-ADCs for decoupling purposes.




Fig. 6.20 Die micrograph of the 28 nm complete IC with a sub-ADC layout view occupying a core area of 0.015 mm²

Each hybrid sub-ADC is laid out as symmetrically as possible, with dummy structures added to guarantee the same environment around critical blocks. The layout practices followed in terms of blocks' arrangement are similar to Chap. 5, to minimize the critical path interconnect and optimize for power and speed. In each sub-SAR stage, the comparator and its clock logic are placed on top of the DAC with the state memory logic part at the two sides interacting with both. The DRAs are placed on top of SAR₁ and SAR₂, respectively, with their control signals coming from the sides. One align block is placed alongside each stage with the final combine block at the top. The boosted input and clock switches are placed at the bottom of the sub-ADC to sample the incoming signals onto the DAC₁.

6.3.1 Measurement Setup

The complete measurement setup used to evaluate the performance of this TI ADC prototype is depicted in Fig. 6.21 (top). A photo of the setup during a one-tone measurement is also shown (bottom-left), together with a closer view of the custom-made boards. For the one-tone measurements, two Agilent E8257D analog signal sources with a low phase noise option are used to generate the input and clock signals. The spectral purity of both is improved by adding appropriate filtering



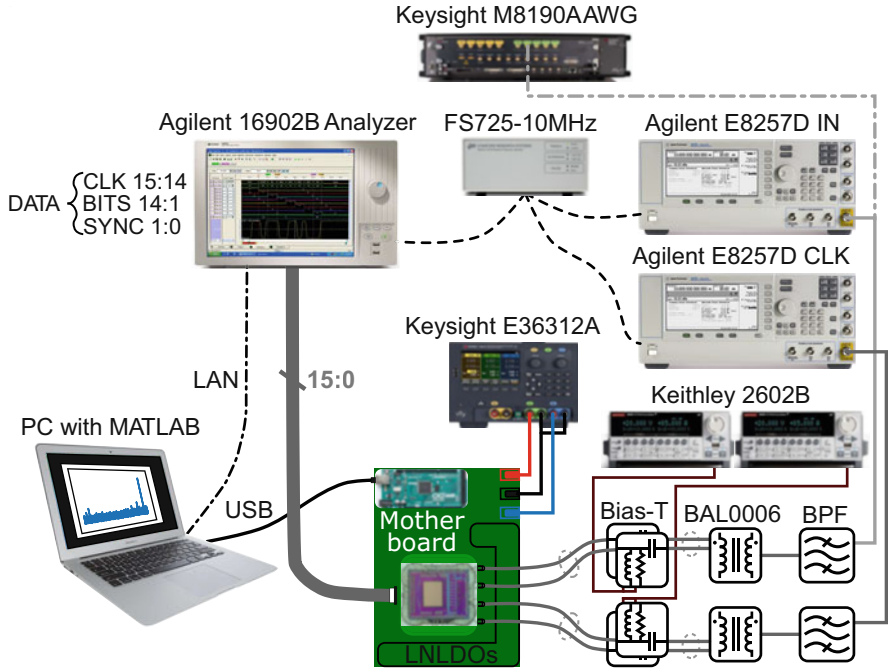


Fig. 6.21 Measurement setup of the 12-bit 5 GS/s TI ADC prototype (top). Photo of the overall setup (bottom-left). Closer view of the motherboard with the four-layer Rogers 4350 chip board mounted on its center through high-speed Samtec connectors (bottom-right). The bare die is placed in a plated cavity

to remove residual noise and spurs. After filtering, both the input and clock are converted into differential signals by two identical wideband baluns with small amplitude (<0.1 dB) and phase ($<1^\circ$) mismatch. Finally, they are AC-coupled to the chip through custom-designed bias-Ts and phase-matched cables as well as identical board traces. For the multi-tone measurements (see Sect. 6.3.2), the input

source is replaced by a Keysight M8190A Arbitrary Waveform Generator (AWG). The differential data, clock and synchronization outputs coming out of the chip at the single-channel sample rate, are captured by an Agilent logic analyzer and reconstructed on a PC in MATLAB. All the equipment is synchronized by a 10 MHz rubidium source.

On the custom high-frequency material chip board, special attention is paid to ensure a high signal integrity. The input and clock traces start from the chip as differential microstrips with outer ground shielding, ending at the high-speed connectors as coplanar waveguides. The transition is made as smooth as possible to guarantee the required characteristic impedance at more than twice the entire band of interest. The output traces are also designed with a characteristic impedance, but due to their digital nature, they are not as critical as the input and clock. To reduce the effect of the critical bondwires, their length is minimized to about 250–300 μm by mounting the bare die on a plated cavity with the same height as the die.

The required supply, bias, and reference voltages for the different chip domains are generated with dedicated low-noise LDOs on the custom motherboard and provided to the ADC after further low-pass filtering with discrete components. Two Keithley sourcemeters are employed to provide the input and clock common-mode voltages, respectively. The motherboard also hosts the microcontroller board that is responsible for interfacing with the calibration through the SPI.

6.3.2 Measurement Results

The functionality and performance of this prototype are characterized by a variety of measurements. These include both one-tone and multi-tone spectral measurements, as the latter are of great importance in RF sampling ADCs [141]. For completeness, static measurements are also performed, as will be shown.

One of the main highlights of this work is introducing circuit and layout techniques that enable a very wide bandwidth by a buffer-less front-end. To demonstrate the effectiveness of the introduced techniques, the measured ADC transfer characteristic is plotted in Fig. 6.22, with the simulated one annotated as well. The measurement includes losses from the board, the connectors, and the bondwires, while the power versus frequency delivered at the board is calibrated to 0 dBFS. The input bandwidth is larger than 6 GHz, allowing a multi-Nyquist zone operation of this ADC. The upper limit of 6 GHz is chosen for the frequency sweep as this is the operation region of the available baluns with tolerable amplitude and phase unbalance. Above about 800 MHz, there is a deviation of about 0.5 dB between the measured and simulated characteristics, which becomes about 1 dB at 6 GHz. This could be attributed to the board and connector parasitics as well as slight underestimation in the extraction tool, which are not captured by our front-end model. The bandwidth could be further extended by reducing the internal termination and source resistors below 50 Ω , in exchange for a larger power burnt




Fig. 6.22 Measured (black solid curve) and simulated (gray dotted curve) ADC transfer characteristic showing a bandwidth in excess of 6 GHz

at the input. Such a property is exploited in the ultra-wideband front-end described in the next chapter.

The measured output spectra after calibration at 5 GS/s for three different input frequencies across two Nyquist zones are shown in Fig. 6.23. The frequency bins of the interleaving tones are annotated as well. At a 75 MHz input (Fig. 6.23a), the SFDR is 75.2 dB, dominated by the HD3 of the bootstrapped input switch. The SFDR drops to 65.4 dB at 2.4 GHz (Fig. 6.23b) and 61.0 dB at 4.8 GHz (Fig. 6.23c). In the latter cases, it is dominated by the increasing with input frequency HD2. This is mainly attributed to the phase imbalance from the off-chip balun and was verified by reproducing the imbalance in simulations. This could be improved by employing two separate identical signal sources and adjusting their relative amplitudes and phases. All remaining sub-ADC/interleaving-related spurs are suppressed by the on-chip calibration to below -75 dBFS across the entire 1st Nyquist zone and below -65 dBFS in the 2nd Nyquist zone. This is due to the DDT being configured to correct timing/bandwidth mismatch errors with a Nyquist test sinusoid. Above the 1st Nyquist zone, these errors are still corrected but to a lesser extent. The SNDR at 75 MHz is 62.4 dB, which drops to 58.5 dB at 2.4 GHz and 53.6 dB at 4.8 GHz. The SNDR across the entire frequency range is mainly limited by noise. When increasing the input frequency toward Nyquist and beyond, frequency-dependent residual interleaving errors and harmonic distortion as well as the increasing effect of signal/clock jitter and loss of signal gain, all contribute their part to the SNDR degradation. The small signal noise floor (NSD) is about -160 dBFS/Hz, which drops to -157 dBFS/Hz at full scale. This specification is key for the sensitivity of RF sampling ADCs.

The measured SFDR/SNDR versus the input frequency at 5 GS/s are plotted in Fig. 6.24a for five different samples. The peak SFDR averaged across the samples

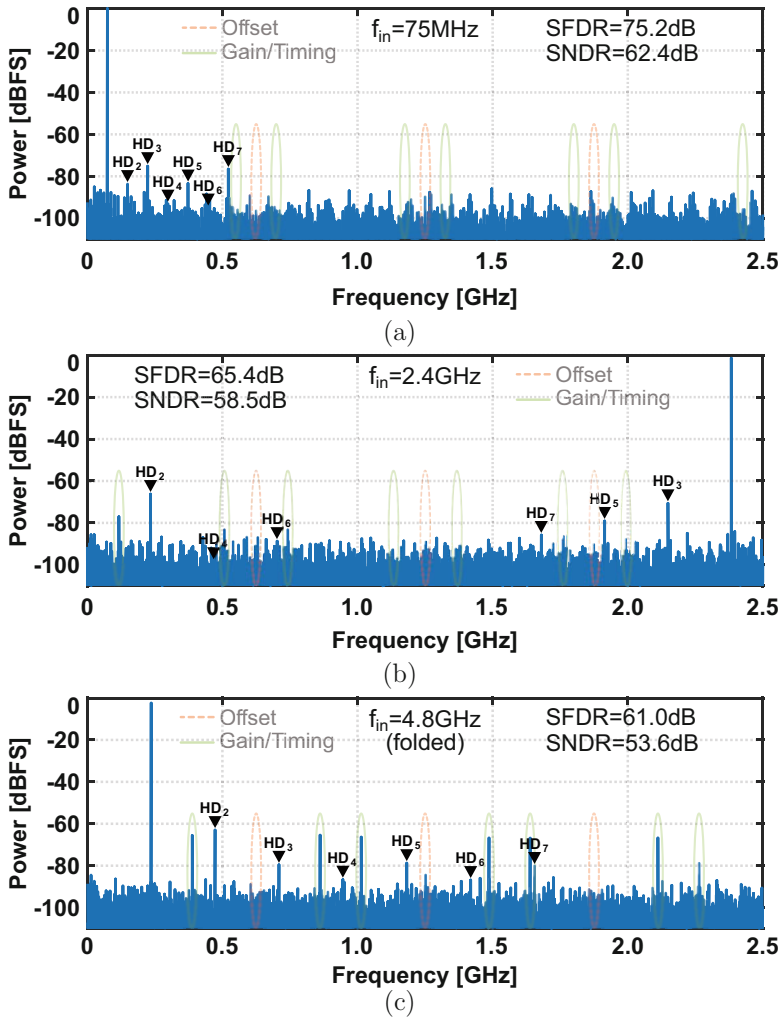


Fig. 6.23 Measured calibrated output spectra at 5 GS/s for (a) 75 MHz, (b) 2.4 GHz, and (c) 4.8 GHz input frequencies

is 73 dB and remains above 60 dB up to the 4.8 GHz frequency. The average low-frequency SNDR is 62.4 dB, which drops to 53.2 dB at the highest frequency. The measured SFDR/SNDR for five samples versus the sample rate at a 2.4 GHz input frequency are also plotted in Fig. 6.24b. Both SFDR and SNDR stay relatively flat up to 5 GS/s, and after that, they start dropping, mainly due to the limited cycle time for the RAs to operate properly. This also verifies that the ADC is running close to its architectural and technology speed limit, which justifies the interleaving factor choice of 8× to achieve the required performance, considering the BEOL.




Fig. 6.24 Measured SFDR/SNDR versus (a) input frequency at 5 GS/s and (b) sample rate for a 2.4 GHz input

The measured DNL and INL characteristics at 5 GS/s for a sinusoidal input are plotted in Fig. 6.25. DNL and INL after calibration lie within $-0.73/+0.86$ LSB and $-1.1/+1.2$ LSB, respectively. Some systematic jumps are noticed in the INL pattern. These could be attributed to any remaining errors in the transitions between the different sub-SARs within each pipeline due to accumulated error sources that the OR needs to smoothly absorb. In the transition between SAR₁ and SAR₂ in particular, any mismatch between the different external references used from their ideal values is contributing to this effect as well.

In Sect. 6.1.1, the superior spectral purity requirement of RF ADCs was stressed, due to the wideband/multi-band nature of the signals to be correctly digitized in the absence of several signal conditioning blocks. For this reason, an important performance characterization of the converter includes its behavior in a multi-tone signal environment. In this case, the intermodulation products closest to the applied signals must be sufficiently suppressed, to avoid out-of-band signals interfering with the useful signals of interest. Figure 6.26 plots the measured output spectrum at 5 GS/s for a -6.1 dBFS two-tone input signal centered at 78 MHz. The most critical metric then is IM3, defined by the harmonics at frequencies $2f_1 - f_2$ and $2f_2 - f_1$. Since these fall closest to the fundamentals, they are the hardest ones to




Fig. 6.25 Measured static performance at 5 GS/s for a sinusoidal input of 7.4 MHz: (a) DNL and (b) INL

filter out, especially at RF frequencies. At a center frequency of 78 MHz, IM3 is limiting the SFDR to -74.0 dB, while the less critical IM2 lies below -82.0 dB. Any remaining sub-ADC/interleaving-related spurs are suppressed by the on-chip calibration to below -80 dBFS. The performance of the ADC for a -6.4 dBFS two-tone input signal centered at 1.76 GHz is also characterized, and the measured spectrum is shown in Fig. 6.27. In this case, IM3 is -73.8 dB, while IM2 is -72.2 dB. However, the SFDR is 71.0 dB, limited by the interleaving spurs due to residual gain/timing errors. These could be reduced by optimizing the correction coefficients and calibrating the ADC for this particular input frequency.

The prototype ADC utilizes multiple core 1.0 V supply voltage domains, and the power partitioning versus the sample rate for a 2.4 GHz input frequency is shown in Fig. 6.28. The total power consumption includes every on-chip contribution, except for the output IO buffers, which have their dedicated domain and are overdesigned to drive the pads. The total power ranges from 72 mW at 1 GS/s to 173 mW at 6 GS/s. The digital calibration power dominates across the entire sample rate range. A considerable part of this power comes from the different calibration modes as well as several extra circuits included for testing/debugging purposes. These circuits even in standby mode are consuming a non-zero amount power and, due to the single supply for the entire calibration block, cannot be turned off. Further, the IO buffers




Fig. 6.26 Measured output spectrum at 5 GS/s for a -6.1 dBFS two-tone input signal at 74.5 MHz 81.7 MHz




Fig. 6.27 Measured output spectrum at 5 GS/s for a -6.4 dBFS two-tone input signal at 1.67 GHz 1.85 GHz

bringing out signals from the calibration block at the single-channel sample rate are also under the same supply as the digital core. Therefore, their contribution, comprising about 30% of the total calibration power according to simulations, is also included in the reported values of Fig. 6.28.




Fig. 6.28 Measured power partitioning versus sample rate for a 2.4 GHz




Fig. 6.29 FoM_S comparison with relevant SotA RF ADCs [36]

6.3.3 State-of-the-Art Comparison

The performance of this work is compared by means of FoM_S to the best existing relevant ADCs (RF and others) at the time of publication [36] and plotted in Fig. 6.29. This work achieves at least 5 dB better FoM_S than any ADC with a sample rate within ± 1 GS/s from the 5 GS/s and at least 6 dB better FoM_S than any actual RF ADC. Furthermore, it achieves an at least 2.5× higher sample rate than the closest competitor with a FoM_S within ± 0.5 dB.

Table 6.2 summarizes the performance of this ADC and compares it with the most noteworthy wideband TI RF ADCs. The majority of these works utilize the pipeline sub-ADC as the architecture of choice [101, 147, 148, 150], while

there are one more pipelined-SAR [89] and one SAR [149]. Further, all these works except for [149] employ a static front-end unity gain buffer to attain a wide bandwidth while driving the large sampling capacitor. The proposed circuit and layout techniques in this work demonstrate a comparable or larger input bandwidth with an on-chip terminated buffer-less front-end. Compared to the next best buffer-less work, the input bandwidth of this ADC is at least $6\times$ higher, mainly due to the smaller sampling capacitor, interleaving factor, and layout parasitics, enabled by the introduced techniques. The absence of the buffer directly impacts the total power consumption, which for this work is at least $2\times$ smaller than the buffered works and about $1.4\times$ smaller than the other buffer-less work. However, the efficiency merits of this work stem also from the sub-ADC architectural and circuit choices, predicted by our architectural analysis. Appendix D estimates the power consumption of this ADC, including a front-end unity gain buffer able to achieve a very wide bandwidth.

The reader will notice that there is one design from ISSCC 2019 with a sample rate of 3.2 GS/s standing above the dashed line [156]. This work employs a $4\times$ -interleaved ringamp-based pipeline ADC in 16 nm FinFET and performs the entire calibration (sub-ADC and TI errors) off-chip. The reported performance is achieved almost entirely due to the off-chip calibration. However, the reported power and FoM_S do not include any estimated calibration overhead. Therefore, it is not considered a valid data point for a fair comparison.

6.4 Conclusion

This chapter elaborated on architectural and circuit capabilities to enable ADC resolutions beyond 10 bits while sampling directly at RF frequencies with multi-GHz rate and bandwidth at maximum power efficiency. Such high-resolution, multi-GHz sample rate and bandwidth, low-power RF ADCs are of utmost interest in next-generation communication, data acquisition, and instrumentation systems, currently being a hot research topic.

First, to provide some context for the increasing importance as well as the challenges of wideband RF ADCs in enabling direct sampling, the direct RF sampling receiver architecture was discussed, highlighting the ADC role. This architecture exploits the DSP advancements with technology scaling, to increase flexibility and integration while reducing complexity and cost. However, this poses considerable challenges for the converter, which has to digitize several GHz of bandwidth with very high spectral purity. Hence, making efficient architectural and circuit choices is key to benefiting from RF sampling.

Architectural choices based on prior art were reviewed, and their trade-offs discussed. Two prevailing design strategies in terms of sub-ADC and interleaving factor were identified: (1) interleave as few as possible faster less efficient pipelines or (2) massively interleave slower highly efficient SARs. The former offers the benefits of easier drivability, signal distribution, and relaxed calibration overhead for a less efficient sub-ADC. The latter trades the superior efficiency of the sub-ADC for

Table 6.2 Performance summary and comparison with state-of-the-art wideband TIRF ADCs

	This work	Straayer [147]	Wu [148]	Ali [150]	Devarajan [104]	Vaz [89]	Nam [149]
Technology	ISSC'19	ISSC'16	ISSC'16	VLSI'16	ISSC'17	VLSI'18	JSSC'18
Architecture	TI-pipe-SAR	TI-pipeline	TI-pipeline	TI-pipeline	TI-pipeline	TI-pipe-SAR	TI-SAR
Interleaving factor	8x	16x	4x	2x	8x	8x	32x
Supply [volts]	1.0	1.0/1.8	0.8/1.8	0.9/1.8/2.5	1.0/2.0	0.9/1.8	1.1/1.2/2.5
Sample rate [GS/s]	5.0	4.0	4.0	5.0	10.0	5.0	1.6/3.2/6.4
Max. f_{in} [GHz] ^a	2.4	1.8	1.9	2.0	4.0	2.4	1.0
Bandwidth [GHz]	> 6.0	4.0	N.A.	5.0	7.4	N.A.	N.A.
Noise floor [dBFS/Hz]	-160	-154	-159	-157	-157	-156	N.A.
SFDR @ low f_{in} [dB]	75.2	N.A.	75.0 ^b	80.0 ^b	69.0 ^b	75.0 ^b	82.5 ^{bd}
SFDR @ max f_{in} [dB]	62.4	N.A.	60.0 ^b	63.0 ^b	59.0 ^b	60.0 ^b	67.3 ^{bd}
SFDR @ max f_{in} [dB]	65.4	64.0	68.0	70.0	64.0	61.9	68.2 ^d
SFDR @ max f_{in} [dB]	58.5	55.5	56.0	58.0	55.0	57.0	58.4 ^d
Power consum. [mW]	158.6	2200.0	300.0	2300.0	2900.0	641.0	225.0 ^d
FoMw ^c [fJ/conv-step]	46.1	1130	145.5	708.7	631.2	221.6	165.5 ^d
FoMs ^c [dB]	160.5	145.1	154.2	148.4	147.4	152.9	154.9 ^d

^a Within the 1st Nyquist

^b Estimated value based on graphs

^c @ $\min\{\text{Max } f_{in}, \text{Nyquist } f_{in}\}$

^d @ 6.4 GS/s

a more complex signal distribution and calibration (unless hierarchy is considered). The first strategy was preferred, due to its better thus far results, combined with an improved efficiency pipelined-SAR hybrid sub-ADC.

Finally, a 5 GS/s 12-bit highly efficient $8\times$ directly interleaved passive-sampling wideband RF ADC was presented. The challenges of efficiently achieving wide input bandwidth and high spectral purity in the absence of a front-end buffer were addressed with a minimized resistance/capacitance network. Sampling purity was ensured by an on-chip clock conditioning/distribution chain with as low as $11\text{ fs}_{\text{rms}}$ additive jitter and a segmented DDT that corrects the timing mismatch among sub-ADCs with a 9 fs step. The power efficiency was enhanced by an asynchronous three-stage pipelined-SAR hybrid sub-ADC with a single comparator per stage and an integrator DRA. A synthesized digital correction block improves the spectral performance over the entire band of interest.

