



HAL
open science

Continuous time signal processing for wake-up radios

Alin Ratiu

► **To cite this version:**

Alin Ratiu. Continuous time signal processing for wake-up radios. Networking and Internet Architecture [cs.NI]. INSA de Lyon, 2015. English. NNT : 2015ISAL0078 . tel-01375171

HAL Id: tel-01375171

<https://theses.hal.science/tel-01375171v1>

Submitted on 13 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ DE LYON

DOCTORAL THESIS

Continuous Time Signal Processing for Wake-Up Radios

Author:

Alin Ratiu

Supervisor:

Dominique Morche

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

at the

Commissariat à l'Énergie Atomique
LETI / DACLE / SCCI / LAIR

October 8, 2015

Abstract

Continuous Time Signal Processing for Wake-Up Radios

by Alin Ratiu

Wake-Up Receivers (WU-RX) have been recently proposed as candidates to reduce the communication power budget of wireless networks. Their role is to sense the environment and wake up the main receivers which then handle the bulk data transfer. Existing WU-RXs achieve very high sensitivities for power consumptions below $50\mu\text{W}$ but severely degrade their performance in the presence of out-of-band blockers. We attempt to tackle this problem by implementing an ultra low power, tunable, intermediate frequency filtering stage. Its specifications are derived from standard WU-RX architectures; it is shown that classic filtering techniques are either not tunable enough or demand a power consumption beyond the total WU-RX budget of $100\mu\text{W}$. We thus turn to the use of Continuous Time Digital Signal Processing (CT-DSP) which offers the same level of programmability as standard DSP solutions while providing an excellent scalability of the power consumption with respect to the characteristics of the input signal. A CT-DSP chain can be divided into two parts: the CT-ADC and the CT-DSP itself; the specifications of these two blocks, given the context of this work, are also discussed.

The CT-ADC is based on a novel, delta modulator-based architecture which achieves a very low power consumption; its maximum operation frequency was extended by the implementation of a very fast feedback loop. Moreover, the CT nature of the ADC means that it does not do any sampling in time, hence no anti-aliasing filter is required. The proposed ADC requires only $24\mu\text{W}$ to quantize signals in the [10MHz 50MHz] bandwidth for an SNR between 32dB and 42dB, resulting in a figure of merit of $3 - 10\text{fJ/conv} - \text{step}$, among the best reported for the selected frequency range.

Finally, we present the architecture of the CT-DSP which is divided into two parts: a CT-IIR and a CT-FIR. The CT-IIR is implemented by placing a standard CT-FIR in a feedback loop around the CT-ADC. If designed correctly, the feedback loop can now cancel out certain frequencies from the CT-ADC input (corresponding to those of out-of-band interferers) while boosting the power of the useful signal. The effective amplitude of the CT-ADC input is thus reduced, making it generate a smaller number of tokens, thereby reducing the power consumption of the subsequent CT-FIR by a proportional amount. The CT-DSP consumes around $100\mu\text{W}$ while achieving more than 40dB of out-of-band rejection; for a bandpass implementation, a 2MHz passband can be shifted over the entire ADC bandwidth.

Acknowledgements

First of all, I would like to thank my advisor, Dominique Morche for providing me with a very interesting research topic. He gave me the liberty to work as I desired but, at the same time, knew when to step in and bring his vast experience in systems analysis. From him, I have learned about the importance of strategy and positioning even in a scientific field such as microelectronics.

I would also like to thank our collaborators from Columbia University, Sharvil Patil for inspiring me with his ambition as well as Prof. Yannis Tsividis for teaching me about the importance of rigor in scientific writing.

I also thank Bruno Allard, Xuefang Lin-Shi and Jacques Verdier for giving me the opportunity of pursuing a PhD under the guidance of INSA de Lyon, as well as for the assistance and feedback they provided over the years.

Thanks also to the members of my defense committee, Prof. Pieter Harpe, Dominique Dallet, Hassan Aboushady, Stéphane Le Tual, Prof. Yannis Tsividis, and Bruno Allard for their interest, time and feedback given to me. A special thanks to Prof. Pieter Harpe and Stéphane Le Tual for their detailed and constructive comments they gave me regarding this manuscript.

I would also like to thank my fellow PhD students, especially Matthieu Verdy, for teaching me about Linux development, Robert Polster, for the interesting discussions regarding circuit design and finally, David Buffeteau for the time he took helping me improve this manuscript. I take this opportunity to send a big thank you to the rest of the LAIR team at CEA for their support and hospitality.

I thank the testing team, especially Sylvain Dumas and Christian Chancel for the support they offered in PCB design and in preparing the testing equipment as well as the testing software.

Finally, I thank my family for the support they have offered me throughout the years as well as for inspiring me to pursue a career which eventually led to this PhD.

INSA Direction de la Recherche - Ecoles Doctorales – Quinquennal 2011-2015

SIGLE	ECOLE DOCTORALE	NOM ET COORDONNEES DU RESPONSABLE
CHIMIE	CHIMIE DE LYON http://www.edchimie-lyon.fr Sec : Renée EL MELHEM Bat Blaise Pascal 3 ^e étage 04 72 43 80 46 Insa : R. GOURDON secretariat@edchimie-lyon.fr	M. Jean Marc LANCELIN Université de Lyon – Collège Doctoral Bat ESCPE 43 bd du 11 novembre 1918 69622 VILLEURBANNE Cedex Tél : 04.72.43 13 95 directeur@edchimie-lyon.fr
E.E.A.	ELECTRONIQUE, ELECTROTECHNIQUE, AUTOMATIQUE http://edeea.ec-lyon.fr Sec : M.C. HAVGODOUKIAN Ecole-doctorale.eea@ec-lyon.fr	M. Gérard SCORLETTI Ecole Centrale de Lyon 36 avenue Guy de Collongue 69134 ECULLY Tél : 04.72.18 60.97 Fax : 04 78 43 37 17 Gerard.scorletti@ec-lyon.fr
E2M2	EVOLUTION, ECOSYSTEME, MICROBIOLOGIE, MODELISATION http://e2m2.universite-lyon.fr Sec : Safia AIT CHALAL Bat Atrium- UCB Lyon 1 04.72.44.83.62 Insa : S. REVERCHON Safia.ait-chalal@univ-lyon1.fr	M. Fabrice CORDEY Laboratoire de Géologie de Lyon Université Claude Bernard Lyon 1 Bât Géode – Bureau 225 43 bd du 11 novembre 1918 69622 VILLEURBANNE Cédex Tél : 04.72.44.83.74 Sylvie.reverchon-pescheux@insa-lyon.fr fabrice.cordev@univ-lyon1.fr
EDISS	INTERDISCIPLINAIRE SCIENCES-SANTE http://www.ediss-lyon.fr Sec : Safia AIT CHALAL Bat Atrium – UCB Lyon 1 04 72 44 83 62 Insa : Safia.ait-chalal@univ-lyon1.fr	Mme Emmanuelle CANET-SOULAS INSERM U1060, CarMeN lab, Univ. Lyon 1 Bâtiment IMBL 11 avenue Jean Capelle INSA de Lyon 696621 Villeurbanne Tél : 04.72.11.90.13 Emmanuelle.canet@univ-lyon1.fr
INFOMATHS	INFORMATIQUE ET MATHEMATIQUES http://infomaths.univ-lyon1.fr Sec : Renée EL MELHEM Bat Blaise Pascal 3 ^e étage infomaths@univ-lyon1.fr	Mme Sylvie CALABRETTO LIRIS – INSA de Lyon Bat Blaise Pascal 7 avenue Jean Capelle 69622 VILLEURBANNE Cedex Tél : 04.72. 43. 80. 46 Fax 04 72 43 16 87 Sylvie.calabretto@insa-lyon.fr
Matériaux	MATERIAUX DE LYON http://ed34.universite-lyon.fr Sec : M. LABOUNE PM : 71.70 –Fax : 87.12 Bat. Direction 1 ^{er} et. Ed.materiaux@insa-lyon.fr	M. Jean-Yves BUFFIERE INSA de Lyon MATEIS Bâtiment Saint Exupéry 7 avenue Jean Capelle 69621 VILLEURBANNE Cedex Tél : 04.72.43 71.70 Fax 04 72 43 85 28 Ed.materiaux@insa-lyon.fr
MEGA	MECANIQUE, ENERGETIQUE, GENIE CIVIL, ACOUSTIQUE http://mega.universite-lyon.fr Sec : M. LABOUNE PM : 71.70 –Fax : 87.12 Bat. Direction 1 ^{er} et. mega@insa-lyon.fr	M. Philippe BOISSE INSA de Lyon Laboratoire LAMCOS Bâtiment Jacquard 25 bis avenue Jean Capelle 69621 VILLEURBANNE Cedex Tél : 04.72. 43.71.70 Fax : 04 72 43 72 37 Philippe.boisse@insa-lyon.fr
ScSo	ScSo* http://recherche.univ-lyon2.fr/scso/ Sec : Viviane POLSINELLI Brigitte DUBOIS Insa : J.Y. TOUSSAINT viviane.polsinelli@univ-lyon2.fr	Mme Isabelle VON BUELTZINGLOEWEN Université Lyon 2 86 rue Pasteur 69365 LYON Cedex 07 Tél : 04.78.77.23.86 Fax : 04.37.28.04.48 isavonb@dbmail.com

*ScSo : Histoire, Géographie, Aménagement, Urbanisme, Archéologie, Science politique, Sociologie, Anthropologie

Contents

Abstract	ii
Acknowledgements	iii
Abbreviations	xvi
1 Introduction: Wake-Up Radios	1
1.1 Power Considerations in Wireless Sensor Networks	2
1.2 Wake-Up Receivers State of the Art	4
1.2.1 Architecture Considerations	4
1.2.2 Performance Analysis of Existing WU-RXs	7
1.3 Proposed WU-RX architecture	9
1.3.1 Proposed System	9
1.3.2 Low Power – Low Sensitivity WU-RX	10
1.3.3 High Sensitivity WU-RX	10
1.3.4 Interferer Resilient WU-RX	15
1.4 Tunable Filter Design	17
1.4.1 Continuous Time Analog Domain Filters	17
1.4.2 Discrete Time Charge Domain Filters	18
1.4.3 Digital Signal Processing	18
1.4.4 Continuous Time Digital Signal Processing	19
1.4.5 Tunable Filter Choice	19
1.5 Study of the Proposed CT Filtering Architecture	20
1.5.1 Single Tone Reception	20
1.5.2 Single Tone Reception with an Interferer	23
1.5.3 Single Tone Reception with Two Interferers	26
1.6 Wake-Up Receiver Design Conclusions	28
2 Continuous Time Processing Chain	29
2.1 Classification of the Different Signal Processing Domains	29
2.2 Description of the CT-DSP Chain	30

2.2.1	CT-ADC	31
2.2.2	CT-DSP	36
2.3	Co-Designing the CT-ADC with the CT-DSP	41
2.4	Proposed CT-ADC Architecture	45
2.4.1	Filtering CT-ADC Principle	46
2.4.2	Reducing the Input Event Rate	47
2.4.3	Effects on the Linearity of the Conversion	49
2.5	CT-ADC/DSP Conclusion	52
3	Energy Efficient CT-ADC	53
3.1	Previous Work	53
3.1.1	Basic Architectures	53
3.1.2	Improved Delta-Modulator Based CT-ADCs	55
3.2	Proposed CT-ADC	56
3.2.1	Improved Commutation Scheme	57
3.2.2	Proposed Architecture	58
3.2.3	Features	61
3.2.4	Possible Errors	62
3.3	Transistor-Level Implementation	66
3.3.1	Comparators	66
3.3.2	Transconductance	72
3.3.3	Threshold Management	76
3.3.4	Breakdown of the CT-ADC Power Consumption	79
3.4	Measurement Results	80
3.4.1	Single Tone Input: Noise	80
3.4.2	Two Tone Input: Linearity	84
3.5	CT-ADC Conclusion	86
4	Power Scalable CT-DSP	87
4.1	CT-DSP Architecture	87
4.1.1	Dual FIR – IIR Implementation	88
4.1.2	CT Digital Filter Design	89
4.1.3	Architecture Simulation	91
4.1.4	CT-DSP Specifications	96
4.2	CT Delay Cell	98
4.2.1	State of the Art for Asynchronous Delay Cells	98
4.2.2	Delay Cell Design	99
4.2.3	Delay Cell Architecture	102
4.2.4	Calibration and Matching	106

4.2.5	Delay Architecture Summary	109
4.3	CT Adder	110
4.3.1	Previous Work	110
4.3.2	Proposed Weighted-CT-Adder	112
4.3.3	Adder Performance	117
4.4	DF-CT-ADC	118
4.4.1	CT-ADC – CT-FIR Integration	119
4.4.2	Dispatcher	120
4.4.3	Voltage Gain and Filtering	121
4.4.4	Feedback G_m	122
4.5	Simulation/Measurements Results	123
4.5.1	DF-CT-ADC Performance	123
4.5.2	CT-FIR Performance	126
4.5.3	Interferer Rejection	128
4.5.4	Power Consumption Scaling	131
4.5.5	Noise	132
4.5.6	Comparison with State of the Art	133
4.6	CT-DSP Conclusion	136
5	Conclusion	137
5.1	Motivations and Contributions of this Work	137
5.2	Improvements of the Proposed Design	139
5.2.1	CT-ADC	139
5.2.2	DF-CT-ADC Feedback Path	141
5.2.3	CT-DSP Delay Cells	143
5.3	Future Work	143
A	Squarer Noise Analysis	145
B	The 28nm UTBB FDSOI CMOS Technology	151
	Bibliography	155

List of Figures

1.1	Iot node available power versus its lifetime when powered by different types of batteries: [2].	2
1.2	Representation of a duty cycled wireless network with synchronization packets as well as a zoomed in view of an arbitrary moment in time where node A initiates communication with node B.	3
1.3	View of a network with wake-up radios; in this representation no duty-cycling is applied.	4
1.4	Standard receiver architectures: heterodyne (a) and homodyne (b).	5
1.5	Architecture of the Uncertain IF receiver [4].	5
1.6	State of the art in frequency references: power consumption versus precision [5], [6], [7], [8], [4], [9], [10], [11], [12].	6
1.7	Spectrum of the signal at different points inside the Uncertain IF architecture.	6
1.8	State of the art low power data receivers and WU-RXs. The radio sensitivities have been normalized to a data-rate of 100kbps.	8
1.9	Top level view of the proposed WU-RX architecture.	10
1.10	Proposed WU-RX architecture in low power – low sensitivity mode.	10
1.11	Proposed WU-RX architecture in high sensitivity mode.	11
1.12	Model of the simulation bench used to characterize prospective sensitivity improvements by reducing the IF bandwidth.	12
1.13	IF signal-to-noise ratio required for a robust demodulation of the IF signal.	12
1.14	Frequency estimated by a zero crossing detector versus the SNR_{in}	13
1.15	Frequency control loop.	14
1.16	Evolution of the local oscillator instantaneous frequency for different values of the input SNR_{in}	14
1.17	Proposed WU-RX architecture.	15
1.18	View of the WU-RX IF with an input signal configuration described by SNR_{in} and SIR_{in}	20
1.19	Single tone reception scenario.	21
1.20	Baseband SNR (SNR_{BB}) versus the input SNR (SNR_{in}) for different values of SNR_{conv}	22
1.21	Required SNR_{in} versus the SNR_{conv} for achieving a BER of 1e-3.	23

1.22	The SNR constraint for the analog to digital conversion	24
1.23	Useful signal & interferer reception scenario.	24
1.24	Required SNR_{in} and SIR_{in} for different values of the conversion precision, SNR_{conv} , and CT-DSP rejection level.	25
1.25	The SFDR constraint for the analog to digital conversion	26
1.26	Useful signal with two interferers reception scenario.	27
1.27	Required SNR_{in} and SIR_{in} for different values of the conversion linearity, $SFDR_{conv}$, and CT-DSP rejection level.	27
2.1	Classification of systems according to the continuous / discrete nature of the time and amplitude [36].	30
2.2	Implementation of an N-level flash CT-ADC.	31
2.3	Input-output characteristic of a 9 level flash CT-ADC.	32
2.4	Representation of the activity dependent power dissipation of a CT-ADC [36].	32
2.5	Comparison between the spectral representation of a signal at the output of a CT-ADC and that of a sampled ADC for a single tone, sinusoid input.	34
2.6	SFDR and number of events generated per period for a CT-ADC with an input sinusoid versus the ADC number of levels, N	34
2.7	SFDR versus the number of events generated per period for a CT-ADC with an input sinusoid for different ADC number of levels, N	35
2.8	Operation principle of a CT adder (left) and of a sampled adder (right).	37
2.9	Architecture of a CT finite impulse response (FIR) filter.	38
2.10	Spectral representation of the transfer function of a low-pass FIR filter which has the values of all of its delay cells equal to $\tau = 1/F_c$	38
2.11	Architecture of a CT infinite impulse response (IIR) filter.	39
2.12	Operation principle of a continuous time delay cell.	40
2.13	Details regarding the equiripple filter design specifications for a lowpass filter implementation.	42
2.14	FIR transfer function example using a bandpass implementation.	43
2.15	FIR transfer function example using a highpass implementation.	43
2.16	FIR transfer function example using spectrum repetitions to filter out unwanted interferers.	43
2.17	Achievable CT FIR filter rejection levels versus the FIR order for the previously presented filter design methods.	44
2.18	Proposed DF-CT-ADC architecture.	46
2.19	Representation of the required feedback FIR transfer function and of the resulting signal transfer function.	47

2.20	Event rate reduction offered by the proposed solution versus the SIR_{in} for different values of A_{amp} and A_{att}	48
2.21	Frequency domain representation of the signals at the input and at the output of the DF-CT-ADC.	49
2.22	Output SIMR for a 22 level ADC with an input signal at -30dB of SIR and different configurations of the proposed DF-CT-ADC.	50
2.23	Output SIMR for different values of A_{att} and A_{amp} versus the number of CT-ADC levels for different configurations of the proposed DF-CT-ADC.	51
2.24	Output SIMR versus the average number of events triggered per signal period for different values of the CT-ADC quantization level.	52
3.1	Architecture of an N -level flash CT-ADC.	54
3.2	Architecture of a delta modulator based CT-ADC.	54
3.3	Evolution of the input, output and some key internal signals of a delta modulator based CT-ADC.	55
3.4	Architecture of a delta modulator based CT-ADC which uses fixed comparator thresholds.	56
3.5	Architecture of a delta modulator based CT-ADC which uses fixed comparator thresholds and has a fast feedback loop.	57
3.6	Evolution of V_C , the signal at the input of the comparators, in a standard delta modulator based CT-ADC.	57
3.7	Evolution of a differential version of V_C , the signal at the input of the comparators, in the proposed delta modulator based CT-ADC.	58
3.8	Front end of the proposed CT-ADC architecture.	58
3.9	Logic required to control the input switches (DAC) of the proposed CT-ADC architecture.	59
3.10	Full view of the proposed CT-ADC architecture.	60
3.11	View of the input, output and some key internal signals of the proposed CT-ADC.	60
3.12	Input and output signals as well as the voltage across the capacitors.	62
3.13	Input and output signals as well as the voltage across the capacitors for a CT-ADC with a non-zero delay in the feedback path.	63
3.14	Spectrum of the output signal containing periods when the CT-ADC goes out of bounds.	64
3.15	Proposed CT-ADC with overflow comparators.	65
3.16	Input, output and some key internal signals around an overflow event.	65
3.17	Spectrum of the CT-ADC output signal with overflow comparator correction.	66
3.18	Transistor level implementation of the continuous time comparators.	67
3.19	First order model of the comparator behavior.	69

3.20	Extracted comparator delay as well as a model for the rise and propagation delays.	69
3.21	SFDR degradation of the proposed CT-ADC architecture with a single input tone at 10MHz versus and an artificially injected threshold mismatch; the absolute value of Δ is 40mV.	70
3.22	Measured drift of the time between two consecutive output pulses for a constant DC input; this time is proportional to the instantaneous value of Δ	71
3.23	Transistor level implementation of the transconductance.	72
3.24	Voltage transfer function of the proposed $G_m - C$ implementation for different values of I_{ref}	73
3.25	Degradation of the voltage gain versus the input peak-to-peak swing at different input frequencies.	74
3.26	g_m normalized to its maximum value versus the input voltage difference for different values of the degeneration resistance.	75
3.27	Sampled PDF of the $G_m - C$ output differential DC offset over 100 MC simulations.	76
3.28	Proposed threshold setting mechanism.	77
3.29	Design of the charge pump used to set comparator thresholds.	78
3.30	Reset and control signals used for the proposed threshold setting mechanism.	79
3.31	Breakdown of the CT-ADC power consumption.	80
3.32	Spectrum of the signal observed at the output of the proposed CT-ADC for different input frequencies	81
3.33	The output SNR and SNDR versus the input tone frequency.	81
3.34	The output SNR, SNDR as well as the ADC power consumption versus the input signal amplitude (normalized to full scale).	82
3.35	The output spectrum for an input consisting of an out-of-band tone located at 60MHz; no aliasing is observed.	82
3.36	Walden and Energy figures of merit of recent state of the art ADC implementations.	84
3.37	SFDR of the proposed CT-ADC versus the input peak-to-peak swing for two values of Δ	85
3.38	SFDR of the proposed CT-ADC versus the input peak-to-peak swing for different values of the back-bias voltage.	86
4.1	Architecture of the proposed CT-DSP.	88
4.2	Example of the IIR transfer function for different values of the adder conversion gain, k_{ad}	89

4.3	Example of a CT-DSP frequency response along with a possible configuration of the IF signal.	90
4.4	Changes observed by adding mismatch to an ideal CT-FIR transfer function.	92
4.5	Effects of delay cell mismatch on the transfer function of a 9 th order FIR filter.	93
4.6	Effects of delay cell jitter on the noise floor of the CT-DSP output.	94
4.7	CT-FIR transfer functions obtained using different representations of its coefficients.	95
4.8	Effects of coefficient mismatch on the transfer function of a 9 th order FIR filter.	97
4.9	RC based, mixed-signal delay cell.	99
4.10	CMOS thyristor based, mixed signal delay cell.	99
4.11	Transistor level implementation of the proposed delay cell.	100
4.12	Time domain evolution of different delay cell signals during a delay event.	100
4.13	Value of the delay cell versus the tail current.	102
4.14	Classic delay cell architecture in the context of our FIR filter.	103
4.15	Delay tap architecture based on a parallel elementary delay arrangement.	103
4.16	Programmable parallel delay tap architecture.	107
4.17	Value of a delay tap versus the control current for different MC runs.	108
4.18	PDFs of the three parameters describing the statistical link between the value of each elementary delay and the control current.	108
4.19	Correlation between b and a (left); c and a (right).	109
4.20	Architecture of a 3 input (4 bits) Carry Save Adder.	111
4.21	Charge pump based analog adder.	112
4.22	Proposed elementary adder cell.	113
4.23	Proposed full adder cell.	115
4.24	Scenario used to illustrate the operation principle of the proposed weighted adder.	115
4.25	Elementary adder with its coefficient equal to 10 and with input V_+ active (at V_{dd}).	116
4.26	Elementary adder with its coefficient equal to 01 and with input V_+ active (at V_{dd}).	116
4.27	Elementary adder with its coefficient equal to 00 and with input V_+ active (at V_{dd}).	116
4.28	Lumped representation of the proposed CT adder	117
4.29	View of the complete DF-CT-ADC implementation.	119
4.30	Schematic of the dispatching circuit used to split the CT-ADC output in 5 parallel streams.	120

4.31	Schematics of the active voltage gain stage used in the feedback loop of the DF-CT-ADC.	121
4.32	Voltage transfer function of the G_V block.	122
4.33	Schematics of the feedback G_m cell.	122
4.34	Highpass, lowpass and bandpass configurations for the DF-CT-ADC transfer function.	124
4.35	Tuning the center frequency of the DF-CT-ADC signal transfer function.	125
4.36	DF-CT-ADC transfer functions for different configurations of the feedback loop voltage gain.	126
4.37	Highpass, lowpass and bandpass configurations for the CT-FIR transfer function.	127
4.38	Tuning the central frequency of the CT-FIR filter transfer function.	128
4.39	Input scenario used to test the performance of the proposed CT-DSP.	128
4.40	Spectrum of the signal at various points in the proposed system (DF-CT-ADC & CT-FIR).	129
4.41	Spectrum of the signal at various points in the proposed system with no feedback loop around the ADC (CT-ADC & CT-FIR).	130
4.42	Breakdown of the complete DF-CT-ADC-DSP power consumption.	131
4.43	CT-DF-ADC and CT-FIR power consumption versus the input amplitude (system event rate). Note that the power consumption of the CT-ADC is included in both previously defined systems.	132
4.44	Output spectrum of a full-scale, 50MHz single tone input, for a transient noise simulation of the entire system.	133
5.1	Concept of the improved CT-ADC design.	140
5.2	Switchable, current mode FIR adder cell.	142
5.3	DF-CT-ADC architecture using the proposed current mode adder.	142
A.1	Scenario used to study the noise performance of a squarer.	145
A.2	Baseband noise generated by the cross-mixing of the IF noise with the signal component.	147
A.3	Baseband noise generated by the self-mixing of the bandlimited IF noise.	148
B.1	Cross sectional view of an UTBB FDSOI transistor [89].	151
B.2	Sectional view of an FDSOI PMOS transistor next to an FDSOI NMOS transistor, along with the allowed backgate bias voltages.	152
B.3	Threshold voltage variation of FDSOI and standard bulk transistors.	152

List of Tables

1.1	WU-RX state of the art. Note that the sensitivities and data-rates are given as reported in the original publications and have not been normalized to 100kbps.	8
1.2	LO precision for different values of the SNR_{in}	14
1.3	Comparison of different tunable filter solutions and the components they require.	19
1.4	Performance limits of the IF interferer rejection stage according to the specifications of the CT-ADC and the CT-DSP.	28
2.1	FIR and IIR filter comparison	39
2.2	Comparison of existing CT-ADC/DSP implementations with the requirements of our application.	44
2.3	Estimation of the power requirements of our CT-ADC/DSP implementation based on results from literature.	45
2.4	CT-ADC operating point for different configurations of the DF-CT-ADC.	52
3.1	Sizing of the comparator components.	67
3.2	Sizing of the $G_m - C$ components.	72
3.3	Sizes of various components used in the design of the charge pump.	79
3.4	Comparison with state of the art discrete time ADCs with bandwidths smaller than 100MHz.	83
3.5	Comparison with existing continuous time ADCs.	83
3.6	Speed of the logic block for different values of the back-bias voltage.	85
4.1	CT-DSP specifications.	96
4.2	Sizing of the elementary delay cell components.	101
4.3	Comparison of the proposed elementary delay cell with previous works.	102
4.4	RMS jitter of a delay tap with different configurations (N_S and N_P).	106
4.5	Performance of the proposed delay tap calibration method.	109
4.6	Summary of the delay cell architecture and its performance.	110
4.7	Component sizes of an elementary adder cell.	114
4.8	Energy required by an elementary adder cell to process a single CT-ADC token.	117

4.9	Sizing of the components used for the dispatcher.	120
4.10	Sizing of the components used in the active voltage gain stage.	121
4.11	Sizing of various components used in the design of the feedback G_m cell.	123
4.12	Coefficients which demonstrate the transfer function reconfigurability of the DF-CT-ADC.	124
4.13	Simulated IIR feedback configurations.	125
4.14	Rejection performance achieved by the proposed DF-CT-ADC.	126
4.15	Highpass, lowpass and bandpass CT-FIR configurations.	127
4.16	Power breakdown of the proposed CT-ADC-DSP system with and without DF-CT-ADC feedback.	131
4.17	Comparison between the proposed DF-CT-ADC-DSP and other, state of the art, CT-DSP implementations.	134
4.18	Comparison between the proposed DF-CT-ADC and existing analog or digital IIR filter implementations.	134
4.19	Comparison between the proposed CT-FIR and existing analog or digital FIR filter implementations.	135

Abbreviations

ACK	ACK nowledge
ADC	A nalog to D igital C onverter
AGC	A utomatic G ain C ontrol
BB	B ase B and
CT-ADC	C ontinuous - T ime A nalog to D igital C onverter
CT-DSP	C ontinuous - T ime D igital S ignal P rocessor
DF-CT-ADC	D igitally F iltering - C ontinuous T ime - A nalog to D igital C onverter
DFF	D ata F lip F lop
DSP	D igital S ignal P rocessing
F2V	F requency T O V oltage
FDSOI	F ully D epleted S ilicon O n I nsulator
FIR	F inite I mpulse R esponse
FoM	F igure of M erit
IF	I ntermediate F requency
IM	I nter M odulation
ISM	I ndustrial S cientific and M edical
IoT	I nternet of T hings
LNA	L ow N oise A mplifier
LO	L ocal O scillator
LPF	L ow P ass F ilter
MAC	M edium A ccess C ontrol
MC	M onte C arlo
NF	N oise F igure
OA	O perational A mplifier
OOK	O N O FF K eying

OTA	O perational T ransconductance A mplifier
RF	R adio F requency
RTC	R ead T ime C lock
RX	R eceiver
SFDR	S purious F ree D ynamic R ange
SIMR	S ignal to I nter M odulation R atio
SIR	S ignal to I nterferer R atio
SNR	S ignal to N oise R atio
TX	T ransmitter
UTBB	U ltra T hin B ody and B uried oxide
WU-RX	W ake U p R adio

Chapter 1

Introduction: Wake-Up Radios

THE Internet of Things (IoT) represents the network which results by extending the Internet Protocol beyond its initial goal of enabling worldwide machine-to-machine communications. A “thing” in the IoT can be any object which has the possibility of sending data over a network. The applications include, but are not limited to, environment monitoring, infrastructure management, manufacturing, energy management, healthcare systems and transportation. According to some technology specialists the IoT is predicted to be the third big industrial revolution which occurred in the past 50 years, after the invention of computers and the Internet. However, due to the projected size of the IoT - between 50 and 100 trillion objects [1] - a key factor in its success is minimizing the implementation and operation costs. The cost of an IoT node can be lowered by increasing the level of integration of its electronic circuits: solutions with a minimum of off-chip components are thus preferred. Second, operation costs can be reduced by minimizing the maintenance required by an IoT network. The *things* should be capable of operating autonomously for extended periods of time without human intervention. Consequently, a key challenge in the development of the IoT is designing low power communication systems thus enabling a long *thing* lifetime on nomad power sources such as batteries or energy harvesters.

The next part of this chapter introduces the energy constraints related to IoT nodes based on different forms of energy harvesting. It is showed that, in order to achieve the targeted IoT node lifetimes, the power requirements of IoT circuits need to be drastically reduced. Therefore, a new concept in receiver design is presented, the Wake-Up Receiver (WU-RX), which thanks to its low power consumption, greatly reduces the total power requirements of IoT transceivers. After an analysis of existing state of the art designs, a new WU-RX architecture is proposed and analyzed. Finally, in the last part of this chapter, the focus shifts to the most critical block of the proposed WU-RX architecture,

the tunable filter. A set of specifications is derived, which further serves as a foundation for its design which is discussed in the rest of this manuscript.

1.1 Power Considerations in Wireless Sensor Networks

The power budget of an IoT node is determined by the energy sources it has access to. Figure 1.1 plots the amount of power available for a node versus its required lifetime for different types of batteries powering it: [2]. By targeting lifetimes of several years, it can be seen that the maximum power budget is around several tens of μW , around two orders of magnitude below the power consumption of the state of the art in low power receivers - 4mW [3]. Currently, the preferred solution for overcoming this power budget gap is by employing receiver duty cycling which consists in only turning ON the receiver for very short periods of time thus saving power.

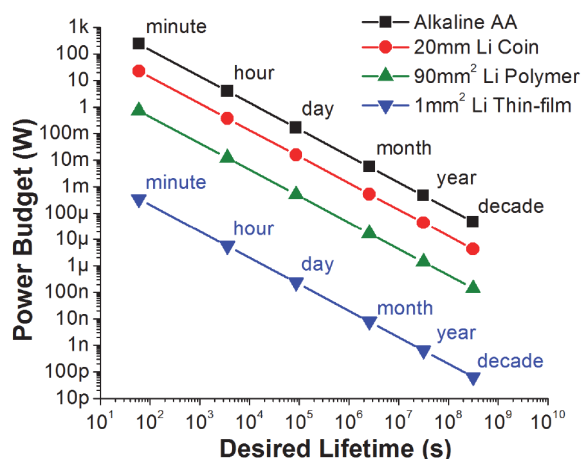


FIGURE 1.1: Iot node available power versus its lifetime when powered by different types of batteries: [2].

Before discussing alternative solutions which lower the power consumption of duty cycled receivers, it is important to understand the main trade-offs related to the design of a classic duty-cycled network (Figure 1.2). One of the basic hypothesis for efficient duty cycling is the fact that the network is perfectly synchronized: if the node A wants to emit a packet to node B, it is supposed that node A has perfect knowledge of the moments when node B is online. This synchronization problem is usually solved at the network level by sending sync packets which reset the internal timer of each node. The maximum time between the synchronization instants is set by the drift of each node's internal time reference which, in turn, is inversely proportional to the amount of power the time reference consumes. The network synchronization mechanism as well as a zoomed in

view of an arbitrary moment in time when node A initiates a communication with node B are represented in Figure 1.2.

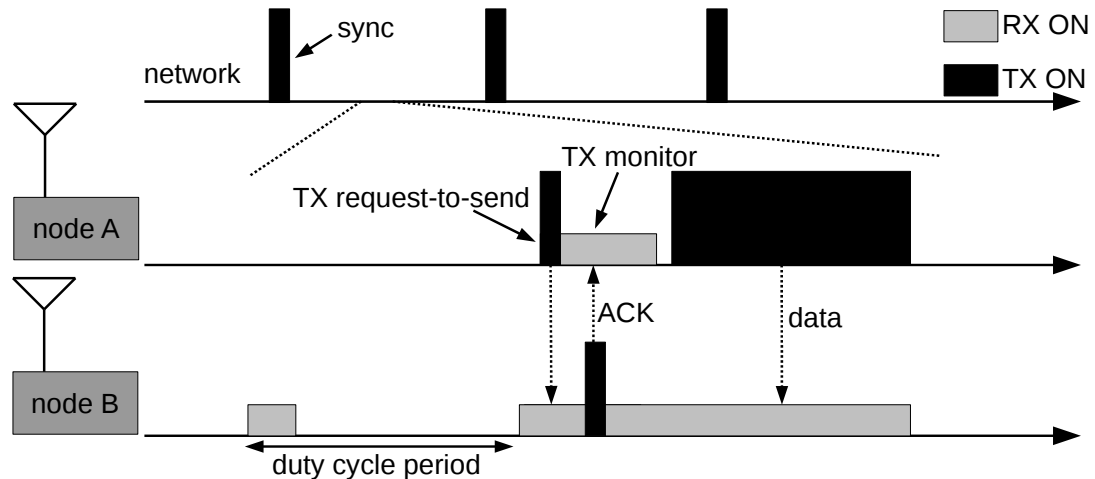


FIGURE 1.2: Representation of a duty cycled wireless network with synchronization packets as well as a zoomed in view of an arbitrary moment in time where node A initiates communication with node B.

This solution works well for networks where the average time between two consecutive exchanges for any node is much smaller than the sync period, as the synchronization adds very little overhead. However, this is not the case for IoT networks: nomad nodes are expected to exchange a very small amount of data, with an average time between consecutive communications in the order of tens of seconds, or even minutes. Consequently, due to synchronization issues, there will be scenarios where duty-cycled IoT networks will exchange more sync packets rather than actual data packets, drastically increasing the synchronization overhead with respect to the actual communication.

Another drawback of duty-cycled networks is linked to the propagation delay of a packet through the network. In order to meet the restrictive power budgets of IoT nodes, imposed by the availability of nomad power sources, existing low power receivers need to be heavily duty-cycled, such that their *average* power consumption becomes sufficiently low. This has the negative effect of increasing the propagation time of the information through the network. This situation becomes even worse for mesh networks, where the worst case node-to-node propagation delay is multiplied by the number of hops between the source of the information and its destination.

Alternatively, a new paradigm for IoT receiver and network design can potentially solve the previously presented problems. The technique consists of adding a supplementary receiver to each IoT node, called henceforth wake-up receiver (WU-RX). The WU-RX monitors the communication channel and only switches ON the main receiver (also referred to as the data receiver) when it receives an actual communication request. The

advantage of such a scheme lies in the fact that the two receivers can now be optimized separately. The WU-RX is switched ON for relatively long periods of time and spends very little of its time actually receiving data, it can thus be optimized for low power consumption while it is in monitoring mode. On the other hand, the data receiver will spend almost the entirety of its ON time actually receiving information and thus can be optimized for a small energy consumption per received bit.

The previously described WU-RX scheme can be implemented either in a completely asynchronous network or in a duty cycled network. Figure 1.3 showcases the asynchronous WU-RX principle: the receiving node B keeps the WU-RX turned ON all the time and is ready to react whenever a TX-request-to-send packet is received. This solution completely obviates the need for real time clocks (RTCs) thus allowing for a very flexible implementation of the network. The disadvantage of this approach is the fact that the budget for the WU-RX is very low making its design potentially troublesome. Alternatively, the WU-RX itself can be duty cycled opening up new possibilities concerning the trade-off between the power consumed by the RTCs (impacting its drift) and the power spent monitoring to the channel. In conclusion, the WU-RX principle represents a substantial expansion of the design space of low power radios, potentially enabling the implementation of more energy efficient trade-offs.




FIGURE 1.3: View of a network with wake-up radios; in this representation no duty-cycling is applied.

1.2 Wake-Up Receivers State of the Art

1.2.1 Architecture Considerations

Choices concerning the design of modern WU-RX can be explained by analyzing the two classic and ubiquitous receiver topologies: the heterodyne and homodyne receiver, which are portrayed in the Figure 1.4(a) and Figure 1.4(b). The key in achieving a low power consumption is the removal of the most power consuming blocks: the RF gain stage and

the frequency synthesis in the heterodyne receiver and the large RF gain stage in the homodyne receiver.




FIGURE 1.4: Standard receiver architectures: heterodyne (a) and homodyne (b).

A true, low power receiver can be implemented by removing the most power consuming blocks from the previously presented architectures and by combining them. What results is an architecture known as the “Uncertain IF” receiver [4], shown in Figure 1.5. The frequency synthesis is completely removed from the heterodyne architecture and replaced by a simple free running LO. This enables a downconversion to an uncertain intermediate frequency which has a frequency precision defined by that of the local oscillator. At this point, it is possible to amplify the signal with a rather low power consumption as the operating frequency is expected to be lower than 100MHz. Finally, since the exact frequency of the signal is unknown, the demodulation is done through energy detection rather than by using a mixer, as in the case of the original homodyne receiver.




FIGURE 1.5: Architecture of the Uncertain IF receiver [4].

Relaxing the frequency precision requirement of the local frequency reference enables a substantial reduction in its power budget. This trade-off is showcased by comparing different state of the art frequency references: Figure 1.6. High precision systems with a frequency uncertainty of smaller than 100ppm require a power of at least several mW - this is equivalent to a frequency precision of 240kHz for an operation in the ISM band at 2.4GHz. On the other hand, ultra low power frequency references, usually implemented

as free running oscillators, achieve power consumptions which fit in the WU-RX budget but are limited in precision: around 1% which corresponds to 24MHz for an oscillating frequency of 2.4GHz.

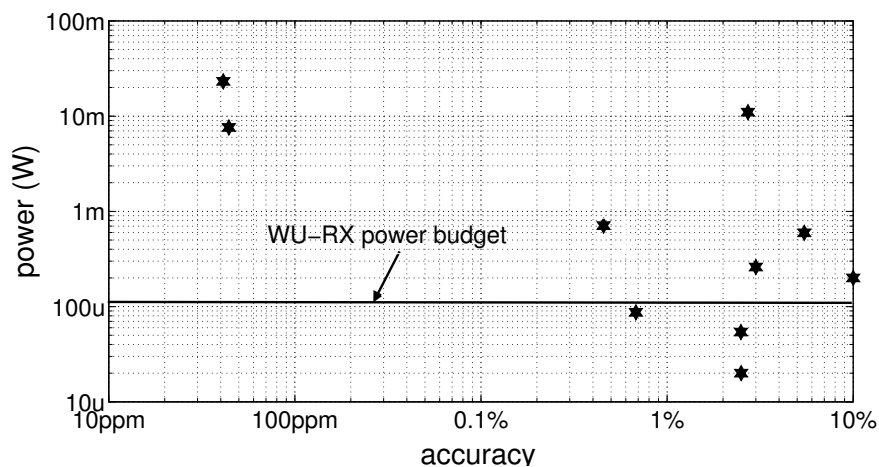


FIGURE 1.6: State of the art in frequency references: power consumption versus precision [5], [6], [7], [8], [4], [9], [10], [11], [12].

The IF filter must be designed with a transfer function wide enough to accommodate the imprecision linked to the free running oscillator. Low power frequency references for the 2.4GHz band thus require the use of a very wide IF filter passband (up to 24MHz), much wider than that required by the WU-RX communication speed (which is usually limited to 100kbps). Consequently, the IF filter can only be used to minimize the noise bandwidth before the demodulation stage and cannot be used for efficient interferer rejection. This situation is illustrated in Figure 1.7: interferers located inside the RF filter passband cannot usually be filtered by the subsequent IF filter stage, as the latter's passband is expected to be widest. These interferers are then downconverted to DC by the energy detector thereby irretrievably corrupting the baseband signal. Consequently, it will be seen that most WU-RX implementations have a very limited interferer resilience.

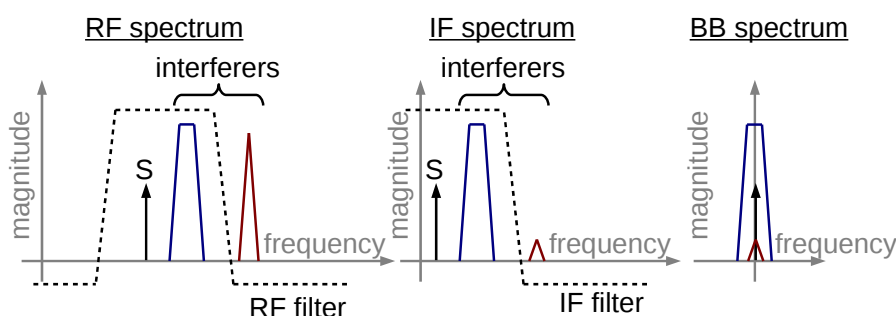


FIGURE 1.7: Spectrum of the signal at different points inside the Uncertain IF architecture.

Another drawback of the Uncertain IF architecture stems from the removal of the front-end high frequency gain stage. The result is a low power radio with a degraded noise figure (NF). Moreover, the demodulator is most of the time implemented using the non-linearity of a diode, in which case, it can be shown [4] that its conversion gain (k_{conv}) is proportional to the voltage swing presented at its input: $k_{conv} = V_{pp}/(4nV_t)$, where V_{pp} is the demodulator input peak-to-peak swing, n – subthreshold factor and V_t – the thermal voltage. Reducing the RF gain results in a lower amplitude signal at the input of the demodulator which decreases its gain thus also degrading its noise figure. To combat this, several strategies have been adopted in literature. As a possible solution, active noise cancellation [13] can be combined with RF envelope detection in order to reduce the noise contribution of the squarer. This solution is however limited by the fact that the envelope detector needs to be used at RF thus requiring it to function at a very low input swing due to the lack of an energy efficient way of providing gain. Another solution, proposed in [14], uses baseband correlators along with information redundancy to improve the effective sensitivity of the receiver at the cost of a reduced data-rate. However, such a scheme can be used at the output of any WU-RX implementation. To provide a fair comparison we are going to focus on the raw bit error rate measured at the output of the demodulator, which is entirely defined by the signal-to-noise ratio of the signal at that point.

1.2.2 Performance Analysis of Existing WU-RXs

The mode of operation described in the introduction of this chapter demands the WU-RX to have a sensitivity as high as that of the data receiver. Any difference in the sensitivities of the two receivers is essentially wasted by either waking up on incoming communication requests which are then too faint for establishing a data transfer (WU-RX is more sensitive) or by missing TX-request-to-send packets which could otherwise be treated by the data receiver (data receiver is more sensitive). Similarly, the performance of the two radios must be identical in the presence of interferers.

The performance of the existing state of the art WU-RXs as well as state of the art data receivers is presented in Figure 1.8 (for WU-RXs, we have normalized the reported sensitivities to a data-rate of 100kbps), a comparison table is also given in Table 1.1. A careful analysis of the data provided in Figure 1.8 shows that next generation WU-RXs should have an improved sensitivity in order to match that of data receivers. Furthermore, as it is shown in Table 1.1, a key feature missing from existing WU-RXs is the robustness to interferers; the interferer rejection achieved currently by WU-RXs is far from that required by modern data receivers.

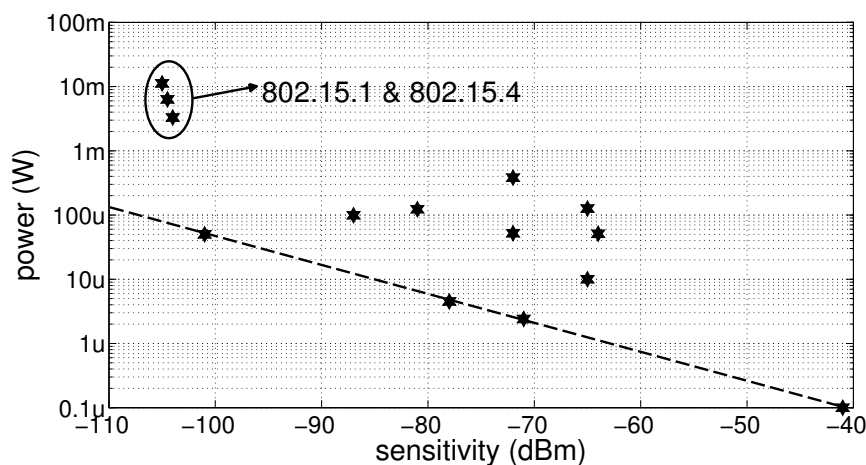


FIGURE 1.8: State of the art low power data receivers and WU-RXs. The radio sensitivities have been normalized to a data-rate of 100kbps.

TABLE 1.1: WU-RX state of the art. Note that the sensitivities and data-rates are given as reported in the original publications and have not been normalized to 100kbps.

reference	frequency	sensitivity	interferer rejection	power	data-rate	process	modulation
[4]	2GHz	-72dBm	n/a	52 μ W	100kbps	90nm	OOK
[15]	915MHz	-83dBm	10dB	121 μ W	10kbps	90nm	OOK
[14]	868MHz	-72dBm	n/a	2.4 μ W	20kbps	130nm	OOK
[16]	2.45GHz	-88dBm	n/a	50 μ W	250kbps	65nm	OOK
[17]	816MHz	-71dBm	< 10dB	382 μ W	125kbps	40nm	FSK
[18]	2.4GHz	-97dBm	22dB	99 μ W	10kbps	65nm	OOK
802.15.1[19]	2.4GHz	-70dBm	27dB @3MHz	n/a	1Mbps	n/a	GFSK
802.15.4[20]	2.4GHz	-85dBm	31dB @10MHz	n/a	250kbps	n/a	QPSK

Moreover, depending on the radiated output power, the ISM band may require transmitters to use spread spectrum techniques such as frequency hopping. Consequently, the WU-RX should also be designed to match these multi-channel requirements in order to avoid limiting choices related to the design of the transmitters.

We can thus define 3 development directions for the improvement of existing WU-RXs:

- power: reducing the power consumption will always improve the quality of the resulting networks.
- sensitivity: despite existing of WU-RX implementations with sensitivities close to those required by low power communication standards, improvements of this performance parameter can increase the communication range of IoT nodes.
- interferer robustness & multichannel capabilities: current WU-RXs have a limited multichannel operation and usually suffer from a very poor interferer rejection.

In conclusion, the current state of the art shows that existing WU-RX implementations severely lack in interferer rejection. As a consequence, this work is focused on the description and implementation of a block designed to improve this performance metric as well as to enable multichannel operation. The proposed system is compatible with most of the existing WU-RX implementations; however, it is important to define a precise setting for our application. Thus, in the next part of this chapter, we propose a new WU-RX architecture which attempts to solve some of the problems presented earlier and is compatible with the system described in the following chapters of this manuscript. Results from a set of architecture simulations of the proposed WU-RX will serve as specifications and as starting point for the design of the proposed IF interferer rejection block.

1.3 Proposed WU-RX architecture

From the previous discussion we have seen that there is a clear trade-off between the sensitivity of a WU-RX and its power consumption: a design achieving around -100dBm requires a power of $50\mu\text{W}$ while a design achieving -40dBm requires only $0.1\mu\text{W}$, corresponding to a relative power reduction of more than two decades. Systems are usually designed to always operate in “worst-case” scenarios and do not scale their power consumption according to the strength of the input signal. This means that the $50\mu\text{W}$ design will always consume the same amount of power regardless of the strength of the input signal (even if it is at -40dBm). We can thus conclude, that a true, low-power WU-RX implementation needs to be capable of scaling its power consumption according to environment conditions such as received signal strength and presence (or lack) of strong interferers.

We propose a multi-stage WU-RX architecture conceived to achieve a scalable power consumption – sensitivity trade-off. Its design, presented in the next part of this chapter, is based on three different operation modes: an ultra-low power, low sensitivity mode; a high sensitivity mode; and finally, an interferer resistant mode. The following architecture analysis is only used to determine the specifications of the tunable filtering system required to enable efficient interferer rejection. The implementation of the latter will then be discussed for the rest of this manuscript.

1.3.1 Proposed System

To achieve a scalable performance level along with a scalable power consumption we propose a WU-RX architecture in which the most power consuming blocks can be

switched ON and OFF according to the current environment conditions. A top level view of the proposed system is presented in Figure 1.9.




FIGURE 1.9: Top level view of the proposed WU-RX architecture.

In the following paragraphs we discuss different WU-RX configurations where we estimate their power consumption as well as their sensitivity and interferer resilience. The control of the architecture configuration can be handled at a higher abstraction level and is beyond the scope of this manuscript.

1.3.2 Low Power – Low Sensitivity WU-RX

The low power – low sensitivity mode is activated only when environment conditions are extremely favourable (strong received signal with no interferers). All intermediate frequency blocks are deactivated, as it consists of a simple energy detector placed at the output of an uncertain receiver front end, as shown in Figure 1.10. The IF bandwidth depends on the precision of the free running oscillator frequency, which is determined by its implementation. No IF amplification is used to boost the signal, as we assume it is sufficiently strong for a correct demodulation.

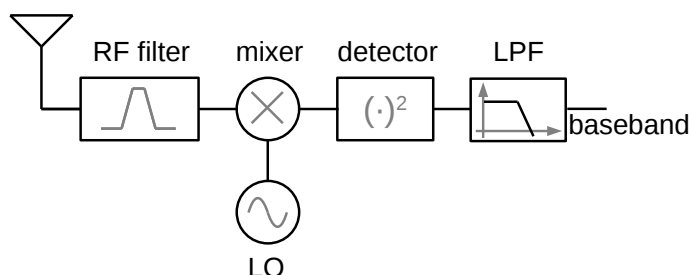


FIGURE 1.10: Proposed WU-RX architecture in low power – low sensitivity mode.

1.3.3 High Sensitivity WU-RX

In the high sensitivity WU-RX mode, all IF blocks of the original architecture, except the tunable IF filter, are switched ON, as seen in Figure 1.11.

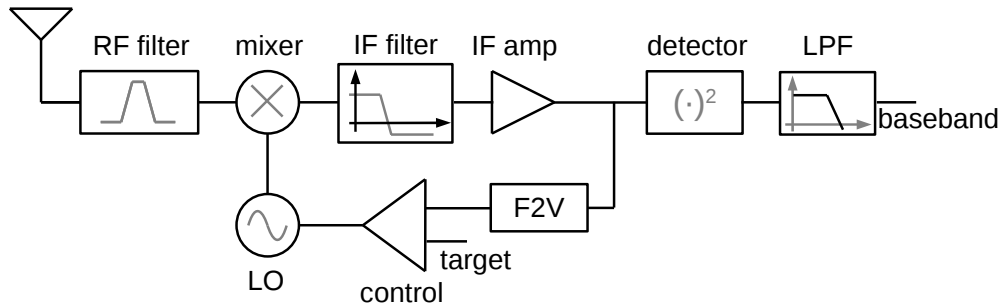


FIGURE 1.11: Proposed WU-RX architecture in high sensitivity mode.

With respect to the previous mode, the sensitivity is increased by the following mechanisms which are described below: increased direct path IF gain and a more narrow IF bandwidth.