The 28 nm bulk CMOS prototype demonstrates a bandwidth in excess of 6 GHz, with a Nyquist SFDR/SNDR of 65.4/58.5 dB and a total power consumption of 158.6 mW at 5 GS/s. This performance advances the SotA among wideband TI RF ADCs in both FoM_S (160.5 dB) and FoM_W (46.1 fJ/conv-step).

Appendix D: TI ADC Power Estimation with On-Chip Input Buffer

For completeness, we provide here a first-order estimation of the increase in the power consumption, in case an on-chip input buffer would be employed to actively drive this ADC with a sufficiently wide input bandwidth. This estimation assumes a push-pull source follower structure [101] with a low enough output impedance ($Z_{\text{out}} \approx 1/g_m \approx 5\ \Omega$). This would necessitate $g_{m,\text{NMOS}}$ and $g_{m,\text{PMOS}}$ of 100 mS from each of the complementary sides, assuming an equal NMOS/PMOS strength. If we design for a g_m/I_D of about 10 S/A, easily achievable in process like 28 nm, this would translate to a total quiescent current from the differential circuit of about 20 mA. It is also necessary to utilize a supply voltage of at least 1.8 V to allow a one-row or two-row cascoding for improving the buffer linearity. This would lead to a power consumption from the differential circuit of about 36 mW. If another 4 mW is added for biasing purposes, a total power consumption from the buffer of 40 mW is estimated. This would result in a total ADC plus input buffer power of 198.6 mW, leading to a FoM_S and FoM_W of 159.5 dB and 58 fJ/conv-step, respectively.

Chapter 7

Ultra-Wideband Direct RF Receiver

Analog Front-End



The challenge to continue increasing the RF sampling ADC sample rate and bandwidth, to enable next-generation ultra-wideband applications, does not lie only with the converter core. While time-interleaving can enhance the sample rate, the same cannot be said about the bandwidth, which should be extended by the front-end preceding the ADC while maximizing the spectral purity and limiting the excess power consumption. Hence, novel front-end solutions toward this direction are essential and highly desirable.

Section 7.1 of this chapter revisits the problem of extending the bandwidth beyond several tens of GHz and discusses the challenges along with a prior art overview. Section 7.2 introduces the prototype ultra-wideband highly integrated analog front-end and discusses its innovative performance-advancing features in detail. The experimental verification, including the complete measurement setup, measured results, and a recent state-of-the-art comparison, are the subject of Sect. 7.3. Finally, Sect. 7.4 draws the conclusion of this chapter.

Parts of this chapter were previously presented at the 2022 *Symposium on VLSI Technology and Circuits (VLSI'22)* in Honolulu, HI, USA [157]. Also, two inventions have been filed with the US Patent and Trademark Office, one of them published and granted [158] and the other one accepted and pending publication.

Special thanks go to Dr. Eng. Gabriele Manganaro, MediaTek USA Inc., Woburn, MA, USA (previously with Analog Devices Inc., Wilmington, MA, USA), for enabling and contributing to the work covered in this chapter.

7.1 Pushing the Bandwidth Beyond 20 GHz

The benefits of direct RF sampling receivers in deep-scaled CMOS, such as the simplification of the receiver analog signal chain, the lower cost and footprint, and the flexibility with the constantly improving DSP, were already highlighted in Chap. 6 (see Sect. 6.1). It is highly desirable to extend the direct RF sampling concept up to and above mm-Wave frequencies (≥ 25 –30 GHz) and several tens of GHz bandwidth, with high spectral purity (HD3, IM3 ≤ -55 dB, NSD ≤ -150 dBFS/Hz) and low power. These requirements are demanded so as to empower leading-edge wireless communications (5G mm-Wave FR2, future 6G), radar, and instrumentation. The analog front-end is paramount in such a receiver, as it must deliver on all the above metrics while seamlessly providing extra integrated functionality (gain/attenuation) to control the signal range.

7.1.1 Revisiting the Analog Front-End Problem

In Chaps. 5 and 6, the challenges of extending the sample rate and bandwidth of both wireline and wireless ADC-based receivers, while preserving high spectral purity with low power, were highlighted. While time-interleaving can extend the sample rate, the same cannot be said about the bandwidth. The main limitation in extending bandwidth stems from the non-negligible time constant at the input of the ADC, which includes the sampling capacitor and the parasitic capacitance from routing and switches, as well as the equivalent input, switch, and wiring resistance. The ESD and pad capacitances further degrade this time constant. The capacitance of the ESD devices for industry-approved protection (1–2 kV Human Body Model (HBM), 250–500 V Charge Device Model (CDM)) is easily on the same order as the sampling capacitor (200–400 fF). To make matters worse, this capacitance is highly non-linear and may already degrade the input linearity considerably¹ while not strictly obeying scaling.

To address this challenge, the solution proposed in the previous chapter [93] optimized this time constant by minimizing each of the above contributions, including employing minimized custom ESD diodes (~ 100 fF with pad), which slightly degrades robustness. Although extremely power efficient, the sample rate and bandwidth of this purely passive solution, with sufficient spectral purity, are limited to below 10 GHz. This is not enough to guarantee a smooth 1st or 2nd Nyquist zone operation into higher frequencies (K- and Ka-bands), found in mm-Wave 5G. The straightforward way of pushing the sample rate with this approach would be to increase the interleaved channel count. However, this increases the parasitic front-end loading, while the noise-limited sampling capacitor and equivalent input

¹ Extracted simulations with industry-approved 16 nm FinFET ESDs show an HD3 of -81 dB for a 500–600 mV_{pp,diff} input at a 5 GHz frequency, which drops to -72 dB at 20 GHz.

resistance remain constant, and the switch resistance cannot reduce indefinitely. Hence, these approaches make it very challenging to improve the bandwidth without compromising the spectral purity.

The most widely adopted approach in literature to enhance the bandwidth is to employ a class-A front-end buffer [72, 88, 89, 147, 148, 150]. This reduces the loading and impedance variations on the ADC preceding circuit and actively drives the noise-limited sampling capacitor and the routing parasitics. This buffer acts as an active impedance transformer, separating the equivalent input resistance and ESD capacitance from the sampling capacitor and switches. However, for the buffer to provide an acceptable linearity, the signal swing must be reduced, necessitating a larger sampling capacitor for the same SNR compared to a buffer-less approach. This is only exacerbated in finer CMOS processes with the constant supply down-scaling. The devices' excess noise imposes an extra overhead on the sampling capacitor up-scaling. To achieve a sufficiently low output impedance ($Z_{\text{out}} \approx 1/g_m$), the front-end buffer dissipates a comparable power to the back-end ADC. The works in [101, 159] partially reduce the power dissipation by employing a class-AB buffer, exploiting the current re-use properties to effectively double the g_m for the same current or half the current for the same g_m . Nevertheless, the large devices needed for a sufficient g_m significantly increase the input-output buffer capacitance, excessively loading the preceding circuit and degrading the isolation at high frequencies. Thus far, [159] reports the highest achievable bandwidth of 18 GHz with an HD3-limited SFDR of 62 dB at 10 GS/s and 4 GHz input frequency, dropping to 54 dB at 18 GS/s and 8 GHz input frequency, thus limiting the bandwidth with acceptable spectral purity below 10 GHz. Further, the 18 GHz bandwidth is at the onset of the K-band (18–27 GHz) and far from the Ka-band (27–40 GHz), while the unity gain re-sampling front-end dissipates 220 mW.

The general take is that an active front-end is imperative to push the bandwidth above 10 GHz, but thorough investigation is needed to further extend it by 2–3 \times . Moreover, the thus far discussed limitations involve solely on-chip contributions. The chip is either flipped or wire bonded on a substrate through a set of copper pillars or bondwires, respectively, as shown in Fig. 7.1, and the critical signal is carried by characteristic impedance traces. For a 20–30 GHz frequency, the flipping option is preferred to minimize the inductance and ensure high signal integrity. The chip-to-substrate interface and the substrate traces should also be taken into account, if possible co-optimized together with the on-chip elements.

To better understand the impact of on-chip/interface/off-chip contributions and identify places for improvement in the chain, the front-end model of Fig. 7.2 is built. The model assumes a 50 Ω terminated buffered front-end, targeting a 30 GS/s operation in a 2 \times ping-pong fashion. Looking from the input, the S-parameters of a high-frequency substrate material trace are included. The copper pillar is modeled as a CLRC Π -network, with the indicated values. The ESD capacitance corresponds to the extracted value of a 16 nm industry-approved RF ESD device pair. The sampling capacitor of 300 fF is chosen for an SNR of 10 bits or above, under typical signal swings of active front-ends (0.5–1 V_{pp,diff}). This capacitance guarantees a NSD below –150 dBFS/Hz for the targeted 30 GS/s operation, considering also a




Fig. 7.1 Two chip-to-substrate illustrations: (a) flip-chip through copper pillars and (b) wire bonding through gold bondwires

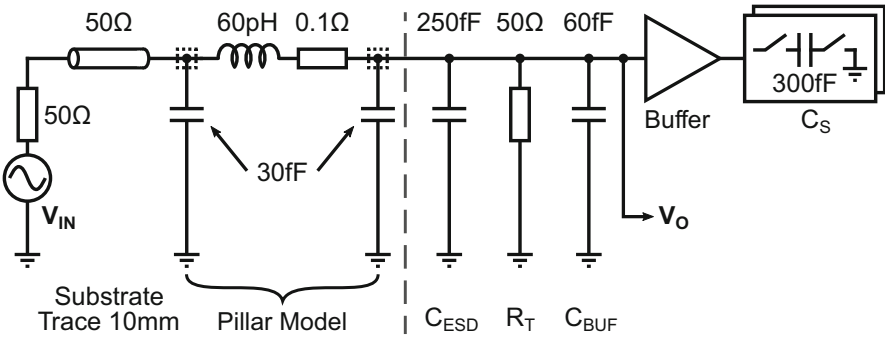


Fig. 7.2 Buffered front-end model including on-chip/interface/off-chip contributions (single-ended shown)

1.5–2× margin for buffer excess noise. Assuming a buffer settling requirement to 1/2 LSB accuracy of 10 bits at a full period of 30 GS/s (due to ping-pong), its input capacitance can be estimated for a chosen f_T , similar to our analysis of Chap. 3

$$C_{BUF} \approx C_S \cdot \frac{f_s}{f_T}. \quad [F] \quad (7.1)$$

For a 30 GS/s f_s and an f_T of about 150 GHz, easily achievable in 16 nm FinFET (see Chap. 3, Fig. 3.4), a minimum C_{BUF} of 60 fF is estimated.

The front-end with the explained values above is simulated, and the results are plotted in Fig. 7.3. The transfer characteristic is assessed at the buffer input (V_O in Fig. 7.2). The different contributions are added one by one, such that the performance degradation from each in the overall transfer characteristic can be more easily quantified. With all the contributions added, a best-case scenario indicates an S_{21} of –3 dB at about 15 GHz. This is dominated by the ESD capacitance, with the buffer capacitance following (Fig. 7.3a). When the buffer capacitance is removed, the –3 dB point is extended to about 18 GHz. S_{11} reaches –10 dB at about 7 GHz and 8 GHz, including or excluding the buffer capacitance, respectively. It quickly becomes evident that the ESD presents among the main bandwidth-limiting capacitances and needs to undergo at least a twofold reduction, if the targeted

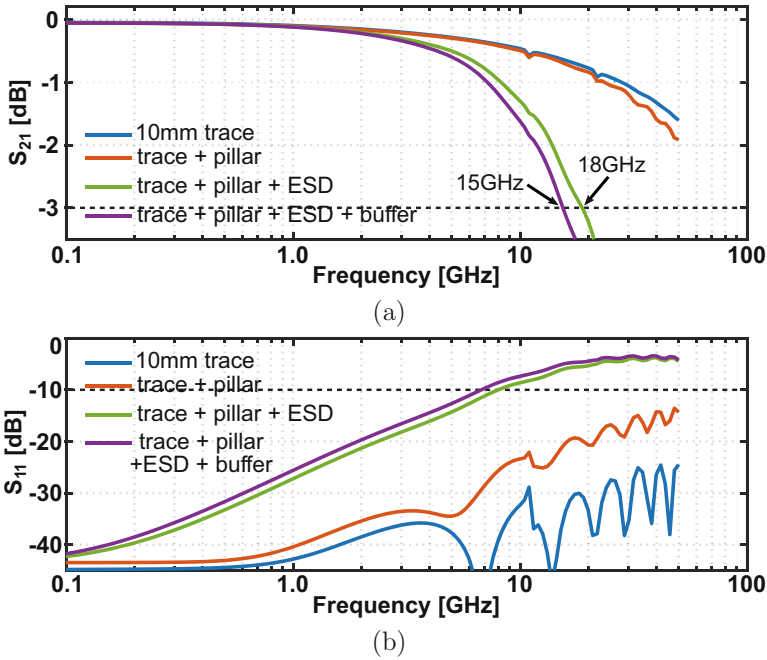


Fig. 7.3 Simulated S-parameters of the front-end model, gradually adding the contributions: (a) S_{21} and (b) S_{11}

30 GHz bandwidth is to be reached and/or surpassed. At the same time, a decent reliability must be preserved with minimum robustness degradation. On top, special attention should be put on the buffer design to ensure a minimum input-output isolation degradation at high frequencies (neglected in this first-order model).

7.1.2 Increasing Integration and Challenges

CMOS integrated RF ADC-based receiver front-ends reported in up-to-date literature are limited to a fixed unity gain buffer.² However, wideband variable gain/attenuation is necessary in the receiver prior to the ADC. The twofold purpose of this function is to prevent SNR degradation by amplifying small amplitude signals to the expected ADC range as well as to attenuate larger signals that would increase distortion and/or cause clipping. It is highly desirable to integrate such a function on the same back-end ADC CMOS chip; to reduce power, area, and cost;

² Whether in a direct interleaving [101] or in a re-sampling scheme [159], the concept and limitations of the fixed unity gain remain the same.

and to increase flexibility. This allows a substantially smaller and better controlled front-end—ADC interface, reducing high-frequency signal integrity degradation due to chip-to-chip transitions. Yet, it is non-trivial to integrate extra variable gain/attenuation in the front-end of Fig. 7.2, without degrading bandwidth and linearity or increasing power. In the simplest case of replacing the buffer with a gain stage having the same settling requirements, its non-linear input capacitance goes up by the gain factor plus the Miller effect, adding both an input bandwidth and power overhead. Hence, it becomes extremely challenging to achieve altogether a highly integrated ultra-wideband operation with high linearity, low noise, and low power. This is further exacerbated in finer CMOS processes due to the supply down-scaling and the increased interconnect parasitic contribution.

ADC-based serial-link receivers have demonstrated the benefits of large front-end integration by combining VGA stages with passive and active high-pass filtering [151, 160–162]. Bandwidths up to 30 GHz have been achieved with the additional support of T -coils [163] and distributed inductive peaking [164]. The T -coil tunes out the ESD capacitance by employing the mutual coupling of two inductors, while the distributed peaking splits it into two or more segments, emulating a pseudo-transmission line. However, to tune out capacitances above 250 fF, the T -coil inductors require large values (> 150 – 200 pH) and a sufficiently large coupling, which might not be easy to achieve on-chip. Further, splitting the ESD into several smaller segments may severely degrade robustness or even damage the protection device itself (e.g., if all the current goes to one segment prior to reaching the others). Besides, the variable gain/attenuation is typically done by controlling device elements (g_m , r_o , C), which results in a non-constant and potentially degraded bandwidth across the different settings. The downside of these solutions is the poor static and dynamic linearity ($\text{HD}_3 \geq -45$ dB), especially across such a wide bandwidth, insufficient to cover the needs of the aforementioned targeted applications. Yet, they hint directions toward ultra-wideband highly integrated front-ends.

7.2 Prototype IC: A 30 GHz-Bandwidth < -57 dB-IM3 Front-End in 16 nm FinFET CMOS

This section presents a 30 GHz-bandwidth analog front-end for a direct RF receiver that achieves a better than 58 dB-SFDR and better than -57 dB-IM3 while supporting 1024-/2048-QAM modulation with excellent spectral purity across its entire bandwidth. A multi-segment Chebyshev LC filter distributes the input ESD protection while providing a 0–11 dB variable attenuation across the entire bandwidth. A new push-pull hybrid amplifier, followed by a push-pull buffer, significantly improves the gain-bandwidth and noise at no extra power consumption, while their linearity across the entire band of interest is enhanced by resistive degeneration and bootstrapped cascoding, respectively. Fabricated in 16 nm FinFET CMOS, the

prototype front-end occupies a compact core area of $540 \times 280 \mu\text{m}$ and draws a total current of 52.5 mA from a dual rail $\pm 2\text{ V}$ supply, for a total power consumption of 210 mW.

7.2.1 High-Level Front-End Chain

The top-level block diagram of the proposed highly integrated front-end is given in Fig. 7.4. The input signal is applied to a differential 50Ω characteristic impedance filter, which distributes a parasitic capacitance-optimized ESD protection circuitry and a digital step attenuation to enhance the frequency response. The attenuation is variable with a between 0 and 11 dB range and a 1 dB step and can be digitally controlled. The filter is able to absorb the parasitic capacitances of the pad, the ESD, the attenuation, the input termination, and the following amplifier, significantly extending the bandwidth at the termination point (see Sect. 7.2.2).

The filter is followed by a two-path push-pull hybrid amplifier with a fixed gain of 6 dB and a unity-gain push-pull cascoded buffer to drive the output. The combined gain/attenuation ranges from +6 dB to -5 dB and is responsible for providing the required fixed swing at the output. Providing both gain and attenuation fulfils the twofold purpose of amplifying small amplitude signals to improve the dynamic range as well as attenuating larger signals to prevent excessive distortion. By designing the variability in the attenuation within the impedance-matched filter, while keeping a fixed gain amplifier, offers the benefit of a constant bandwidth, linearity, and noise across the settings, compared to the typical variable gain amplifier. Further, since the variable attenuation mainly comprises passive elements, any degradation in its superior linearity and noise negligibly impacts the overall front-end performance.

The output network, containing the series termination resistor $R_{T,O}$, the ESD, and the pad, presents the load for the standalone evaluation of this front-end. This network would not be needed with an integrated ADC back-end, which would instead present a capacitive load, with typical values ranging between 200




Fig. 7.4 Top-level block diagram of the proposed front-end (single-ended shown for simplicity)

and 400 fF, depending on swing and SNR requirements. Therefore, for a realistic bandwidth assessment, the component values of this network were chosen for a time constant that roughly matches the one, where the buffer would directly drive a 300 fF ADC capacitive load.

7.2.2 Filter with Distributed ESD and Variable Attenuation

As already mentioned, to extend the input bandwidth in ADC-based receivers, the T -coil and distributed inductive peaking have been largely employed. The goal of these approaches is to tune out ESD and other device capacitances to keep the impedance of the chain closer to the desired characteristic impedance Z_O for a targeted range of frequencies. This, in turn, minimizes the reflection coefficient (S_{11}) and maximizes the transmission coefficient (S_{21}) for the targeted range of frequencies. Assuming a perfectly terminated and lossless network, S_{21} and S_{11} can be linked as follows:

$$|S_{11}|^2 = 1 - |S_{21}|^2. \quad [\text{dB}] \quad (7.2)$$

For capacitances on the order of 300 fF and above, a single T -coil usually does not suffice to guarantee a bandwidth larger than 30 GHz; hence, the distributed inductive peaking is more effective. The size of the implemented inductors depends on the capacitance to be tuned out on each segment and the desired characteristic impedance of the chain

$$Z_O = R_T = \sqrt{\frac{L}{C}} \Rightarrow L = R_T^2 \cdot C \quad [\Omega] \quad (7.3)$$

where R_T is the termination resistance, typically 50 Ω . The above expression reveals that the smaller the capacitance in each segment, the smaller the required inductance. To closely resemble the behavior of a transmission line, a large number of segments are required to minimize the ripple in the S_{21} and S_{11} due to the lumped nature of the components. This can significantly increase the area and the design complexity. The latter holds true when the total capacitance cannot be evenly split; hence, a different inductance needs to be realized in every segment, increasing the optimization effort. Finally, the above expression hints that if possible, a smaller characteristic impedance would be beneficial to reduce the inductor values and minimize unwanted coupling to surrounding blocks.

A more attractive way of applying the distributed inductive peaking would be to distribute the total capacitance and size the inductors following filter design theory [165]. Different types of filters exist (Chebyshev, Butterworth, Bessel, Legendre). Depending on the total capacitance to be tuned out and the targeted bandwidth, the filter order and the values of the components can be determined [166]. To properly absorb a capacitance above 300 fF (Fig. 7.2) and maintain a larger than 30 GHz

bandwidth, a minimum fifth-order filter is necessary, considering an ESD splitting of no more than $2\times$ for a minimum robustness degradation. In reality, an even higher order is needed to account for other capacitances in different parts of the chain, including termination, pad, routing, and other functional blocks (e.g., attenuation). Further, the filter cut-off frequency $f_{-3\text{dB}}$ should be designed considerably higher than the targeted 30 GHz, such that it does not limit the aggregate bandwidth of the entire front-end chain. On the other hand, too high of a filter order may significantly increase the area and affect the wideband matching due to the finite quality factor of series practical inductors and routing.