Increasing the IF Gain

As shown in [21], boosting the power of the IF signal increases the sensitivity of the demodulator thereby improving the sensitivity of the WU-RX. In fact, the demodulator, which is usually a squarer block, has a conversion gain which depends on the input signal amplitude: increasing the signal amplitude increases the conversion gain which improves the receiver noise figure. Based on results from literature, we estimate the sensitivity of such a system can be improved from -40dBm to beyond -70dBm for a total power consumption under $50\mu\text{W}$ ([22] and [4]).

Decreasing the IF Bandwidth

The sensitivity of the previously presented WU-RX architecture is limited by the wide IF bandwidth which can be as high as 100MHz [4] while the useful IF signal has a bandwidth of only several hundreds of kHz (imposed by the communication data-rate). The demodulator, implemented as an energy detector, then integrates its input noise over the entire bandwidth and downconverts it to baseband. The sensitivity can thus be improved by limiting the width of the IF stage passband, but this requires an improved precision of the local oscillator frequency. To illustrate prospective sensitivity improvements achieved by reducing the IF bandwidth, a set of simulations is done, in which we determine the minimum IF input signal-to-noise ratio (SNR_{in}) for a given bandwidth BW_{IF} which results in 12dB of baseband SNR (SNR_{BB}), corresponding to a bit error rate (BER) of $1e-3$. A representation of the simulation test bench is given in Figure 1.12; note that the first spectral representation of the signal is used only to normalize the SNR to a constant bandwidth and is obtained by extending the IF noise power from BW_{IF} to 100MHz – it never actually occurs in the proposed system.




FIGURE 1.12: Model of the simulation bench used to characterize prospective sensitivity improvements by reducing the IF bandwidth.

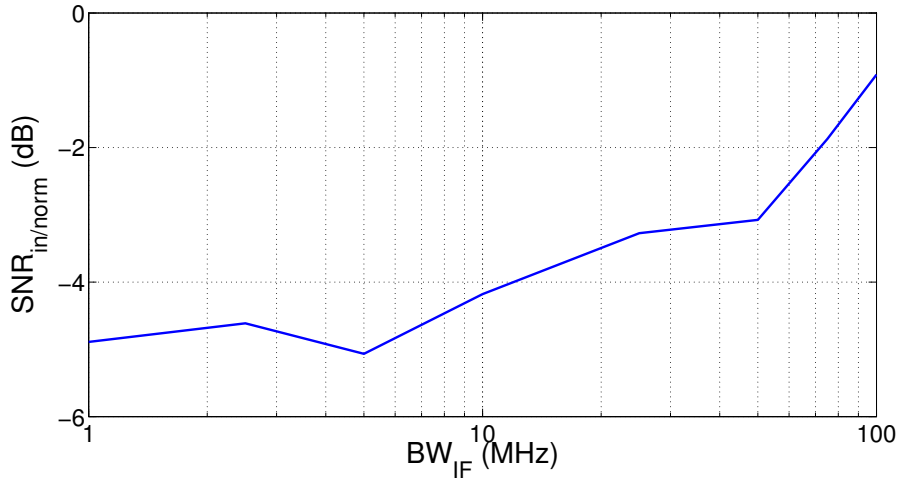


FIGURE 1.13: IF signal-to-noise ratio required for a robust demodulation of the IF signal.

Results are plotted in Figure 1.13; the SNR_{in} has been computed by normalizing the integrated noise to a bandwidth of 100MHz, according to equation 1.1. It can be seen that by reducing the bandwidth from 100MHz to about 10MHz improves the sensitivity of the receiver by about 3dB.

$$SNR_{in/norm} = SNR_{in} - 10 \log_{10} \frac{BW_{IF}}{100MHz} \quad (1.1)$$

We now propose a new scheme aimed at improving the local oscillator precision by constructing a feedback loop which senses the frequency of the IF signal (through the use of the F2V – frequency to voltage – block) and controls the local oscillator, as showed initially in Figure 1.11 on the preceding page. At this point, we operate under the assumption that the IF signal does not contain any interferers. Compared to a standard PLL, the proposed LO control scheme does not employ any blocks running at RF other than the local oscillator itself. Consequently, its power requirements are not expected to increase dramatically [23].

A brief overview of the prospective performance of the proposed system is presented next.

The proposed frequency locking mechanism achieves a precision which is limited by the estimation of the central frequency at IF level. At lower input signal levels, corresponding to a lower SNR_{in} , a bias in the estimation of the signal frequency is induced by the IF noise. To illustrate this we consider a scenario where the useful signal is situated at 100MHz, the IF bandwidth is [0Hz–100MHz] and the input SNR is defined as SNR_{in} . As an estimator, we use the average 0 crossing frequency computed over an interval ranging from $0.5\mu s$ to $10\mu s$; the estimated frequency is thus given by equation 1.2 with N – the number of 0 crossings occurring in the time interval ΔT . Results are plotted in Figure 1.14: the solid lines represent the “average” result while the error bars represent the most extreme cases occurring over the course of 100 MC (Monte Carlo) simulations.

$$E(f) = \frac{N - 1}{2 \cdot \Delta T} \quad (1.2)$$

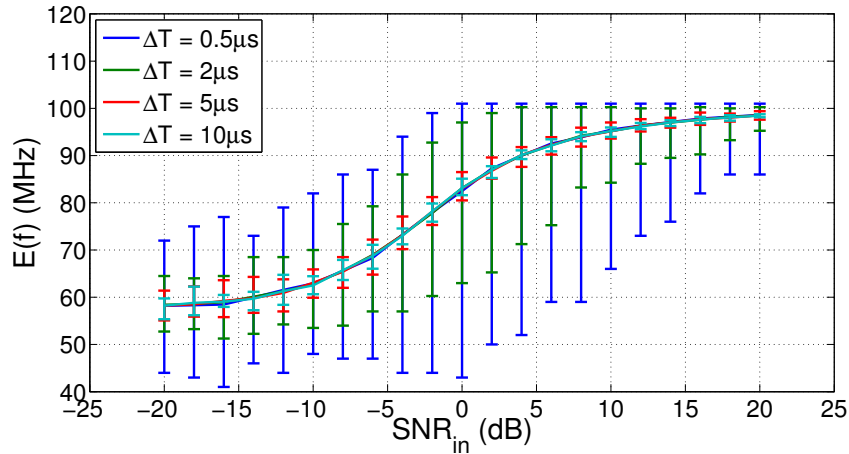


FIGURE 1.14: Frequency estimated by a zero crossing detector versus the SNR_{in} .

It can be seen that, for low signal strengths, the estimated frequency of the IF signal is very close to the middle of the IF noise bandwidth, 50MHz rather than to the frequency of the useful signal (100MHz). Moreover, the random nature of the IF noise increases the standard deviation observed between two consecutive estimations of the central frequency for a given estimation window. Having constructed the central frequency estimator, we now inject its output into a proportional – derivative (PD) controller which is used to tune the local VCO frequency. A more detailed view of the frequency control loop is given in Figure 1.15. Note that $V_{F2V}(f)$ represents the voltage output of the F2V block when presented a signal of frequency f .

The frequency precision achieved by the proposed LO calibration loop is assessed through a series of simulations in which, initially, the useful signal is located at 100MHz and a 50MHz target is given to the control loop (f_{target}). The frequency estimation is done over





FIGURE 1.15: Frequency control loop.

a duration $\Delta T = 10\mu s$. The evolution of the instantaneous frequency of the oscillator over a duration of $250\mu s$ is plotted in Figure 1.16; as previously, the solid lines represent the evolution of the “average” result, while the error bars represent the most extreme cases occurring over the course of 100 MC simulations.

FIGURE 1.16: Evolution of the local oscillator instantaneous frequency for different values of the input SNR_{in} .

The coefficients chosen for the control loop have been optimized for an SNR_{in} of -10dB ; the under-damped behavior of the control loop for other values of the SNR_{in} can be easily corrected by adjusting the corrector coefficients. In steady state, depending on the value of SNR_{in} , the frequency precision of the local oscillator is reduced from an initial value of 100MHz to values given in Table 1.2.

TABLE 1.2: LO precision for different values of the SNR_{in}

SNR_{in}	ΔF_{LO}
0dB	7MHz
-2dB	10MHz
-5dB	15MHz
-10dB	28MHz

Comparing results from Table 1.2 with those from Figure 1.13 on page 12, we conclude that the proposed frequency calibration mechanism operates correctly at around -5dB of SNR_{in} which reduces the IF uncertainty to about 15MHz thereby improving the sensitivity of the WU-RX by 3dB with respect to the initial configuration (BW of 100MHz). Despite the modest improvement in the sensitivity, this result is very useful as it drastically limits the IF bandwidth which relaxes the requirements of the interferer rejection stage.

1.3.4 Interferer Resilient WU-RX

In this operation mode, all blocks of the initially proposed architecture are switched ON. For completeness, a view of the resulting architecture is given in Figure 1.17. Interferer rejection can only be achieved using filters which have bandwidths similar to that of the useful signal – up to 1MHz . This cannot be done at RF, since it would demand a filter with a quality factor which is too large. On the other hand, descending the signal to IF would solve the filter quality issue but would introduce another problem related to the frequency uncertainty of the useful signal. To solve this, we propose the use of a tunable IF filtering stage which is used to scan the entire IF band and attempt a demodulation at each frequency point in the search for the wake-up signal.

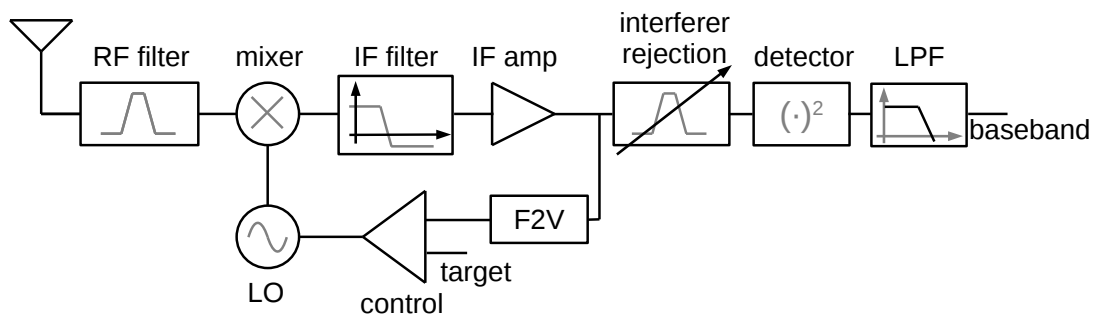


FIGURE 1.17: Proposed WU-RX architecture.

The total bandwidth which requires scanning is defined by the precision achieved by the proposed signal tracking loop. Contrary to the previous situation, here we study a scenario in which the input signal is present alongside much stronger interferers, making the tracking loop lock onto one of the interferers rather than the useful signal. Considering that the total bandwidth containing signals at IF is 10MHz – equal to that of the RF filter – we can conclude that the useful signal will always be at most 10MHz away from the frequency the tracking loop locks onto. This gives us a total IF bandwidth of $2 \cdot 10\text{MHz} + 15\text{MHz} = 35\text{MHz}$, corresponding to the two extreme cases when strong interferers are located at either ends of the RF filter passband plus the absolute uncertainty achieved by the frequency tracking loop presented previously.

In conclusion, assuming an implementation of a WU-RX with a data-rate of 100kbps, a tunable IF filter needs to be designed, with a passband equal to several times the data-rate, 1MHz for instance, which can be shifted over the 35MHz of effective IF. Depending on the communication scenario and on the constraints presented by the application, several search algorithms are possible:

- *Linear search:* The central frequency of a 1MHz bandpass filter is swept across the entire 35MHz frequency range; demodulation by energy detection is attempted at every frequency point. Despite being slow, this strategy guarantees the best interferer rejection for a given filter transfer function as it guarantees that, at some point in time, the useful signal will be situated inside the filter passband with the interferer located in the stopband.
- *Binary search:* The useful bandwidth is sequentially split in two parts, using complementary lowpass / highpass filters; energy detection and demodulation are then attempted on both the resulting half-spectrums. If the signal is successfully separated from the interferer, the demodulator can intercept the resulting wake-up signal and trigger the wake-up function of the main receiver. On the other hand, if the interferer is still located in the same half spectrum as the useful signal, the filters are reconfigured to enable the analysis of the resulting two quarter-spectrum. This operation is repeated until a successful separation of the signal from the interferer has been achieved.

The advantage of this solution lies in its speed compared to the that of the linear search. However, more complex filtering transfer functions are necessary, since dynamic reconfiguration from lowpass to highpass is required by the proposed search algorithm.

- *Custom search:* Depending on the communication scenario and the statistical properties of the signal at IF, a custom search algorithm can be designed to optimize either response time or power consumption.

The optimization of the previously described search problem is dependent on the communication scenario and is usually handled at a higher abstraction layer such as the Medium Access Control (MAC) level. In order to provide a generic solution which can be adapted to any sensor network, the IF filter is required to be perfectly tunable in terms of transfer function type and achieved bandwidth. Finally, it is important to note that the previously presented interferer resistant radio back-end has, by design, multi-channel capabilities. Simply reprogramming the transfer function of the IF filter along with changing the target of the frequency tracking loop allows us to cover the entire ISM bandwidth of 80MHz.

In the next part of this chapter we discuss the existing possibilities for the implementation of the previously described filtering stage.

1.4 Tunable Filter Design

Given the design constraints of the previously presented IF filtering stage (very low power budget and high tunability), we show that classic solutions cannot be employed.

1.4.1 Continuous Time Analog Domain Filters

This class of filters usually relies on active gain elements either in the form of operational amplifiers (OAs), operational transconductance amplifiers (OTAs) or transconductance cells (G_m). Lowpass transfer functions are constructed using capacitors and resistors, while low frequency zeros, required for the highpass or bandpass designs are implemented using active circuits mimicking the behavior of inductances [24] (which are usually impractical for below GHz frequencies). $G_m - C$ filters are preferred for low power operation due to their open loop implementation [25], [26], but suffer from linearity issues; on the other hand OAs are preferred for low frequency, high precision applications, however, the closed loop nature of OAs make them consume more power [27].

The tunable element can be either the capacitor (switched capacitor bank or a varicap), resistor (transistor in triode region or a switched resistor bank) or gain element [28], such as a transconductance (by varying the bias current or adding supplementary elements in parallel [24]). The most popular methods currently employed are either relying on active resistors, thanks to their high tuning range and small silicon footprint, or transconductance variation through its bias current which enables a reduction in power consumption when lowering the operation frequency.

All in all, continuous time analog filters provide a high precision, high performance solution for applications requiring filtering. Implementations targeting the frequency range mentioned in the previous section tend to require a power in the milliwatt range and hence do not satisfy the small power budget of a wake-up radio. The interested reader is referred to [29] and [25] for more insight into the design of continuous time analog filters.

1.4.2 Discrete Time Charge Domain Filters

The proposed tunable IF stage can also be implemented using charge domain filtering in which the information, represented as an electrical charge (hence it's analog), is transferred from the plates of one capacitor to another by means of clock controlled switches. An example of such systems is represented by the switched capacitor circuits which implement filters similar to the ones previously presented using only OAs, switches and capacitors [30], [31], [32], [33] and [34]. Compared to classic, resistor-based solutions, switched capacitor circuits simplify the design of OAs' output stage, which does not need to be low impedance, and also have the potential of reducing the silicon area, as resistors are replaced by capacitors and switches.

Despite enabling the implementation of high precision filters, these circuits suffer from several drawbacks regarding their power requirements. The sampled nature of the charge in the time domain demands the use of an anti-aliasing filter as well as a circuit for generating the sampling clock. In reality, the sampling clock frequency is usually much higher than the frequency of the input signals, thereby simplifying the implementation of the anti-aliasing filter. Moreover, charge domain filters usually employ OAs which require a non-negligible amount of static power.

1.4.3 Digital Signal Processing

Alternatively, the proposed filtering function can be implemented using a Nyquist analog-to-digital converter, followed by a digital signal processor (DSP). The major advantage of this solution lies in the outstanding programmability offered by the DSP as it can generate almost arbitrary transfer functions as well as any baseband logic required for signal demodulation. Nevertheless, this solution comes with several drawbacks as it requires extra circuitry in order to operate correctly. The sampled nature of the approach imposes the use of a local oscillator generating the ADC sampling signal as well as the DSP clock. Furthermore, due to aliasing constraints, an anti-aliasing filter (usually in the form of a continuous time analog filter) needs to be added to the input of the ADC, thereby increasing the static power consumption of the system.

In conclusion, the power overhead incurred by using a discrete time ADC & DSP makes this solution impractical for wake up radio back-ends.

1.4.4 Continuous Time Digital Signal Processing

CT-DSP is a new paradigm in signal processing which has been extensively studied over the past years. CT-DSPs are obtained by removing clocks from classic DSP solutions [35]. A CT digital processing chain usually starts with a continuous time analog to digital conversion (CT-ADC) which converts the input analog signal into a signal which encodes the information two-fold: in an output binary word as well as in the arrival time between consecutive output events. Contrary to a Nyquist ADC, a CT-ADC can generate an event at any point in time, rather than just at integer multiples of the sampling period. The lack of sampling in the time domain also means that the conversion does not create any aliases, thus no anti-aliasing filter is required [36].

A classic DSP cannot be used to process the information from a CT-DSP because it does not have the capabilities of preserving the exact timing between events. As such, a CT (asynchronous) DSP, made up of asynchronous delay cells and adders, must be used. In terms of programmability, CT-DSPs have been proven to be as agile as their sampled counterparts [37], without requiring any clocks. The only drawback related to this solution lies in the type of processing which can be implemented in the CT-DSP: state of the art designs are limited to simple filtering functions. A more detailed description of such systems is given in Chapter 2 on page 29.

1.4.5 Tunable Filter Choice

A comparison table summarizing the previous discussion is given in Table 1.3.

TABLE 1.3: Comparison of different tunable filter solutions and the components they require.

filter type	active gain	clocks	anti-aliasing
continuous time analog	yes	no	no
discrete time charge domain	yes	yes	relaxed
discrete time DSP	no	yes	yes
CT-DSP	no	only for calibration	no

For the implementation of the proposed tunable filter stage, the CT-ADC-DSP solution has been chosen since it can accomplish the required functionality without adding any supplementary circuitry such as a sampling clock or anti-aliasing filters or requiring active gain components, such as OAs, OTAs or G_m cells.

1.5 Study of the Proposed CT Filtering Architecture

For the rest of this chapter, we shift the focus of this discussion on the specifications of the CT-ADC-DSP IF filtering stage. These results will then serve as a starting point for the design and implementation of the tunable IF filter.

A detailed view of the proposed IF filtering stage inside the proposed WU-RX architecture is given in Figure 1.18. The IF signal has a bandwidth of 40MHz (a 5MHz margin has been taken with respect to the results from the previous section) and is described by its input signal-to-noise ratio (SNR_{in}) and input signal-to-interferer ratio (SIR_{in}). In order to avoid flicker noise, we choose 10MHz as the start of the IF bandwidth. The quantization process then adds a certain amount of noise, corresponding to the component noise (SNR_{conv}), and harmonics defined by the conversion spurious free dynamic range, $SFDR_{conv}$. As stated in Section 1.3, the primary goal of the tunable filter solution is to enable an interferer-cancelling IF stage; consequently, most of its performance metrics are related and defined by the behaviour of the system in the presence of strong out-of-band interferer-like signals. Most of these requirements are linked to the CT-ADC since its conversion noise or linearity may limit the performance of the subsequent CT-DSP.

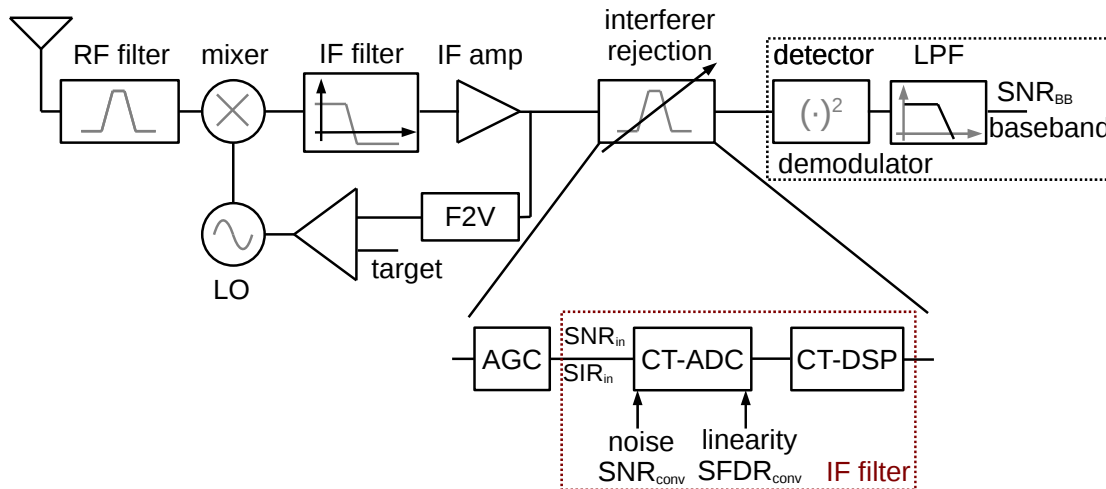


FIGURE 1.18: View of the WU-RX IF with an input signal configuration described by SNR_{in} and SIR_{in} .

1.5.1 Single Tone Reception

First, the performance of the proposed IF filtering stage (Figure 1.18) is tested in a scenario where its input contains only the useful signal and no interferers. As in the case of most existing WU-RX implementations, the input signal is considered to be on-off-keying (OOK) modulated with a data rate of 100kbps. Several assumptions are

also made: we suppose that the automatic gain control (AGC) block situated before the CT-ADC amplifies the input signal to the full scale of the ADC; the CT-DSP has a bandwidth of 1MHz and is configured such that the useful signal falls inside its pass-band; the bandwidth of the last low-pass filter used is 500kHz. The configuration of the system as well as the frequency domain representation of the signal at various intermediate points is provided in Figure 1.19

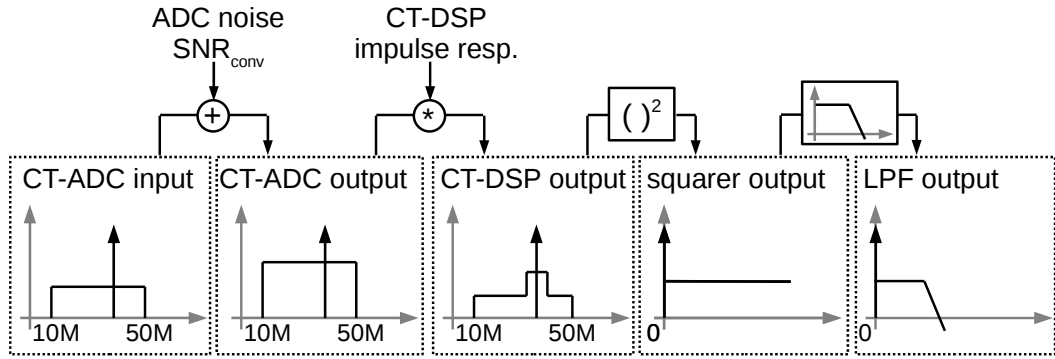


FIGURE 1.19: Single tone reception scenario.

The sensitivity of a radio is defined as the minimum input signal power which yields a given bit error rate (usually equal to 1e-3). The front-end of the proposed receiver (located before the CT-ADC, Figure 1.18) is a linear system which follows a simple rule equation 1.3. Here, P_{RF} represents the input power, k – the Boltzmann’s constant, T – the temperature in Kelvin degrees and B – the input bandwidth.

$$SNR_{in} = SNR_{RF} - NF_{FE} = 10 \log_{10} \frac{P_{RF}}{kTB} - NF_{FE} \quad (1.3)$$

We can see that the signal-to-noise ratio at the antenna output (SNR_{RF}) is entirely defined by the power of the input signal and an environment parameter, the absolute temperature. It follows that the signal-to-noise ratio at the input of our IF stage (SNR_{in}) depends only on the front-end noise figure (NF_{FE}) and on the input signal power. Thus, optimizing the sensitivity of the radio for a given NF_{FE} , is equivalent to designing an IF stage which achieves a bit error rate (BER) of 1e-3 for the lowest SNR_{in} . Consequently, a primary metric of performance of our system is the minimum SNR_{in} for which the baseband demodulator achieves an SNR_{BB} of 12dB (equivalent to a BER of 1e-3). The SNR_{in} is defined as the ratio of the useful IF signal power divided by the noise power integrated over the entire IF: from 10MHz to 50MHz; the SNR_{BB} is defined as the power of the useful signal in baseband divided by the noise power integrated from 0 to 500kHz.

Contrary to the RF front-end, the demodulation stage has a non-linear transfer function due to the use of the squarer before the last low-pass filter. Thus, the relation between

its input SNR and output SNR cannot be described compactly by its noise figure. There are several papers in the literature which attempt to derive the analytical expression linking the noise at the output of a squarer to the noise at its input [21], [38]. A full analytical derivation of these equations requires several assumptions regarding the signal at the input of the demodulator which, unfortunately, are not true in our case. For more insight in this derivation, the reader is referred to Appendix A on page 145.

For a more realistic interpretation of the system proposed in Figure 1.19 we perform system level simulations to plot the SNR_{BB} as a function of the SNR_{in} for different levels of noise added by the CT-ADC (SNR_{conv}): Figure 1.20. The intersections of these plots with the horizontal line located at 12dB represent the required values for SNR_{in} which yield a BER of $1e-3$. This operation is repeated for two values of rejection achieved by the CT-DSP: 30dB in Figure 1.20a and 40dB in Figure 1.20b.




FIGURE 1.20: Baseband SNR (SNR_{BB}) versus the input SNR (SNR_{in}) for different values of SNR_{conv} .

We can see that systems using low precision CT-ADCs, which achieve SNR_{conv} of 0dB, never reach a baseband SNR of 12dB regardless of the input signal power. In such configurations the noise added by the CT-ADC is simply too high and completely masks the noise coming from the RF front-end even for high values of SNR_{in} . On the other hand, we can conclude that ADCs achieving SNR_{conv} of around 20dB are sufficient, as increasing the SNR_{conv} to infinity does not yield any substantial improvements. Finally, increasing the CT-DSP rejection levels from 30dB to 40dB does not improve the sensitivity of our IF stage. This can be explained by the fact that the baseband noise floor is determined by the noise levels around the signal (inside the CT-DSP pass-band) and not by the noise situated far from the signal, which gets attenuated by the CT-DSP transfer function. These results can be alternatively viewed by intersecting the previous plots with the line corresponding to 12dB of SNR_{BB} , giving us the required SNR_{in} for different values of the SNR_{conv} : Figure 1.21.




FIGURE 1.21: Required SNR_{in} versus the SNR_{conv} for achieving a BER of $1e-3$.

We conclude that, in order to limit the impact of the ADC noise on the sensitivity of the receiver, SNR_{conv} should be at least 20dB.

1.5.2 Single Tone Reception with an Interferer

According to results from the previous section, correct demodulation of OOK signals requires only around 20dB of SNR_{conv} for a bit error rate of $1e-3$. However, when receiving the useful signal along with interferers the ADC is required to operate at a much higher SNR_{conv} . Such a configuration is presented in Figure 1.22. In the first example, the useful signal is received without any interferers, and the automatic gain control (AGC) is configured such that the amplitude of the input signal corresponds to the ADC full-scale. The resulting ADC output signal SNR is equal to the conversion SNR (SNR_{conv}), given by its effective number of bits. In the second scenario Figure 1.22 (right) an interferer X dB above the useful signal has been added. The AGC now equalizes the power of the interferer to the ADC full-scale. Since the noise level at the output of the ADC is only determined by the conversion performance, the useful SNR decreases by the same amount as the useful signal and becomes $SNR_{conv} - X$ dB. Thus if, for instance, a 30dB interferer rejection is targeted, this means that the conversion precision (SNR_{conv}) should be at least 30dB above the initial demodulation requirements, assuming a perfect filtering of the unwanted signal.

This scenario supposes a single tone interferer, so no intermodulation products are generated. Consequently, the conversion process can be considered perfectly linear ($SFDR_{conv} \rightarrow \infty$), without having an impact on the results of this study. As previously, the width of the CT-DSP passband as well as the bandwidth of the low pass filter in the demodulation block have a passband equal to 5 times the communication speed which




FIGURE 1.22: The SNR constraint for the analog to digital conversion

is 100kbps. The frequency domain representation of the signal at various intermediate points of the system is given in Figure 1.23.

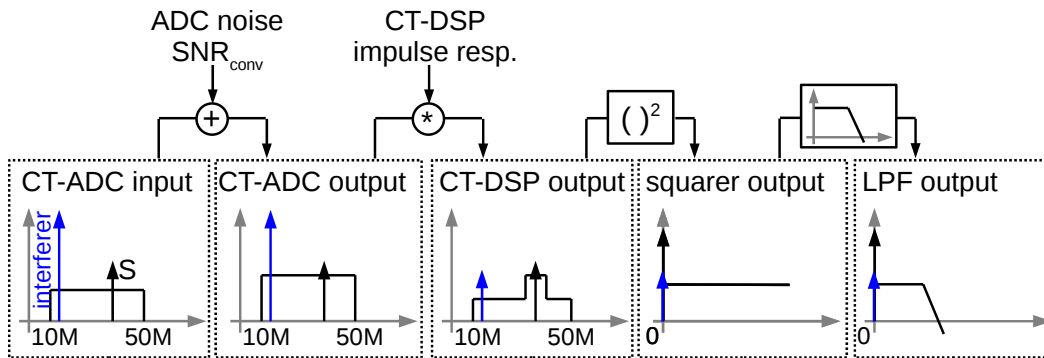


FIGURE 1.23: Useful signal & interferer reception scenario.

Depending on the rejection levels achieved by the CT-DSP, a different minimum conversion precision (SNR_{conv}) is required for the CT-ADC. In order to avoid over design, and thus save power, the goal of this section is to determine the minimum CT-ADC precision which does not compromise the potential results offered by the rejection levels of the CT-DSP. Consequently, for different values of SNR_{in} , SIR_{in} has been swept in order to extract the minimum value which yields 12dB of baseband SNR. This operation has been repeated for different CT-DSP rejection levels (30dB in Figure 1.24a and 40dB in Figure 1.24b) as well as for several CT-ADC precisions (ranging from $SNR_{conv} = 20$ dB to an ideal ADC, which does not add noise).

The colored lines split the (SNR_{in}, SIR_{in}) two dimension plane in two parts: points situated above the respective curves yield baseband SNRs above 12dB equivalent to BERs less or equal to $1e-3$. Input scenarios with less than 0dBs of SNR_{in} contain too much noise, and regardless of the power of the interferers, the baseband SNR is always below 12dB. In both plots, the dotted line represents the results obtained using an ideal, noiseless ADC. It can be seen that, as a rule of thumb, the SNR_{conv} of the CT-ADC has to be equal to the rejection level achieved by the CT-DSP. This represents the minimum conversion precision, which does not compromise the system's ability of robustly demodulating input signals.




FIGURE 1.24: Required SNR_{in} and SIR_{in} for different values of the conversion precision, SNR_{conv} , and CT-DSP rejection level.

Mathematical Interpretation

This result can be alternatively interpreted by considering that the signal at the input of the ADC is primarily made up of the interferer. Given the supposition that the AGC amplifier equalizes the input to the full-scale of the ADC, this means that, at the output of the ADC the ratio between the interferer and the total noise is bounded by the ADC precision (equation 1.4).

$$10\log_{10} \frac{P_{interferer}}{P_{noise-total}} = 10\log_{10} \frac{P_{interferer}}{P_{noise-in} + P_{noise-ADC}} < 10\log_{10} \frac{P_{interferer}}{P_{noise-ADC}} = SNR_{conv} \quad (1.4)$$

Given the definition of the SIR_{in} (equation 1.5) it follows that the ratio between the useful signal and the total noise at the output of the ADC is bounded by equation 1.6.

$$SIR_{in} = 10\log_{10} \frac{P_{signal}}{P_{interferer}} \quad (1.5)$$

$$\begin{aligned} 10\log_{10} \frac{P_{signal}}{P_{noise-total}} &= 10\log_{10} \frac{P_{signal}}{P_{noise-in} + P_{noise-ADC}} < \\ &< 10\log_{10} \frac{P_{signal}}{P_{noise-ADC}} = SNR_{conv} + SIR_{in} \end{aligned} \quad (1.6)$$

This means that for a given ADC precision, the useful SNR at its output is bounded by a function which depends on the SIR_{in} . When targeting higher rejection ratios, the SNR_{conv} needs to be adjusted accordingly, in order to maintain a constant signal to total noise ratio at the output of the CT-ADC.

1.5.3 Single Tone Reception with Two Interferers

Finally, the last parameter required to completely specify the conversion performance is its linearity, measured by its spurious free dynamic range – $SFDR_{conv}$. The main limitation related to this performance metric is showed in Figure 1.25: the interferer is made up of a two-tone signal. Since the analog to digital conversion is a non-linear operation, intermodulation spurs are expected to appear throughout the spectrum [39]. Such intermodulation products may fall arbitrary close to the useful signal frequency and thus become impossible to filter out at the CT-DSP level. Consequently, the intermodulation spur now acts as “noise” with respect to the received signal and has its power added to that of the conversion noise thus diminishing the useful signal-to-noise ratio. Consequently, the two tone spurious free dynamic range (SFDR) is used as the main linearity metric for the rest of this manuscript.




FIGURE 1.25: The SFDR constraint for the analog to digital conversion

In this section a two-tone interferer scenario is analyzed in which the frequency difference between the interferers is chosen such that their third order intermodulation product falls at a frequency very close to that of the useful signal and inside the CT-DSP pass-band. The power of this intermodulation product is strictly determined by the power of the two interferers (measured through the input signal to interferer ratio – SIR_{in}) and by the linearity of the CT-ADC ($SFDR_{conv}$). The evolution of the signal throughout the system is presented in Figure 1.26. To differentiate the effects of ADC noise from those of the non-linear behavior of the CT-ADC, it is supposed that the SNR_{conv} is infinite. As previously, it is desired that the CT-ADC is sufficiently linear not to compromise the robust reception of input signals without wasting power on over designing its performance parameters.

The (SNR_{in}, SIR_{in}) two dimensional space is explored for determining the performance limits of the architecture. Two CT-DSP scenarios are considered (30dB rejection in Figure 1.27a and 40dB rejection in Figure 1.27b) for several values of the ADC linearity ($SFDR_{conv}$ ranging from 20dB to a perfectly linear ADC).

A perfectly linear ADC ($SFDR_{conv} \rightarrow \infty$) is limited primarily by the input noise level (SNR_{in}) and also by the amount of rejection brought by the CT-DSP. As the linearity of

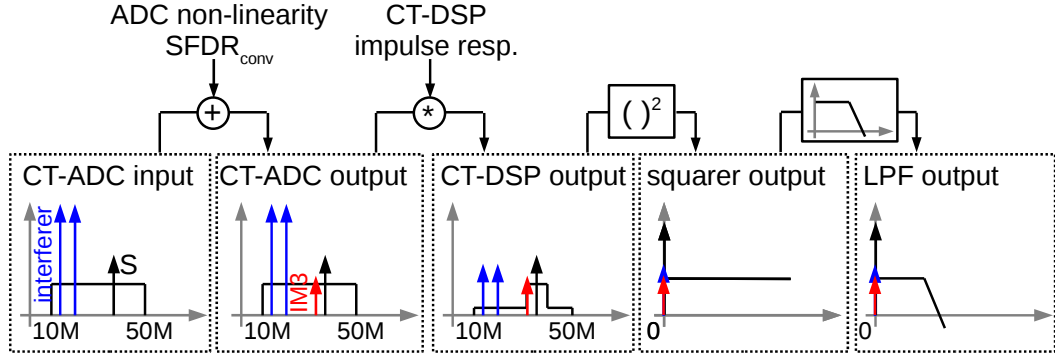



FIGURE 1.26: Useful signal with two interferers reception scenario.

FIGURE 1.27: Required SNR_{in} and SIR_{in} for different values of the conversion linearity, $SFDR_{conv}$, and CT-DSP rejection level.

the ADC degrades, more powerful spurs are created inside the CT-DSP passband which go through, unattenuated, into the demodulation block. If a 3dB degradation of the ideal sensitivity of the WU-RX is accepted, then according to Figure 1.27, $SFDR_{conv}$ needs to be at least 40dB when the CT-DSP achieves 30dB of rejection and 50dB when 40dB of rejection are targeted by the CT-DSP.

Mathematical Interpretation

For the signal configuration described previously, it can be proven that useful SNR at the output of the ADC is limited by the SIR_{in} and by the $SFDR_{conv}$. Given that SIR_{in} is given by equation 1.5 on page 25 and $SFDR_{conv}$ by equation 4.28, it follows that the useful signal to 3rd order intermodulation term, at the output of the ADC is given by equation 1.8.

$$SFDR_{conv} = 10 \log_{10} \frac{P_{interferer}}{P_{IM3}} \quad (1.7)$$

$$10 \log_{10} \frac{P_{signal}}{P_{IM3}} = 10 \log_{10} \frac{P_{signal}}{P_{interferer}} + 10 \log_{10} \frac{P_{interferer}}{P_{IM3}} = SIR_{in} + SFDR_{conv} \quad (1.8)$$

The useful SNR at the output of the ADC is given by the first term in equation 1.9 and is bounded by the second term in the same equation. This represents the signal reception limit for a given ADC linearity ($SFDR_{conv}$): increasing SIR_{in} and the rejection levels of the DSP does not yield better results since the third order intermodulation term is too strong compared to the input useful signal.

$$10\log_{10}\frac{P_{signal}}{P_{noise-total}} = 10\log_{10}\frac{P_{signal}}{P_{noise-in} + P_{IM3}} < 10\log_{10}\frac{P_{signal}}{P_{IM3}} = SIR_{in} + SFDR_{conv} \quad (1.9)$$

1.6 Wake-Up Receiver Design Conclusions

We have seen that the performance of existing wake-up receivers does not match that of standard data receivers for either the 802.15.1 or 802.15.4 standard. Performance areas which can be targeted for improvement include but are not limited to interferer robustness and multi-channel capabilities. As a solution to this we have proposed the use of an IF filtering stage based on a CT-ADC-DSP which, thanks to its highly tunable transfer function, enables the separation of the useful signal from excess noise and interferers.

The proposed stage is compatible with most existing WU-RX implementations, however, for completeness, we have also proposed a new WU-RX architecture which has a scalable performance – power trade-off. This architecture has only been studied at a system level in order to define a set of specifications for the proposed IF filter (summarized in Table 1.4); its implementation is beyond the scope of this manuscript and is not discussed in any of the following chapters.

TABLE 1.4: Performance limits of the IF interferer rejection stage according to the specifications of the CT-ADC and the CT-DSP.

CT-DSP rejection	CT-ADC		input signal requirements	
	SNR_{conv}	$SFDR_{conv}$	SNR_{in}	SIR_{in}
30dB	30dB	40dB	4dB	-20dB
40dB	40dB	50dB	6dB	-30dB

This set of specifications will now serve as a starting point for the implementation of the tunable IF filter which is presented in the rest of this manuscript.

Chapter 2

Continuous Time Processing Chain

IN this chapter we introduce the continuous time digital signal processing systems. Their principle of operation as well as critical system-level trade-offs for minimizing the power consumption are highlighted and analyzed. In the second part of this chapter, based on previously derived system-level specifications, we estimate the power consumption of our system, by extrapolating state of the art CT-ADC-DSP implementations. We show that the reuse of concepts from existing systems results in a power consumption well beyond the $100\mu\text{W}$ budget of our implementation. Consequently, in the last part of this chapter we present a new CT-ADC architecture designed to greatly reduce the power requirements of the subsequent CT-DSP. This architecture is particularly adapted for efficiently processing IF signals in the back-end of low power radios where signals are characterized by having a high likelihood of containing strong, out-of-band components which need to be removed for successful demodulation.

2.1 Classification of the Different Signal Processing Domains

According to the discrete / continuous nature of the time and amplitude of signals, we can distinguish four signal processing domains, as shown in Figure 2.1:

- *Continuous in Amplitude and Time:* The circuits which fall in this category are usually build around operational transconductance amplifiers (OTAs) or $G_m - C$ cells. These truly analog systems have been widely studied from the beginning years of microelectronics and even in the beginning years of electronics.

		time	
		continuous	discrete
amplitude	discrete	CT-ADCs CT-DSPs	DSPs
	continuous	analog filters	switched capacitor circuits

FIGURE 2.1: Classification of systems according to the continuous / discrete nature of the time and amplitude [36].

- *Continuous in Amplitude and Discrete in Time:* For this class of systems the signal is defined as having an arbitrary value at specific moments in time, usually defined as multiples of a master clock. An example of such systems are switched capacitor circuits (in which case the signal is defined as the charge moving from and to capacitors).
- *Discrete in Amplitude and Time:* Classic digital signal processors fall in this category. The signal is quantized by an ADC which associates a discrete binary word to the analog signal presented at its input at time instants defined by the sampling clock.
- *Discrete in Amplitude and Continuous in Time:* Such an operation mode is obtained by removing the clocks from the previously described class of systems.

For this manuscript we focus on the implementation of a system falling in the last category. The reasons behind this choice are mainly related to the specific requirements of our application: a high degree of programmability for a very low power consumption. It will be seen in the following sections that CT-ADC/DSP systems are very good candidates for achieving such goals.

2.2 Description of the CT-DSP Chain

A CT-DSP chain can be divided in two parts: one for the CT analog-to-digital conversion and the other containing the CT digital signal processor unit.

2.2.1 CT-ADC

The role of a CT-ADC is to convert the input analog signal into a signal which is discrete in amplitude and continuous in time.

Operation Principle

Like sampled ADCs, there exist several CT-ADC architectures which can be differentiated based on their principle of operation as well as on the encoding of the data at their output. So far, two CT-ADC architectures have been proposed in literature: the flash CT-ADC [37], [40] and the delta-modulator CT-ADC [41], [42], [43], [44], [45], [46], [47], [48]. For the purpose of this introduction, the operation principle of the flash CT-ADC is analyzed, however, the conclusions based on this analysis apply to *all* CT-ADCs presented in the literature so far.

The basic operation of a flash CT-ADC is presented in Figure 2.2: a series of discrete quantization levels are defined and assigned to a digital word ranging from 0 to the full scale of the ADC. The input, analog signal, is compared to each of these levels using a bank of continuous time (non clocked) comparators. Whenever one of these levels is crossed by the input signal, the corresponding comparator triggers its output. This change is reflected in the digital output word through a decoder (in this case a thermometer to binary decoder is required). This digital output is fed to the CT-DSP as the digital representation of the signal at the input.




FIGURE 2.2: Implementation of an N-level flash CT-ADC.

The input-output characteristic of a 9 level mid-thread flash CT-ADC is presented in Figure 2.3.




FIGURE 2.3: Input-output characteristic of a 9 level flash CT-ADC.

Activity Dependent Power Consumption

The previously described system has the interesting property of automatically scaling its power consumption according to the input signal properties [36]. The ideal configuration for such a system is achieved when the input is sparse in the time domain, i.e. it contains long periods when the input is constant. The power consumption of a CT-ADC quantizing such a signal is portrayed in Figure 2.4. When the input is constant, no comparators are triggered, hence no dynamic power is drawn. The power consumption of the CT-ADC thus becomes equal to the static power consumption of its building blocks which can be very low. As the input signal starts to evolve more rapidly (i.e. its first derivative is large), the frequency of events at the output of the CT-ADC also increases thereby increasing its dynamic power consumption.




FIGURE 2.4: Representation of the activity dependent power dissipation of a CT-ADC [36].

We can thus conclude that, in terms of power consumption, the CT-ADC performs favorably when quantizing signals with long periods of “silence”. According to the Nyquist criterion, the sampled ADC is required to sample at a frequency greater than twice the highest frequency expected at the input, even when the signal to be quantized is constant. On the other hand, this advantage is lost for time varying input signals such as pure sinusoids: Nyquist ADCs require two samples per period while CT-ADCs can trigger all of their quantization levels (if the input is full-scale) over one period. This increases activity at the output of the CT-ADC which translates into an increased power consumption of both the CT-ADC and of the CT-DSP. Consequently, it may seem

tempting to reduce the number of quantization levels of a CT-ADC in order to optimize its power consumption; unfortunately, this comes at a cost as its linearity also degrades. This trade-off occupies a central role in CT-ADC design and is discussed next.

Linearity - Activity Trade-Off

For a mid thread quantizer with a step of Δ and a total number of levels, N , as defined in the Figure 2.3, it can be proven that a single tone input ($V_{in}(t)$) of frequency f yields an output $q(t)$ which can be decomposed in a Fourier series ($c(n)$), as given by equation 2.2 - equation 2.4 ([39] and [49]). Thus we can conclude that, in the frequency domain, the signal at the output of an ideal CT-ADC contains only the input tone, its odd order harmonics and no quantization noise in-between. These spectral properties have been derived and experimentally proven in [50]. For more mathematical properties of continuous time digital signals, the interested reader can refer to [51], [52], [53], [54].

$$v_{in}(t) = A \cdot \sin(2\pi ft) \quad (2.1)$$

$$q(t) = \sum_{n=1}^{+\infty} c(n) \cdot \sin(2\pi fnt) \quad (2.2)$$

$$c(n) = \frac{4\Delta}{\pi n} \cdot \sin\left(\frac{\pi}{2}n\right) \sum_{i=1}^{(N-1)/2} \sin(d_i \pi n) \quad (2.3)$$

$$d_i = \frac{1}{2} - \frac{1}{\pi} \sin^{-1} \left(\frac{(2i-1)\Delta}{2A} \right) \quad (2.4)$$

As in the case of sampled ADCs, it is desirable that the digital interpretation of the output resembles the most the analog input, in the time domain as well as in the frequency domain. A comparison between the spectral representation of a signal at the output of a CT-ADC and that of a sampled ADC is presented in Figure 2.5. In terms of conversion noise, CT-ADCs are only affected by component noise, as they do not add any quantization noise (the signal is never sampled in the time domain). The only unwanted effects of the continuous time conversion are the odd harmonics appearing at the output. As seen in equation 2.3, these harmonics have a power proportional to Δ . This means that, for a sinusoid input of a given amplitude, the only way of reducing the harmonic power, thereby increasing the conversion spurious free dynamic range (SFDR) and signal to noise and distortion ratio (SNDR), is by decreasing Δ . However, this has the unwanted effect of causing an increased average commutation frequency at the output of the CT-ADC, f_{avg} – given by equation 2.5, increasing the power consumption of both the CT-ADC and of the CT-DSP.

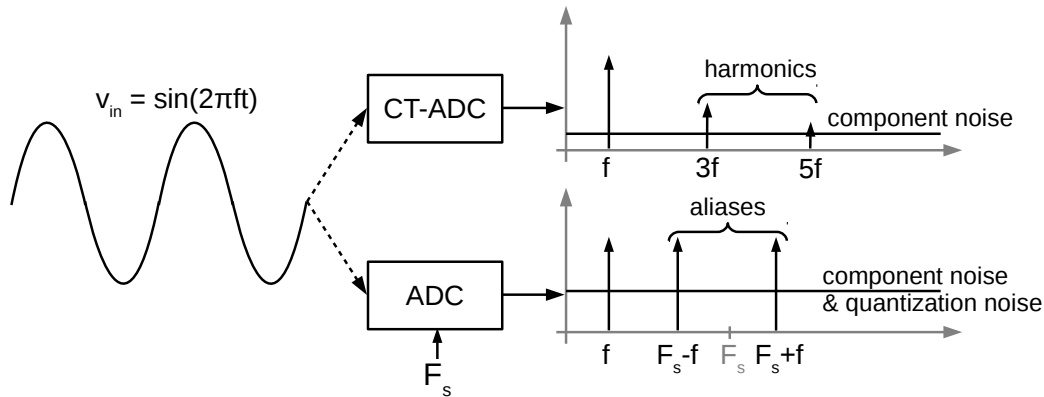


FIGURE 2.5: Comparison between the spectral representation of a signal at the output of a CT-ADC and that of a sampled ADC for a single tone, sinusoid input.

$$f_{avg} = \frac{1}{\Delta T_{ev}} = 2Nf \quad (2.5)$$

When designing power efficient CT-DSP systems, a critical trade-off in the design of the CT-ADC is between its linearity and its power consumption which is reflected in the output average event frequency. Figure 2.6 plots both the SFDR (expressed in dB) of a flash ADC and the average output event frequency versus the number of quantization levels (defined as the integer part of the ratio between the input peak-to-peak swing and the quantization step). Figure 2.7 highlights this trade-off directly, by plotting the SFDR versus the event rate.

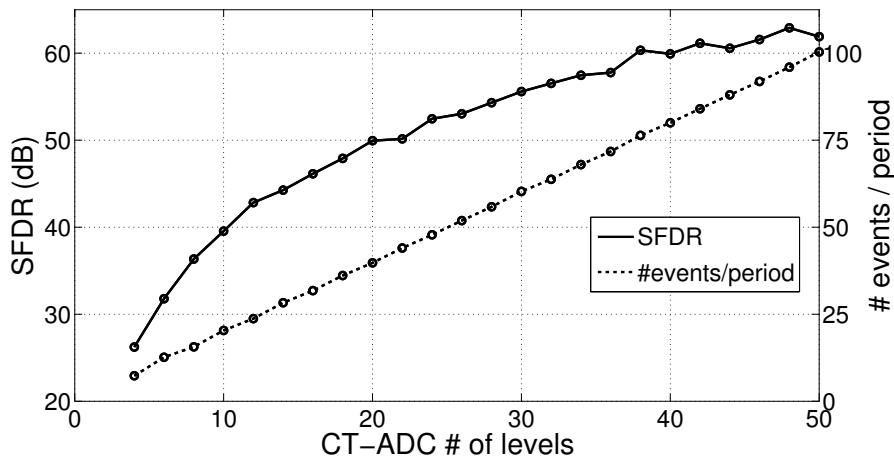


FIGURE 2.6: SFDR and number of events generated per period for a CT-ADC with an input sinusoid versus the ADC number of levels, N .

Given the importance of the previously described trade-off, several attempts have been made to reduce the event frequency without sacrificing the linearity of the respective converters. These attempts can be classified in five categories:




FIGURE 2.7: SFDR versus the number of events generated per period for a CT-ADC with an input sinusoid for different ADC number of levels, N .

- *Transfer function optimization:* For a given statistical description of the input signal and a required conversion bandwidth, the position of the quantization levels can be mathematically derived to optimize the trade-off between the CT-ADC output event rate and its linearity. This solution is similar to the μ -Law companding mechanism used to improve the dynamic range of sampled ADCs. A continuous time implementation has been presented in [55].
- *Dynamic quantization step:* Another option consists in dynamically tuning the quantization step according to the local input signal characteristics. In [56] the authors have implemented a system which continuously tracks the first derivative of the input in order to increase (or decrease) the quantization step whenever the derivative is large (or small). Consequently, large quantization steps are used for fast varying signals thereby reducing the average output event rate. The authors have also shown that the error corresponding to this treatment is usually high frequency and does not affect the in-band SNDR. The disadvantage lies in the extra circuitry used for the derivative tracking; the required analog implementation does not scale its power consumption according to the signal activity which is an imperative characteristic of *true* event driven systems.
- *Hysteresis:* Adding hysteresis to the continuous time comparators can reduce the output event frequency while having a small impact on the quantization error. This technique can efficiently suppress events which consist of the input signal crossing a quantization level for only very brief moments of time before going back to the initial level. This technique is, to a certain extent, used in all existing CT-ADC implementations, as all continuous time comparators exhibit some hysteresis. The main limitation of this solution is that the potential event rate reduction it offers is minimal and cannot be scaled effectively.

- *Decimation*: This technique is inspired by the decimation filters used in $\Sigma - \Delta$ modulators and consists in defining a time window during which all events issued by the ADC are accumulated inside a decimation block which only outputs at most an event per timing window. The CT-ADC still functions at a high frequency, but the CT-DSP sees a reduced event rate, defined by the value chosen for the timing window of the decimation block [57]. The disadvantage of this solution is that, instead of increasing the power of the harmonics, the decreased event rate causes an increase in the conversion noise floor. Given that CT-ADCs are only limited by component noise, it is possible that such a penalty can be tolerated while still achieving an efficient conversion.
- *CT-ADCs with embedded filters*: a *novel* technique proposed and detailed at a later point in this chapter consists of constructing a feedback loop around the CT-ADC such that the overall response of the system *before* the CT-ADC favors certain frequencies while attenuating others ([RMA⁺15] and [RMPT15]). This solution is particularly adapted for implementations which have a high likelihood of containing out-of-band signals. Attenuating these components enables an activity reduction at the level of the CT-ADC. This solution is limited by the fact that in-band signals do not benefit from a reduced output event rate as they are not attenuated by the embedded CT-ADC transfer function.