Considering the above, this work adopts a ninth-order filter, and its component values follow a symmetric Chebyshev type I sizing (Fig. 7.5a). To reduce the inductor count, a shunt-first topology with capacitors $C_1 = C_9$ at the edges is adopted, which absorbs better the pad and termination capacitances. This results in four inductors (instead of five for series-first), with $L_2 = L_8$, $C_3 = C_7$, and $L_4 = L_6$ due to the filter symmetry around its center tap C_5 . The Chebyshev type I is preferred since, for the same order and $f_{-3\text{dB}}$, it can absorb a higher capacitance compared to other practical filter types [165]. The component values for two termination options are also shown. An $f_{-3\text{dB}}$ greater than 70 GHz does not impose an aggregate bandwidth limitation, while a passband ripple smaller than 0.1 dB is a good compromise between capacitance absorption and flatness in S_{21} and S_{11} . As implied by Eq. (7.3), for the same $f_{-3\text{dB}}$, an R_T of $25\ \Omega$ reduces the inductor values by $2\times$ for an equivalent increase in the capacitor values, compared to the typical $50\ \Omega$, both of which are highly desirable. Due to this dual benefit, a $25\ \Omega$ on-chip termination ($50\ \Omega$ differentially) is adopted, realized with high-R polysilicon resistors. It should be noted that this puts a burden on the external circuitry driving this front-end, which for the same voltage swing should be able to deliver twice the current. The simulated S-parameters and group delay of this filter with top-metal EM-extracted inductors are shown in Fig. 7.5b. An S_{11} below -15 dB is maintained up to about 70 GHz with the S_{21} dropping by 3 dB from its low-frequency value at about 75 GHz. The simulated group delay deviation is less than 1 ps up to the aggregate front-end bandwidth of 30 GHz. A constant group delay is desirable to minimize the phase distortion across the entire bandwidth of interest.

The specific implementation and distribution of the components in the proposed filter is shown in Fig. 7.6. The ESD protection is split into only two segments, occupying the first two segments of the filter and avoiding the robustness and reliability issues of multiple segments with too small devices [166]. Both segments comprise stacked devices to reduce the parasitic capacitance and improve the high-frequency linearity. The stacking and device sizing in each segment are different and optimized for equal current paths, including the wiring resistance through L_2 . This configuration minimizes altogether the bandwidth and linearity penalty while still offering a sufficient protection (250 V CDM, 1 kV HBM). The 0–11 dB/1 dB-step attenuation is distributed across the filter in the form of four impedance-matched attenuator cells with ascending values 1 dB, 2 dB, 4 dB, and 4 dB. The ESD and the capacitance of the attenuators and the termination are almost entirely absorbed as part of the filter capacitance. The first segment is entirely occupied by the pad and

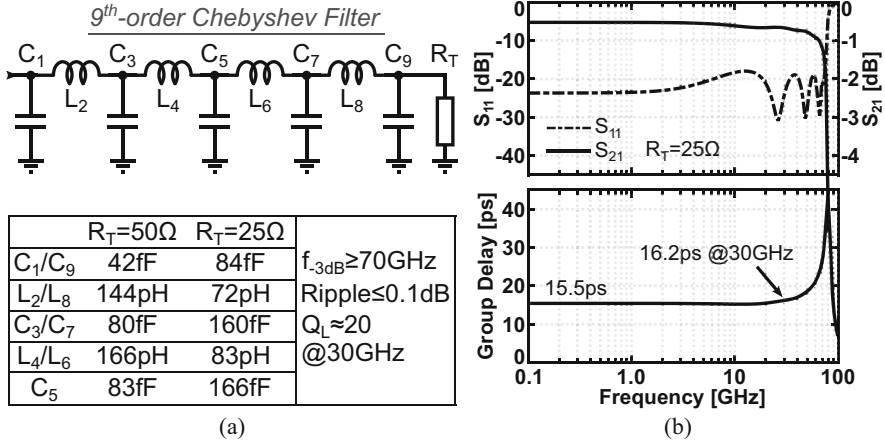


Fig. 7.5 (a) Ideal ninth-order Chebyshev filter with its component values for two R_T values and (b) simulated S-parameters and group delay




Fig. 7.6 Implementation of the proposed filter with the distributed ESD and variable attenuation (single-ended shown)

first ESD part. The capacitances of the second ESD part, the attenuators, and the termination are accommodated by subtracting from the filter MOM capacitors an amount that maximizes the bandwidth after EM-extracting the entire filter.

The attenuator cells employed in this work are shown in Fig. 7.7. Two different hybrid polysilicon resistor/NMOS device cells are adopted, a Π -cell (Fig. 7.7a) and a T-cell (Fig. 7.7b). Both cells include a one-device or two-device series path R_{ser} , one or two shunt paths R_{sh} , and a bypass path R_{byp} for the 0dB setting. The series paths are realized with triode region NMOS devices M_1 and M_{1A}/M_{1B} , to achieve the very small R_{ser} values (Fig. 7.7c) with small area overhead. These are always-on devices to lower the resistance for the 0dB setting due to the parallel combination with M_3 devices. The shunt paths comprise a combination of polysilicon resistor and NMOS devices M_{2A}/M_{2B} and M_2 . All devices are realized

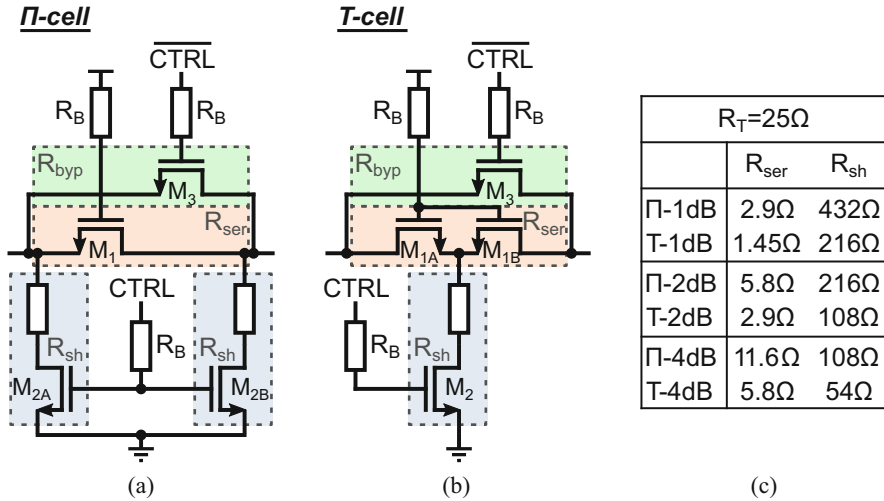


Fig. 7.7 Attenuator cells employed in this work: (a) Π -cell, (b) T-cell, and (c) their resistance values rounded to include multiple matched units

using multiples of the same unit, connected on the top metal layers, for optimum matching and minimum performance degradation due to the interconnect-intensive FinFET process. The linearity of both cells is enhanced by bootstrapping their gates through the large resistors R_B . Their linearity is improved by gate-bootstrapping through large resistors R_B .

The arrangement of the attenuator cells in the filter requires special attention to yield maximum bandwidth and linearity benefits. Intuitively speaking, if a fixed swing is needed at the output of the filter, then the input is adjusted according to the attenuation needed to reach that swing. Looking at Fig. 7.7c, it becomes clear that the cell with the smallest attenuation should be placed at the filter input since it experiences the smallest voltage drop across a non-linear on-resistance; thus, it yields a better linearity at larger swings. To verify the intuition and determine the best arrangement and cell type for the different settings, Fig. 7.8 compares their bandwidth and linearity (IM3) for a swing at the output of the filter of about $300\text{ mV}_{pp,diff}$.³ The Π -cell offers a higher bandwidth, thanks to the smaller capacitance of its $1\times$ relative size single M_1 , compared to the two series M_{1A} - M_{1B} of the T-cell, with $2\times$ relative size each. The benefit is particularly pronounced for the 1 dB setting, while for the higher settings, the gap is gradually bridging. However, the T-cell exhibits a superior linearity by virtue of its $2\times$ relative size M_{1A} - M_{1B} , since each offers half the non-linear on-resistance. To achieve the best performance and make the best use of the available capacitance for absorption across

³ The swing at the input then goes as high as about $1.1\text{ V}_{pp,diff}$ ($\sim 3.6\text{-}0.3\text{ V}_{pp,diff}$) for 11 dB attenuation.

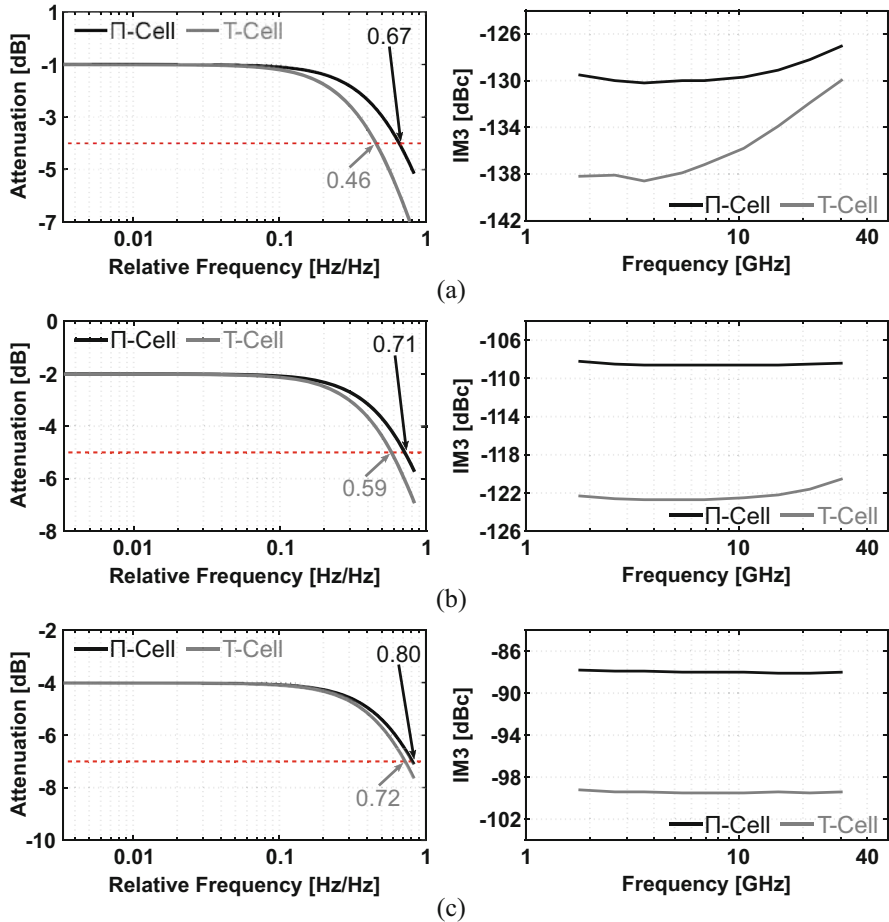


Fig. 7.8 Simulated bandwidth (relative) and linearity of the two attenuator cells for the different attenuation settings: (a) 1 dB, (b) 2 dB, and (c) 4 dB

the filter chain, the Π -cell is adopted for the 1 dB setting with sufficient linearity for the utilized swing. This also adds the minimum extra capacitance on the existing one from the second ESD part sharing the same filter segment (Fig. 7.6). The 2 dB and 4 dB settings utilize the T-cell for its superior linearity, with sufficient room in the other filter segments to easily absorb its extra capacitance.

The complete filter with a combined RC extraction of the active cells and EM extraction of the passive elements is simulated, and the S_{21} and S_{11} for all the settings are plotted in Fig. 7.9. In Fig. 7.9a, S_{21} shows a uniform step across the settings with a deviation less than 0.15 dB from the nominal 1 dB step. It also shows a very low loss for all settings of about 1.2 dB or less. This is attributed to the accumulated series resistance of the metal inductors and the attenuators' bypass




Fig. 7.9 Simulated S-parameters of the implemented filter across the different attenuation settings: (a) S_{21} and (b) S_{11}

devices, which is maximum for the 0 dB setting. The worst-case $f_{-3\text{dB}}$ is about 70 GHz for the 0 dB setting, while the best-case $f_{-3\text{dB}}$ for the 11 dB setting is about 75 GHz. In Fig. 7.9b, S_{11} shows a very wideband well-behaved characteristic, with values below -14 dB up to about 70 GHz for all the attenuation settings. This simulation indicates that for the targeted front-end bandwidth of 30 GHz, the proposed filter and design techniques do not impose a limitation, while adding the variable attenuation, not present in previously reported deep-scaled CMOS RF sampling ADC-based receivers. Finally, since the filter comprises mainly passive elements and switches, its power consumption is negligible.

The linearity of the complete filter chain is also simulated by applying a two-tone with fixed frequency spacing and sweeping the carrier frequency. The power of each tone is chosen at -6 dBFS compared to a one-tone test, and the amplitude is adjusted to yield a fixed output swing of about $300 \text{ mV}_{\text{pp,diff}}$ (-6 dBFS each tone). The IM3 vs. frequency for the best- and worst-case attenuation settings is plotted in Fig. 7.10. The best-case IM3 occurs for the 0 dB setting, when all attenuator cells are in the bypass mode, and stays below -110 dB up to the targeted 30 GHz. This is mainly due to the split ESD protection and the on-resistance of the attenuators' bypass devices. The worst-case IM3 is found for the 11 dB setting, when all attenuator cells are active, and remains below -85 dB across the entire 30 GHz band of interest. The




Fig. 7.10 Simulated filter two-tone IM3 vs. frequency for the best (0dB) and worst (11dB) attenuation settings

degradation is attributed to the accumulated distortion from all the attenuator cells, with the 4 dB cells heavily dominating, as already indicated by Fig. 7.8.

7.2.3 Two-Path Push-Pull Hybrid Amplifier

The proposed fully differential hybrid amplifier is detailed in Fig. 7.11. The amplifier plays a critical role in an ultra-wideband highly integrated front-end. Its performance may limit the achievable bandwidth, noise, and linearity as well as dominate the front-end power consumption in the given 16 nm process. This work introduces a new hybrid amplifier topology that combines several key aspects in order to maximize the performance and reduce the power consumption. To meet the stringent bandwidth requirements, an open-loop topology is preferred over a traditional amplifier with feedback. Limiting the description to the left side, it comprises a push-pull Common Gate (CG) pair $M_{1NA}-M_{1PA}$ and a push-pull Common Source (CS) pair $M_{2NA}-M_{2PA}$ stacked on the same branch. The push-pull topology is chosen over class-A, due to its current re-use properties, to reduce the power and noise for a certain bandwidth. The input V_{I+} connects to the CG pair through the termination $R_{T,I}$, which in series with the $1/(g_{m1,N}+g_{m1,P})$ composes the very wideband input termination of $25\ \Omega$. $R_{T,I}$ also makes a readily available source degeneration resistor for the CG pair, improving its g_m distortion [21]. In parallel, the complementary input V_{I-} is AC-coupled to the CS pair through capacitors C_C . This novel hybridism enables a two-path parallel signal amplification on the load R_D . To improve the g_m distortion of the CS pair, the explicit source degeneration resistors R_S are added. The low-frequency total gain on each R_D can be approximated by

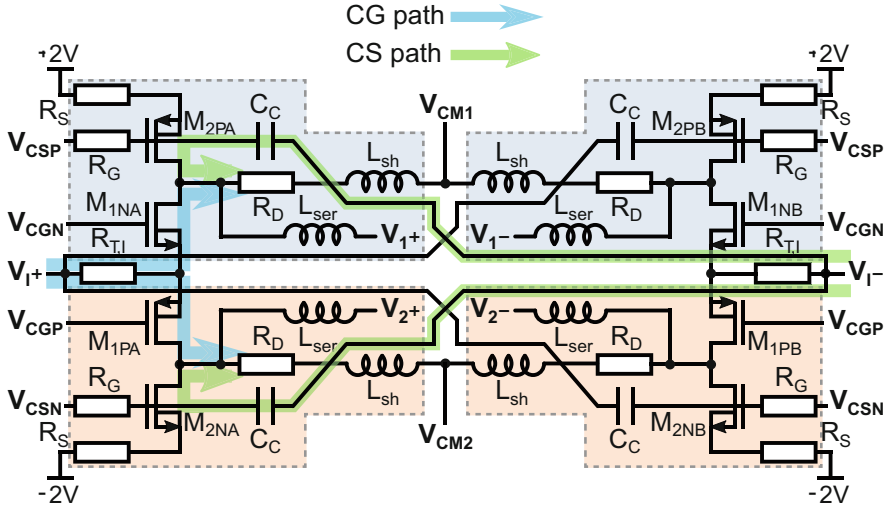


Fig. 7.11 Proposed push-pull hybrid CG-CS amplifier with resistive source degeneration and series-shunt peaking

$$\begin{aligned}
 |Gain| &= \left(\frac{g_{m1}}{1 + 2g_{m1}R_{T,I}} + \frac{g_{m2}}{1 + g_{m2}R_S} \right) \cdot R_D \\
 &\approx \frac{R_D}{2R_{T,I}} + \frac{R_D}{R_S}, \text{ if } g_{m1}R_{T,I}, g_{m2}R_S \gg 1, \quad (7.4)
 \end{aligned}$$

where the devices' intrinsic output impedance r_o contribution is excluded and the $2\times$ takes into account the splitting of $R_{T,I}$ due to the push-pull CG path.

The above expression reveals, either explicitly or implicitly, several performance enhancements the proposed hybridism offers. First, for a fixed load resistor, which together with the output capacitance determines the dominant pole, the gain is increased by an amount of $g_{m2}R_D/(1+g_{m2}R_S)$ (w.r.t the CG pair) or $g_{m1}R_D/(1+2g_{m1}R_{T,I})$ (w.r.t the CS pair). This improves both the gain-bandwidth and the SNR compared to a CG- or CS-only stage, where the load devices would add only noise without any signal amplification. Further, for a certain gain-bandwidth, the input loading is also reduced. Most of the M_{1NA} – M_{1PA} source capacitance is “hidden” behind the degeneration $R_{T,I}$ making this node close to an AC ground. For the M_{2NA} – M_{2PA} , the Miller effect is reduced by the amount of the gain increase due to the hybridism ($g_{m1}+g_{m2}/+g_{m2}$), while their remaining gate capacitance is easily accommodated by the preceding filter (see Sect. 7.2.2). The linearity is also improved, by de-sensitizing the gain to the devices' g_m , making the gain largely a ratio of well-defined resistors implemented with the same unit cell. The other term affecting linearity is the devices' r_o , which is in parallel to the R_D

$$R_{o,tot} = [(1 + g_{m1}r_{o1}) \cdot 2R_{T,I} + r_{o1}]/[(1 + g_{m2}r_{o2}) \cdot R_S + r_{o2}]/R_D. \quad (7.5)$$

To maximize r_{o1} and r_{o2} and reduce their contribution to the overall gain and linearity, a large V_{DS} of 750 mV is allocated to both M_1 and M_2 devices. This results in the unavoidably large supply rails of $-2\text{ V} \rightarrow +2\text{ V}$ to properly accommodate the device stacking. However, no device experiences a voltage difference above the nominal process supply across any two terminals. Proper sequencing is employed during powering on the chip by first applying the input and biasing voltages with the supply rails low and then gradually increasing the rails to their final values, always keeping V_{GS} , V_{GD} , and V_{DS} below the nominal process supply.

To achieve the very wide targeted bandwidth, the gain of this hybrid amplifier is limited to about 6 dB with equal contribution from both pairs. All the devices are biased in S-I with a relatively low g_m/I_D ($\sim 12\text{ S/A}$) for the 16 nm process capabilities (see Chap. 3, Sect. 3.2.3), to minimize their parasitics on the critical output nodes. This helps in improving altogether the bandwidth, noise, and high-frequency linearity. The total current drawn from the dual rail $\pm 2\text{ V}$ supply is about 28.5 mA, which is reasonable given the stringent specifications. A higher gain without sacrificing the bandwidth may be achieved by reducing R_S and/or $R_{T,1}$ in exchange for a higher current from this stage and/or the preceding circuit driving this front-end with a lower than $25\ \Omega$ termination. Finally, to enhance the output impedance and reduce the gain drop at frequencies around the targeted 30 GHz, a series-shunt peaking combination [167] through L_{ser} and L_{sh} is employed. These inductors are kept small (below 100 pH) to avoid excessive overshoot and large area overhead, making use of the already existing top-metal interconnects on these nodes.

Although the single-branch device stacking requires higher supply rails, it is preferred over alternative multi-branch configurations (e.g., folded cascode). Folded cascode configurations would unavoidably increase the interconnect parasitics and degrade the high-frequency performance due to extra poles in the folding nodes while not necessarily reducing the power consumption. Therefore, to achieve the stringent bandwidth and analog performance needed for this front-end in a digital process like 16 nm FinFET, it is important to stress the importance of minimizing the current branches, where the interconnect parasitic contribution can easily dominate one of the devices.