2.2.2 CT-DSP

Continuous time digital signal processors (CT-DSPs) are a class of systems designed to process the data produced by CT-ADCs [58].

Event Driven Nature

The information contained in the data at the output of a CT-ADC is encoded in the succession of digital output words as well as in the precise timing between them. Consequently, any system processing such an output (in our case the CT-DSP itself) must preserve the entirety of this information otherwise risk adding an error to the quantized CT-ADC output. This means that any block in the CT-DSP must also be event driven, rather than being triggered by a synchronous clock, which is the case in classic DSP implementations. A comparison between these two systems is presented in Figure 2.8: a continuous time adder needs to react immediately to a change presented at its input (therefore δt becomes a major constraint in CT systems and must be kept as low as possible) while a synchronous adder only begins its operation on the rising edge

of the clock (clk) and has to resolve its output before the next rising clock edge occurs (at which point the addition operation is repeated).

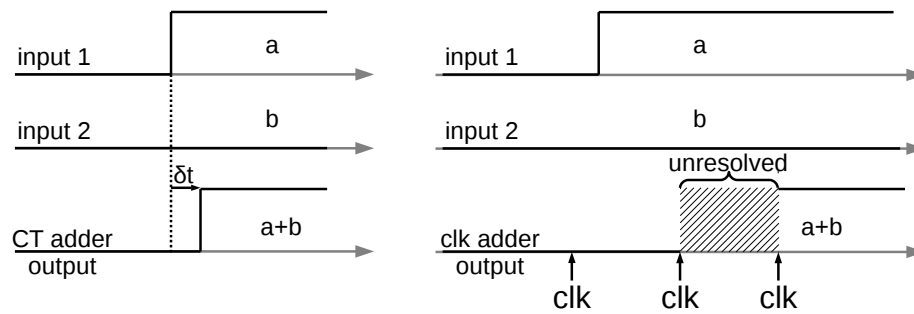


FIGURE 2.8: Operation principle of a CT adder (left) and of a sampled adder (right).

To avoid the design of a new class of DSPs, certain papers in the literature, [59], suggest quantizing the time between consecutive events issued by the CT-ADC thereby rendering its output data compatible with classic DSPs. This facilitates the implementation of more complex processing functions in the digital domain at the cost of introducing a supplementary error in the CT-ADC signal given by the noise of the time quantization. Furthermore, the event-driven nature of a true CT-DSP implementation is lost when opting for the clocked approach: clocked DSP blocks repeat their operation at every clock rising edge even when the input signal remains constant. Time quantization at the CT-ADC output is not discussed here as it is beyond the scope of this manuscript.

Functionality

The functionality of existing CT-DSP implementations is severely limited by the continuous time nature of the CT-ADC output which renders it incompatible with classic memory cells. Consequently, existing CT-DSPs usually have narrow functionalities and are limited to finite impulse response (FIR) filters [43], [60], [61] or infinite impulse response (IIR) filters [62]. The architecture and the equations behind such a system are developed next, followed by a discussion about the most important trade-offs at building block level.

The architecture of a continuous time FIR filter is presented in Figure 2.9: it consists of CT delay cells, CT multipliers, CT adders and no clocks! The equations governing this system can be derived by expressing the time representation of the output as a sum of delayed versions of the input, equation 2.6, and taking the Laplace transform on both sides, equation 2.7, [35]. The system transfer function $H(s)$ can then be derived by replacing s with $j\omega$ and taking the ratio $Y(s)/X(s)$. We thus recognize the transfer function of an K -th order filter with a_0 to a_K as coefficients.





FIGURE 2.9: Architecture of a CT finite impulse response (FIR) filter.

$$y(t) = \sum_{i=0}^K a_i q(t - i\tau) \quad (2.6)$$

$$Y(s) = Q(s) \cdot \sum_{i=0}^K a_i e^{-s\tau i} \quad (2.7)$$

$$H(j\omega) = \frac{Y(s)}{Q(s)} = \sum_{i=0}^K a_i e^{-j\omega\tau i} \quad (2.8)$$

Similarities exist between the frequency representation of the previously derived transfer function and that of a sampled FIR implementation. As it can be seen in equation 2.8, the transfer function is periodic in the frequency domain, with a period equal to $F_c = 1/\tau$. Supposing a low-pass implementation of $H(j\omega)$, pass-bands are expected to appear throughout the spectrum with a period equal to F_c , as showed in Figure 2.10.

FIGURE 2.10: Spectral representation of the transfer function of a low-pass FIR filter which has the values of all of its delay cells equal to $\tau = 1/F_c$.

For completeness, the architecture of a CT-IIR, [62], is given in Figure 2.11: the output of an FIR filter (FIR_1) needs to be fed back to its input through another FIR filter (FIR_2). The corresponding transfer function is given by equation 2.9. As in the case of a sampled IIR implementation, the adder needs to have its output digitally truncated. Supposing the output word, $y(t)$, is M bits wide, this means that the output of the second FIR filter, $p(t)$, has a width of P bits with $P > M$. However, this bus represents the input of an adder which we already supposed to be only M bits wide, hence digital truncation is required.




FIGURE 2.11: Architecture of a CT infinite impulse response (IIR) filter.

$$H(j\omega) = \frac{FIR_1(s)}{1 - FIR_2(s)} = \frac{\sum_{i=0}^K a_i e^{-j\omega\tau i}}{1 - \sum_{i=1}^L b_i e^{-j\omega\tau i}} \quad (2.9)$$

FIR vs. IIR

Despite being implemented in continuous time, FIR and IIR filters face the same issues as they do in discrete time. Table 2.1 presents a short comparison between the two solutions.

TABLE 2.1: FIR and IIR filter comparison

	FIR	IIR
phase	linear	not defined
order	high	low
stability	always stable	can be unstable
limit cycles	none	can occur

These characteristics are closely linked to the history of these filters, as the IIR filters have been derived, and can be considered as digital approximations of analog filters (such as Butterworth filters). Compared to sampled implementations, the biggest challenge regarding CT-IIRs is related to minimizing delay cell mismatch. Clocked IIR filters are controlled by a master clock and treat each cycle as a constant time unit, regardless of any jitter present around the fundamental clock frequency. In CT-IIR filters, any mismatch which occurs in the value of a delay cell can be treated as a small amount of noise which is injected and possibly reamplified by the loop structure of the filter.

Consequently, most CT-DSP implementations rely on the use of FIR filters rather than IIR filters.

Building Blocks

As seen previously, the three basic building blocks required for the implementation of FIR or IIR transfer functions are the delay cells, the multipliers and the adders. A quick overview of these blocks is provided below:

- *Continuous Time Delay Cells:* The basic operation principle of a continuous time delay cell is presented in Figure 2.12: inputs, which can arrive at any moment, have to be delayed by a fixed amount of time, τ , before being presented at the output [63], [64], [65], [66]. Furthermore, if we define the minimum time between two consecutive events (ΔT_{ev}) at the output of the CT-ADC as T_{gran} , delay cells are usually required to have $\tau > T_{gran}$. This means that the capacity (maximum number of events which have to be delayed *at the same time*) of the delay cell is larger than 1, and defined by the ratio τ/T_{gran} . This problem is usually addressed by opting for a series implementation of the delay macro-block using several elementary delay cells with a delay value of T_{gran} or smaller. Elementary delay cells have a mostly analog implementation: usually a pre-charged capacitor is discharged by a current source until a certain voltage is reached at which point a positive feedback mechanism is triggered and the capacitor is completely discharged.




FIGURE 2.12: Operation principle of a continuous time delay cell.

The two most important parameters required to specify a delay cell are its delay value, τ , and its event capacity (or granularity T_{gran}). It will be seen later that these two parameters, due to choices regarding the implementation, play a vital role in determining the energy required per delay as well as the area of the elementary delay cell.

- *Continuous Time Multipliers and Adders:* Most CT-DSP implementations combine the multiplication and the addition functionalities into a single block called weighted adder. Depending on the speed of the input signal, the weighted addition can be done either in the analog or in the digital domain. Low frequency, voice processing systems [43], [61] have been proven to operate at speeds low enough to allow for a digital implementation of the weighted adder. On the other hand, the CT conversion of high frequency signals, such as those of RF applications [37], generates events which have very hard timing constraints and requires a faster, analog implementation of the adder.

2.3 Co-Designing the CT-ADC with the CT-DSP

In this section we try to assess the power requirements of the proposed CT-ADC-DSP system based on the specifications derived previously (Section 1.6 on page 28) and by extrapolating results from the literature. We show that in order to achieve a power consumption compatible with the budget of a WU-RX ($< 100\mu\text{W}$) the CT-ADC and the CT-DSP need to be designed and optimized together rather than separately.

In the previous chapter we have seen that the CT-ADC must achieve an $SFDR_{conv}$ of 40dB (50dB), depending on the amount of rejection offered by the CT-DSP: 30dB (40dB). The linearity – activity trade-off presented previously (Section 2.2.1, Figure 2.7 on page 35) allows us to conclude that the respective SFDR is achieved if the CT-ADC has 12 levels (22 levels). In the worst case scenario, the interferers are located around the high end of the IF bandwidth, 50MHz, thus imposing an average event frequency at the output of the CT-ADC of 1.2GHz (2.2GHz) (given by equation 2.5 on page 34), supposing a full-scale signal. The average time between two consecutive events is thus 830ps (455ps).

To avoid strict matching requirements for the delay cells, we are going to suppose, for now, that we are opting for a CT-FIR filter to remove unwanted IF interferers. Having specified the ADC requirements, the next step is to determine the minimum FIR filter order which enables us to achieve the required rejection of 30dB (40dB). The actual filter design is done using automated tools which compute the FIR coefficients required to achieve a transfer function respecting a set of requirements defined by the user. In this section we choose a minimum order equiripple design method; the user-imposed constraints in the case of lowpass filter design are presented in Figure 2.13: a ripple of A_{pass} is allowed in the passband, the minimum stopband rejection is defined by A_{stop} , the stopband starts at F_{stop} while the passband ends at F_{pass} respectively. Due to the frequency repetition property of the filter, we only need to specify the stopband until

$F_c/2$, with $F_c = 1/\tau$; the transfer function over the rest of the spectrum can be computed from that. This filter design method is implemented in various design and analysis tools; its description is beyond the scope of this manuscript. For more information, the interested reader can refer to [67].




FIGURE 2.13: Details regarding the equiripple filter design specifications for a lowpass filter implementation.

Out of the five design parameters (A_{pass} , A_{stop} , F_{pass} , F_{stop} and F_c), four are imposed by the application: F_{pass} is defined by the communication speed, we choose a value of 1MHz; F_{stop} needs to be as close as possible to F_{pass} , we thus choose a value of 2MHz; A_{stop} has been derived in the previous chapter, we require between 30dB and 40dB of rejection; finally, for A_{pass} we choose the standard value which defines the cutoff frequency of a filter, 3dB. In general, the order of an FIR filter is determined by its quality factor (Q_{filter} – equation 2.10) and by the amount of rejection it achieves: since we fix F_{pass} , increasing F_c or the stopband rejection (A_{stop}) also increases the order of the filter and hence its power consumption.

$$Q_{filter} = \frac{F_c}{F_{pass}} \quad (2.10)$$

One solution for the implementation of the FIR filter is to choose $F_c/2$ greater than the maximum ADC frequency and then implement a bandpass filter, with a central frequency which can be tuned by changing the ADC coefficients. As an example we choose $F_c = 200\text{MHz}$, the corresponding FIR transfer function is plotted in Figure 2.14. Applying the equiripple design method for a bandpass filter around a central frequency (F_{mid}) around 50MHz we deduce that a 165 (220) order filter is necessary to achieve 30dB (40dB) of rejection.

The previous solution is not optimal, since it relies on a bandpass FIR which requires a set of poles to implement the highpass half of the transfer function and another set of poles for the lowpass part. To combat this, we could use an FIR filter in highpass configuration; thanks to the transfer function repetition property, a passband is created around $F_c/2$, as shown in Figure 2.15. The position of the passband can be tuned by




FIGURE 2.14: FIR transfer function example using a bandpass implementation.

changing F_c through τ . Contrary to the previous solution, here we require only a single set of poles which defines the highpass shape of our filter. For an F_c of 100MHz, the equiripple design method results in filters of order 90 (120) for 30dB (40dB) of rejection.




FIGURE 2.15: FIR transfer function example using a highpass implementation.

Both previously presented solutions require filters with very high orders, making them impractical. However, they can be further improved by taking advantage of the properties of the input signal. The RF filter limits the total bandwidth of the IF signal to 10MHz. We can thus guarantee an attenuation of all interferers, despite having several passbands in the [10MHz–50MHz] band, as long as the frequency difference between these passbands is greater than the total RF front-end bandwidth, 10MHz. The filter quality factor can thus be minimized by choosing $F_c = 10\text{MHz}$ and opting for a lowpass implementation: Figure 2.16. Using this strategy, we only require a filter of order 8 (10) to achieve the targeted 30dB (40dB) of rejection.




FIGURE 2.16: FIR transfer function example using spectrum repetitions to filter out unwanted interferers.

A comparison of the three filter design strategies is given in Figure 2.17 where we plot the rejection achieved by each filter versus its order. We can see that, by taking advantage of the spectral properties of the input signal, the FIR order, and hence its

power consumption, can be reduced by a factor of 20, with respect to the initial bandpass design.




FIGURE 2.17: Achievable CT FIR filter rejection levels versus the FIR order for the previously presented filter design methods.

In Table 2.2 we now compare the requirements of our design with existing CT-ADC/DSP realizations that we chosen as benchmarks ([68] and [37]). Despite being published only four years later, the second design ([37]) is about 500 times more energy efficient than the first ([68]). This can be explained by looking at the operating frequency of both solutions: in general, high frequency designs are more energy efficient but require much more power.

TABLE 2.2: Comparison of existing CT-ADC/DSP implementations with the requirements of our application.

CT-DSP	Schell [68]	Kurchuk [37]	our implementation
frequency	10kHz	3GHz	50MHz
ADC # of levels	256	8	12 (22)
event rate – average (T_{avg})	195ns	20.8ps	830ps (455ps)
DSP order	16	6	8 (10)
ADC energy / event ($E_{ADC/ev}$)	10pJ	30fJ	n/a
DSP energy / event /tap ($E_{DSP/ev}$)	20pJ	40fJ	n/a

In terms of operating frequency, our design is situated midway between [68] and [37]; meaning that the expected energy efficiency should be estimated accordingly. However, for completeness, we extrapolate the power efficiencies from both previous designs to determine the power requirements of our application. The equations linking $E_{ADC/ev}$ (energy required by the CT-ADC to produce an event) and $E_{DSP/ev}$ (energy required by one DSP tap to process one CT-ADC event) to the actual ADC and DSP power consumptions are given in equation 2.11 and equation 2.12, with T_{avg} – the average time between two consecutive events generated by the CT-ADC and N_{FIR} – the number of

taps of the CT-DSP. A summary of the results is presented in Table 2.3.

$$P_{ADC} = \frac{E_{ADC/ev}}{T_{avg}} \quad (2.11)$$

$$P_{DSP} = \frac{E_{DSP/ev}}{T_{avg}} \cdot N_{FIR} \quad (2.12)$$

TABLE 2.3: Estimation of the power requirements of our CT-ADC/DSP implementation based on results from literature.

estimation bench	Schell [68]	Kurchuk [37]
ADC estimated power	12mW (22mW)	36 μ W (66 μ W)
DSP estimated power	193mW (440mW)	385 μ W (879 μ W)

Given our limited power budget of 100 μ W, the comparison with the state of the art enables us to draw two conclusions. First, despite our efforts to limit the order of the CT-FIR, an improvement in the power efficiency of CT-DSPs is required, since existing works require simply too much power. Second, given the disproportionate power requirements between existing CT-ADCs and CT-DSPs, we can conclude that for an optimal solution, these systems should be designed together. To achieve this, we propose a novel system which combines a CT-ADC with a CT-FIR in the feedback path thereby achieving an CT-IIR transfer function; this structure is presented next.

2.4 Proposed CT-ADC Architecture

The CT-ADC linearity / event-rate trade-off can be optimized based on prior knowledge regarding the signals which need to be quantized, in a manner similar to the one used to reduce the FIR filter order. In the presented wake-up receiver case, unwanted, interferer-like signals are expected to appear at any frequency inside the [10MHz–50MHz] bandwidth, with a maximum total input signal bandwidth of 10MHz as defined by the input RF bandpass filter.

The power consumption of our system is dramatically increased by interferers which are expected to have an amplitude much greater than that of the useful signal, in some cases as high as 30dB above it. This increases the input signal swing, making the CT-ADC “spend” most of its output events quantizing the interferers and injecting them in the CT-DSP which then filters them out, rather than using this energy to quantize the signal of interest. In the end, this process leads to an increased power consumption of the entire CT-ADC/DSP system. However, this can be avoided by removing unwanted tones *before* the CT-ADC rather than *after* it. Since the input of the CT-ADC is analog, one

possibility is to use analog filtering in a manner similar to the anti-aliasing filters used before sampled ADCs. This poses a problem with respect to our implementation, since tunability is also a required characteristic of the filtering stage.

2.4.1 Filtering CT-ADC Principle

To solve the previously presented problem we propose a new, filtering, CT-ADC architecture which relies on feedback from the output, through a tunable, CT-FIR filter [RMA⁺15]. The digital nature of the feedback offers a very good programmability, equivalent to that of any CT-DSP. The signal issued from this digital stage is then injected into the input of the CT-ADC through an analog adder which enables the cancelling of unwanted input signals. This cancellation mechanism is similar to the one employed in predictive modulators [69]. The proposed digital filtering CT-ADC (DF-CT-ADC) architecture is presented in Figure 2.18. The only restriction placed on the behavior of the CT-ADC is that it has a unity, or constant, transfer function over the entire bandwidth of interest, [10MHz–50MHz] in our case. The transfer functions between different points in the system are expressed in equation 2.13 - equation 2.15. It is interesting to remark that the useful signal transfer function (equation 2.14) now resembles that of an IIR filter.

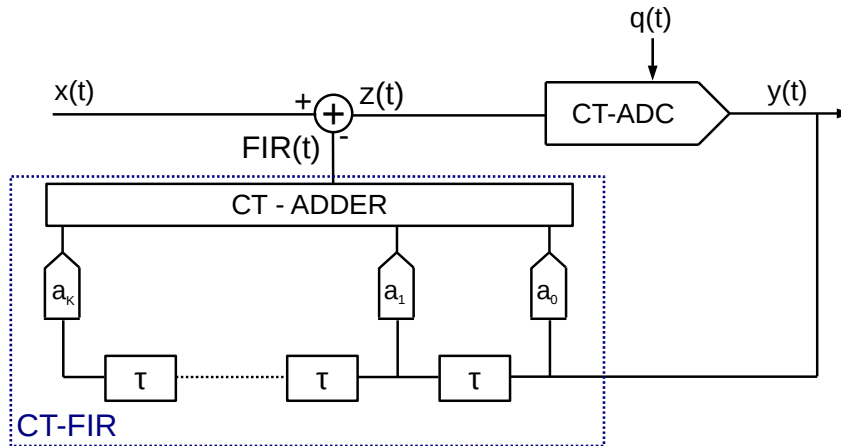


FIGURE 2.18: Proposed DF-CT-ADC architecture.

$$Y(s) = X(s) \frac{ADC(s)}{1 + ADC(s)FIR(s)} + Q(s) \frac{1}{1 + ADC(s)FIR(s)} \quad (2.13)$$

$$STF(s) = \frac{Y(s)}{X(s)} = \frac{ADC(s)}{1 + ADC(s)FIR(s)} \quad (2.14)$$

$$\frac{Z(s)}{X(s)} = \frac{1}{1 + ADC(s)FIR(s)} \quad (2.15)$$

The CT-FIR characteristics required to remove out-of-band signals can be easily derived from the analysis of the previously defined expressions. For in-band signals, $z(t)$ should be equal to $x(t)$, hence $FIR(s)$ should have a low absolute value such that the right hand side of equation 2.15 becomes close to 1. On the other hand, out-of-band components should be attenuated from $z(t)$, meaning that the magnitude of $FIR(s)$ at these frequencies should be high, thereby making the right hand side of equation 2.15 decrease towards zero. Consequently, $FIR(s)$ should have a band-pass behavior amplifying out-of-band signals while attenuating in-band components, as showed in Figure 2.19.




FIGURE 2.19: Representation of the required feedback FIR transfer function and of the resulting signal transfer function.

Given our supposition that the CT-ADC has a unity transfer function over the entire bandwidth, it follows that the transfer function between $Z(s)$ and $X(s)$ is also equal to $STF(s)$. The CT-ADC thus *sees* the input signal containing attenuated interferers. This has the major advantage of reducing the effective amplitude of the input signal before it is quantized by the CT-ADC; thus resulting in less events generated by the CT-ADC and a lower power consumption in the CT-ADC/DSP system.

The CT-FIR filter required for the proposed DF-CT-ADC architecture can be designed following the procedure presented in the previous section: highpass implementation with $F_c = 10\text{MHz}$. Stopbands are thus created around all integer multiples of F_c (Figure 2.19).

2.4.2 Reducing the Input Event Rate

The main performance criterion related to the DF-CT-ADC is the event rate reduction versus the SFDR compared to an implementation without FIR feedback. To study this, we suppose our system has an input signal containing an out-of-band interferer, located at 50MHz which is attenuated by A_{att} , and a single sinusoid in-band, located around 45MHz, which is amplified by A_{amp} . We sweep the input signal-to-interferer ratio (SIR_{in}) and plot the event rate reduction factor (Γ_r) achieved for different values of (A_{att}, A_{amp}) . The event rate reduction factor is computed according to equation 2.16 with T_{avg} – the average event rate observed at the output of the CT-ADC and $T_{avgDFADC}$ – the average

event rate observed at the output of the DF-CT-ADC; its value is plotted in Figure 2.20. Again, we suppose that the AGC amplifies the input signal to the full-scale of the ADC.

$$\Gamma_r = \frac{T_{avg}}{T_{avgDFADC}} \quad (2.16)$$




FIGURE 2.20: Event rate reduction offered by the proposed solution versus the SIR_{in} for different values of A_{amp} and A_{att} .

At -30 dB of SIR_{in} we can assume that the input signal consists mainly of the interferer; attenuating it by 20dB (or 10dB) without amplifying the in-band signal ($A_{amp} = 0$ dB) results in a diminished input swing, as seen by the ADC, of a factor of 10 (or 3.2). At the same time, if we increase A_{amp} to 10dB, the swing of the ADC input signal, $z(t)$, is slightly greater than previously, thereby reducing Γ_r . If we now apply the same reasoning to an input scenario with a lower power interferer (corresponding to a higher SIR_{in}), the initial supposition of the input being primarily made up of the interferer, becomes less and less valid. The DF-CT-ADC attenuation (A_{att}) starts showing diminishing returns since it is filtering out components which were not that strong in the first place; we thus observe $\Gamma_r \rightarrow 1$. If, on top of the previous case, we affect a positive gain (A_{amp}) to the in-band signal, the DF-CT-ADC will actually generate *more* events than the standalone ADC. However, this is not problematic, since we make the supposition that this WU-RX configuration is only active when the useful signal is corrupted by strong interferers.

The previously derived results show us that, in theory, the event rate at the output of the ADC can be reduced by factors up to 10. However, this reduction is useful only if the CT-ADC linearity is sufficient to enable a robust demodulation of input signals. This problem is discussed next.

2.4.3 Effects on the Linearity of the Conversion

To assure a robust reception of the useful signal, even in the presence of more than one interferer, the CT-ADC needs to be sufficiently linear to limit the power of intermodulation products (IM) falling in-band. Interferers, themselves, can be filtered by the CT-DSP, however IM can fall arbitrarily close to the useful signal and thus cannot always be removed by post-filtering. In a two interferer scenario, located at 48MHz and 50MHz, results from Section 1.6 on page 28 show that 50dB of ADC SFDR is required when targeting 30dB of interferer rejection. This means that, for robust reception, the ratio between the useful signal (located at 45MHz for example) and the third order intermodulation term (46MHz) is greater than 20dB. This value, shifted by a 3dB margin, will serve as target and benchmark in the following study.

The proposed DF-CT-ADC architecture constructs a feedback loop around the CT-ADC which reduces the power of out-of-band signals before presenting them to the ADC input. This has two effects on the value of the ratio between the signal and intermodulation product. On the one hand the power of the interferers at the ADC input is reduced by A_{att} ; since the power of the IM term is proportional to the power of the signals causing it, it follows that the IM term's absolute power also reduces. On the other hand, the ADC achieves 50dB of SFDR *only* if the signal at its input is full-scale; in standalone configuration we assume that the input AGC fulfills this role, however, the central purpose of the filtering structure is to reduce the ADC event rate, meaning that its input signal is no longer full-scale. In the end, this makes the ADC effectively operate at a lower SFDR, since not all of its quantization levels are triggered. Consequently, the goal of the next part of this section is to find a balance between A_{att} and A_{amp} such that the effect of interferer attenuation is greater than the SFDR degradation caused by the ADC not operating at its optimum point, thereby leaving us with a better signal to IM ratio (SIMR) for an ADC generating less events. A view of the signals *before* and *after* the DF-CT-ADC is given in Figure 2.21.




FIGURE 2.21: Frequency domain representation of the signals at the input and at the output of the DF-CT-ADC.

We start by doing a set of architecture-level simulations for the configuration presented previously: an input consisting of two interferers (48MHz and 50MHz) and a useful signal (45MHz) with an SIR_{in} of -30 dB and a 22 level ADC achieving 50dB of SFDR. When applying no filtering, this configuration is at the limit of robust reception, any degradation of the useful signal to third order intermodulation ratio results in an irrecoverable degradation of the received baseband signal. In Figure 2.22 we plot the output SIMR for the proposed DF-CT-ADC versus A_{att} for different values of the in-band gain, A_{amp} .




FIGURE 2.22: Output SIMR for a 22 level ADC with an input signal at -30 dB of SIR and different configurations of the proposed DF-CT-ADC.