7.2.4 Push-Pull Bootstrapped Cascoded Buffer

Due to the hybrid CG-CS push-pull operation, there are two outputs per branch V_1 and V_2 coming from the amplifier stage. These are transferred to the unity gain buffer stage shown in Fig. 7.12, to lower the output impedance and drive the final load. The buffer adopts a push-pull Common Drain (CD) topology [74, 101] to effectively double the g_m for a given current. A two-level cascoding is implemented, and the gates of both the input and the cascode devices are bootstrapped to the inputs through AC-coupling. This significantly reduces the drain modulation of both the input and the inner cascode devices, yielding a superior linearity compared to a single-level cascoding [159], at the expense of higher supply rails. Parallel

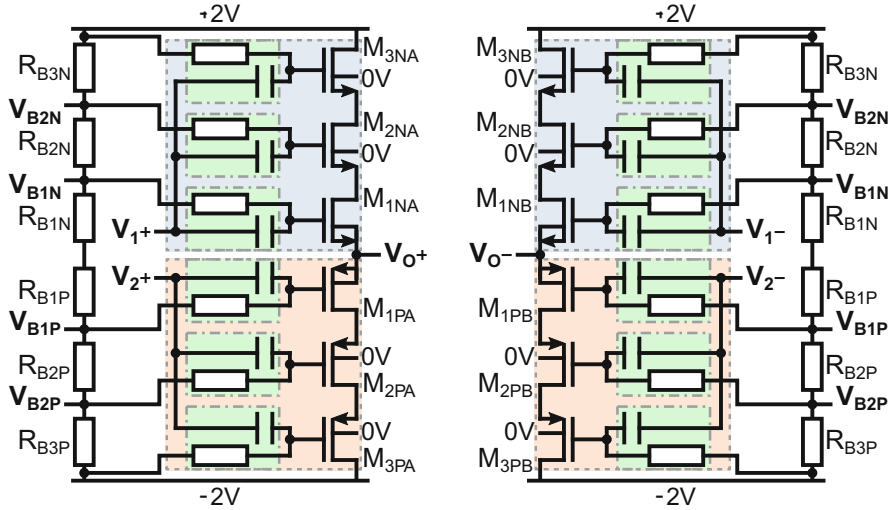


Fig. 7.12 Push-pull buffer with two-level bootstrapped cascoding

bootstrapping is adopted for all the buffer devices by connecting their gates directly to the amplifier outputs. In contrast to both [101, 159], where bootstrapping occurs sequentially through series combination of capacitors, the proposed approach makes the signals at all the gates more equal to each other and to the input, minimizing the loss due to sequential capacitive divisions and improving the effectiveness of the bootstrapping. The extra capacitive loading is not an issue due to the separate inputs coming from the amplifier and each driving three devices of the same type, in contrast to [101, 159], where the same input drives four complementary devices. Further, both the inner and outer cascodes are sized smaller than the input devices, and all devices are biased at relatively low g_m/I_D to minimize their parasitic contribution. Finally, while the input devices are tied to their sources for a gain as close to unity as possible, the bulks of the cascodes are reverse biased to 0 V to reduce well-diode parasitics in their sources.

The buffer conveniently shares the ± 2 V supply rails with the amplifier. Regarding the biasing voltages for the different devices, these are generated through resistive ladders comprising the same unit resistor and DC-coupled to the different gates through large resistors (Fig. 7.12). The buffer draws a total current of about 24 mA including the biasing. For prototype evaluation and tuning purposes, the ladder biasing voltages V_B are also taken off-chip through large series resistors (not shown in Fig. 7.12). In case a more robust biasing against PVT variations is needed, current sources can be used across the ladder [101], controlled through a common-mode feedback loop.

The simulated transfer characteristic of the amplifier and buffer chain is plotted in Fig. 7.13. Two output load scenarios are compared, one with a 300 fF capacitive load and one with the actual implemented matched load (Fig. 7.4). Both load cases




Fig. 7.13 Simulated amplifier-buffer transfer characteristic for a capacitive load and the implemented matched load




Fig. 7.14 Simulated two-tone IM3 vs. frequency of the amplifier-buffer chain

exhibit a similar bandwidth and about 30 GHz, as intended. In the case of the capacitive load, an always-on series switch with an on-resistance of $3\ \Omega$ is also included. A slight overshoot below 0.5 dB is also noticed, which is not detrimental and also not uncommon for source follower circuits driving a capacitively dominated load. In the matched load case, the characteristic is normalized for a more straightforward comparison. Also, the extra parasitics of the ESD, the resistor, as well as routing to this resistor and toward the pad slightly degrade the bandwidth versus the capacitive loaded case. These extra parasitics are to a large extent non-existent in the end application of this front-end, within an ADC-based receiver.

To estimate the linearity of the amplifier-buffer chain, a similar two-tone test as the one for the filter is applied, and the simulation results with a fixed 300 fF load are plotted in Fig. 7.14. The power of each input tone is chosen at $-6\ \text{dBFS}$ of $300\ \text{mV}_{\text{pp,diff}}$, so as to yield about $600\ \text{mV}_{\text{pp,diff}}$ at the output after the 6 dB gain. The simulated IM3 is about $-73\ \text{dB}$ close to 2 GHz, dominated by the amplifier.




Fig. 7.15 AC-noise simulation of the amplifier-buffer chain with 300 fF load

The resistive degeneration and large V_{DS} allocation help in improving the linearity at the expense of the higher supply rails required. The buffer alone achieves at least 10 dB better IM3 at this frequency due to the bootstrapping; therefore, it is not the dominant source. IM3 degrades smoothly due to the increasing effect of the non-linear device parasitics on the amplifier stage as well as the gradual diminishing of the bootstrapping benefit and device parasitic contribution on the buffer stage. However, at a 30 GHz frequency, it is still below -60 dB, thanks to the small device sizes due to the relatively low g_m/I_D values employed. At these high frequencies, this simulated value is already more than 15 dB lower than any highly integrated wireline receiver front-end and more than 5 dB lower than the unity gain buffer-only front-ends in literature while consuming a similar or lower power. Calibration as in [168] could be applied to further improve performance and/or lower the power.

The noise of the amplifier and buffer is also important, as it can dominate the total noise of this front-end and, eventually, determine the sensitivity of the entire receiver it aims to be integrated into. The simulated noise at the amplifier-buffer chain output, as it would be delivered to a back-end ADC with a 300 fF sampling capacitor on each side (150 fF differentially), is shown in Fig. 7.15. The integrated noise voltage is about $350 \mu V_{rms}$, with the biggest device contribution coming from the CS devices of the amplifier, followed by the input devices of the buffer stage. For the utilized output swing of $600 mV_{pp,diff}$, this simulated noise results in an SNR of 55.6 dB delivered to the back-end. This leads to a NSD of about -157 dBFS/Hz within the 1st Nyquist of a 30 GS/s ADC and about -160 dBFS/Hz within the 1st Nyquist of a 60 GS/s one, according to Eq. (2.33) from Chap. 2.

The linearity of the complete front-end is simulated and plotted in Fig. 7.16 for three different load scenarios, (1) 300 fF fixed capacitive load; (2) implemented matched load; and (3) 300 fF switched at a 30 GS/s $2\times$ ping-pong S&H fashion, as it would occur with an integrated ADC back-end. The last scenario is highly interesting as it allows evaluating the kickback during the sampling instants to the front-end output and its potential linearity degradation. For this simulation, the worst-case (for linearity) 11 dB attenuation setting is considered. The amplifier-buffer chain limits the IM3 in all cases, even with all the attenuator cells active (Fig. 7.10). The two fixed load scenarios show a comparable IM3 within 1.5 dB difference. This degrades smoothly as the frequency increases, remaining below




Fig. 7.16 Simulated two-tone IM3 vs. frequency of the complete front-end chain for three different load cases

–60 dB for the capacitive load and below –61 dB for the matched load at 30 GHz. For the switched load scenario, the IM3 degrades slightly more above 5 GHz, resulting in a value of about –57.5 dB at 30 GHz. This is due to the generated kickback into the front-end complex output impedance, producing non-linear current glitches and slightly increasing the distortion.

7.3 Experimental Verification

The prototype analog front-end is fabricated in a single-poly thirteen-metal (1P13M) 16 nm FinFET CMOS process. Figure 7.17 shows a die micrograph, which measures a pillar-limited total area of 1.7 mm². The core front-end area, including the LC filter with ESD, attenuator cells, control, and in-between decoupling grid, the hybrid amplifier, the series-shunt peaking inductors, and the buffer, is only 540 × 280 μm. Although this area is already comparable to a well-optimized sub-ADC of a TI RF sampling converter, it could be further reduced by removing the grid between the two sides of the filter and bringing them closer together. However, the effect of the mutual coupling between the differential side inductors on the filter performance has to be properly assessed. Further, since the filter does not limit the overall bandwidth, it may be reduced from ninth to seventh order by removing one tap, for a slight bandwidth degradation.

The preferred signal flow in this chip is horizontal, prioritizing the design of the filter inductors and minimizing their interaction with the surrounding blocks. The layout of each device and the arrangement of the blocks are carefully optimized to minimize parasitics, which is key to maximizing performance in an interconnect-intensive process such as 16 nm as FinFET. Further, differential symmetry is kept as much as possible across the entire front-end chain, with dummy structures and interconnect, replicating the same environment around the critical blocks. Finally, the flip-chip option with copper pillars is preferred over traditional wire bonding,




Fig. 7.17 Die micrograph of the 16 nm FinFET IC with front-end occupying a core area of about 0.15 mm^2

as it minimizes the interface and routing inductance of critical signals while also providing a better current return path by placing multiple ground pillars close to their corresponding signals.

7.3.1 Measurement Setup

The complete setup for the performance evaluation of the prototype analog front-end via different measurements is depicted in Fig. 7.18, along with photos. On the input side, the setup includes two Agilent E8257D low phase noise analog signal sources, which can be used either separately for the one-tone measurements or combined via a wideband power combiner to accommodate the two-tone measurements. For the modulated signal measurements, a R&S SMW200A vector signal generator with an up to 32 GHz modulated signal generation capability is employed. The spectral purity of the incoming signal is enhanced through appropriate band-pass filtering and converted into a differential signal by an ultra-wideband balun with small amplitude and phase mismatch. Finally, the differential signal is AC-coupled to the chip through wideband bias-Ts and identical phase-matched cables. At the output, the differential signal, through phase-matched cables and a second ultra-wideband balun, is converted back into single-ended and collected for spectral analysis. Two R&S FSW spectrum analyzers are used, one for the one-tone and two-tone measurements with up to 67 GHz frequency capability and another one with a




Fig. 7.18 Measurement setup of the 30 GHz bandwidth front-end prototype

dedicated vector signal analysis software for the modulated signal measurements. Small-signal measurements are also performed by means of a differential input—differential output 67 GHz R&S vector network analyzer (VNA).

The bare die is flip-chipped through copper pillars on a custom-designed LGA package with a low-loss high-frequency laminate material. Special attention is paid on the package, to ensure a minimum signal integrity degradation across the entire band of interest. The input and output traces are carefully designed as differential microstrips with the appropriate outer ground shielding. To guarantee a controlled characteristic impedance at more than twice the bandwidth of interest, they are extensively simulated with a commercial 3D EM-solver. The package is covered by an epoxy molding compound, which is etched away from the input and the output sides to allow the signals to be probed with high-frequency differential probes at the top side of the package. Prior to characterizing the prototype front-end performance, the complete measurement setup losses, including probes, baluns, cables, and connectors, are fully characterized with the VNA and de-embedded in the results hereafter, for proper signal swing adjustment and harmonic content.

The packaged chip is soldered onto a custom daughterboard. The required supply and bias voltages for the different blocks are generated with dedicated low-noise LDOs (both positive and negative) on the custom larger motherboard. These are transferred to the daughterboard through coaxial cables, and they are finally provided to the chip after further low-pass filtering with both off-chip components and on-chip decoupling capacitors. The control signals for the different attenuation settings are also coming from the motherboard. Finally, a dual-channel Keithley sourcemeter is employed to independently adjust the common-mode voltages of the complementary inputs, for debugging and characterization purposes.

7.3.2 Measurement Results

The key demonstration goals of this front-end prototype are a bandwidth in excess of 20–25 GHz, with beyond state-of-the-art spectral purity and low power consumption. As such, several measurements are performed, including both small-signal and large-signal (one-tone, two-tone, modulated) measurements. The measured small-signal performance across different attenuation settings and six samples are illustrated in Fig. 7.19. The gain demonstrates an ultra-wide bandwidth of about 30 GHz for all the settings while maintaining a very uniform step within ± 0.1 dB from its nominal value (Fig. 7.19a). Across six different samples (Fig. 7.19b), the performance is very consistent, with the gain showing a spread within ± 0.3 dB in amplitude and ± 0.5 GHz in bandwidth. The input return loss is about -10 dB up to 10 GHz and remains < -7 dB up to 20 GHz and < -6 dB up to 30 GHz (Fig. 7.19c). A possible explanation for this gradual degradation is the manual etching to remove the epoxy molding compound and expose the traces for probing, deteriorating somewhat the input matching. Finally, the measured group delay shows a variation within 8 ps for all settings up to 35 GHz, while between the different settings, the relative group delay variation remains within 3 ps for the entire band of interest (Fig. 7.19d).

Two measured one-tone output spectra, at 2.5 and 5 GHz, and HD2/HD3 vs. the input frequency are shown in (Fig. 7.20). The spectra include the setup losses; however, these are de-embedded from the HD2/HD3 measurements, so as to correspond to a swing of about $600 \text{ mV}_{\text{pp,diff}}$ at the buffer output. In both spectra, the 2nd and 3rd harmonics are the only evident tones, other than the fundamental. From Fig. 7.20c, the measured HD3 dominates SFDR across the entire bandwidth. The median across six samples remains below -67 dB up to 5 GHz and still better than -58 dB at 20 GHz. It is not possible to characterize HD3 close to 30 GHz because the 3rd harmonic for a 30 GHz fundamental falls out of the 67 GHz spectrum analyzer capability. HD2 is about 10 dB lower than HD3 across the entire bandwidth, and it is better than -65 dB up to 30 GHz.

The measured two-tone spectra at two carrier frequencies, 5 and 20 GHz, and a power of -6 dBFS for each tone are shown in Fig. 7.21, for 0 dB (best case, left) and 11 dB (worst case, right) attenuation settings. As expected, the front-end

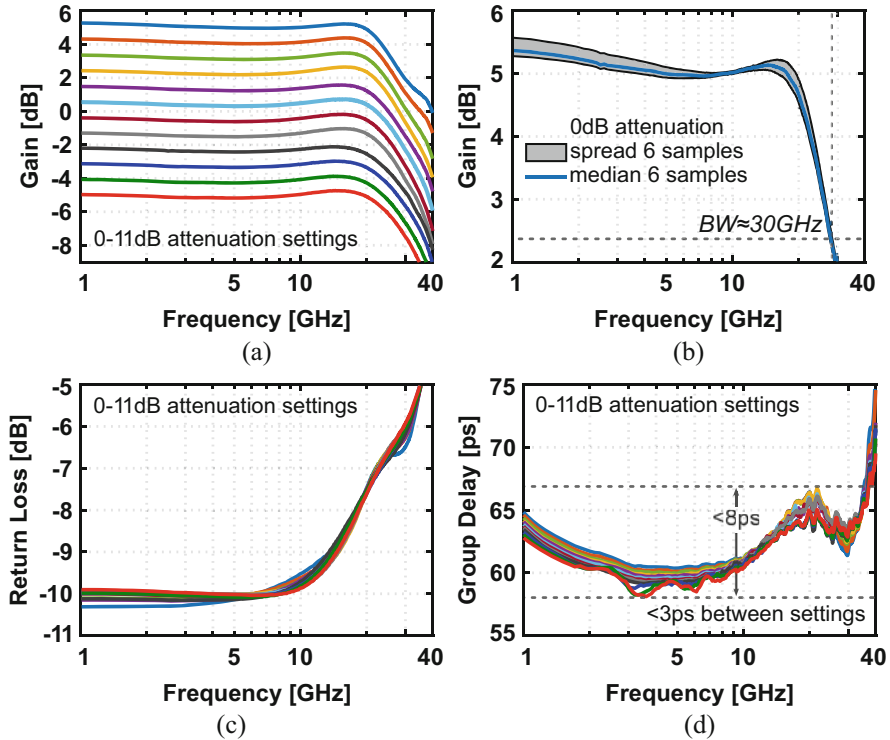


Fig. 7.19 Measured front-end small-signal performance for different attenuation settings and six samples

IM3 is primarily limited by the amplifier and buffer, which are progressively more dominant as the frequency increases, due to their faster rising distortion profile at higher frequencies compared to the filter. This was also indicated by the linearity simulations of Figs. 7.10 and 7.14. At the 11 dB setting, with all the attenuator cells active, the additional distortion contribution, with respect to the 0 dB setting, is 1.3 dB at 5 GHz and less than 0.5 dB at 20 GHz.

The measured IM3 vs. the carrier frequency of a 40 MHz spaced two-tone signal and vs. tone spacing at different carrier frequencies are plotted in Fig. 7.22. The six samples included exhibit very similar performance, with a spread within 4 dB across both carrier frequency and tone spacing. The median IM3 is below -70 dB up to 5 GHz and remains below -57 dB close to 30 GHz (Fig. 7.22a), with the degradation attributed mainly to the increasing dynamic distortion of the amplifier and buffer. When sweeping the tone spacing between 10 and 200 MHz, IM3 stays reasonably flat for a large range of carrier frequencies (Fig. 7.22b). A maximum variation of about 4 dB is noticed at a 20 GHz frequency. This might be attributed to the synchronization and combination accuracy of the two analog signal sources (Fig. 7.18) across such a wide frequency range.

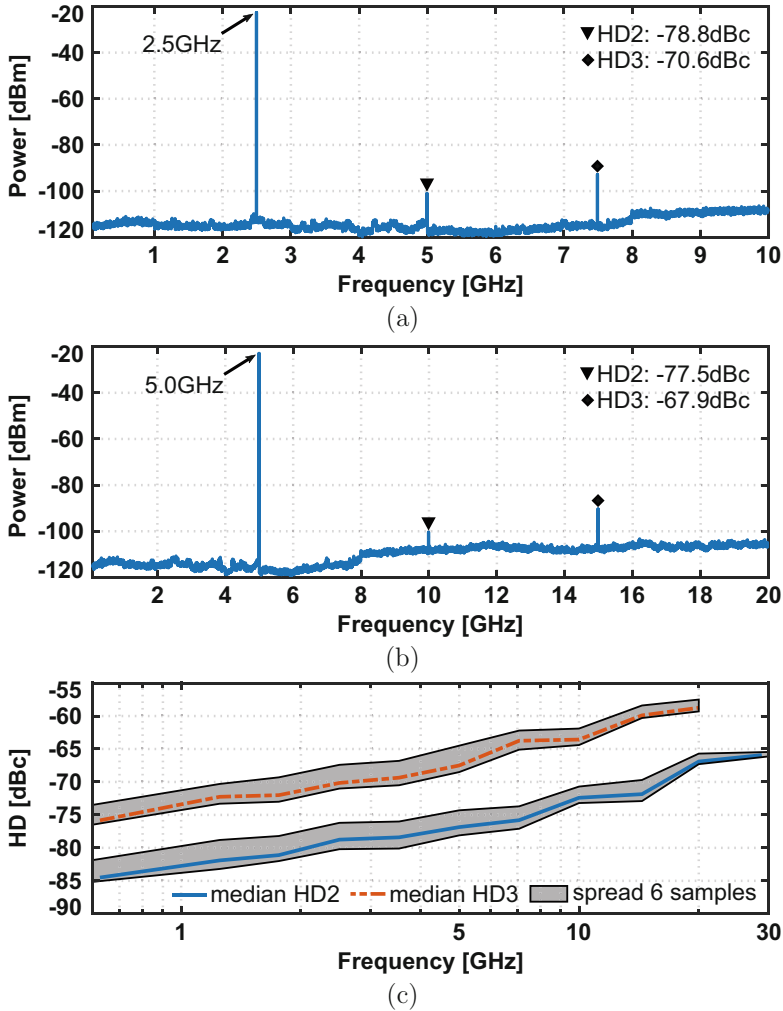


Fig. 7.20 Measured one-tone output spectra for (a) 2.5 GHz and (b) 5 GHz input frequencies and (c) measured HD2/HD3 vs. input frequency

Apart from the aforementioned one-tone and two-tone measurements, modulated signal measurements are also performed. Almost all wireless networks use digital modulation to transmit and receive signals, and the most common modulation in modern radios, including the 5G mm-Wave FR2, is Quadrature Amplitude Modulation (QAM). Very high data rates with a very high spectral efficiency are possible with QAM by setting a suitable constellation size. However, these are ultimately limited by the linearity and noise of the circuits in the transmitter and receiver chains. One key metric to evaluate is the Error Vector Magnitude (EVM), which is a noise-to-signal ratio that quantifies the distance between the




Fig. 7.21 Measured two-tone spectra at 5 and 20 GHz for 0 dB (left) and 11 dB (right) attenuation settings

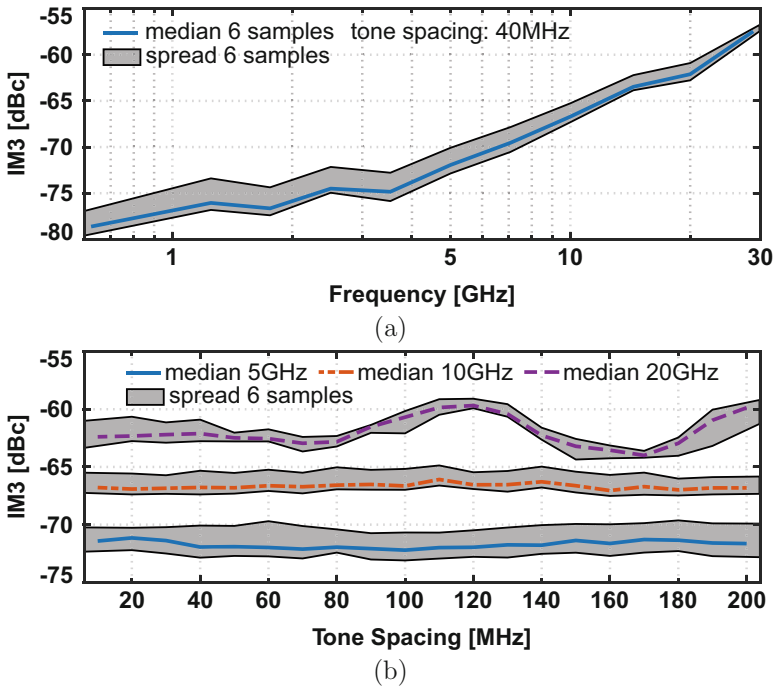


Fig. 7.22 Measured two-tone IM3 for six samples vs. (a) carrier frequency and (b) tone spacing

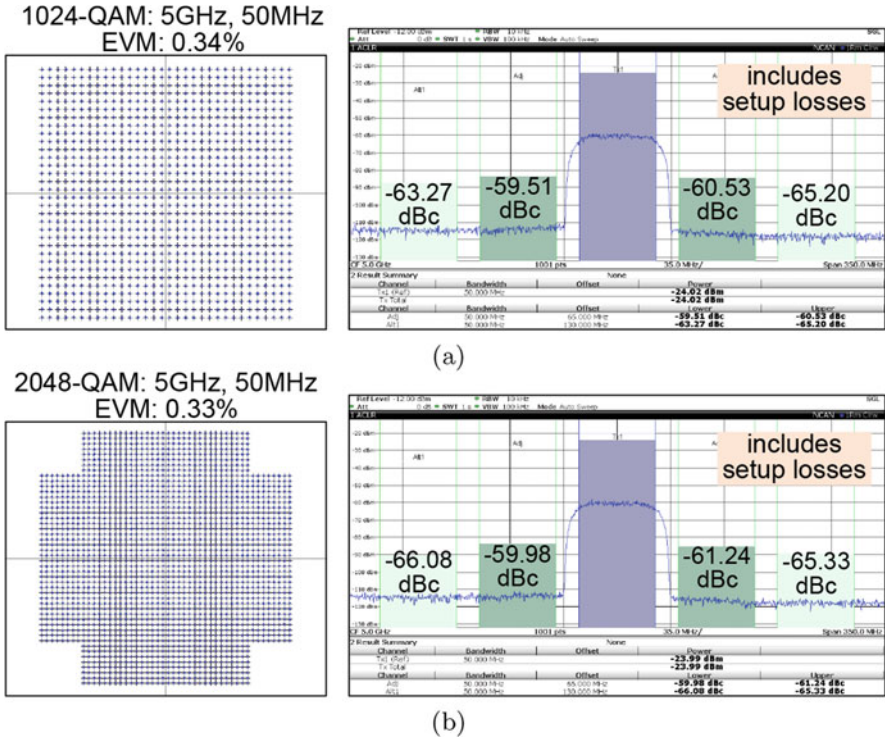


Fig. 7.23 Measured constellations and spectra at 5 GHz frequency for (a) 1024-QAM and (b) 2048-QAM signals

measured symbol points and the ideal reference points in the constellation. Another particularly important metric is the Adjacent Channel Leakage Ratio (ACLR), which quantifies the amount of leakage power into adjacent channels and is defined as the ratio of the mean power centered on a certain channel frequency to the mean power centered on an adjacent channel frequency.

The measured constellations and spectra of 5 GHz-carrier frequency 50 MHz-modulation bandwidth signals are shown in Fig. 7.23, both for 1024-QAM (Fig. 7.23a) and for 2048-QAM (Fig. 7.23b) schemes (>8.5 dB PAPR). Both modulation schemes demonstrate clean measured constellations with clearly distinct symbol points. The spectra for both schemes are free of any visible spectral regrowth or any out-of-the-ordinary intermodulation products. The measured EVM and ACLR vs. the carrier frequency across the entire 30 GHz bandwidth are plotted in Fig. 7.24. For both 1024- and 2048-QAM, EVM and ACLR demonstrate the beyond state of the art < -0.35 % and < -59.5 dB up to 5 GHz, which remain < -1.6 % and < -48 dB all the way up to 30 GHz. No calibration or additional distortion cancellation were used for any of the large-signal measurements of this front-end.




Fig. 7.24 Measured EVM and ACLR vs. frequency for 1024-QAM and 2048-QAM modulated signals

It is worth noting that the modulation bandwidth is limited by the demodulation capability of the spectrum analyzer.

7.3.3 State-of-the-Art Comparison

A performance summary of this work is given in Table 7.1, together with a recent SotA comparison against the best published most relevant ultra-wideband ADC-based receiver front-ends. The proposed front-end with the distributed variable attenuation LC filter and the push-pull hybrid CG—CS amplifier significantly advances the state of the art in terms of bandwidth and linearity, with comparable or lower power, area, and noise, while providing additional integrated functionality of variable attenuation and gain, compared to the relevant prior art. Relative to the next best alternative [74], on top of the extra functionality, this front-end attains a larger than $1.6\times$ net bandwidth, with at least $3\times$ larger bandwidth for similar SFDR/IM3 and at least 8 dB better SFDR/IM3 at similar frequencies, with similar noise, power, and area. Finally, supporting up to 2048-QAM modulation with excellent spectral purity, this work is the first to demonstrate the SoC integration in deep-scaled CMOS of the front-end required to enable direct RF sampling up to mm-Wave frequencies.

Table 7.1 Performance summary and comparison with state-of-the-art ADC-based receiver front-ends

	This work VLSI'22	Straayer [147] ISSCC'16 ^{a,b}		Devarajan [104] ISSCC'17 ^{a,b}		Ali [74] ISSCC'20 ^{a,b}	
Front-end topology	Variable attenuation and hybrid	Dual push-pull	Push-pull	Push-pull	Push-pull	T&H push-pull	
	Push-pull amplifier-buffer	Buffer	Buffer	Buffer	Buffer	Buffer	
Technology	16 nm FinFET	65 nm CMOS	28 nm CMOS	28 nm CMOS	16 nm FinFET		
Area [mm ²]	0.151	0.360	0.610	0.610	0.148		
Power [mW]	210.0	207.0	400.0	400.0	220.0		
Bandwidth [GHz]	30.0	4.0	7.4	7.4	18.0		
Input frequency [GHz]	2.5	5.0	10.0	20.0	29.0	4.0	8.0
One-tone SFDR [dBc] ^d	70.6	67.9	64.3	58.7	66.0 ^e	66.0	61.0
Two-tone IM3 [dBc]	-74.5	-72.1	-67.1	-61.9	-57.2	-72.0	N.A.
NSD [dBFS/Hz]	-155.7	-155.7	-155.5	-155.1	-155.0	-154.0	-157.0
SNR [dB] ^e	53.9	53.9	53.7	53.3	53.2	56.1	N.A.
SNDR [dB]	53.8	53.7	53.3	52.1	N.A.	55.5	53.0
Modulation scheme	1024-QAM	2048-QAM		N.A.	N.A.	N.A.	N.A.
EVM [%] ^f	0.34/1.58	0.33/1.57		N.A.	N.A.	N.A.	N.A.
ACLR [dBc] ^f	-59.5/-48.1	-60.0/-48.3		N.A.	N.A.	N.A.	N.A.

^aData as reported^bFront-end dominates performance^cOnly HD2, 3rd-harm. is above the 67 GHz analyzer range^dHD2 or HD3 limits SFDR^eIntegration over 0.5-bandwidth^fFor 5 GHz/29 GHz frequencies

7.4 Conclusion

This chapter addressed the analog front-end challenges in pushing the sample rate and bandwidth of RF ADC-based receivers to several tens of GHz while delivering high spectral purity with low power. These challenges stem from the large ADC input load the front-end has to deal with as well as the constant pursuit for higher functionality and integration in deep-scaled CMOS processes. After introducing the problem and overviewing some noteworthy prior art, a buffered front-end model was developed, including on-chip/interface/off-chip contributions altogether, to investigate the bandwidth limits and identify potential points for improvement. The ESD and buffer input capacitances were found among the dominant bandwidth-limiting factors. Existing solutions to deal with these capacitances, such as T -coils and distributed inductive peaking, were briefly discussed, highlighting their advantages and drawbacks.

The proposed front-end was presented, along with the introduced solutions in this work to extend the bandwidth while ensuring a higher integration. The introduced solutions start with an impedance-matched ninth-order Chebyshev filter that distributes the input ESD while providing a 0–11 dB variable attenuation by means of distributed switched attenuator cells. The filter is able to absorb the parasitic capacitances of the pad, the ESD, the attenuation, the input termination, and the following amplifier, significantly extending the bandwidth. The design, arrangement, and type of the attenuator cells were motivated, and simulation results were provided, justifying the motivations. The filter is followed by a new two-path hybrid CG-CS 6 dB-gain amplifier and a CD buffer stage, adopting push-pull topologies. These allow for a significant gain-bandwidth improvement at no extra noise or power consumption, while their linearity across the entire band is improved by resistive degeneration and bootstrapped cascoding, respectively. The design choices on the amplifier were also detailed and backed up by simulations.

The prototype analog front-end with the proposed innovations, fabricated in a 16 nm FinFET CMOS process, demonstrates a 30 GHz bandwidth with a better than 58 dB-SFDR and better than -57 dB-IM3 across its entire bandwidth and settings. It also supports 1024-/2048-QAM modulation with beyond state-of-the-art spectral purity while occupying a compact core area of $540 \times 280 \mu\text{m}$ and drawing a total current of 52.5 mA from a dual rail ± 2 V supply. This work is the first to prove the viability of power-efficient highly linear ADC-based receiver analog front-ends in deep-scaled CMOS, to enable direct RF sampling up to mm-Wave frequencies.

Chapter 8

Conclusions, Contributions, and Future Work



Data converters are in the catbird's seat during the ever-growing digitization trend of the last decades, as they are the heart of any application/device that exchanges information between the real analog world and the digital signal processor. With this privileged position within modern electronics comes a set of distinctive challenges, since attaining the required specifications is paramount for the target application's correct functionality. Furthermore, their roles as interfaces of analog and digital signals mean that they must deal with the imperfections of the former and keep up with the advancements of the latter.

The overarching goal of this book was to devise innovations toward maximal A/D conversion accuracy and power efficiency at the multi-GHz sample rate and bandwidth regime. The solutions proposed to this end yielded beyond state-of-the-art performance at the circuit, architectural, system, and technology levels, proving that modern ultrahigh-speed ultra-wide-bandwidth ADCs demand innovations on all fronts. Most importantly, the work in this book has shown that holistically fulfilling the proposed innovations is paramount to shattering the rigid performance boundaries.

This closing chapter recapitulates the steps taken to fulfill this goal. Section 8.1 first provides an overview of the developed analyses and the proposed architectural and circuit techniques of this work. Section 8.2 then describes some of the major contributions brought by the work in this book and their advancing the field of high-speed low-power ADCs. Finally, suggestions for future research directions are listed in Sect. 8.3.

8.1 Overview and General Conclusions

Chapter 1 discussed the fundamental role and applicability of data converters, specifically ADCs, in an increasingly digital dominated era. The three main

performance parameters, accuracy, speed, and power, were introduced, and the multiple level challenges, circuit, architecture, system, and technology, to achieve an optimum accuracy-speed-power set were reviewed. These discussions guided the definition of this book's research goal and objectives: to propose scalable friendly architectural and circuit solutions to address the challenges at all levels and maximize the $accuracy \cdot speed \div power$ of next-generation multi-GHz sample rate and bandwidth ADCs in deep-scaled CMOS.

Chapter 2 covered fundamental A/D conversion principles, performance metrics, and limitations. After reviewing the basic functions of sampling and quantization, the major error sources from the circuit blocks in a practical converter chain were identified and analyzed. A solid understanding of these error sources leads to the establishment of fundamental limits in terms of accuracy-speed-power on a converter's performance. The derived limits, imposed by the circuits and ultimately by physics, built an insight as to what may be theoretically achievable from the elementary building blocks in a converter and what has to be traded off to maximize the $accuracy \cdot speed \div power$ ratio.

The fundamental limits' analysis was extended to the architectural level in Chap. 3. By studying state-of-the-art high-performance architectures, such as flash, SAR, pipeline, and pipelined-SAR, models were derived to estimate and compare their accuracy-speed-power limits. In addition to offering a complete decomposition of the individual blocks' contributions, the models were further enhanced by including the process effects of four deep-scaled CMOS nodes. It was found that for a low-to-medium resolution, the simplicity of the SAR is hard to beat even approaching GHz sample rates. With the comparator its single analog block, it makes a very good candidate at this resolution range, either standalone or as a base for a larger system integration to enhance speed (through time-interleaving), resolution (through pipelining), or both. For a medium-to-high resolution, the pipelined-SAR hybrid with more than two stages showed high promise, being able to compete with the traditional pipeline even at GHz sample rates for the same or similar (or even lower) stage count. This was a natural consequence of the low-resolution sub-SAR stages within the pipelined-SAR being more efficient than the sub-flash in the traditional pipeline.

This chapter also discussed time-interleaving as the most popular and potentially the only way to boost the sample rate beyond a standalone converter with acceptable efficiency. Key aspects, such as the interleaving errors due to offset, gain, timing, and bandwidth mismatches between the channels, were detailed and modeled to quantitatively compare their accuracy degradation impact. Errors due to bandwidth mismatch were found to be among the toughest to cope with as the input frequency over the analog bandwidth increases, hinting to the importance of guaranteeing a wide bandwidth by design. Finally, the interleaver architecture was detailed, as it determines to a great extent the converter's analog bandwidth and sampling accuracy, as well as presents a considerable power/area overhead. The main interleaver architectures, namely, direct, demux, and resamp, were discussed, modeled, and compared in terms of achievable bandwidth and sampling accuracy providing insight in determining the optimum interleaver depending on the design.

Direct interleavers were found to achieve the best results for a channel count ≤ 8 , while for ≥ 16 channels, resamp were found superior to demux in terms of accuracy and bandwidth. Further insight was gained by comparing the interleavers across the different process nodes. Both the standalone ADC and interleaving analyses constituted the basis of the design choices for the prototypes of the following chapters.

Chapter 4 focused on the comparator block, whose speed, sensitivity, and power consumption significantly impact the sample rate, accuracy, and efficiency of the ADC. The choice for dynamic regenerative comparators was motivated by the pursuit of maximizing speed and minimizing power in deep-scaled CMOS, as per the theoretical analyses in the previous chapters. The delay of two most widely adopted topologies, the strong-ARM and the double-tail, was analyzed, and their drawbacks in minimizing the delay and its variations pin-pointed. The insight gained led to the proposal of a dynamic comparator, able to minimize its delay and its variations through the combination of a high-gain three-stage configuration and an extra parallel feed-forward path. Additionally, the cascaded triple-latch arrangement with reduced device stacking significantly reduced the delay across a wide common-mode and supply voltage range. The prototype comparator was fabricated in 28 nm CMOS, and its performance was verified through extensive measurements. When compared to the state of the art, it reported the highest data rate as well as the smallest delay slope and variation with similar input-referred noise and competitive energy/comparison.

Chapter 5 moved toward extending the sample rate and the bandwidth of low-to-medium-resolution single-channel SAR ADCs in the GHz range. While doing so, it was of utmost importance to minimize their accuracy degradation across the entire band and not compromise their digital nature, excellent efficiency, and simplicity, such that they could be both easily used as standalone blocks and integrated into larger systems with minimum overhead and complexity. The main speed-limiting factors in the conventional SAR loop were identified as the comparator evaluation time, the DAC settling time, and the digital logic delay, which have to be accommodated within each cycle. On top, the input sampling switch was identified to directly impact the achievable bandwidth and high-frequency sampling linearity. After reviewing noteworthy prior art speed-boosting techniques, the prototype single-channel SAR ADC was presented, with proposed techniques to tackle all the aforementioned in a single shot. On the architectural level, a semi-asynchronous processing was introduced with a dynamically allocated internal timing, eliminating the logic delay from the critical path by overlapping it with the comparator evaluation. On the circuit level, a dual-loop bootstrapped input switch was proposed to improve the input bandwidth and high-frequency linearity. A USPC CDAC topology and a triple-tail dynamic comparator were also proposed, to reduce the settling and evaluation times, respectively. The prototype converter was fabricated in a 28 nm CMOS process, and its performance was verified through several measurements across multiple samples. Compared to the state of the art, it showed among the highest sample rates and the lowest accuracy degradation from the designed aggregate resolution across the whole band while attaining a very wide bandwidth with a power dissipation, area, and efficiency on par or better than the

state of the art. This power dissipation is about $3\times$ higher than the predicted one from our models with the same sample rate and effective resolution. The difference stems partly from the increased logic power compared to the models, due to the higher number of gates in an actual design. Further, some deviations are also to be expected, given that our model captures mainly first-order effects. However, it can still be within one order of magnitude with actual designs, provided that the assumptions are adjusted accordingly to the specific design under comparison.

Chapter 6 delved deeper into investigating circuit, architectural, and system capabilities to enable a higher ADC resolution (>10 bits), while sampling directly at RF frequencies with multi-GHz sample rate and bandwidth, and maximum efficiency. To achieve the required sample rate, time-interleaving would be the way to go. Upon prior art review, two prevailing design strategies in terms of sub-ADC and interleaving factor were identified: (1) interleave ≤ 8 , faster less efficient pipelines or (2) massively interleave ≥ 32 slower highly efficient SARs. The former was found to offer the benefits of easier drivability, signal distribution, and relaxed calibration overhead for a less efficient sub-ADC. The latter traded the superior efficiency of the sub-ADC with an increased front-end loading and a more complex signal distribution and calibration. The widely used front-end buffer to ensure a wide bandwidth was found the primary performance and efficiency bottleneck of such ADCs, especially at the highest frequencies. Combining these observations with the architectural and interleaving analyses of Chap. 3, the prototype passive-sampling $8\times$ -interleaved hybrid RF ADC was presented, and the proposed techniques to improve on the prior art were discussed in detail. This prototype utilized the insight from the theoretical analyses to collectively address challenges at all the different levels. A fully dynamic three-stage pipelined-SAR sub-ADC was employed that maximized the efficiency for the given resolution, sample rate, and technology. A wide input bandwidth with high spectral purity and efficiency in the absence of a front-end buffer was achieved altogether with an optimized input network. Sampling purity was ensured by an on-chip clock conditioning/distribution chain with negligible additive jitter. On top, a combined custom analog-synthesized digital calibration improved the spectral performance over the entire band of interest. The prototype converter, fabricated in 28 nm CMOS and characterized with measurements, demonstrated a comparable or larger input bandwidth than existing works with an input buffer and at least $6\times$ higher compared to the next best buffer-less approach. While preserving similar accuracy and spectral purity levels, it also showed at least $2\times$ smaller power than buffered works and about $1.4\times$ smaller power than the next best buffer-less work, demonstrating the validity of the analyses and proposed system, architectural, and circuit solutions in this book. When comparing the power dissipation of one sub-ADC with a three-stage integrator-based pipelined-SAR with the same sample rate and effective resolution (including its sub-calibration estimation), the estimated power from the models is about $2\times$ – $4\times$ lower. This is also to be expected due to the underestimated logic and calibration power. Further deviations are caused by the fact that the f_T (or g_m/I_D) is not the same for every block in the sub-ADC, which is assumed by the models for simplicity. However, it is also within one order of magnitude, verifying the usefulness of the proposed models.