It can be seen that, for a standard CT-ADC, as the interferer power is reduced (without adjusting the ADC full-scale) the power of the IM products actually increases making the SIMR decrease. Adding direct path amplification, through A_{amp} , readjusts the signal swing, as seen by the ADC, making it closer to its full-scale, thus decreasing the power of IM products and increasing the SIMR. It is important to note that A_{amp} does not only increase the power of the useful signal, but also of all the intermodulation products falling in-band, as derived from equation 2.14 on page 46. However, these results need to be put into perspective: any improvements in SIMR offered by the proposed DF-CT-ADC architecture are useful *only* if they come along with a reduced output event rate.

It is clear, from the previous discussion, that despite greatly reducing the output event rate, adding any kind of attenuation to the DF-CT-ADC actually reduces the SIMR, thus corrupting the correct reception of input signals. This can be avoided by using a CT-ADC with a smaller quantization step, Δ . The DF-CT-ADC output SIMR is thus plotted versus the number of CT-ADC quantization steps for different values of A_{att} and A_{amp} in Figure 2.23.

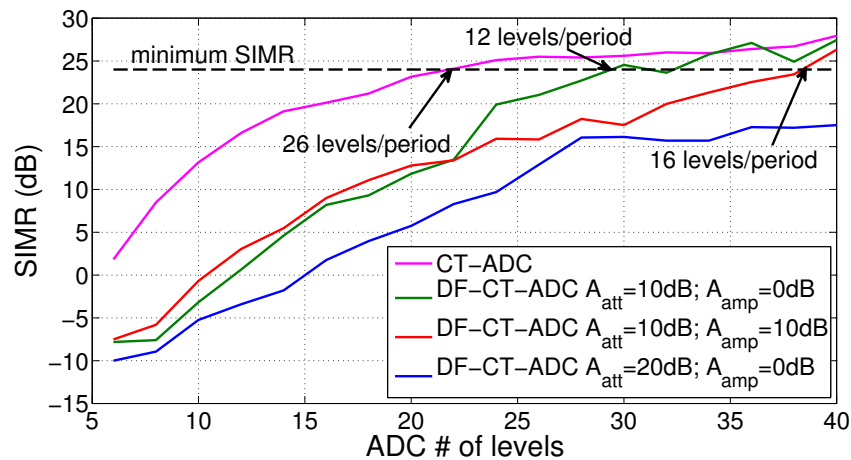


FIGURE 2.23: Output SIMR for different values of A_{att} and A_{amp} versus the number of CT-ADC levels for different configurations of the proposed DF-CT-ADC.

The pink line corresponds to the performance of a 22 level standalone CT-ADC while the dotted line represents the minimum SIMR which does not degrade the demodulator performance. Configurations with 20dB of attenuation require either more than 10dB of gain or more than 40 CT-ADC levels to achieve the previously defined performance specification. On the other hand, configurations with 10dB of attenuation achieve the SIMR threshold for an ADC with 29 levels when $A_{amp} = 0\text{dB}$ and 38 levels when $A_{amp} = 10\text{dB}$. The average activity observed in these two configurations is 12 and 16 events per period compared to the standalone average activity of 26 events per period. Note that the input signal is two tone which creates a slow beating in its amplitude, thus making it, in average, not trigger all its quantization levels twice over the course of one period (which would have yielded 44 events per period).

Another way of looking at the previous results is by plotting the output SIMR versus the average number of events triggered in a signal period, as shown in Figure 2.24. As explained at the beginning of this chapter, standard CT-ADC solutions are bounded to the corresponding trade-off line in Figure 2.24. The proposed DF-CT-ADC architecture enables us to achieve a better linearity while, at the same time, generate less output events, thereby reducing the power consumption of the subsequent CT-DSP.

In the end, depending on the configuration chosen for the DF-CT-ADC, the ADC will see a signal, $z(t)$, with an attenuated version of the interferers and an amplified useful signal which will trigger a finite subset of its quantization levels. Thus, based on the number of levels triggered during a period, we can compute the SFDR at which the CT-ADC operates, by referring back to Figure 2.7 on page 35; results are summarized in Table 2.4.

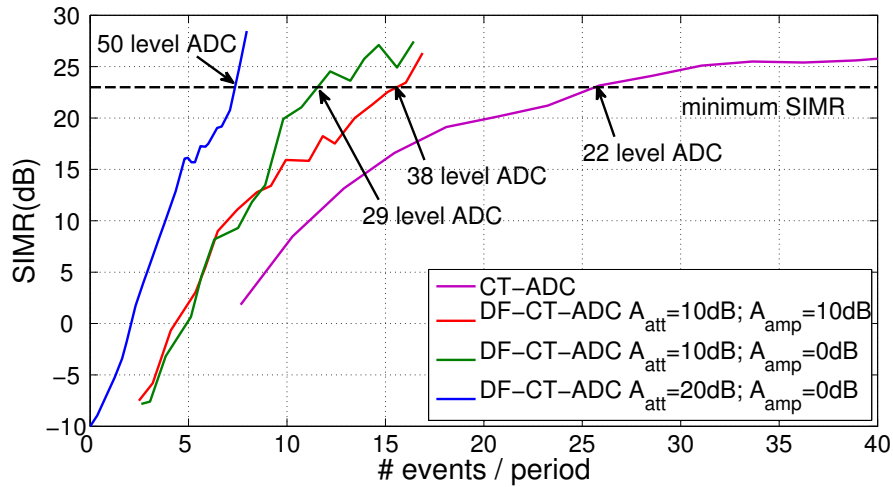


FIGURE 2.24: Output SIMR versus the average number of events triggered per signal period for different values of the CT-ADC quantization level.

TABLE 2.4: CT-ADC operating point for different configurations of the DF-CT-ADC.

A_{att}	A_{amp}	ADC # of levels	# events / period	ADC SFDR
10dB	0dB	29	12	32dB
10dB	10dB	38	16	37dB
20dB	0dB	50	7	26dB

2.5 CT-ADC/DSP Conclusion

In this chapter we have introduced the continuous time signal processing digital domain. We have showed that CT digital systems are usually composed of a CT-ADC and a CT-DSP; the different trade-offs related to these systems have also been presented. In the second part we have specified the requirements of our implementation and have compared them to existing CT-DSP implementations. It has been shown that in order to respect the $100\mu\text{W}$ power budget, improvements in the power efficiency of existing CT-DSPs were necessary. Finally, in the last part of the chapter we have proposed two improvements to existing digital CT processing chains aimed to reduce their power consumption. First, a novel FIR architecture has been proposed, which uses the repetition property of its impulse response to drastically lower the FIR order without sacrificing the out-of-band attenuation. Second, we have presented a new filtering CT-ADC architecture, which we call the DF-CT-ADC, which besides fulfilling the role of a CT digital filter, also reduces the average event rate generated by the CT-ADC without sacrificing its linearity, thereby reducing the power requirements of the entire system.

Chapter 3

Energy Efficient CT-ADC

HAVING studied different trade-offs related to the implementation of low power CT-ADC/DSP systems, we now start the discussion regarding the design of the tunable IF filtering stage with its most critical block, the CT-ADC. For the implementation, we have chosen ST's 28nm Ultra Thin Body and Buried oxide Fully Depleted Silicon On Insulator (UTBB FDSOI) CMOS technology with a V_{dd} of 0.65V; more information regarding this technology node can be found in [Appendix B on page 151](#). The rest of this chapter is organized as follows: first, existing CT-ADC architectures are presented along with their strengths and weaknesses. In the second part of this chapter, a new CT-ADC architecture is proposed which aims to improve the power efficiency of previous works. A detailed description of the transistor-level design of our solution is provided in [Section 3.3 on page 66](#) of this chapter. Finally, in the last part of this chapter we compare the performance of our proposed CT-ADC [[PRMT15a](#)] extracted from measurement data with that of state of the art ADC and CT-ADC implementations.

3.1 Previous Work

Existing continuous time ADCs can be divided into two categories: delta modulator based CT-ADCs and flash CT-ADCs.

3.1.1 Basic Architectures

The architecture of a flash-based CT-ADC has already been presented in [Section 2.2 on page 30](#). For completeness, a brief description of its operating principle is given below. A bank of comparators continuously compares the input signal to a set of predefined levels.

Whenever one of these levels is crossed, the corresponding comparator triggers and the change is reflected in the value of the output word (Figure 3.1).

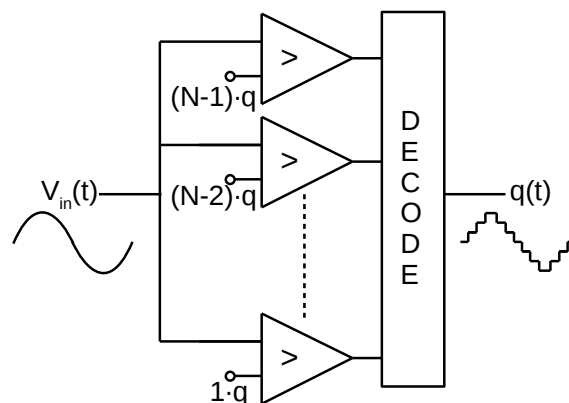


FIGURE 3.1: Architecture of an N -level flash CT-ADC.

The delta modulator CT-ADC also has its operating principle based on the use of continuous time comparators. However, contrary to the flash CT-ADC, only two comparators are required regardless of the number of levels associated with the conversion transfer function. To achieve this, the input signal (V_{in}) is tracked using two outputs (inc and dec) coming from a digital to analog converter (DAC) located in the feedback loop. Whenever one of these comparators is triggered, an increase or decrease instruction is generated. This instruction is then sent to the output and also sent to the feedback DAC in order to generate a shift of one quantization step in the level of the two tracking signals, inc and dec . The input signal is then reconstructed by injecting the output instructions into an accumulator. Figure 3.2 presents the architecture while Figure 3.3 presents the evolution of the input, output and some key internal signals.




FIGURE 3.2: Architecture of a delta modulator based CT-ADC.

Previous works, [60], have shown us that for high frequency implementations, the flash CT-ADC is the architecture of choice. Due to the lack of a feedback path, the input frequency for such systems is limited only by the speed of the continuous time comparators employed. The disadvantages related to this choice are linked to its complex implementation: the architecture necessitates a high number of comparators with thresholds requiring a precision which increases linearly with the number of quantization levels.




FIGURE 3.3: Evolution of the input, output and some key internal signals of a delta modulator based CT-ADC.

On the other hand, existing implementations of delta-modulator based CT-ADCs [68] usually use very small quantization steps (corresponding to a total number of up to 256 quantization levels) but are limited to the quantization of low frequency signals, such as voice (up to 20kHz). The high frequency operation of delta modulator CT-ADCs is limited by the reaction speed of the feedback loop which is usually slow due to the large number of components it consists of: the CT comparators themselves, the logic as well as the DAC.

Since our application requires up to 22 quantization levels, it seems that the delta modulator based CT-ADC architecture is more appropriate. Besides having a complex implementation, a flash based solution demands the generation of a large number of quantization levels. Any systematic error in the generation of these levels would degrade the linearity of the conversion which is the main constraint of this particular block. However, our application requires the CT-ADC to operate at frequencies up to 50MHz, well beyond the current limit of delta modulator based CT-ADCs which is several tens of kHz. Consequently, architecture level improvements are required.

3.1.2 Improved Delta-Modulator Based CT-ADCs

In this section we explore several delta modulator implementations, both in continuous time and in discrete time, in order to analyze their different strengths and weaknesses and thus propose a fast, energy-efficient CT-ADC architecture.

Possible improvements of the architecture presented in Figure 3.2 on the facing page can have two goals: decreasing the delay in the feedback loop to allow for a higher frequency operation or improving the overall energy efficiency of the CT-ADC.

The first component, adding delay to the feedback path, is the continuous time comparator: improving its comparison time will enable a higher frequency operation. A common characteristic of existing continuous time comparator implementations is the fact that

the time required to resolve the comparison depends on the level used for the comparison. Thus, due to reduced overdrive voltages, comparisons with levels close to the supply rails (V_{dd} and V_{ss}) take more time to be resolved than comparisons with levels situated around the middle of the supply rails. We can thus conclude that in order to improve the feedback delay in a delta modulator based CT-ADC, it is desirable to use *fixed* comparison levels which can be optimized for speed. This can be achieved by subtracting from V_{in} the reconstructed output value, as shown in Figure 3.4.




FIGURE 3.4: Architecture of a delta modulator based CT-ADC which uses fixed comparator thresholds.

The input signal is tracked by an analog image of the quantized output so that V_C is always kept between $-\Delta/2$ and $\Delta/2$. The input adder can be designed such that the common mode of V_C is situated around $V_{dd}/2$, thereby minimizing the comparison time of the two comparators. This delta modulator architecture has been implemented in [45]. Despite reducing the comparator contribution to the loop delay, the previously presented architecture introduces the accumulator and its corresponding delay to the feedback loop. Furthermore, the architecture still requires an N level DAC, making the linearity of the conversion sensitive to any static mismatch in the level generation.

A careful analysis of Figure 3.4 reveals that the values at the output of the DAC *always* change by one least significant bit (LSB), either $+\Delta$ or $-\Delta$. This hints to the possibility of using a DAC with only one bit, which is, by construction, perfectly linear. Furthermore, the information required to pilot this DAC is already contained in the *inc* and *dec* signals, therefore the accumulator doesn't need to be included in the feedback path. A view of the modified architecture is presented in Figure 3.5: the *inc* and *dec* signals are integrated through an analog integrator placed in the feedback path. In [70] the input subtraction operation, as well as the feedback integrator are all implemented around an operational transconductance amplifier (OTA), thereby minimizing the delay in the feedback loop.

3.2 Proposed CT-ADC

In the previous section we have shown that the key in achieving a fast, energy efficient delta modulator based CT-ADC is by using comparators with fixed thresholds coupled




FIGURE 3.5: Architecture of a delta modulator based CT-ADC which uses fixed comparator thresholds and has a fast feedback loop.

with a simple, fast feedback path. The architecture presented in Figure 3.5 seems to fulfill these requirements, however we show that the operations performed on the input signal are not optimal. To illustrate this, we take as an example an input signal (V_{in}) which is a ramp; the evolution of V_C is presented in Figure 3.6: V_{cm} is the common mode at the output of the subtractor while $+\Delta/2$ and $-\Delta/2$ represent the comparator thresholds.




FIGURE 3.6: Evolution of V_C , the signal at the input of the comparators, in a standard delta modulator based CT-ADC.

After each level crossing, the analog integrator and the input subtractor must shift the input signal by an amount equal to Δ . Depending on the value of Δ and on the slew-rate of the OTA this operation might take a significant amount of time with respect to the average time between two consecutive events. Increasing the slew-rate of the OTA minimizes this “dead-time” at the cost of an increased power consumption. Furthermore, the electronic charge moved to achieve this level shift also inevitably incurs a penalty on the energy required to produce an event.

3.2.1 Improved Commutation Scheme

The problems presented previously can be solved by using a new commutation scheme, in which, instead of shifting the signal by Δ for every level crossing we propose to simply flip the direction of the signal. This is illustrated in Figure 3.7: V_C now looks like a saw-tooth wave and no longer undergoes any discontinuity in the voltage domain. Consequently, both the commutation time and the energy required to achieve the previously described operation are reduced. This has the effect of increasing the maximum input frequency and improving the energy efficiency of the conversion.




FIGURE 3.7: Evolution of a differential version of V_C , the signal at the input of the comparators, in the proposed delta modulator based CT-ADC.

3.2.2 Proposed Architecture

In this section we are going to construct the CT-ADC architecture required to achieve the previously described commutation scheme. The idea of this work originated at Columbia University, however, the design, layout and tests have been done as part of a collaboration between the CEA and Columbia University.

In the previous example we have supposed that the input signal is represented by a ramp, which is monotonously rising. However, the scheme proposed in Figure 3.7 requires the generation of rising and falling signals. The input signal cannot be used to directly generate the alternating rising and falling elements of the required saw-tooth wave signal. To achieve this, we propose the use of a differential integrator in the direct signal path: the direction of the integration can be switched by simply swapping the inputs of the differential integrator. The “front-end” of the proposed CT-ADC architecture is shown in Figure 3.8.

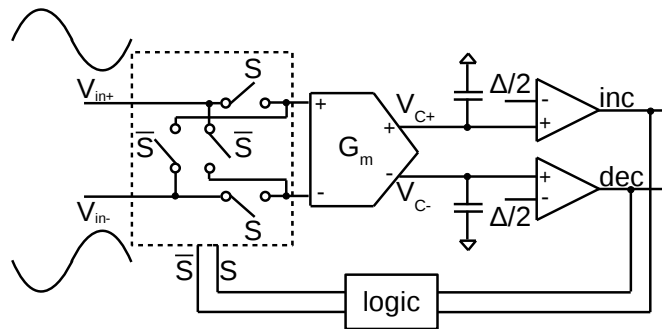


FIGURE 3.8: Front end of the proposed CT-ADC architecture.

Here, the switches present in the dotted box enable us to switch the polarity of the $G_m - C$ input to generate the saw-tooth wave described in Figure 3.7. Whenever one of the two comparators is triggered, a narrow pulse is generated and sent to the output of the CT-ADC as well as to a digital block, referred to as *logic*, which flips the position of the input switches, S and \bar{S} . It is interesting to remark that the proposed architecture does not require an actual DAC. In fact, the input switches controlling the polarity of the G_m , behave like a simple 1-bit DAC. This ensures a good linearity as well as a very fast operation since the DAC delay has been reduced to the that of turning ON or OFF a CMOS switch.

The logic required to control the polarity of S and \bar{S} is very simple: an OR gate between inc and dec generates the clock signal for an inverter feedback D flip-flop (DFF), as showed in Figure 3.9. Thus, the output of the DFF is flipped at the arrival of each pulse on the inc or dec nets. The total delay of the logic required for the proposed CT-ADC is thus equal to that of only two logic gates (an OR and a DFF).




FIGURE 3.9: Logic required to control the input switches (DAC) of the proposed CT-ADC architecture.

The CT-ADC front-end presented in Figure 3.8 achieves the goal of converting a signal from the analog domain (V_{in}) to the digital domain (inc and dec). However, the conversion is “incomplete” since it is still required to correctly interpret inc and dec in order to reconstruct the input signal. In the case of a classic delta modulator, this reconstruction process is represented by the output accumulator. In the proposed architecture, the digital domain integrator has been replaced by one in the analog domain, the $G_m - C$ block. This integrator sees at its input (V_{GMin}) a scrambled version of V_{in} ; this process can be represented by multiplying the input with a sequence of 1s and -1 s which change in time based on a function of the input signal, $f(V_{in}(t))$, as given in equation 3.1.

$$V_{GMin}(t) = V_{in} \cdot (-1)^{f(V_{in}(t))} \quad (3.1)$$

The original signal can be reconstructed by “unscrambling” the comparator outputs. To do this, we multiply inc and dec with the same signal, $(-1)^{f(V_{in}(t))}$, in a manner similar to the principle used in chopping circuits. Therefore, a new set of switches is required at the comparator outputs; the proposed CT-ADC architecture thus becomes as shown in Figure 3.10.

Due to timing constraints, the second pair of switches is controlled by S_d and \bar{S}_d , delayed versions of S and \bar{S} . This can be explained by the fact that the inc and dec signals are represented by narrow pulses which are triggered by level crossings and have a very narrow width. These pulses must be completely sent to out_+ and out_- before the output switches are flipped. In other words the output switches must never trigger while either inc or dec are equal to V_{dd} . To achieve this, another logic block (called \overline{logic}) is used, which instead of being triggered by *rising* edges on inc and dec which is the case of the

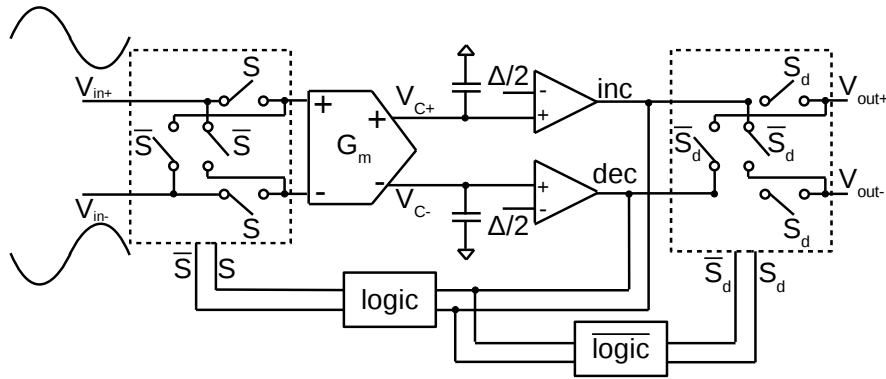


FIGURE 3.10: Full view of the proposed CT-ADC architecture.

block called *logic*, is triggered by *falling* edges. This ensures that pulses are entirely sent to the output *before* any of its switches are triggered. Since \overline{logic} is not located in the feedback path, there are no hard timing constraints for the propagation delay related to it.

A more detailed view of different internal signals as well as of the input and output is presented in Figure 3.11. V_{out} is obtained by subtracting V_{out-} from V_{out+} ; the density of pulses obtained at this level depends on the instantaneous value (amplitude) of the input signal V_{in} .




FIGURE 3.11: View of the input, output and some key internal signals of the proposed CT-ADC.

3.2.3 Features

The signal going through the proposed CT-ADC architecture undergoes an integration, performed by the G_m and a differentiation, performed by the two output comparators which behave like a delta modulator. Overall, it comes out at the output unattenuated with the event frequency directly proportional to the input signal amplitude. However, the succession of the operations has two interesting effects on the average output sample rate as well as on the spectrum of the quantization error at the output.

Output Event Rate

The ADC produces a delta modulated version of the integral of the input. Therefore, the output pulse rate, and hence the power dissipation, is proportional to the slope of the integral of the input, hence to the input signal instantaneous value. Contrary to existing CT-ADCs which exhibit an increased event rate as the frequency increases, as shown in 2.5 on page 34, our implementation does not scale its power consumption with the input signal frequency. High frequency input signals generate *fewer* output events per period while low frequency signal generate *more* events in a period such that the *average* event rate remains constant. We can interpret this as having an ADC which scales its quantization step based on the input frequency: low frequencies, susceptible to generate low frequency in-band error, are quantized with a fine quantization step, while high frequency out-of-band signals which generate mostly high frequency error [56] are quantized with a large quantization step, thereby reducing output activity and saving power.

Error Shaping

The quantization error, responsible for the harmonics in the output, is introduced at the interface between the analog and the digital domain, at the comparator level. Before arriving at the output, it undergoes a differentiation transfer function which attenuates low frequency, in-band harmonics, while applying out-of-band high frequency harmonics. It is interesting to note that the component noise introduced *after* the G_m block also benefits from the first order error shaping bestowed by the back-end of the proposed CT-ADC.

3.2.4 Possible Errors

The previously described operation principle works *if* we assume that the delay in the feedback loop is much smaller than the time between two consecutive events. On the one hand, the loop delay is determined by the transistor level implementation and will be discussed at a later point in this chapter. On the other hand, the minimum time between events is determined by the characteristics of the input signal which are discussed next.

Analyzing the time domain operation of the proposed CT-ADC, showed in Figure 3.12, and using an ideal model without any delay in the feedback loop, we can distinguish two situations which generate events close to each other. First, when the input signal has a high absolute value, it generates ramps with a high slope on the capacitors situated before the comparators thereby reducing the time between consecutive events (this is happening around 12.5ns, when the input signal is close to 1V). Supposing the amplitude of the input (A_{in}) stays constant for the time elapsed between two consecutive events (ΔT_{event}), then this time can be written as given in equation 3.2, with C_1 – the capacitance value, g_m – the transconductance value and Δ – the quantization step. Given the proposed commutation scheme, it is unlikely that this type of events generates a glitch in the behavior of the CT-ADC as this would require very high values for the input amplitude, well beyond the full scale of the converter.




FIGURE 3.12: Input and output signals as well as the voltage across the capacitors.

$$\Delta T_{event} = \frac{\Delta \cdot C_1}{A_{in} \cdot g_m} \quad (3.2)$$

The second type of potentially problematic events occurs when the input signal undergoes a zero crossing as this induces a change in the integration direction. If we couple this with a recent comparator trigger, it follows that output events can occur at arbitrarily close intervals. This can be seen in Figure 3.12: the first zero crossing of the input (around 25ns) generates events very close to each other, while the second zero crossing (around 50ns) generates comparator triggers that are spaced by a large amount of time.

If we now artificially introduce a delay in the feedback loop, the voltage across the capacitors starts to slightly overshoot the $[-\Delta/2 \Delta/2]$ window set by the two comparators, as seen in Figure 3.13. This overshoot corresponds to the charge integrated on the capacitors for the time corresponding to the loop delay and can be modeled as an artificial increase of Δ . Since the loop delay depends on the input signal, it follows that each level crossing *sees* a different local value of Δ which turns out to have, as an effect, a reduction of the CT-ADC linearity. If we now couple this overshoot with the event described in the previous paragraph, a glitch in the behavior of the ADC can occur: input zero crossings happening while capacitor voltages are above (or below) the $\Delta/2$ level cause the V_{c+} and V_{c-} to stay outside the $[-\Delta/2 + \Delta/2]$ window for an undetermined amount of time, as seen in Figure 3.13. We choose to call this event as out of bounds operation.




FIGURE 3.13: Input and output signals as well as the voltage across the capacitors for a CT-ADC with a non-zero delay in the feedback path.

Having showed that for an ideal feedback loop the minimum time between consecutive events can be arbitrarily small, it follows that it is impossible to design a feedback loop




FIGURE 3.14: Spectrum of the output signal containing periods when the CT-ADC goes out of bounds.

which is always faster. The out of bounds behavior will thus always *eventually* occur. This event can be modeled by a time domain multiplication of the ideal “correct” output (assuming no glitches occur) with a binary signal equal to 0 when V_{c+} and V_{c-} are out of bounds and 1 otherwise, as given in equation 3.4. This can be translated in the frequency domain by a convolution with sinc function; for a sinusoid input, the corresponding output spectrum is plotted Figure 3.14. The conversion SNR thus degrades proportionally to the width of the sinc “lobe” situated around the main tone which, in turn, depends on how long the binary function $b(t)$ is equal to 0. Consequently, to minimize the noise added by these glitches, we must minimize the time the capacitor signals are out of bounds. Without changing the architecture, this time is entirely defined by the input signal and thus can be arbitrarily long. Thus, this problem needs to be dealt with by changing the very behavior of the proposed architecture.

$$V_{out}^{real} = V_{out}^{ideal} \cdot b(t) \quad (3.3)$$

$$b(t) = \begin{cases} 0 & V_{c+} \text{ out of bounds} \\ 0 & V_{c-} \text{ out of bounds} \\ 1 & \text{otherwise} \end{cases} \quad (3.4)$$

The solution we propose consists of introducing an extra set of comparators to which we refer, henceforth, as overflow comparators. Their role is to compare V_{c+} and V_{c-} with an overflow threshold and reset (electrically short) the capacitors every time an overflow event is detected. The overflow threshold needs to be high enough not to get triggered by the overshoots occurring during normal operation but also low enough to minimize the time the comparator voltage is out of bounds. For this implementation, we choose an overflow threshold of Δ . The architecture with the proposed overflow comparators is

presented in Figure 3.15 while the evolution of internal signals around an overflow event is shown in Figure 3.16.