Finally, Chap. 7 shifted the focus toward addressing the front-end challenges in pushing the sample rate and bandwidth of direct RF sampling ADC-based receivers to multiple tens of GHz while delivering high spectral purity with low power consumption. These challenges stemmed from the large ADC input load and the constant pursuit for higher integration in deep-scaled CMOS. After revisiting the bandwidth problem and discussing its root causes, an enhanced front-end chain model was built to better understand the impact of on-chip/interface/off-chip contributions and identify places for improvement in the chain. The ESD capacitance was found among the dominant bandwidth-limiting factors, followed by the buffer input capacitance. Integrating more functionality, such as wideband variable gain/attenuation, only makes matters worse in terms of bandwidth degradation. To solve these challenges on a system, architectural, and circuit level, the proposed prototype front-end introduced an on-chip impedance-matched multi-segment LC Chebyshev filter with a two-segment split ESD, and a variable stepped attenuation, by distributing attenuator cells across the filter taps. The capacitances of the pad, ESD, attenuators, and termination were almost entirely absorbed by filter. The design, arrangement, and type of the attenuator cells to achieve the largest bandwidth and spectral purity were motivated, and simulation results were provided, validating the motivations. The filter was followed by a new two-path hybrid CG-CS 6dB-gain amplifier and a CD buffer stage, adopting push-pull topologies. These allowed for a significant gain-bandwidth improvement at no extra noise or power consumption, while their linearity across the entire band was enhanced by resistive degeneration and bootstrapped cascoding, respectively. The prototype analog front-end was fabricated in a 16 nm FinFET process, and its performance was verified through different measurements and across multiple samples. The proposed front-end significantly advanced the state of the art in terms of bandwidth and linearity, with comparable or lower power, area, and noise, while providing additional integrated functionality of variable attenuation and gain, compared to the relevant prior art. It demonstrated a larger than $1.6\times$ net bandwidth, with at least $3\times$ larger bandwidth for similar SFDR/IM3 and at least 8 dB better SFDR/IM3 at similar frequencies, with similar noise, power, and area, compared to any prior work.

8.2 Original Scientific Contributions

Despite the field of high-speed ADCs being immense and well established, the work developed in this book brought some new insights in advancing the field with the key contributions listed below:

Architectural Limits' Models with Block Decomposition and Technology Effects from Four Deep-Scaled CMOS Processes

The study of the internal operation of state-of-the-art ADC architectures led to the proposal of mathematical models to estimate and compare their accuracy-speed-power limits across four different process nodes with enhanced process effects.

These models offered a unique insight on the architecture comparison across different resolutions and sample rates and motivated design choices in the implementation prototypes. They can also serve as a “cookbook” when starting a new design and/or moving to a new process node. Architecture coverage includes flash, SAR, pipeline (1,2,3,4-bit/stage, extendable), and pipelined-SAR (2,3,4,5-stage, extendable), while block coverage includes sampler, comparator, open-loop residue-amplifier, resistor ladder, DAC, and first-order digital logic. Special features include redundancy allocation when pipelining, supply utilization, settling allocation, target BER, DAC reference overhead, and residue amplifier linearity overhead. It should be noted that in literature there can be found similar attempts to estimate ADC power consumption bounds [169], but they are limited in terms of architectures and architecture variants and included block coverage. This is the first proposed framework to collectively capture architectures (and variants) and contributions that have not been demonstrated previously in literature, in addition to including technology parasitic effects across different deep-scaled CMOS process nodes.

Interleavers’ Models to Estimate Bandwidth and Accuracy Across Four Deep-Scaled CMOS Processes

The interleaver being one of the most important design considerations in a TI ADC, it was important to quantitatively analyze and compare different architectures in making the optimum choice for given specifications. This led to the proposal of models to compare the main interleaver architectures, direct, demux, and resamp, in terms of achievable bandwidth and sampling accuracy, providing insight in their trade-offs. These were compared across different number of channels, sample rates, and interleaver variants. Also, similar to the ADC architectural models, they were extended to capture process effects from four different deep-scaled CMOS nodes. They greatly motivated the choice for a direct interleaver in the TI RF ADC prototype, as being able to achieve the required bandwidth and sampling accuracy with small complexity given the design targets.

A Three-Stage Triple-Latch Feed-Forward Comparator

To improve the delay across a large input range, reduce device stacking for lower supply operation, remove series turn on of the latch devices, and create a higher signal gain prior to latching, a three-stage triple-latch comparator topology with a reduced stacking and parallel direct/feed-forward paths was introduced. The multi-stage nature with cascaded latches enabled a very high total gain prior to the final latching. The concurrent turn on of the latch devices with a large overdrive voltage increased the effective regeneration rate by increasing. The horizontal cascading, instead of the vertical latch stacking, further reduced the required headroom compared to the two currently widely adopted topologies, allowing a favorable operation at lower supply voltages. These merits were demonstrated with measurements, leading to the highest reported data rate, with the smallest delay variations across supply and common mode for similar noise and competitive power with the state of the art.

A Semi-asynchronous Processing with Comparator-DAC-Logic Delay Overlapping

A new semi-asynchronous processing was introduced to combine the merits of simple logic and cycle control from synchronous processing with the dynamically allocated internal timing of asynchronous processing. Within each fixed cycle, the time is dynamically shared between the comparator and the DAC while triggering the logic in parallel to the comparator and majorly hiding its time from the critical loop. Considering that the logic may occupy as much as 30–40% of an internal cycle delay of a GHz sample rate SAR, this led to a considerable sample rate improvement and contributed to the prototype single-channel SAR achieving among the highest sample rates at the time of publication.

A Dual-Loop Bootstrapped Input Switch

The linearity of the input switch was identified to directly impact and even dominate the total converter spectral purity with increasing gravity at GHz sample rates and bandwidths. Bootstrapping is the popular way to improve the sampling switch linearity, but existing bootstrap circuits have an slow internal loop, rendering them largely ineffective. An improved bootstrap circuit was proposed with the introduction of a separate loop to control the critical devices and significantly speed up the boosting mechanism, leading to a considerably higher and relatively constant sampling linearity over the entire band of interest. The validity of the proposed circuit was demonstrated with measurements employed in both prototype ADCs, showcasing its effectiveness for both medium and high resolutions.

A Unit-Switch-Plus-Cap DAC

The DAC settling time was identified as one of the main speed-limiting factors within the SAR loop. Except for the “clean” switch on-resistance and capacitance, the parasitic resistance and capacitance due to interconnect were found to significantly increase this time and even dominate for very small unit capacitors. To minimize the effect of these parasitics, the reference switches were merged with the unit capacitors, introducing the Unit-Switch-Plus-Cap DAC topology, which minimized the interconnect between them, in contrast to the conventional Unit-Cap topology. The benefits of this topology were also demonstrated with measurements by employing it in both prototype ADCs and significantly contributed to their favorable state-of-the-art standings.

A Symmetrical Intertwisted Input/Clock Y-Tree Structure

The routing of the input in a TI ADC and its necessary ground shielding to prevent coupling with the clock were identified to be a significant factor to the front-end loading. To minimize the routing parasitic contribution, a symmetrical differential intertwined input/clock Y-tree was introduced. The input and clock were intentionally routed side by side (but in different metals) in a tree-like structure, with the clock being intertwined around the input in every turn of the tree. The intertwining cancelled out the mutual differential input/clock coupling, eliminating the ground shielding and reducing the loading by more than 2×. It contributed a

big part to the drivability and large input bandwidth of the buffer-less TI RF ADC prototype, comparable or larger than existing works with an input buffer.

A Multi-Stage (>2) Pipelined-SAR Hybrid Architecture

The multi-stage (>2) pipelined-SAR hybrid was proposed, whose potential and promise were predicted by the introduced architectural limits' models. From a theoretical standpoint, analogous to the regular pipeline with flash sub-ADCs, there did not seem to be a fundamental reason preventing a higher order than two-stage pipelined-SAR, which is the vast majority found in literature. On the contrary, it seemed the way forward to achieve a higher absolute sample rate while increasing the resolution and competing with or surpassing the efficiency of the regular pipeline at GHz range, where it is currently the architecture of choice. This was motivated by the fact that a low-resolution SAR is more efficient than a low-resolution flash, which was also proven with the models. Also, the SAR has been shown to progressively improve its sample rate with technology scaling at a faster pace than the flash, considering also secondary effects (e.g., increased BEOL at finer nodes). This proposed hybrid architectural extension showed to already advance the state of the art with the three-stage pipelined-SAR converter implemented as a sub-ADC of the TI RF ADC prototype.

A Two-Path Hybrid Amplifier with Improved Gain-Bandwidth and Noise

The amplifier was determined to be among the most critical blocks in an ultra-wideband front-end, dominating the achievable bandwidth, linearity, noise, and power. To overcome the limitations of traditional power-hungry amplifiers, a new hybrid amplifier was introduced, which uniquely combined two parallel paths, a common gate and a common source, by connecting one input to the one path and the complementary input to the other path, so as to both process and amplify the signal by adding their contributions. However, being stacked on a single current branch, the two paths did not add any extra current. This novel hybridism enabled a superior gain-bandwidth and noise compared to any traditional single amplifier with the same current. On top, due to the additional gain-bandwidth, the proposed hybrid amplifier allowed a size reduction in its devices, resulting in an improved dynamic linearity due to a reduced parasitic capacitance. The benefits of this new amplifier were demonstrated with measurements in the proposed analog front-end, which significantly advanced the state of the art in terms of bandwidth and spectral purity, with comparable or lower power, and area.

8.3 Suggestions for Future Work

The work in this book demonstrated the benefits of several deep-scaled CMOS-compatible solutions in the circuit, architectural, and system levels to improve *accuracy · speed ÷ power* and advance the state of the art in multi-GHz sample rate and bandwidth ADCs. Yet, there are unexplored and/or new directions toward

understanding and advancing the field further, with some important and promising ones listed below:

Additional Architectures in the Fundamental Limits' Models

The introduced architectural limits' models were proven very useful in opting for the best architectural choice for a certain resolution, sample rate, process, and underlying assumptions. However, they are limited to four architectures and some variants of theirs. Including additional architectures, such as $\Sigma\Delta$ converters and recently emerging time-based converters, will enhance the usefulness and applicability of these models. This will, in turn, lead to a better understanding in making the optimum choice, given the design space, rather than blindly following the most recent trend.

Higher-Order Pipelined-SAR (>3, 4, 5)

The introduced models already hinted to the fact that a higher-order pipelined-SAR should be able to achieve a higher absolute sample rate due to further reduction of the sequential cycles in each sub-SAR. The energy efficiency of a higher-order pipelined-SAR should also increase when opting for a high-resolution and a GHz-range sample rate due to reduced parasitic loading. This efficiency might potentially surpass the one of the regular pipeline provided that the sub-SAR is more efficient than the sub-flash for the chosen per stage sample rate and resolution. The validity of these predictions was already demonstrated with the prototype three-stage converter of Chap. 6. However, since these models do not include all the potential second-order and practical effects, it is unclear to what extent the actual trends moving forward will match the predicted ones. A silicon realization of such a higher-order pipelined-SAR at GHz-range operation could provide some more clear answers.

More Hybridization Within the Pipeline

Similar to the case with the flash and the SAR, any architecture can be theoretically pipelined to enhance its speed and/or its resolution. Also, there is no fundamental reason for the sub-stages within a pipeline to be of the same architecture. Depending on the design requirements, different architectures may fit different pipeline stages better, leading to an ADC with a better overall efficiency. One good example of such a hybridization is [170] that combined a coarse SAR and a fine digital-slope ADC for a 12-bit aggregate resolution with superior efficiency compared to traditional pipelines and pipelined-SARs. However, due to the slow nature of the digital-slope (step-at-a-time), the sample rate was limited to 100 MS/s. From a research point of view, it will be of great interest to investigate the efficiency benefits (if any) of such enhanced hybrids at GHz-range sample rates.

Clock Power Estimation in the Optimum Interleavers' Models

The introduced interleavers' models provided a lot of insight in determining the most suitable architecture or combination of architectures to achieve the best *accuracy · speed* for a given set of specifications. Yet, these models are missing one essential contribution in every TI-ADC, the power of the clock generation and

distribution circuitry and interconnect. This power is a complex function of the total number of channels as well as the hierarchical split and the channels in each rank. The reason is that clock edges at different nodes can have different requirements in terms of jitter and/or mismatch, making a correct power estimation less trivial than just counting number of gates driving a capacitor. Enhancing the interleavers' models with such an inclusion will tremendously increase their applicability and result in better optimized ADCs.

Jitter-Tolerant Converters

In the prototype ADC of Chap. 6, it was stressed that the high quality especially of the sampling pulses with respect to jitter was key to achieving the desired SNR at high input frequencies. To minimize the on-chip jitter contribution, a synchronous clock conditioning/distribution chain with re-timing was implemented. Although this clock chain minimized the on-chip jitter, that of the external generator ended up dominating the high-frequency SNR. Extending the bandwidth further is expected to only make matters worse. A CMOS integrated clock generation solution cannot compete currently in phase noise with external crystal oscillators, and those cannot easily improve by tenfold without an enormous amount of power [171]. One interesting approach could be to investigate techniques that make the converter more jitter tolerant, similar to some continuous-time ADCs [45]. Yet, something like this has not been shown in silicon, at GHz-range bandwidths.

Bibliography

1. G. Moore, Cramming more components onto integrated circuits. *Electron. Mag.* **38**(8), 114–117 (1965)
2. G. Moore, No exponential is forever: but "forever" can be delayed! in *2003 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2003), pp. 20–23
3. G. Manganaro, *Advanced Data Converters* (Cambridge University Press, Cambridge, 2011)
4. M. Mota, [Online Available], *Very High-Speed Data Converters for 5G Analog Front-End, White Paper* (Synopsis Incorporated, 2020)
5. S. Condra, K. Mäki, A. Purmonen, [Online Available], *Extended Spectrum Docsis: A Pragmatic Approach, White Paper* (Teleste Corporation, 2020)
6. T. Neu, *Direct RF Conversion: From Vision to Reality* (Texas Instruments Incorporated, 2015).
7. J. Mitola, The software radio architecture. *IEEE Commun. Mag.* **33**(5), 26–38 (1995)
8. S. Palermo, S. Hoyos, S. Cai, S. Kiran, Y. Zhu, Analog-to-digital converter-based serial links: an overview. *IEEE Solid-State Circuits Mag.* **10**(3), 35–47 (2018)
9. J. Im, K. Zheng, A. Chou, L. Zhou, J.W. Kim, S. Chen, Y. Wang, H. Hung, K. Tan, W. Lin et al., A 112Gb/s PAM-4 long-reach wireline transceiver using a 36-way time-interleaved SAR-ADC and inverter-based RX analog front-end in 7 nm FinFET, in *2020 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2020), pp. 116–118
10. W. Sansen, Analog CMOS from 5 micrometer to 5 nanometer, in *2015 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2015), pp. 1–6
11. M. Babaie, Power Efficient RF/mm-wave Oscillators and Power Amplifiers for Wireless Applications—Ph.D. Dissertation. TU Delft, Delft, NL, 2016
12. M.J.M. Pelgrom, *Analog-to-Digital Conversion*, 3rd edn. (Springer, New York, 2017)
13. H. Nyquist, Certain factors affecting telegraph speed. *Trans. Am. Inst. Electr. Eng.* **43**, 412–422 (1924)
14. H. Nyquist, Certain topics in telegraph transmission theory. *Trans. Am. Inst. Electr. Eng.* **47**(2), 617–644 (1928)
15. R.J. Van de Plassche, *CMOS Integrated Analog-to-Digital and Digital-to-Analog Converters* (Springer, Berlin, 2013)
16. J. Max, Quantizing for minimum distortion. *IRE Trans. Inform. Theory* **6**(1), 7–12 (1960)
17. W.R. Bennett, Spectra of quantized signals. *Bell Syst. Tech. J.* **27**(3), 446–472 (1948)
18. W. Kester, Analog Devices Technical Staff, *The Data Conversion Handbook* (Newnes, 2005)

19. M.S.O. Alink, A.B. Kokkeler, E.A. Klumperink, K.C. Rovers, G.J. Smit, B. Nauta, Spurious-free dynamic range of a uniform quantizer. *IEEE Trans. Circuits Syst. II: Exp. Briefs* **56**(6), 434–438 (2009)
20. R. Schreier, G.C. Temes et al., *Understanding Delta-Sigma Data Converters*, vol. 74 (IEEE Press Piscataway, 2005)
21. B. Razavi, *Design of Analog CMOS Integrated Circuits*, 2nd edn. (McGraw-Hill Education, 2017)
22. T. Sepke, P. Holloway, C.G. Sodini, H.-S. Lee, Noise analysis for comparator-based circuits. *IEEE Trans. Circuits Syst. I: Regul. Papers* **56**(3), 541–553 (2008)
23. I. Opris, Noise estimation in strobed comparators. *Electron. Lett.* **33**(15), 1273–1274 (1997)
24. A. Roy, C. Enz, Compact modeling of thermal noise in the MOS transistor. *IEEE Trans. Electron Devices* **52**(4), 611–614 (2005)
25. M. Van Elzакker, E. van Tuijl, P. Geraedts, D. Schinkel, E.A. Klumperink, B. Nauta, A 10-bit charge-redistribution ADC consuming $1.9 \mu\text{W}$ at 1MS/s . *IEEE J. Solid-State Circuits* **45**(5), 1007–1015 (2010)
26. B. Murmann, *EE315B: VLSI Data Conversion Circuits—Lecture Notes* (Stanford University, Stanford, 2020)
27. W. Kester, The Importance of Data Converter Static Specifications—Don't Lose Sight of the Basics. Analog Devices, Tutorial MT-010, 2009
28. M.J.M. Pelgrom, A.C. Duinmaijer, A.P. Welbers, Matching properties of MOS transistors. *IEEE J. Solid-State Circuits* **24**(5), 1433–1439 (1989)
29. S. Smith, *Digital Signal Processing: A Practical Guide for Engineers and Scientists* (Elsevier, Amsterdam, 2013)
30. W. Kester, Understand SINAD, ENOB, SNR, THD, THD + N, and SFDR so You Don't Get Lost in the Noise Floor. Analog Devices, Tutorial MT-003, 2009
31. N.I. Staff, [Online available], OFDM and multi-channel communication systems, white paper (2011), pp. 1–11
32. W. Kester, Op Amp Distortion: of HD, THD, THD + N, IMD, SFDR, MTPR. Analog Devices, Tutorial MT-053, 2009
33. R.H. Walden, Analog-to-digital converter survey and analysis. *IEEE J. Sel. Areas Commun.* **17**(4), 539–550 (1999)
34. B. Murmann, A/D converter trends: power dissipation, scaling and digitally assisted architectures, in *2008 IEEE Custom Integrated Circuits Conference* (IEEE, Piscataway, 2008), pp. 105–112
35. A.M. Ali, A. Morgan, C. Dillon, G. Patterson, S. Puckett, P. Bhoraskar, H. Dinc, M. Hensley, S. Bardsley, D. Lattimore et al., A 16-bit 250-MS/s IF sampling pipelined ADC with background calibration. *IEEE J. Solid-State Circuits* **45**(12), 2602–2612 (2010)
36. B. Murmann, ADC performance survey 1997–2022. <http://web.stanford.edu/~murmanna/adcsurvey.html>
37. P. Kinget, *Analog VLSI Integration of Parallel Signal Processing Systems—Ph.D. Dissertation*. KU Leuven, Leuven, BE, 1996
38. E. A. Vittoz, Future of analog in the VLSI environment, in *1990 IEEE International Symposium on Circuits and Systems (ISCAS)* (IEEE, Piscataway, 1990), pp. 1372–1375
39. R. Kapusta, H. Zhu, C. Lyden, Sampling circuits that break the kT/C thermal noise limit. *IEEE J. Solid-State Circuits* **49**(8), 1694–1701 (2014)
40. L. Shen, Y. Shen, X. Tang, C.-K. Hsu, W. Shi, S. Li, W. Zhao, A. Mukherjee, N. Sun, A 0.01mm^2 $25 \mu\text{W}$ 2MS/s 74dB-SNDR continuous-time pipelined-SAR ADC with 120fF input capacitor, in *2019 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2019), pp. 64–66
41. J. Liu, X. Tang, W. Zhao, L. Shen, N. Sun, A 13b 0.005mm^2 40MS/s SAR ADC with kT/C noise cancellation, in *2020 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2020), pp. 258–260
42. W. Sansen, *Advanced Analog Circuit Design—Training Slides*. IMEC Academy, Leuven, BE, 2014

43. L.E. Larson, High-speed analog-to-digital conversion with GaAs technology: prospects, trends and obstacles, in *1988 IEEE International Symposium on Circuits and Systems-(ISCAS)* (IEEE, Piscataway, 1988), pp. 2871–2878
44. H. Shibata, V. Kozlov, Z. Ji, A. Ganesan, H. Zhu, D. Paterson, A 9GS/s 1GHz-BW oversampled continuous-time pipeline ADC achieving—161dBFS/Hz NSD, in *2017 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2017), pp. 278–279
45. H. Shibata, G. Taylor, B. Schell, V. Kozlov, S. Patil, D. Paterson, A. Ganesan, Y. Dong, W. Yang, Y. Yin et al., An 800MHz-BW VCO-based continuous-time pipelined ADC with inherent anti-aliasing and on-chip digital reconstruction filter, in *2020 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2020), pp. 260–262
46. W. Heisenberg, *Die Physikalischen Prinzipien der Quantentheorie* (Verlag S. Hirzel, Leipzig, 1930)
47. A. Ramkaj, M.J. Pelgrom, M.S. Steyaert, F. Tavernier, In the pursuit of the optimal accuracy–speed–power analog-to-digital converter architecture: a mathematical framework. *IEEE Solid-State Circuits Mag.* **14**(1), 45–53 (2022)
48. A. Varzaghani, A. Kasapi, D.N. Loizos, S.-H. Paik, S. Verma, S. Zogopoulos, S. Sidiropoulos, A 10.3-GS/s, 6-bit flash ADC for 10G ethernet applications. *IEEE J. Solid-State Circuits* **48**(12), 3038–3048 (2013)
49. V.H.-C. Chen, L. Pileggi, An 8.5 mW 5 GS/s 6b flash ADC with dynamic offset calibration in 32 nm CMOS SOI, in *2013 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2013), pp. C264–C265
50. V.H.-C. Chen, L. Pileggi, A 69.5 mW 20 GS/s 6b time-interleaved ADC with embedded time-to-digital calibration in 32 nm CMOS SOI. *IEEE J. Solid-State Circuits* **49**(12), 2891–2901 (2014)
51. S. Zhu, B. Wu, Y. Cai, Y. Chiu, A 2GS/s 8b flash ADC based on remainder number system in 65 nm CMOS, in *2017 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2017), pp. C284–C285
52. P.M. Figueiredo, J.C. Vital, Kickback noise reduction techniques for CMOS latched comparators. *IEEE Trans. Circuits Syst. II: Exp. Briefs* **53**(7), 541–545 (2006)
53. S. Callender, W. Shin, H.-J. Lee, S. Pellerano, C. Hull, FinFET for mmwave-technology and circuit design challenges, in *2018 IEEE BiCMOS and Compound Semiconductor Integrated Circuits and Technology Symposium (BCICTS)* (IEEE, Piscataway, 2018), pp. 168–173
54. R. Kapusta, J. Shen, S. Decker, H. Li, E. Ibaragi, H. Zhu, A 14 b 80 MS/s SAR ADC with 73.6 dB SNDR in 65 nm CMOS. *IEEE J. Solid-State Circuits* **48**(12), 3059–3066 (2013)
55. M.J. Kramer, E. Janssen, K. Doris, B. Murmann, A 14 b 35 MS/s SAR ADC achieving 75 dB SNDR and 99 db SFDR with loop-embedded input buffer in 40 nm CMOS. *IEEE J. Solid-State Circuits* **50**(12), 2891–2900 (2015)
56. H.S. Bindra, A.-J. Annema, S.M. Louwsma, B. Nauta, A 0.2-8MS/s 10b flexible SAR ADC achieving 0.35-2.5fJ/conv-step and using self-quenched dynamic bias comparator, in *2019 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2019), pp. C74–C75
57. X. Tang, Y. Shen, X. Xin, S. Liu et al., A 10-bit 100MS/s SAR ADC with always-on reference ripple cancellation, in *2020 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2020), pp. C72–C73
58. A.T. Ramkaj, *Analysis and Design of High-Speed Successive Approximation Register ADCs – M.Sc. Thesis*. TU Delft, Delft, NL, 2014
59. J.L. McCreary, P.R. Gray, All-MOS charge redistribution analog-to-digital conversion techniques—Part I. *IEEE J. Solid-State Circuits* **10**(6), 371–379 (1975)
60. A.T. Ramkaj, M. Strackx, M.S. Steyaert, F. Tavernier, A 1.25-GS/s 7-b SAR ADC with 36.4-dB SNDR at 5 GHz using switch-bootstrapping, USPC DAC and triple-tail comparator in 28-nm CMOS. *IEEE J. Solid-State Circuits* **53**(7), 1889–1901 (2018)
61. H. Wei, C.-H. Chan, U.-F. Chio, S.-W. Sin, U. Seng-Pan, R.P. Martins, F. Maloberti, An 8-b 400-MS/s 2-b-per-cycle SAR ADC with resistive DAC. *IEEE J. Solid-State Circuits* **47**(11), 2763–2772 (2012)

62. B. Sedighi, A.T. Huynh, E. Skafidas, D. Micusik, Design of hybrid resistive-capacitive DAC for SAR A/D converters, in *2012 19th IEEE International Conference on Electronics, Circuits, and Systems (ICECS 2012)* (IEEE, Piscataway, 2012), pp. 508–511
63. K. Doris, E. Janssen, C. Nani, A. Zanakopoulos, G. Van der Weide, A 480 mW 2.6 GS/s 10b time-interleaved ADC with 48.5 dB SNDR up to Nyquist in 65 nm CMOS. *IEEE J. Solid-State Circuits* **46**(12), 2821–2833 (2011)
64. B.P. Ginsburg, A.P. Chandrakasan, An energy-efficient charge recycling approach for a SAR converter with capacitive DAC, in *2005 IEEE International Symposium on Circuits and Systems (ISCAS)* (IEEE, Piscataway, 2005), pp. 184–187
65. Y.-K. Chang, C.-S. Wang, C.-K. Wang, A 8-bit 500-kS/s low power SAR ADC for bio-medical applications, in *2007 IEEE Asian Solid-State Circuits Conference (ASSCC)* (IEEE, Piscataway, 2007), pp. 228–231
66. C.-C. Liu, S.-J. Chang, G.-Y. Huang, Y.-Z. Lin, A 10-bit 50-MS/s SAR ADC with a monotonic capacitor switching procedure. *IEEE J. Solid-State Circuits* **45**(4), 731–740 (2010)
67. L. Kull, T. Toifl, M. Schmatz, P.A. Francese, C. Menolfi, M. Braendli, M. Kossel, T. Morf, T.M. Andersen, Y. Leblebici, A 3.1 mW 8b 1.2 GS/s single-channel asynchronous SAR ADC with alternate comparators for enhanced speed in 32 nm digital SOI CMOS. *IEEE J. Solid-State Circuits* **48**(12), 3049–3058 (2013)
68. V. Hariprasath, J. Guerber, S.-H. Lee, U.-K. Moon, Merged capacitor switching based SAR ADC with highest switching energy-efficiency. *Electron. Lett.* **46**(9), 620–621 (2010)
69. D. Stepanovic, B. Nikolic, A 2.8 GS/s 44.6 mW time-interleaved ADC achieving 50.9 dB SNDR and 3 dB effective resolution bandwidth of 1.5 GHz in 65 nm CMOS. *IEEE J. Solid-State Circuits* **48**(4), 971–982 (2013)
70. V. Tripathi, B. Murmann, Mismatch characterization of small metal fringe capacitors. *IEEE Trans. Circuits Syst. I: Regul. Papers* **61**(8), 2236–2242 (2014)
71. L. Singer, S. Ho, M. Timko, D. Kelly, A 12b 65MSample/s CMOS ADC with 82dB SFDR at 120 MHz, in *2000 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2000), pp. 38–39
72. A.M. Ali, H. Dinc, P. Bhoraskar, C. Dillon, S. Puckett, B. Gray, C. Speir, J. Lanford, J. Brunsilius, P.R. Derounian et al., A 14 bit 1 GS/s RF sampling pipelined ADC with background calibration. *IEEE J. Solid-State Circuits* **49**(12), 2857–2867 (2014)
73. J. Mulder, D. Vecchi, Y. Ke, S. Bozzola, M. Core, N. Saputra, Q. Zhang, J. Riley, H. Yan, M. Introini et al., An 800MS/s 10b/13b receiver for 10GBASE-T ethernet in 28 nm CMOS, in *2015 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2015), pp. 1–3
74. A.M. Ali, H. Dinc, P. Bhoraskar, S. Bardsley, C. Dillon, M. Kumar, M. McShea, R. Bunch, J. Prabhakar, S. Puckett, A 12b 18GS/s RF sampling ADC with an integrated wideband track-and-hold amplifier and background calibration, in *2020 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2020), pp. 250–252
75. Y. Chiu, P.R. Gray, B. Nikolic, A 14-b 12-MS/s CMOS pipeline ADC with over 100-dB SFDR. *IEEE J. Solid-State Circuits* **39**(12), 2139–2151 (2004)
76. B.-G. Lee, B.-M. Min, G. Manganaro, J.W. Valvano, A 14-b 100-MS/s pipelined ADC with a merged SHA and first MDAC. *IEEE J. Solid-State Circuits* **43**(12), 2613–2619 (2008)
77. S. Devarajan, L. Singer, D. Kelly, S. Decker, A. Kamath, P. Wilkins, A 16-bit, 125 MS/s, 385 mW, 78.7 dB SNR CMOS pipeline ADC. *IEEE J. Solid-State Circuits* **44**(12), 3305–3313 (2009)
78. B. Murmann, B.E. Boser, A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification. *IEEE J. Solid-State Circuits* **38**(12), 2040–2050 (2003)
79. B. Verbruggen, J. Craninckx, M. Kuijk, P. Wambacq, G. Van der Plas, 'A 2.6 mW 6b 2.2GS/s 4-times interleaved fully dynamic pipelined ADC in 40 nm digital CMOS, in *2010 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2010), pp. 296–297

80. J. Wu, A. Chou, C.-H. Yang, Y. Ding, Y.-J. Ko, S.-T. Lin, W. Liu, C.-M. Hsiao, M.-H. Hsieh, C.-C. Huang et al., A 5.4GS/s 12b 500 mW pipeline ADC in 28nm CMOS, in *2013 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2013), pp. C92–C93
81. M. Kramer, *High-Resolution SAR A/D Converters with Loop-Embedded Input Buffer – Ph.D. Dissertation*. Stanford University, Stanford, CA, 2015
82. S.M. Louwsma, A.J.M. van Tuijl, M. Vertregt, B. Nauta, A 1.35 GS/s, 10 b, 175 mW time-interleaved AD converter in 0.13 μm CMOS. *IEEE J. Solid-State Circuits* **43**(4), 778–786 (2008)
83. C.C. Lee, M.P. Flynn, A SAR-assisted two-stage pipeline ADC. *IEEE J. Solid-State Circuits* **46**(4), 859–869 (2011)
84. F. van der Goes, C.M. Ward, S. Astgimath, H. Yan, J. Riley, Z. Zeng, J. Mulder, S. Wang, K. Bult, A 1.5 mW 68 dB SNDR 80 MS/s $2\times$ interleaved pipelined SAR ADC in 28 nm CMOS. *IEEE J. Solid-State Circuits* **49**(12), 2835–2845 (2014)
85. V. Tripathi, B. Murmann, A 160 MS/s, 11.1 mW, single-channel pipelined SAR ADC with 68.3 dB SNDR, in *2014 IEEE Custom Integrated Circuits Conference* (IEEE, Piscataway, 2014), pp. 1–4
86. Y. Zhou, B. Xu, Y. Chiu, A 12 bit 160 MS/s two-step SAR ADC with background bit-weight calibration using a time-domain proximity detector. *IEEE J. Solid-State Circuits* **50**(4), 920–931 (2015)
87. M. Brandolini, Y.J. Shin, K. Raviprakash, T. Wang, R. Wu, H.M. Geddada, Y.-J. Ko, Y. Ding, C.-S. Huang, W.-T. Shih et al., A 5 GS/s 150 mW 10 b SHA-less pipelined/SAR hybrid ADC for direct-sampling systems in 28 nm CMOS. *IEEE J. Solid-State Circuits* **50**(12), 2922–2934 (2015)
88. B. Vaz, A. Lynam, B. Verbruggen, A. Laraba, C. Mesadri, A. Boumaalif, J. Mcgrath, U. Kamath, R. De Le Torre, A. Manlapat et al., A 13b 4GS/s digitally assisted dynamic 3-stage asynchronous pipelined-SAR ADC, in *2017 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2017), pp. 276–277
89. B. Vaz, B. Verbruggen, C. Erdmann, D. Collins, J. Mcgrath, A. Boumaalif, E. Cullen, D. Walsh, A. Morgado, C. Mesadri et al., A 13bit 5GS/s ADC with time-interleaved chopping calibration in 16 nm FinFET, in *2018 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2018), pp. 99–100
90. W. Jiang, Y. Zhu, M. Zhang, C.-H. Chan, R.P. Martins, A 7.6mW 1GS/s 60dB SNDR single-channel SAR-assisted pipelined ADC with temperature-compensated dynamic Gm-R-based amplifier, in *2019 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2019), pp. 60–62
91. A. Ramkaj, J.C.P. Ramos, Y. Lyu, M. Strackx, J.M. Pelgrom, M. Steyaert, M. Verhelst, F. Tavernier, A 5GS/s 158.6mW 12b passive-sampling $8\times$ -interleaved hybrid ADC with 9.4 ENOB and 160.5 dB FoMs in 28 nm CMOS, in *2019 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2019), pp. 62–64
92. V. Giannini, P. Nuzzo, V. Chironi, A. Baschiroto, G. Van der Plas, J. Craninckx, An $820\mu\text{W}$ 9b 40MS/s noise-tolerant dynamic-SAR ADC in 90 nm digital CMOS, in *2008 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2008), pp. 238–240
93. A.T. Ramkaj, J.C.P. Ramos, M.J. Pelgrom, M.S. Steyaert, M. Verhelst, F. Tavernier, A 5-GS/s 158.6-mW 9.4-ENOB passive-sampling time-interleaved three-stage pipelined-SAR ADC with analog-digital corrections in 28-nm CMOS. *IEEE J. Solid-State Circuits* **55**(6), 1553–1564 (2020)
94. E. Janssen, K. Doris, A. Zanicopoulos, A. Murrioni, G. Van der Weide, Y. Lin, L. Alvado, F. Darthenay, Y. Fregeais, An 11b 3.6GS/s time-interleaved SAR ADC in 65 nm CMOS, in *2013 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2013), pp. 464–465
95. C.-H. Chan, Y. Zhu, S.-W. Sin, U. Seng-Pan, R.P. Martins, A 5.5 mW 6b 5GS/s $4\times$ -interleaved 3b/cycle SAR ADC in 65 nm CMOS, in *2015 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2015), pp. 1–3

96. J.P. Keane, N.J. Guilar, D. Stepanovic, B. Wuppermann, C. Wu, C.W. Tsang, R. Neff, K. Nishimura, An 8GS/s time-interleaved SAR ADC with unresolved decision detection achieving -58dBFS noise and 4 GHz bandwidth in 28 nm CMOS, in *2017 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2017), pp. 284–285
97. M. Zhang, Y. Zhu, C.-H. Chan, R.P. Martins, A $4\times$ interleaved 10GS/s 8b time-domain ADC with $16\times$ interpolation-based inter-stage gain achieving 37.5 dB SNDR at 18 GHz input, in *2020 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2020), pp. 252–254
98. B. Razavi, Design considerations for interleaved ADCs. *IEEE J. Solid-State Circuits* **48**(8), 1806–1817 (2013)
99. A. Buchwald, High-speed time interleaved ADCs. *IEEE Commun. Mag.* **54**(4), 71–77 (2016)
100. N. Kurosawa, H. Kobayashi, K. Maruyama, H. Sugawara, K. Kobayashi, Explicit analysis of channel mismatch effects in time-interleaved ADC systems., *IEEE Trans. Circuits Syst. I: Fund. Theory Appl.* **48**(3), 261–271 (2001)
101. S. Devarajan, L. Singer, D. Kelly, T. Pan, J. Silva, J. Brunsilius, D. Rey-Losada, F. Murden, C. Speir, J. Bray et al., A 12-b 10-GS/s interleaved pipeline ADC in 28-nm CMOS technology. *IEEE J. Solid-State Circuits* **52**(12), 3204–3218 (2017)
102. D. Stepanovic, *Calibration Techniques for Time-Interleaved SAR A/D Converters – Ph.D. Dissertation*. UC Berkeley, Berkeley, CA, 2012
103. N. Le Dortz, J.-P. Blanc, T. Simon, S. Verhaeren, E. Rouat, P. Urard, S. Le Tual, D. Goguet, C. Lelandais-Perrault, P. Benabes, A 1.62GS/s time-interleaved SAR ADC with digital background mismatch calibration achieving interleaving spurs below 70dBFS, in *2014 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2014), pp. 386–388
104. S. Devarajan, L. Singer, D. Kelly, S. Kosic, T. Pan, J. Silva, J. Brunsilius, D. Rey-Losada, F. Murden, C. Speir et al., A 12b 10GS/s interleaved pipeline ADC in 28 nm CMOS technology, in *2017 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2017), pp. 288–289
105. L. Kull, T. Toifl, M. Schmatz, P.A. Francese, C. Menolfi, M. Braendli, M. Kossel, T. Morf, T.M. Andersen, Y. Leblebici, A 90GS/s 8b 667 mW $64\times$ interleaved SAR ADC in 32 nm digital SOI CMOS, in *2014 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2014), pp. 378–379
106. L. Kull, D. Luu, C. Menolfi, M. Braendli, P.A. Francese, T. Morf, M. Kossel, A. Cevrero, I. Ozkaya, T. Toifl, A 24–72-GS/s 8-b time-interleaved SAR ADC with 2.0–3.3-pJ/conversion and $> 30\text{dB}$ SNDR at Nyquist in 14-nm CMOS FinFET. *IEEE J. Solid-State Circuits* **53**(12), 3508–3516 (2018)
107. S. Le Tual, P.N. Singh, C. Curis, P. Dautriche, A 20GHz-BW 6b 10GS/s 32mW time-interleaved SAR ADC with master T&H in 28 nm UTBB FDSOI technology, in *2014 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2014), pp. 382–383
108. A.T. Ramkaj, M.S. Steyaert, F. Tavernier, A 13.5-Gb/s 5-mV-sensitivity 26.8-ps-CLK–OUT delay triple-latch feedforward dynamic comparator in 28-nm CMOS, in *2019-IEEE 45th European Solid State Circuits Conference-(ESSCIRC)* (IEEE, Piscataway, 2019), pp. 167–170
109. A.T. Ramkaj, M.J.M. Pelgrom, M.S.J. Steyaert, F. Tavernier, A 28 nm CMOS triple-latch feed-forward dynamic comparator with $<27\text{ ps}/1\text{ V}$ and $<70\text{ ps}/0.6\text{ V}$ delay at 5 mV-sensitivity. *IEEE Trans. Circuits Syst. I: Regul. Papers* **69**(11), 4404–4414 (2022)
110. B. Wicht, T. Nirschl, D. Schmitt-Landsiedel, Yield and speed optimization of a latch-type voltage sense amplifier. *IEEE J. Solid-State Circuits* **39**(7), 1148–1158 (2004)
111. D. Schinkel, E. Mensink, E. Klumperink, E. Van Tuijl, B. Nauta, A double-tail latch-type voltage sense amplifier with 18 ps setup+hold time, in *2002 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2007), pp. 314–605

112. P. Nuzzo, F. De Bernardinis, P. Terreni, G. Van der Plas, Noise analysis of regenerative comparators for reconfigurable ADC architectures. *IEEE Trans. Circuits Syst. I: Regul. Papers* **55**(6), 1441–1454 (2008)
113. T. Jiang, W. Liu, F.Y. Zhong, C. Zhong, K. Hu, P.Y. Chiang, A single-channel, 1.25-GS/s, 6-bit, 6.08-mW asynchronous successive-approximation ADC with improved feedback delay in 40-nm CMOS. *IEEE J. Solid-State Circuits* **47**(10), 2444–2453 (2012)
114. P.M. Figueiredo, Comparator metastability in the presence of noise. *IEEE Trans. Circuits Syst. I: Regul. Papers* **60**(5), 1286–1299 (2013)
115. B. Razavi, *Principles of Data Conversion System Design*, vol. 126 (IEEE Press, New York, 1995)
116. B. Razavi, The design of a comparator [the analog mind]. *IEEE Solid-State Circuits Mag.* **12**(4), 8–14 (2020)
117. B. Goll, H. Zimmermann, A 65 nm CMOS comparator with modified latch to achieve 7 GHz/1.3 mW at 1.2 V and 700 MHz/47 μ W at 0.6V, in *2009 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2009), pp. 328–329
118. L. Kull, T. Toifl, M. Schmatz, P.A. Francese, C. Menolfi, M. Braendli, M. Kossel, T. Morf, T.M. Andersen, Y. Leblebici, A 3.1mW 8b 1.2GS/s single-channel asynchronous SAR ADC with alternate comparators for enhanced speed in 32 nm digital SOI CMOS, in *2013 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2013), pp. 468–469
119. H.S. Bindra, J. Ponte, B. Nauta, A 174 μ V_{RMS} input noise, 1GS/s comparator in 22 nm FDSOI with a dynamic-bias preamplifier using tail charge pump and capacitive neutralization across the latch, in *2022 IEEE International Solid-State Circuits Conference-(ISSCC)*, vol. 65 (IEEE, Piscataway, 2022), pp. 1–3
120. A. Ramkaj, F. Tavernier, M. Steyaert, Fast switch bootstrapping for GS/s high-resolution analog-to-digital converter, in *2015 11th Conference on Ph. D. Research in Microelectronics and Electronics (PRIME)* (IEEE, Piscataway, 2015), pp. 73–76
121. A. Ramkaj, M. Strackx, M. Steyaert, F. Tavernier, A 36.4dB SNDR @ 5GHz 1.25GS/s 7b 3.56mW single-channel SAR ADC in 28 nm bulk CMOS, in *2017-43rd IEEE European Solid State Circuits Conference-(ESSCIRC)* (IEEE, Piscataway, 2017), pp. 167–170
122. A. Ramkaj, M. Strackx, M. Steyaert, F. Tavernier, An 11 GHz dual-sided self-calibrating dynamic comparator in 28 nm CMOS. *Electronics* **8**(1), 13 (2019)
123. Y. Zhou, B. Xu, Y. Chiu, A 12b 160MS/s synchronous two-step SAR ADC achieving 20.7fJ/step FoM with opportunistic digital background calibration, in *2014 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2014), pp. 1–2
124. S.-W.M. Chen, R.W. Brodersen, A 6-bit 600-MS/s 5.3-mW asynchronous ADC in 0.13- μ m CMOS. *IEEE J. Solid-State Circuits* **41**(12), 2669–2680 (2006)
125. J. Yang, T.L. Naing, R.W. Brodersen, A 1 GS/s 6 bit 6.7 mW successive approximation ADC using asynchronous processing. *IEEE J. Solid-State Circuits* **45**(8), 1469–1478 (2010)
126. Y.-C. Lien, A 4.5-mW 8-b 750-MS/s 2-b/step asynchronous subranged SAR ADC in 28-nm CMOS technology, in *2012 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2012), pp. 88–89
127. D. Bankman, A. Yu, K. Zheng, B. Murmann, Understanding metastability in SAR ADCs: part I: synchronous. *IEEE Solid-State Circuits Mag.* **11**(2), 86–97 (2019)
128. A. Yu, D. Burkman, K. Zheng, B. Murmann, Understanding metastability in SAR ADCs: part II: asynchronous. *IEEE Solid-State Circuits Mag.* **11**(3), 16–32 (2019)
129. H. Wei, C.-H. Chan, U.-F. Chio, S.-W. Sin, U. Seng-Pan, R. Martins, F. Maloberti, A 0.024 mm² 8b 400MS/s SAR ADC with 2b/cycle and resistive DAC in 65 nm CMOS, in *2011 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2011), pp. 188–190
130. C.-H. Chan, Y. Zhu, I.-M. Ho, W.-H. Zhang, U. Seng-Pan, R.P. Martins, A 5 mW 7b 2.4GS/s 1-then-2b/cycle SAR ADC with background offset calibration, in *2017 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2017), pp. 282–283