With the proposed solution, the pulse generated by the overflow event can also be integrated into the ADC output. In the example given in Figure 3.16, ideally, the CT-ADC would have generated two pulses around the 25ns mark. However, after generating one of them, an out of bounds behavior is observed and the second pulse is lost. This loss can be recovered by injecting the overflow pulse in the V_{out-} output; the second pulse is thus observed at the output with a slight delay compared to its ideal position, corresponding to the time required by V_{c+} to reach Δ from its initial value of $\Delta/2$. A small error is thus added to the ADC output whenever an out of bounds behavior occurs. This error has a negligible effect on the spectrum of the output signal, as seen in

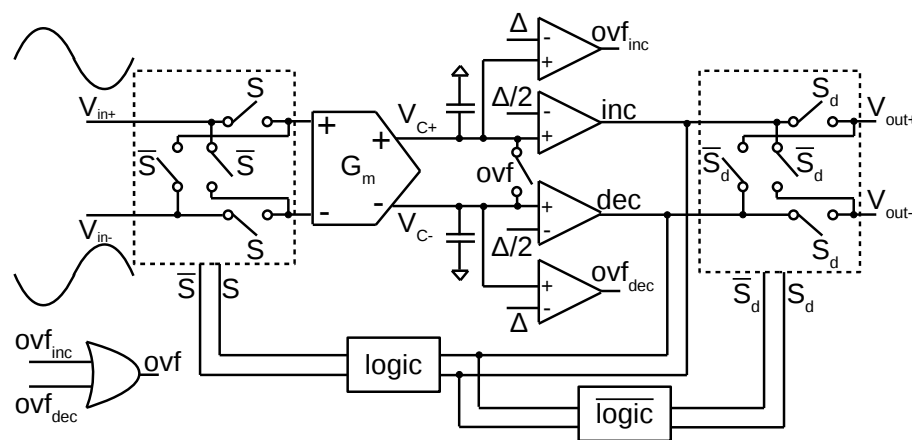


FIGURE 3.15: Proposed CT-ADC with overflow comparators.

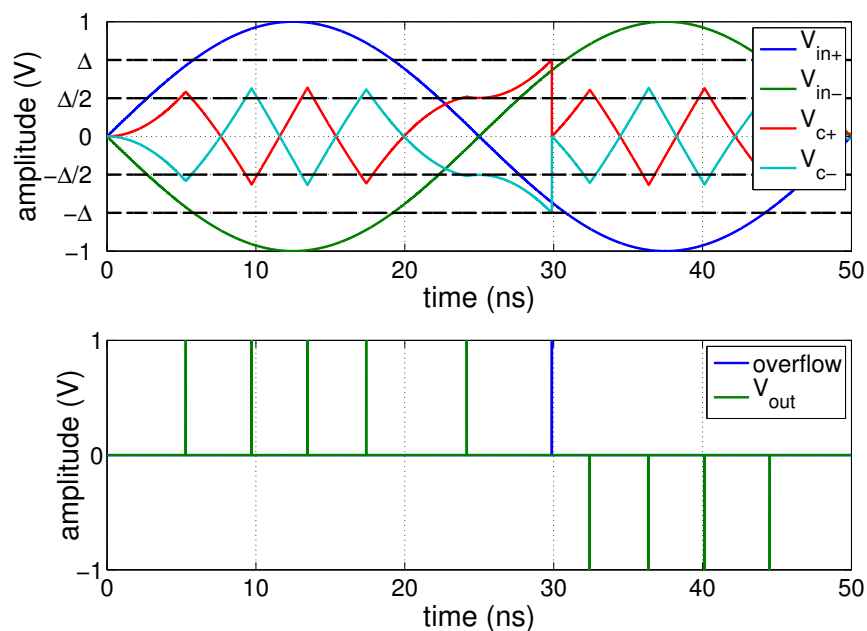


FIGURE 3.16: Input, output and some key internal signals around an overflow event.

Figure 3.17. Moreover, due to the choice of the overflow threshold as Δ , the reset event restores the signal on top of the capacitor to its normal value (as if no out of bounds behavior had occurred); no charge is thus lost.




FIGURE 3.17: Spectrum of the CT-ADC output signal with overflow comparator correction.

3.3 Transistor-Level Implementation

In this section we will explore different design choices related to the transistor-level implementation of the blocks required for the proposed CT-ADC architecture. In order to further minimize power consumption we choose a V_{dd} of 0.65V, lower than that of the ST 28nm UTBB FD-SOI technology, which is 1V.

3.3.1 Comparators

Despite dealing with signals up to 50MHz, the continuous time comparators are required to operate at a much higher speed since their comparison delay is directly reflected in the total delay of the feedback loop. Therefore, as comparator architecture we choose the one used in [37], which required comparing signals against fixed thresholds at GHz frequency. The comparator implementation is presented in Figure 3.18. We can see that only two transistors are required between V_{dd} and V_{ss} , making this architecture particularly adapted for operating at the low supply voltage chosen for our implementation (0.65V).

The transistor sizes as well as the biasing chosen for their back-gate are presented in Table 3.1. CMOS switches S_2 and S_3 have been optimized to minimize leakage while S_1 has been optimized to minimize resistive losses. Comparator inverters are back-biased at




FIGURE 3.18: Transistor level implementation of the continuous time comparators.

0.75V to reduce the threshold voltage of the transistors thereby increasing the speed of the comparison.

TABLE 3.1: Sizing of the comparator components.

component	type	value	back-biasing
S_1	NMOS	$1\mu\text{m}/30\text{nm}$	2V
S_1	PMOS	$1\mu\text{m}/30\text{nm}$	-2V
S_2	NMOS	$100\text{nm}/30\text{nm}$	0V
S_2	PMOS	$100\text{nm}/30\text{nm}$	0V
S_3	NMOS	$180\text{nm}/120\text{nm}$	0.75V
S_3	PMOS	$180\text{nm}/120\text{nm}$	-0.75V
M_1	PMOS	$1.4\mu\text{m}/30\text{nm}$	-0.75V
M_2	NMOS	$1.4\mu\text{m}/60\text{nm}$	0.75V
M_3	PMOS	$480\text{nm}/30\text{nm}$	-0.75V
M_4	NMOS	$240\text{nm}/30\text{nm}$	0.75V
M_5	PMOS	$480\text{nm}/30\text{nm}$	-0.75V
M_6	NMOS	$240\text{nm}/30\text{nm}$	0.75V
M_7	PMOS	$480\text{nm}/30\text{nm}$	-0.75V
M_8	NMOS	$240\text{nm}/30\text{nm}$	0.75V
C_1	capacitance	81fF	n/a

In normal operation mode, CMOS switches S_2 and S_3 are open while S_1 is closed. The comparator threshold value is stored on the plates of the capacitor C_1 as described by equation 3.5. The second constraint we impose to achieve the desired behavior is that when $V_{in} = V_{th} + V_{cm}$ then $V_{C-} = V_{trip}$, with V_{cm} – the common mode at the input of the comparator and V_{trip} – the trip point of the first inverter. Thus if $V_{in} > V_{th} + V_{cm}$ then the first inverter is tripped in one direction, while if $V_{in} < V_{th} + V_{cm}$ then the first inverter trips in the opposite direction.

$$V_C = V_{C+} - V_{C-} = V_{th} + V_{cm} - V_{trip} \quad (3.5)$$

This initial condition required for a correct operation of the comparator is achieved by using a setup time in which switches S_2 and S_3 are closed and S_1 is open. V_{C-} naturally becomes V_{trip} since the first inverter is in feedback; meanwhile, a voltage equal to $V_{th} + V_{cm}$ is applied to the V_{prog} terminal to program the required threshold. This

scheme is used just to illustrate the operation principle of the comparator. For our implementation, due to constraints related to the $G_m - C$ block, we have chosen another threshold setup mechanism which achieves the same result using different means and which is detailed later on in this chapter.

The most important performance parameters related to the proposed comparator implementation are its power consumption, its comparison delay, its offset with respect to the required threshold value and its leakage of the charge representing the threshold.

Power Consumption

The power consumption of the proposed comparator can be divided into two parts: one corresponding to the first inverter which, by construction, operates near its trip point thereby drawing a significant crowbar current and another one corresponding to the last three inverters which have an almost purely dynamic power consumption depending only on the switching frequency.

For a correct assessment of the first inverter power consumption, the proposed comparator needs to be tested inside the CT-ADC architecture in order to assure correct input conditions. Furthermore, depending on the value of its threshold, $\Delta/2$, the voltage on the net connecting the gates of M_1 and M_2 spends more (or less) time around the trip point of the inverter thus drawing more (or less) crowbar current. The power consumption of this first inverter stage decreases from $5\mu\text{W}$ to $4\mu\text{W}$ as Δ increases from 40mV to 80mV despite decreasing the average toggle frequency from 250MHz to 125MHz . On the other hand, the last three inverters require a power of only $1.8\mu\text{W}$ to 900nW which scales linearly with the output event frequency.

Comparison Delay

Continuous time comparators usually have a comparison delay (δT_{comp}) which can be decomposed into two elements: a part depending on the evolution of the input signal (δT_{rise}) along with a propagation delay corresponding to any digital block present at their output (δT_{prop}). The signal-dependent behavior can be modeled as a requirement for a certain amount of energy (E_{trig}) to trigger the “analog” part of the comparator, as shown in Figure 3.19. Fast rising signals provide this energy faster and therefore trigger a change in the comparator output sooner. The trigger energy provided to the comparator is given in equation 3.6, with m – the slope of the input signal.

$$E_{trig} = \frac{1}{2} \cdot m \cdot \delta T_{rise}^2 \quad (3.6)$$




FIGURE 3.19: First order model of the comparator behavior.

For the chosen comparator topology, δT_{rise} represents the time required to trigger the first inverter stage, while δT_{prop} corresponds to the propagation time through the last three inverters designed to increase the slew rate of the comparator output signal. The comparison delay is determined by a set of post-layout simulations, in which we apply a ramp as an input and we sweep its slope from 5MV/s to 50MV/s. The corresponding extracted comparator delay as well as a model for δT_{rise} and δT_{prop} fitted on simulation results are plotted in Figure 3.20. We can see that, as the slope of the input increases, the comparison delay decreases accordingly, as predicted by our model, from 480ps to around 190ps.




FIGURE 3.20: Extracted comparator delay as well as a model for the rise and propagation delays.

Threshold Mismatch

Threshold mismatch is important because our CT-ADC architecture uses two comparators which are supposed to have the same, exact thresholds. Any mismatch in the value of these thresholds will result in a CT-ADC which has half of the quantization levels slightly different from the other half thereby generating extra non-linearity in the digital output. To quantify this mismatch, a set of MC simulations is conducted from which the actual

comparator threshold is extracted. The standard deviation of the actual threshold value with respect to the 40mV target is 0.4mV. To conclude on the impact of this variation on the SFDR of the CT-ADC, a set of architecture level simulations is done in which mismatch is intentionally injected in the value of the two thresholds. Results are plotted in Figure 3.21; the SFDR starts to degrade for threshold differences above 5mV, more than 12 times greater than the standard deviation of the observed threshold mismatch in our implementation. We can thus conclude that comparator threshold mismatch has an insignificant effect on the CT-ADC linearity.




FIGURE 3.21: SFDR degradation of the proposed CT-ADC architecture with a single input tone at 10MHz versus and an artificially injected threshold mismatch; the absolute value of Δ is 40mV.

Threshold Leakage

In the proposed comparator architecture, the required threshold value is stored, as a charge, on the plates of the input capacitance (C_1). Leakage through any of the gates of the first inverter, through S_2 or S_3 or even charge injection incurred by the first inverter switching from one state to another may modify the total charge on the plates of the capacitor thereby shifting the comparator threshold. This process is hard to model and heavily dependent on the input signal characteristics which determine the comparator switching frequency.

Since we want to have a certain amount of control over the comparator quantization step Δ , we limit this threshold drift by resetting the comparator periodically in time. The reset frequency has been determined experimentally by applying a DC input to the CT-ADC and measuring the time between consecutive output events. At this point we suppose that both comparators are behaving in the same way. Ideally, no charges are ever injected or removed from the plates of the input capacitances and the time between

consecutive output events stays constant. However, measurements show that, apart from the random variations observed in the time duration between consecutive pulses due to noise, there is also a static drift caused by the phenomena predicted in the previous paragraph. This is shown in Figure 3.22 where the initial period observed between two pulses drifts from 4.6ns to about 5.4ns over 4.5ms of operation. This corresponds to a slow drift of the quantization step or a reduction in the CT-ADC gain which may have a negative impact on CT-ADC performance parameters such as SNR, when measuring the output over long periods of time. In the case of low power wireless receivers however, thanks to duty-cycling, the circuit is not expected to remain powered for long enough durations, necessary for the previously described phenomenon to occur.




FIGURE 3.22: Measured drift of the time between two consecutive output pulses for a constant DC input; this time is proportional to the instantaneous value of Δ .

We thus conclude that, assuming we can tolerate a threshold drift of 10% (which has a negligible effect on the CT-ADC characteristics such as average event rate, power consumption etc.), the initial threshold corresponding to a period of 4.6ns is allowed to drift up to a value corresponding to a period of 5ns equivalent to 2ms of operation between two consecutive phases when the threshold is refreshed. This gives us a timing window in which 200 bits can be demodulated, assuming an input data-rate of 100kbps. Furthermore, it will be seen later that the total time the CT-ADC is offline due to threshold resetting is much smaller than the effective operation time (of 2ms): less than 2 bits are expected to be lost.

For an application requiring continuous operation with a fixed threshold, the previously presented comparator threshold drift issue can be easily solved at the expense of an increased power consumption. One solution is to use two capacitors for each comparator: one could be used in the main comparator path while the other is charged. Alternatively, the capacitor voltage drop-off could be sensed and corrected continuously, using an analog control loop.

3.3.2 Transconductance

The role of the transconductor is to transform the input voltage difference in a proportional current over a bandwidth equal to that of the input, [10MHz–50MHz]. We choose a G_m architecture based on a differential pair with an active load (M_1 and M_2) controlled by a common mode feedback circuit. The output common mode is sensed through the net V_{cmfb} and injected in a source follower circuit (M_8 and R_4) to control the gates to the active load. The transistor level implementation is presented in Figure 3.23, while the component values are detailed in Table 3.2.




FIGURE 3.23: Transistor level implementation of the transconductance.

TABLE 3.2: Sizing of the $G_m - C$ components.

component	type	value	back-biasing
M_1	PMOS	$1.28\mu\text{m}/400\text{nm}$	-0.75V
M_2	PMOS	$1.28\mu\text{m}/400\text{nm}$	-0.75V
M_3	NMOS	$1.8\mu\text{m}/190\text{nm}$	0.75V
M_4	NMOS	$1.8\mu\text{m}/190\text{nm}$	0.75V
M_5	NMOS	$1.5\mu\text{m}/500\text{nm}$	0.75V
M_6	NMOS	$1.5\mu\text{m}/500\text{nm}$	0.75V
M_7	NMOS	$1.5\mu\text{m}/500\text{nm}$	0.75V
M_8	NMOS	$4\mu\text{m}/100\text{nm}$	0.75V
C_1	capacitance	30fF	n/a
C_2	capacitance	30fF	n/a
R_1	resistor	350k Ω	n/a
R_2	resistor	350k Ω	n/a
R_3	resistor	60k Ω	n/a
R_4	resistor	220k Ω	n/a

For the input signal characteristics we have chosen a common mode of 500mV with a maximum peak-to-peak swing of 200mV. The common mode allows us to bias the input differential pair, made up of M_3 and M_4 , in the subthreshold regime. This gives us a high g_m/I_d at the cost of a reduced linearity. The linearity is increased using source degeneration: separating the tail current of the differential pair with the resistance

R_3 . Concerning the input peak-to-peak swing, it has been chosen based on the voltage transfer function of the $G_m - C$: at maximum input swing, V_{C+} and V_{C-} should achieve high enough slopes to minimize the comparator delay, given in Figure 3.20 on page 69.

Transfer Function

To determine the transfer function of the $G_m - C$ block we start by stating that the comparator input is high impedance, therefore, connecting it to the output of the $G_m - C$ cell does not alter the latter's transfer function. Furthermore, node V_{cmfb} acts as an AC ground, therefore the AC output current of the differential pair sees a load made up of a capacitor (C_1) in parallel with a resistor (R_1), neglecting the output impedance of the transconductance stage. Consequently, the voltage transfer function of the $G_m - C$, defined in equation 3.7 and plotted in Figure 3.24, has a low-pass behavior with a DC gain determined by the value of the reference current injected in the I_{ref} terminal. For all values of I_{ref} the cutoff frequency is at 15MHz.

$$TF_{gmc} = 20 \cdot \log_{10} \left(\frac{V_{c+} - V_{c-}}{V_{in+} - V_{in-}} \right) \quad (3.7)$$




FIGURE 3.24: Voltage transfer function of the proposed $G_m - C$ implementation for different values of I_{ref} .

Linearity

Depending on the input frequency which sets the voltage gain, according to Figure 3.24, two different mechanisms can limit the linearity of the transconductor. On the one hand, low frequency signals undergo a high voltage gain and consequently have a low input

compression point: the elevated output swing can push the input differential pair out of saturation ($V_{ds} < V_{ds-sat}$). This is unlikely to occur for high frequency signals as the corresponding voltage gains are much lower. This allows for a higher input swing which can trigger the second non-linearity inducing mechanism: the subthreshold behavior of the input differential pair which has $V_{gs} < V_{th}$. Figure 3.25 plots the voltage gain of the proposed transconductor with its capacitive load versus the input peak-to-peak amplitude for different frequencies. As expected, the input compression point, corresponding to a gain degradation of 1dB, increases with frequency.




FIGURE 3.25: Degradation of the voltage gain versus the input peak-to-peak swing at different input frequencies.

It is important to note that, for the previous analysis, the $G_m - C$ has been used in stand-alone conditions. However, the flipping operation of the CT-ADC guarantees that its output swing is limited to $\pm\Delta/2$ around its common mode. This means that, for values of Δ in the tens of millivolts, the drain-source voltage of the input differential pair will *always* be greater than the saturation margin, V_{ds-sat} . Consequently, the $G_m - C$ linearity can only be limited by the subthreshold behavior of M_3 and M_4 . To measure the linearity of the $G_m - C$ in more realistic conditions we choose a different simulation setup: we apply at the input a unit step voltage difference and we measure the current injected in the output capacitors immediately after applying the change at the input. This mimics the flipping behavior of the CT-ADC “front-end”. We thus guarantee that V_{C+} and V_{C-} are around the common mode when the current is measured. The resulting g_m , normalized to its maximum value (equation 3.8), is plotted in Figure 3.26 versus the input unit step amplitude for different values of the degeneration resistance R_3 .

$$g_m(dB) = 20 \cdot \log_{10} \left(\frac{g_m(V_{inp-p})}{\max(g_m(V_{inp-p}))} \right) \quad (3.8)$$

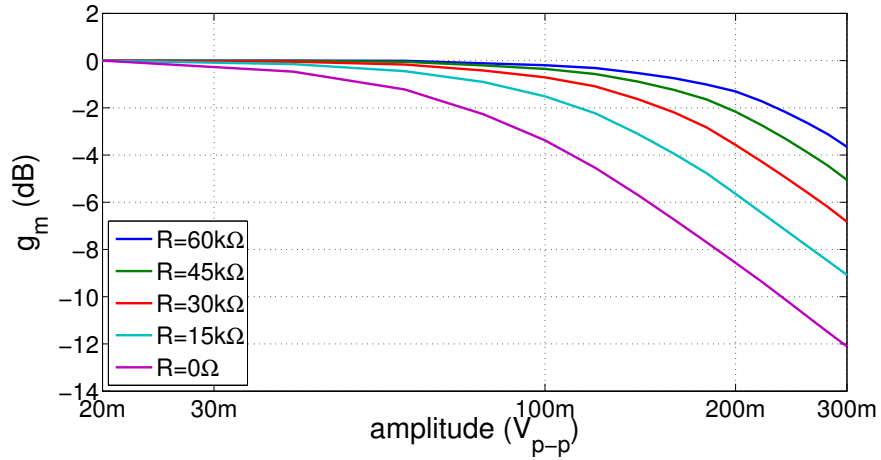


FIGURE 3.26: g_m normalized to its maximum value versus the input voltage difference for different values of the degeneration resistance.

In the final implementation we have chosen an R_3 of $60\text{k}\Omega$, which allows us to achieve a 1dB input compression point at 190mV of input peak-to-peak swing, only 10mV off the original target of 200mV.

Output Common Mode

Finally, we study the differential DC offset observed on the nets V_{C+} and V_{C-} , in order to determine the number of reference voltages required to program the two threshold levels $\Delta/2$ and Δ . Ideally, the differential DC offset observed at the output of the $G_m - C$ is much smaller than Δ , thus enabling us to use a single reference voltage V_{prog1} for both core comparators and another one, V_{prog2} , for both overflow comparators, as given in equation 3.9 and equation 3.10.

$$V_{prog1} = V_{cm+} + \frac{\Delta}{2} \simeq V_{cm-} + \frac{\Delta}{2} \quad (3.9)$$

$$V_{prog2} = V_{cm+} + \Delta \simeq V_{cm-} + \Delta \quad (3.10)$$

The histogram of the output differential DC offset ($V_{cm+} - V_{cm-}$) observed over 100 MC simulations is plotted in Figure 3.27; the standard deviation observed is 37mV. We thus conclude that the initial supposition of approximately equal common modes is false as the differential DC offset is usually higher than a standard Δ value (40mV). Consequently, the previously described threshold setting scheme requires four different reference voltages, as defined by equation 3.11 – equation 3.14, thereby greatly increasing the complexity of the CT-ADC control circuit.

$$V_{prog1} = V_{cm+} + \frac{\Delta}{2} \quad (3.11)$$

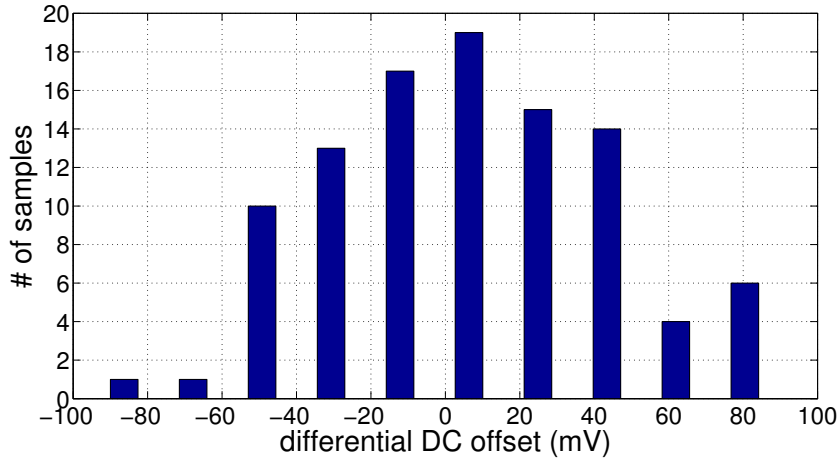


FIGURE 3.27: Sampled PDF of the $G_m - C$ output differential DC offset over 100 MC simulations.

$$V_{prog2} = V_{cm-} + \frac{\Delta}{2} \quad (3.12)$$

$$V_{prog3} = V_{cm+} + \Delta \quad (3.13)$$

$$V_{prog4} = V_{cm-} + \Delta \quad (3.14)$$

To simplify the testing of the CT-ADC, a new, simpler threshold management circuit is proposed.

3.3.3 Threshold Management

In this section we propose a new threshold setting mechanism which enables us to automatically set-up all comparators using a single reference current which is injected from the exterior and which is directly proportional to Δ . Furthermore, this allows us to easily tune the value of Δ thereby offering more flexibility to the CT-ADC.

The proposed threshold setting mechanism is presented in Figure 3.28 and is based on a sequential process consisting of four steps. For simplicity, in this paragraph we only follow the evolution of different internal nodes corresponding to the two core comparators, which we call *comp1* and *comp2*; however this can be easily extended to the two overflow comparators as well.

- In the first step the two $G_m - C$ inputs are disconnected and a voltage corresponding to their common mode is applied instead. The two $G_m - C$ outputs thus become equal to their common mode which we call V_{cm+} and V_{cm-} . During this step CMOS switches S_1 and S_3 are closed while S_2 is open; thus the output common modes of the $G_m - C$ are sampled on the input capacitance of the two comparators.




FIGURE 3.28: Proposed threshold setting mechanism.

The voltage across these capacitors (V_C^{comp1} and V_C^{comp2}) becomes equation 3.15 and equation 3.16.

$$V_C^{comp1}(T_1) = V_{cm+} - V_{trip}^{comp1} \quad (3.15)$$

$$V_C^{comp2}(T_1) = V_{cm-} - V_{trip}^{comp2} \quad (3.16)$$

- In the second phase the $G_m - C$ output is disconnected from the comparators by opening S_1 while closing S_2 and S_3 . Instead of applying a static voltage V_{prog} to set the threshold of the comparators, we inject a current I_{th} for a duration of δT . For now, this current is injected only on the plates of the capacitor from the first comparator ($comp1$). The voltage on the capacitors can thus be rewritten as equation 3.17 and equation 3.18.

$$V_C^{comp1}(T_2) = V_{cm+} - V_{trip}^{comp1} + \frac{I_{th} \cdot \delta T}{C_1} \quad (3.17)$$

$$V_C^{comp2}(T_2) = V_{cm-} - V_{trip}^{comp2} \quad (3.18)$$

- The third phase consists of repeating the operation from phase two, but this time on $comp2$. The key in achieving matching thresholds is reusing the same charge pump for I_{th} and the same delay cell used to generate δT_{th} . The two comparator voltages now become: equation 3.19 and equation 3.20.

$$V_C^{comp1}(T_3) = V_{cm+} - V_{trip}^{comp1} + \frac{I_{th} \cdot \delta T}{C_1} \quad (3.19)$$

$$V_C^{comp2}(T_3) = V_{cm-} - V_{trip}^{comp2} + \frac{I_{th} \cdot \delta T}{C_1} \quad (3.20)$$

- The last phase consists in connecting V_{in+} and V_{in-} to the inputs of the $G_m - C$ cell and also connecting the $G_m - C$ output to the comparator input (S_1 closed and

S_2 and S_3 open). By identifying the different parts of equation 3.5 on page 67 with equation 3.19 and equation 3.20 we deduce that the threshold of both comparators is now $I_{th} \cdot \delta T_{th}/C_1$ above *whatever* common mode is presented at their input.