131. D.-R. Oh, K.-J. Moon, W.-M. Lim, Y.-D. Kim, E.-J. An, S.-T. Ryu, An 8b 1GS/s 2.55mW SAR-Flash ADC with complementary dynamic amplifiers, in *2020 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2020), pp. 1–2
132. M. Miyahara, Y. Asada, D. Paik, A. Matsuzawa, A Low-noise self-calibrating dynamic comparator for high-speed ADCs, in *2008 IEEE Asian Solid-State Circuits Conference-(ASSCC)* (IEEE, Piscataway, 2008), pp. 269–272
133. W. Liu, P. Huang, Y. Chiu, A 12-bit, 45-MS/s, 3-mW redundant successive-approximation-register analog-to-digital converter with digital calibration. *IEEE J. Solid-State Circuits* **46**(11), 2661–2672 (2011)
134. C.-C. Liu, S.-J. Chang, G.-Y. Huang, Y.-Z. Lin, C.-M. Huang, C.-H. Huang, L. Bu, and C.-C. Tsai, A 10b 100MS/s 1.13mW SAR ADC with binary-scaled error compensation, in *2010 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2010), pp. 386–387
135. Y. Duan, E. Alon, A 6b 46GS/s ADC with >23GHz BW and sparkle-code error correction, in *2015 Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2015), pp. C162–C163
136. A.M. Abo, P.R. Gray, A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter. *IEEE J. Solid-State Circuits* **34**(5), 599–606 (1999)
137. C.-H. Chan, Y. Zhu, S.-W. Sin, R.P. Martins et al., A 6 b 5 GS/s 4 interleaved 3 b/cycle SAR ADC. *IEEE J. Solid-State Circuits* **51**(2), 365–377 (2016)
138. E. Swindlehurst, H. Jensen, A. Petrie, Y. Song, Y.-C. Kuan, M.-C.F. Chang, J.-T. Wu, S.-H.W. Chiang, An 8-bit 10-GHz 21-mW time-interleaved SAR ADC with grouped DAC capacitors and dual-path bootstrapped switch. *IEEE Solid-State Circuits Lett.* **2**(9), 83–86 (2019)
139. K.D. Choo, J. Bell, M.P. Flynn, Area-efficient 1GS/s 6b SAR ADC with charge-injection-cell-based DAC, in *2016 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2016), pp. 460–461
140. F. Kuttner, A 1.2V 10b 20MSample/s non-binary successive approximation ADC in 0.13 μ m CMOS, in *2002 IEEE International Solid-State Circuits Conference-(ISSCC)*, vol. 1 (IEEE, Piscataway, 2002), pp. 176–177
141. X. Staff, [Online Available], *Understanding Key Parameters for RF-Sampling Data Converters*, White Paper (Xilinx Incorporated, Feb 2019)
142. D. Kozischek, J. Burton, [Online available], get ready—'cause here it comes: DOCSIS 4.0, white paper, in *Broadband Success Partners in collaboration with Corning Optical Communications* (2020)
143. P. Delos, *A Review of Wideband RF Receiver Architecture Options* (Analog Devices, Incorporated, 2017)
144. E.H. Armstrong, A new system of short wave amplification. *Proc. Inst. Radio Eng.* **9**(1), 3–11 (1921)
145. A.A. Abidi, Direct-conversion radio transceivers for digital communications. *IEEE J. Solid-State Circuits* **30**(12), 1399–1410 (1995)
146. J. Crols, M.S. Steyaert, A single-chip 900 MHz CMOS receiver front-end with a high performance low-IF topology. *IEEE J. Solid-State Circuits* **30**(12), 1483–1492 (1995)
147. M. Straayer, J. Bales, D. Birdsall, D. Daly, P. Elliott, B. Foley, R. Mason, V. Singh, X. Wang, A 4GS/s time-interleaved RF ADC in 65 nm CMOS with 4GHz input bandwidth, in *2016 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2016), pp. 464–465
148. J. Wu, A. Chou, T. Li, R. Wu, T. Wang, G. Cusmai, S.-T. Lin, C.-H. Yang, G. Unruh, S.R. Dommaraju et al., A 4GS/s 13b pipelined ADC with capacitor and amplifier sharing in 16 nm CMOS, in *2016 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2016), pp. 466–467
149. J.-W. Nam, M. Hassanpourghadi, A. Zhang, M.S.-W. Chen, A 12-bit 1.6, 3.2, and 6.4 GS/s 4-b/cycle time-interleaved SAR ADC with dual reference shifting and interpolation. *IEEE J. Solid-State Circuits* **53**(6), 1765–1779 (2018)

150. A.M. Ali, H. Dinc, P. Bhoraskar, S. Puckett, A. Morgan, N. Zhu, Q. Yu, C. Dillon, B. Gray, J. Lanford et al., A 14-bit 2.5GS/s and 5GS/s RF sampling ADC with background calibration and dither, in *2016 IEEE Symposium on VLSI Circuits-(VLSI)* (IEEE, Piscataway, 2016), pp. 1–2
151. T. Ali, E. Chen, H. Park, R. Yousry, Y.-M. Ying, M. Abdullatif, M. Gandara, C.-C. Liu, P.-S. Weng, H.-S. Chen et al., A 460mW 112Gb/s DSP-based transceiver with 38dB loss compensation for next-generation data centers in 7 nm FinFET technology, in *2020 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2020), pp. 118–120
152. S. Lee, A. P. Chandrakasan, H.-S. Lee, A 1 GS/s 10b 18.9 mW time-interleaved SAR ADC with background timing skew calibration. *IEEE J. Solid-State Circuits* **49**(12), 2846–2856 (2014)
153. B. Razavi, The strongARM Latch [a circuit for all seasons]. *IEEE Solid-State Circuits Mag.* **7**(2), 12–17 (2015)
154. J. Lin, M. Miyahara, A. Matsuzawa, A 15.5 dB, wide signal swing, dynamic amplifier using a common-mode voltage detection technique, in *2011 IEEE International Symposium of Circuits and Systems-(ISCAS)* (IEEE, Piscataway, 2011), pp. 21–24
155. T. Astgimath, *A Low-Noise Low-Power Dynamic Amplifier with Common Mode Detect and a Low-Power Low-Noise Comparator for Pipelined SAR-ADC – M.Sc. Thesis.* TU Delft, Delft, NL, 2012
156. B. Hershberg, D. Dermit, B. van Liempd, E. Martens, N. Markulic, J. Lagos, J. Craninckx, A 3.2GS/s 10 ENOB 61 mW ringamp ADC in 16 nm with background monitoring of distortion, in *2019 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2019), pp. 58–60
157. A. Ramkaj, A. Cantoni, G. Manganaro, S. Devarajan, M. Steyaert, F. Tavernier, A 30GHz-BW <math><57\text{dB-IM3}</math> direct RF receiver analog front end in 16 nm FinFET, in *2022 IEEE Symposium on VLSI Technology and Circuits (VLSI)* (IEEE, Piscataway, 2022), pp. 100–101
158. G. Manganaro, A. Ramkaj, F. Tavernier, Amplifiers for RF ADCs. Mar. 24 2022, U.S. Patent App. 17/031,426
159. A.M.A. Ali, H. Dinc, P. Bhoraskar, S. Bardsley, C. Dillon, M. McShea, J.P. Periathambi, S. Puckett, A 12-b 18-GS/s RF sampling ADC with an integrated wideband track-and-hold amplifier and background calibration. *IEEE J. Solid-State Circuits* **55**(12), 3210–3224 (2020)
160. K. Zheng, Y. Frans, S.L. Ambatipudi, S. Asuncion, H.T. Reddy, K. Chang, B. Murmann, An inverter-based analog front-end for a 56-Gb/s PAM-4 wireline transceiver in 16-nm CMOS. *IEEE Solid-State Circuits Lett.* **1**(12), 249–252 (2018)
161. M. Pisati, F. De Bernardinis, P. Pascale, C. Nani, N. Ghittori, E. Pozzati, M. Sosio, M. Garampazzi, A. Milani, A. Minuti et al., A 243-mW 1.25–56-Gb/s continuous range PAM-4 42.5-dB IL ADC/DAC-based transceiver in 7-nm FinFET. *IEEE J. Solid-State Circuits* **55**(1), 6–18 (2020)
162. J. Im, K. Zheng, C.H.A. Chou, L. Zhou, J.W. Kim, S. Chen, Y. Wang, H.W. Hung, K. Tan, W. Lin, A.B. Roldan, D. Carey, I. Chlis, R. Casey, A. Bekele, Y. Cao, D. Mahashin, H. Ahn, H. Zhang, Y. Frans, K. Chang, A 112-Gb/s PAM-4 long-reach wireline transceiver using a 36-way time-interleaved SAR ADC and inverter-based RX analog front-end in 7-nm FinFET. *IEEE J. Solid-State Circuits* **56**(1), 7–18 (2021)
163. B. Razavi, The bridged T-Coil [A circuit for all seasons]. *IEEE Solid-State Circuits Mag.* **7**(4), 9–13 (2015)
164. S. Cao, J.-H. Chun, S.G. Beebe, R.W. Dutton, ESD design strategies for high-speed digital and RF circuits in deeply scaled silicon technologies. *IEEE Trans. Circuits Syst. I: Regul. Papers* **57**(9), 2301–2311 (2010)
165. A.B. Williams, F.J. Taylor, *Electronic Filter Design Handbook* (McGraw-Hill Education, 2006)
166. M.-S. Chen, C.-K.K. Yang, A 50–64 Gb/s serializing transmitter with a 4-Tap, LC-ladder-filter-based FFE in 65 nm CMOS technology. *IEEE J. Solid-State Circuits* **50**(8), 1903–1916 (2015)

167. S. Shekhar, J.S. Walling, D.J. Allstot, Bandwidth extension techniques for CMOS amplifiers. *IEEE J. Solid-State Circuits* **41**(11), 2424–2439 (2006)
168. N. Rakuljic, C. Speir, E. Otte, J. Bray, C. Petersen, G. Manganaro, In-situ nonlinear calibration of a RF signal chain, in *2018 IEEE International Symposium on Circuits and Systems-(ISCAS)* (IEEE, Piscataway, 2018), pp. 1–5
169. T. Sundstrom, B. Murmann, C. Svensson, Power dissipation bounds for high-speed Nyquist analog-to-digital converters. *IEEE Trans. Circuits Syst. I: Regul. Papers* **56**(3), 509–518 (2009)
170. C.-C. Liu, M.-C. Huang, Y.-H. Tu, A 12 bit 100 MS/s SAR-assisted digital-slope ADC. *IEEE J. Solid-State Circuits* **51**(12), 2941–2950 (2016)
171. B. Razavi, Lower bounds on power consumption of clock generators for ADCs, in *2020 IEEE International Symposium on Circuits and Systems-(ISCAS)* (IEEE, Piscataway, 2020), pp. 1–5

Index

A

Accuracy, vii, viii, 4, 5, 7, 11, 16, 30, 33, 35, 37–39, 41, 42, 44, 47, 48, 50, 53–55, 57, 59, 60, 70, 71, 84–87, 94, 99, 102, 103, 109, 113, 116–120, 122, 147, 150, 154, 158, 161, 163, 164, 166, 169, 175, 189, 204, 220, 240, 247–252

Accuracy-energy, 57, 58

Accuracy-speed, vii, viii, 40, 43, 46, 49–53, 57, 58, 187

Accuracy-speed-power limits, vii, 10, 30, 37–55, 57, 62–65, 68, 78–81, 83–90, 93–101, 118, 248, 251

ADC-based receiver, viii, 4, 12, 13, 218, 221, 224, 229, 234, 244–246

Aliasing, 18, 20, 21

Analog-digital calibration, viii, 12, 188, 202

Analog front end, viii, 7, 11–13, 26, 27, 32, 54, 217–246, 251, 254

Analog-to-digital (A/D)

Analog-to-digital conversion, 1, 2, 12, 15–56, 247, 248

Analog-to-digital converter (ADC), vii, 1

Aperture jitter, 10, 26, 29–30, 47–49, 53–55

Architecture, vii, viii, 3, 4, 11, 12, 21, 57–88, 98–104, 108–119, 122, 155–157, 180, 183–186, 188–190, 199, 214, 215, 248, 251, 252, 254, 255

B

Bandwidth, vii, viii, 3, 18, 57, 125, 150, 183

Book structure, 12–13

Bootstrapped input switch, viii, 12, 156, 159–162, 173, 181, 208, 253

C

Capacitive digital-to-analog converter (CDAC), 70–79, 93–95, 162–165, 168–176, 181, 249

Cascoded buffer, 223, 232–236

Circuit, vii, 1, 15, 78, 121, 149, 183

Clock jitter, 176, 209

Common-gate (CG), 230, 231, 254

Common-source (CS), 230, 231, 235, 251, 254

Comparator, 6, 27, 59, 121, 149, 188

Complementary metal-oxide-semiconductor (CMOS), vii, viii, 4, 5, 7–13, 28, 33, 41, 57, 58, 65–67, 70, 79, 115, 118, 125, 128–137, 145–147, 149, 156–173, 178, 180, 181, 184, 188–204, 216, 218–236, 244–246, 248–252, 254, 256

Conclusion, 13, 15, 54–55, 57, 116, 118–119, 121, 147, 149, 179, 181, 183, 214, 216, 217, 246–256

D

Deep-scaled CMOS, vii, viii, 4, 7–13, 57, 65, 66, 145, 147, 184, 218, 229, 244, 246, 248, 249, 251–252, 254

Digital signal processing (DSP), 1–5, 184, 214, 218

- Digital-to-analog converter (DAC), viii, 1, 12, 69–79, 81, 82, 88, 106, 150–152, 154–159, 161–165, 179, 181, 202–205, 249, 252, 253
- Digital-to-analog (D/A), 1
- Direct RF sampling, viii, 4, 184, 185, 214, 218, 244, 246, 251
- Distributed filter, 224–230
- Double-tail, 122, 125–128, 169, 171, 249
- Dynamic comparator, viii, 12, 62, 121–147, 156, 166–171, 181, 249
- Dynamic integrator, 201
- E**
- Electrostatic discharge (ESD), 156, 188, 191, 218–220, 222–230, 234, 236, 246, 251
- F**
- Feed-forward, viii, 12, 128–137, 142, 147, 249, 252
- Figure-of-merit (FoM), 36–37, 104, 177–181
- FinFET CMOS, viii, 12, 222, 246
- Flash, vii, 3, 11, 12, 57–70, 78, 79, 81, 83, 84, 86–88, 92, 93, 98, 102, 118, 121, 152, 248, 252, 254, 255
- Flash limits, 68, 81
- Fully dynamic, 121, 128, 166, 250
- Future research, 13, 119, 247
- G**
- GHz-sample rate, vii, 5, 10, 11, 40, 150, 159, 183, 214, 248, 250, 253, 255
- 5G mm-Wave, 218, 241
- H**
- High-resolution, viii, 12, 59, 81, 93, 98, 156, 183–256
- High-sensitivity, 121–147
- High-speed, vii, viii, 4, 7, 12, 42, 62, 81, 121, 122, 147, 149–181, 189, 200, 201, 206, 207, 251
- Hybrid, viii, 3, 4, 6, 11–13, 59, 70, 88–97, 118, 139, 183, 186, 188–205, 216, 226, 244–246, 248, 250, 251, 254, 255
- Hybrid amplifier, viii, 12, 222, 223, 230–232, 236, 254
- I**
- Interleaver architectures, viii, 11, 12, 57, 108–118, 186, 248, 252
- Interleaver model, 117
- Interleaving errors, 12, 57, 104–109, 117, 119, 185, 188, 189, 202, 209, 248
- L**
- Linearity, 3–5, 70, 85–87, 91, 94, 95, 113, 155, 156, 159–161, 163, 177, 181, 188, 191, 199–201, 216, 218, 219, 222, 223, 225, 227–232, 234, 235, 240, 241, 244, 246, 249, 251–254
- Low-power, 4, 13, 99, 121, 184, 200, 214, 218, 222, 239, 246, 247
- M**
- Metastability, 10, 42–46, 53–55, 61, 121, 132, 133, 136, 152, 154, 158, 166, 168, 169
- Mismatch, 5, 9, 11, 29, 33, 38, 48–51, 53–55, 85, 92, 102, 104–109, 111–113, 119, 165, 185, 188, 204, 207, 208, 210, 216, 237, 248, 256
- N**
- Non-linearity, 26, 30–33, 38, 54, 118, 152, 154, 156, 187, 190
- Nyquist criterion, 19–21, 54
- O**
- Original contributions, 13, 251–254
- Overview, 12, 13, 15, 54, 57, 60–62, 69–71, 81–83, 88–93, 102–104, 183, 217, 247–251
- P**
- Passive front end, 190
- Performance metrics, 15, 248
- Pipeline, 3, 45, 57, 121, 186
- Pipelined, 255
- Pipelined-SAR hybrid, 11, 118, 216, 248, 254
- Power, 2, 21, 57, 121, 152, 183
- Push-pull, 216, 222, 223, 230–233, 244–246, 251

Q

Quantization, 12, 15, 16, 21–23, 25, 26, 30, 31, 34, 35, 38, 44, 52, 54, 82, 154, 163, 178, 187, 189, 194, 199, 248
 Quantization noise, 39, 42, 45, 102, 176, 189

R

Reduced-stacking, 125, 136, 142, 147, 252
 Research goal, 10–12, 248
 Residue amplifier (RA), 11, 59, 81, 84–88, 91–95, 98, 99, 106, 119–120, 155, 199, 201, 202, 204, 252

S

Sampling, 4, 15, 57, 144, 150, 184
 Sampling rate, 16
 SAR limits, 78, 79
 Single-channel, viii, 12, 13, 77, 149–181, 183, 207, 212, 249, 253
 Sixth-generation (6G), 4, 149, 212, 218
 Spectral purity, 10, 13, 21, 36, 61, 159, 174, 184, 185, 187, 190, 193, 207, 210, 214, 216–219, 222, 237, 239, 246, 250, 251, 253
 Speed, 3, 16, 57, 121, 149, 183
 State-of-the-Art (SoTA), 57, 59, 102, 118, 147, 178, 179, 185, 213, 216, 244
 Strong-ARM, 122–125, 169, 171, 249

Successive approximation register (SAR), 3, 57, 121, 149, 183
 Switched attenuator cells, 246
 System, vii, viii, 1–7, 10–13, 21, 36, 39, 61, 83, 121, 124, 125, 136, 147, 149, 150, 156, 177–179, 181, 183, 214, 247–251, 254

T

Technology, vii, viii, 1, 2, 5, 7, 9–12, 15, 33, 44, 50, 59, 61, 65, 66, 116, 118, 146, 180, 185, 188, 210, 214, 215, 245, 247, 248, 250–252, 254
 Thermal noise, 8, 10, 26–29, 38, 41, 42, 48, 54, 61, 133, 156, 158, 176, 191, 198
 Three-stage, viii, 12, 97, 98, 118, 119, 121, 128–137, 145, 147, 169, 188, 190, 199, 200, 204, 216, 249, 250, 252, 254, 255
 Time-interleaved (TI) ADC, 183–256
 Time-interleaving, 3, 5, 6, 11–13, 57, 58, 102–119, 185, 217, 218, 248, 250
 Triple-latch, viii, 12, 128–137, 249, 252

V

Variable gain/attenuation, 222, 251
 Very large scale integration (VLSI), viii, 2, 57, 215, 217, 245