The threshold of the overflow comparators can be set using the same mechanism as before and by mirroring I_{th} with the correct ratio of 2, thus obtaining I_{ovf} . Process and temperature variations might affect I_{th} , I_{ovf} and δT giving us an uncertainty on the absolute value of Δ . However the proposed mechanism guarantees matching between comparator thresholds as well as obtaining the correct ratio between the overflow and core comparator thresholds. Noise on the supply rails can be a source of mismatch, however, its effects can be drastically limited through careful routing and decoupling. Alternatively, all thresholds could be set up in the same δT window using four current sources instead of two. In this case, we are trading mismatch caused by supply noise for the limited matching achieved by current mirrors.

As seen in Figure 3.28 the threshold setting phase takes only $1.5\mu s$ compared to the total time between two consecutive resets of the comparators which occurs every 2ms.

Charge Pump

The implementation of a switchable charge pump, required for generating and injecting I_{th} and I_{ovf} , is presented in Figure 3.29; the corresponding transistor sizes are also given in Table 3.3. Transistor M_6 is used to switch ON and OFF the charge pump, while current mirrors (M_1, M_2) and (M_3, M_4, M_5) are used to generate I_{th} and I_{ovf} from a reference current, I_{ctrl} injected from outside the chip.




FIGURE 3.29: Design of the charge pump used to set comparator thresholds.

Sequential Logic

The last element missing in the implementation of our proposed CT-ADC is the sequential logic required for driving the threshold setting mechanism. An active on low pulse

TABLE 3.3: Sizes of various components used in the design of the charge pump.

component	type	value	back-biasing
M_1	NMOS	$1\mu\text{m}/1\mu\text{m}$	0V
M_2	NMOS	$1\mu\text{m}/1\mu\text{m}$	0V
M_3	PMOS	$5\mu\text{m}/1\mu\text{m}$	0V
M_4	PMOS	$2\mu\text{m}/1\mu\text{m}$	0V
M_5	PMOS	$4\mu\text{m}/1\mu\text{m}$	0V
M_6	PMOS	$200\text{nm}/60\text{nm}$	0V

injected from the exterior, which we call ZCD_{rst} , is used to set into motion the entire threshold setting scheme; the period between these pulses thus defines the CT-ADC reset period. A schematic representation of the control signals thus generated is presented in Figure 3.30. The common mode of the $G_m - C$ is sampled while CM_{sample} is active, the upper comparator thresholds (core and overflow) are set when TH_1 is active, the lower comparator thresholds (core and overflow) are set while TH_2 is active; the CT-ADC is in “active” mode while the signal labeled as *convert* is at 1.




FIGURE 3.30: Reset and control signals used for the proposed threshold setting mechanism.

The detailed implementation of this specific block is not discussed since it consists only of trivial sequential logic blocks and delay cells. For the implementation of the delay cells we used a series of inverters loaded by capacitors. The achieved delay value is not important and there are no matching requirements since the actual CT-ADC threshold will be controlled by a reference current injected into the chip, I_{ctrl} .

3.3.4 Breakdown of the CT-ADC Power Consumption

For a single tone full-scale input, with $\Delta = 80\text{mV}$ and $I_{ref} = 5\mu\text{A}$ the simulated CT-ADC power consumption is $24\mu\text{W}$. A breakdown of this power consumption according to the

different CT-ADC blocks is given in Figure 3.31.

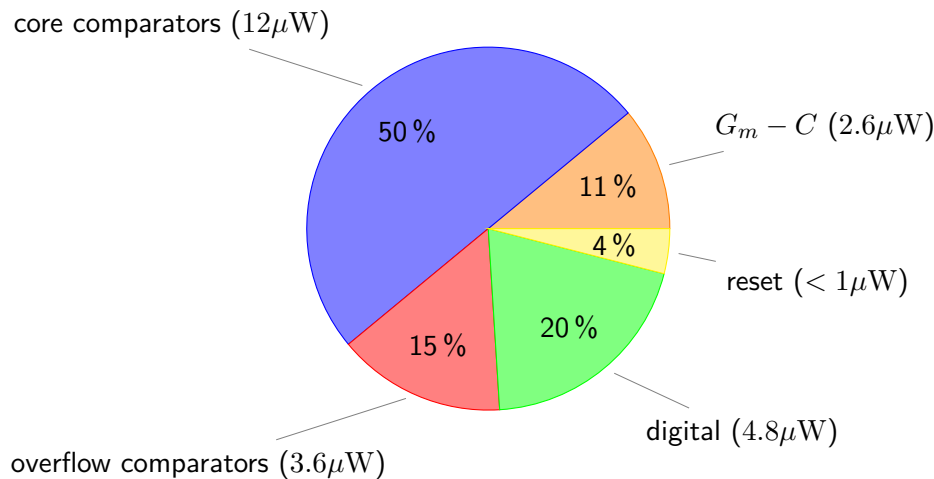


FIGURE 3.31: Breakdown of the CT-ADC power consumption.

3.4 Measurement Results

In this section we present the CT-ADC measurement results as well as a comparison with other CT-ADCs and sampled ADCs with similar bandwidths.

3.4.1 Single Tone Input: Noise

Based on results from Section 1.6 on page 28 the signal-to-noise ratio (SNR) required for our CT-ADC is above 30dB over the bandwidth [10MHz–50MHz].

Spectrum

An example of the output spectrum is given in Figure 3.32a for an input tone at 10MHz and Figure 3.32b for a 50MHz tone.

As expected, the output spectrum contains the input tone along with first order shaped harmonics and component noise, as predicted by Section 3.2.3 on page 61.

SNR and SNDR

To find the SNR (SNDR) we compute the ratio of power between the useful signal and the noise (and distortion) integrated between 10MHz and 50MHz. The resulting SNR




FIGURE 3.32: Spectrum of the signal observed at the output of the proposed CT-ADC for different input frequencies

and SNDR for $\Delta = 80\text{mV}$ and $I_{ref} = 2\mu\text{A}$ are plotted versus the frequency of the input tone in Figure 3.33. For low frequencies, input signal harmonics fall inside the ADC bandwidth and slightly degrade the SNDR with respect to the SNR. We can see that the initial specifications are met, since the SNR stays above 33dB over the entire bandwidth. The observed power consumption is constant, around $24\mu\text{W}$ regardless of the input frequency, as explained in Section 3.2.3 on page 61.




FIGURE 3.33: The output SNR and SNDR versus the input tone frequency.

Alternatively, we can plot the SNR and SNDR versus input signal amplitude normalized to the full scale (in dB). In this case, by reducing the amplitude of the input signal, we also reduces the frequency of flipping events observed at the output which reduces the power consumption of the ADC. This is plotted in Figure 3.34 for an input tone of 10MHz, $\Delta = 80\text{mV}$ and $I_{ref} = 5\mu\text{A}$. Despite reducing the input amplitude by 14dB with respect to the full scale of 200mV, the power reduces only by a factor of 1.6. This can be explained by the fact that only the digital part of the CT-ADC scales its power

with the input amplitude; the analog part draws a constant DC power regardless of the characteristics of the input.

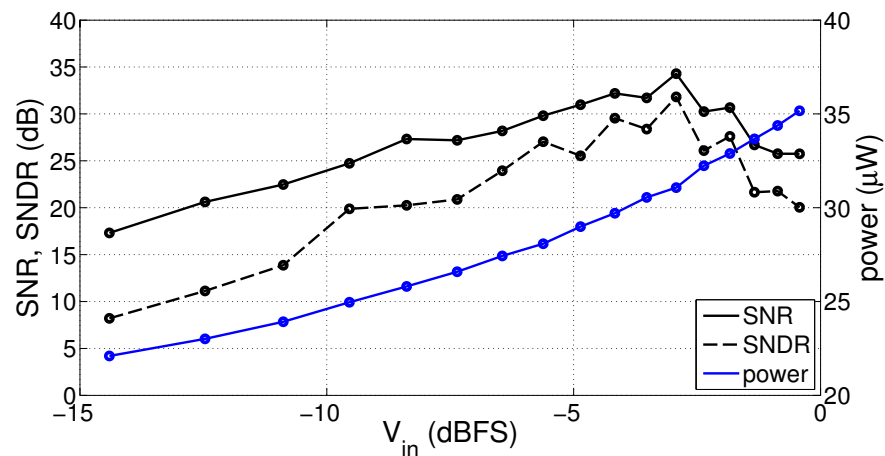


FIGURE 3.34: The output SNR, SNDR as well as the ADC power consumption versus the input signal amplitude (normalized to full scale).

Aliasing

To verify the alias free operation of the proposed CT-ADC, we test it with an out-of-band tone located at 60MHz. The resulting spectrum is plotted in Figure 3.35; no degradation of the in-band spectrum is observed (noise or aliases).




FIGURE 3.35: The output spectrum for an input consisting of an out-of-band tone located at 60MHz; no aliasing is observed.

Comparison with State of the Art

We start by comparing our CT-ADC with existing discrete time implementations which have similar bandwidths. The results are presented in Table 3.4. As a figure of merit (FoM), we have chosen the Walden FoM since it is the most adapted for comparing low power ADC implementations; its equation is given in equation 3.21 with P – the ADC power, ENOB – the effective number of bits and f_{BW} – the ADC bandwidth (40MHz in our case).

TABLE 3.4: Comparison with state of the art discrete time ADCs with bandwidths smaller than 100MHz.

	Yoshioka[71]	Tsai[72]	Van der Plas[73]	Brooks[74]	this work[PRMT15a]
technology	40nm CMOS	90nm CMOS	90nm CMOS	180nm CMOS	28nm UTBB FDSOI CMOS
supply	0.7V	1V	1V	1.8V	0.65V
input bandwidth	12.3MHz	20MHz	75MHz	100MHz	40MHz (10MHz–50MHz)
implementation	discrete time	discrete time	discrete time	discrete time	continuous time
sampling rate	24.6MS/s	40MS/s	150MS/s	200MS/s	signal-dependent
core area	0.0058mm ²	0.055mm ²	0.0625mm ²	0.05mm ²	0.0032mm ²
SNDR	44.2dB	44.5dB	40dB	40.3dB	32dB–42dB
total power	54.6μW	113μW	133μW	8.5mW	24μW
FoM_W	17fJ/ conv-step	20fJ/ conv-step	10.9fJ/ conv-step	503.3fJ/ conv-step	3–10fJ/ conv-step
antialiasing filter required?	yes	yes	yes	yes	no

$$FoM_W = \frac{P}{2^{ENOB} \cdot 2 \cdot f_{BW}} \quad (3.21)$$

The comparison with existing CT-ADCs is presented in Table 3.5. The reader is advised to take these results with a grain of salt as previous implementations have targeted different bandwidths as well as different ENOBs.

TABLE 3.5: Comparison with existing continuous time ADCs.

	Schell[68]	Kurchuk[37]	Weltin-Wu[45]	this work[PRMT15a]
technology	90nm CMOS	65nm CMOS	130nm CMOS	28nm UTBB FDSOI CMOS
supply	1V	1.2V	1V	0.65V
input bandwidth	10kHz	2.4GHz (0.8GHz–3.2GHz)	20kHz	40MHz (10MHz–50MHz)
core area	0.06mm ²	0.0036mm ²	0.36mm ²	0.0032mm ²
SNDR	58dB	20.3dB	47 – 54dB	32dB–42dB
total power	50μW	2.7mW	2 – 8μW	24μW
FoM_W	3.14pJ/ conv-step	66fJ/ conv-step	200 – 850fJ/ conv-step	3–10fJ/ conv-step

The figures of merit achieved by our implementation are among the best reported: Figure 3.36.

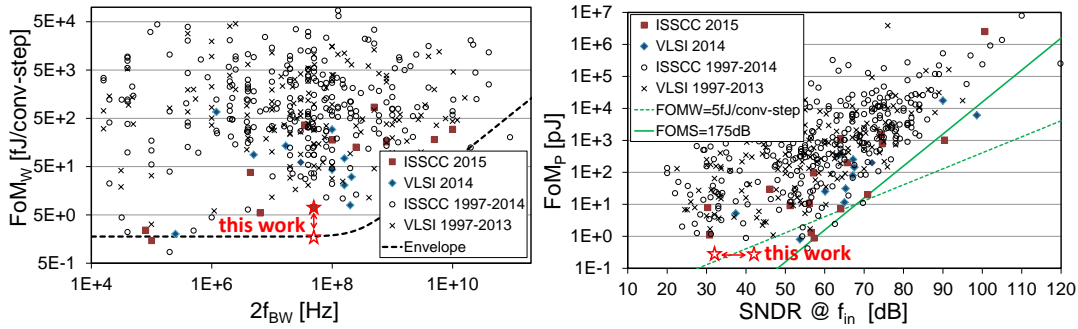


FIGURE 3.36: Walden and Energy figures of merit of recent state of the art ADC implementations.

3.4.2 Two Tone Input: Linearity

In Section 1.6 on page 28 we have showed that depending on the amount of out-of-band attenuation achieved by the CT-DSP (30dB or 40dB), the CT-ADC must have an SFDR of 40dB or 50dB. However, the loop architecture of the proposed DF-CT-ADC (Section 2.4 on page 45) enables us to “boost” the linearity of the CT-ADC. Given it is reasonable to assume an SFDR boost of at least 10dB, we conclude that our CT-ADC must have a linearity between 30dB and 40dB.

SFDR

The CT-ADC SFDR is measured by injecting an input signal consisting of two tones situated at the high end of the ADC bandwidth (48MHz and 50MHz). Due to loop delay constraints, this configuration is the most susceptible of showing a degraded SFDR caused by the overflow behavior described in Section 3.2.4 on page 62. The SFDR is thus plotted in Figure 3.37 versus the input signal amplitude for two values of Δ : 80mV and 140mV; for these measurements we have used an I_{ref} of $5\mu\text{A}$. We can see that, by tuning Δ through I_{ctrl} , the peak SFDR is achieved at different values of the input peak-to-peak swing. Effectively, the inherent tunability of the quantization step, which is controlled through only one current, fulfills a role similar to that of an automatic gain control block: instead of adapting the signal to the characteristics of the ADC, we adapt the ADC to the characteristics of the signal. Compared to a standard implementation, the proposed CT-ADC thus requires less AGC thereby saving power.




FIGURE 3.37: SFDR of the proposed CT-ADC versus the input peak-to-peak swing for two values of Δ .

Influence of the Back-bias

In this section we study the influence of the delay in the feedback loop on the linearity achieved by the ADC. Since this delay cannot be controlled directly, we choose to change it indirectly, by tuning the back-biasing of the *logic* block in Figure 3.15 on page 65. This enables us to increase the threshold voltages of the transistors used in the various logic blocks from an initial value corresponding to a back-biasing of 2V, to a higher value, corresponding to a back-biasing going as low as 0V. This translates to an increase of the propagation delay, observed in simulation and given in Table 3.6.

TABLE 3.6: Speed of the logic block for different values of the back-bias voltage.

back-bias voltage	NMOS threshold	PMOS threshold	delay
2V	186mV	166mV	40ps
1V	260mV	249mV	55ps
0V	335mV	332mV	83ps

SFDR variations, plotted in Figure 3.38, are obtained by applying a two tone input (48MHz and 50MHz), sweeping the input peak-to-peak amplitude and applying the previously defined controls to the back-bias terminals. Here, we use $\Delta = 80\text{mV}$ and $I_{ref} = 5\mu\text{A}$. For lower input amplitudes, the time between consecutive ADC events is large and a higher delay in the feedback loop can be tolerated: the SFDR obtained for all back-biases is the same. However, as the amplitude increases the time between consecutive ADC events diminishes and we observe slow feedback configurations (low back-biasing – high delay) start failing. We can thus conclude that the technology back-biasing option allowed us to boost the SFDR of the CT-ADC from a maximum of 28dB achieved when 0V are applied to the back-gate to 34dB corresponding to a back-bias of 2V.




FIGURE 3.38: SFDR of the proposed CT-ADC versus the input peak-to-peak swing for different values of the back-bias voltage.

3.5 CT-ADC Conclusion

The proposed CT-ADC achieves an improvement, in terms of FoM_W , of over 300 times, 3 to 12 times and 20 times respectively, better than previous CT-ADC implementations ([43], [60] and [75]). The CT-ADC is thus no longer the bottleneck in implementing IF stages for ultra low power radios. This paves the way for new CT-DSP approaches, with no sampling in time, to implement the receiver back-end. Such an architecture is presented in the next chapter.

Chapter 4

Power Scalable CT-DSP

RESULTS from the previous chapter show that it is possible to implement an energy efficient CT-ADC covering the entire band of interest [10MHz–50MHz]. However, this energy efficiency needs to be carried over to the CT-DSP for a successful implementation of an ultra low power IF stage, for applications like wake-up radios. In this chapter we present such a CT-DSP, consisting of two digital filters designed specifically to operate with the previously presented CT-ADC. The rest of this chapter is organized as follows: first, the CT-DSP architecture as well as its building blocks are specified. Next, we discuss the transistor level implementation of the CT delay cells, CT adder and various extra blocks required for the DF-CT-ADC. Finally, in the last part of this chapter we present transistor level simulation results for the entire system. At the time of writing, the chip has unfortunately not come back from the foundry (the delivery date has been delayed by 4 months) therefore no measurement results are available.

4.1 CT-DSP Architecture

The goal of this section is to completely specify the requirements of the CT-DSP system. According to the proposed WU-RX architecture the IF signals have a frequency support limited to 10MHz, corresponding to the width of the RF passband filter, which can be situated anywhere in the [10MHz–50MHz] band. Furthermore, results from the previous chapter show that the IF CT-ADC achieves an SNR of 32dB–42dB. Therefore, according to Section 1.6 on page 28, the CT-ADC’s dynamic range is sufficient for a robust reception of input signals characterized by an SIR between –20dB and –30dB, given a CT-DSP which achieves a rejection of 30dB to 40dB.

4.1.1 Dual FIR – IIR Implementation

The chosen CT-DSP architecture, presented in Figure 4.1, is based on a dual solution relying on the use of a CT-FIR coupled with a CT-IIR filter. On the one hand, the CT-IIR filter (DF-CT-ADC), described in Section 2.4 on page 45, attenuates out-of-band components from the input of the CT-ADC thereby boosting its linearity and reducing the average number of events it generates. On the other hand, the CT-FIR filter situated at the output of the CT-ADC has the role of increasing the rejection levels achieved by the CT-IIR to levels compatible with the previously derived specifications.

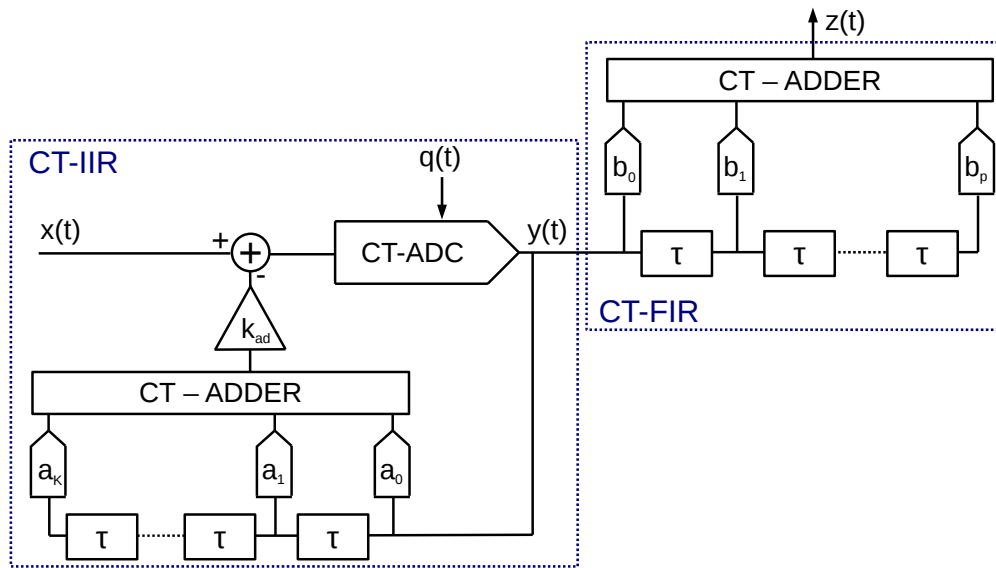


FIGURE 4.1: Architecture of the proposed CT-DSP.

The power consumption of the previously proposed filtering stage is strictly dependent on the order of the two filters (p and k), which, in turn, depends on the amount of rejection they achieve. In theory, an IIR filter is capable of achieving much higher rejection levels than an equivalent order FIR. It would thus seem that the best strategy is to maximize the rejection offered by the IIR in order to minimize (and maybe even remove) the FIR output filter. However, we show in the following paragraph that due to limitations concerning the IIR transfer function implementation, the previous argument is not always valid.

To illustrate this, we take as an example a fourth order IIR with the following coefficients: 0.75, 0.5, -0.5 , 0.25; we also consider that the adder has a conversion gain of k_{ad} and that the ADC has a unity transfer function over the frequency band of interest. The IIR transfer function can thus be written as given in equation 4.1. The frequency response related to this transfer function is now plotted in Figure 4.2 for different values of the adder conversion gain k_{ad} . As expected, the total rejection, defined as the difference

between the maximum and the minimum of the frequency response, increases with k_{ad} and, in theory, can be arbitrarily high. However, this happens by boosting the in-band signal gain rather than attenuating out-of-band components, which has the negative effect of increasing the signal swing seen by the CT-ADC thereby, increasing its power consumption. We therefore conclude that, for a given IIR filter order, there exists a limited amount of “useful“ rejection the filter is capable of achieving. In this case, simulations show that the filter is limited to about 10dB of out-of-band attenuation. Beyond that point, increasing the gain in the feedback loop (k_{ad} in our case) has the effect of only slightly improving the out-of-band attenuation while greatly increasing the in-band signal gain at the expense of the digital circuitry’s power consumption.

$$\frac{Y(s)}{X(s)} = \frac{1}{1 + k_{ad} \cdot FIR(s)} = \frac{1}{1 + k_{ad} \cdot (0.75 + 0.5e^{\tau s} - 0.5e^{2\tau s} + 0.25e^{3\tau s})} \quad (4.1)$$




FIGURE 4.2: Example of the IIR transfer function for different values of the adder conversion gain, k_{ad} .

For this implementation, we settle for a 3^{rd} order IIR followed by a 9^{th} order FIR. Together, these two filters should provide us with sufficient rejection required by the targeted interferer levels.

4.1.2 CT Digital Filter Design

We minimize the power consumption of the CT-DSP by using the spectrum repetition property of CT digital filters’ transfer function, as explained in Section 2.3 on page 41. The transfer function of either of the two filters thus becomes similar to the one given in Figure 4.3: lowpass FIR implementation creates passbands around every integer multiple of $1/\tau$.




FIGURE 4.3: Example of a CT-DSP frequency response along with a possible configuration of the IF signal.

The repeated spectrum solution also simplifies the design in terms of tunability: LO uncertainties result in an IF signal which can be anywhere inside the [10MHz–50MHz] band, thereby requiring sufficient tunability of the CT-DSP to cover an equivalent frequency band. However, using the proposed solution, it can be seen that by changing τ from 100ns to 66ns the passband frequency of the first lobe is shifted from 10MHz to 15MHz, the second lobe – from 20MHz to 30MHz, the third lobe – 30MHz to 45MHz and lastly, the fourth lobe from 40MHz to 60MHz. In the end, this small change in the value of the tap delay cells, enables us to cover the entire ADC bandwidth, except the [15MHz–20MHz] band. This can be solved by reconfiguring the coefficients of the filter in order to obtain a highpass behavior instead of the lowpass one. Consequently, passbands are created around $(1 + 2K)/(2\tau)$, with K – an integer. It can be easily seen that varying τ from 100ns to 66ns, the first spectrum repetition passband shifts from 15MHz to 20MHz; thus the entire ADC bandwidth is covered by the previously defined delay cell variation.

Furthermore, the repeated spectrum solution allows us to scan several frequency bands at one time, thereby reducing the total scan time. Instead of varying τ so that $1/\tau$ covers the entire 40MHz bandwidth, we only have to change τ such that $1/\tau$ varies from 10MHz to 15MHz twice (once for the lowpass scan and once for the highpass scan). The effective scan time is thus divided by a factor of 4. Note that this solution works because we know that, thanks to the RF front-end filter, the frequency support of the signal at IF is limited to 10MHz. The uncertain IF architecture then downconverts it to an unknown IF, situated between 10MHz and 50MHz.

The only drawback of the low $1/\tau$ solution presented previously is related to the fact that noise from all IF passbands is injected into the output envelope detector (Figure 1.18 on page 20) thereby slightly degrading the overall sensitivity of the radio. However, results from [38] detailed in Appendix A on page 145, show that for low data-rate systems, the demodulation sensitivity is mostly limited by noise situated around the signal rather than by wideband noise downconverted from the IF and integrated in the baseband. The sensitivity degradation related to this phenomenon is thus negligible.

4.1.3 Architecture Simulation

Before discussing the transistor level implementation of the previously proposed system, a set of specifications is defined for the different blocks of the CT-DSP.

Delay Cells

The delay cells deviate from their ideal behavior by two mechanisms:

- mismatch – random, constant deviations from the ideal value
- jitter – random, time varying deviations from the ideal value

Delay mismatch does not add any noise to the output, but has the effect of changing the transfer function of the filter. On the other hand, jitter does not affect the transfer function of the filter, but adds extra noise to the output. By design, we must ensure that this extra noise is small enough compared to the noise of the CT-ADC, otherwise risk degrading the overall sensitivity of the WU-RX.

To study the effects of both jitter and mismatch, we conduct a series of behavioral simulations of the output FIR filter, since its higher order is more susceptible of negatively impacting the characteristics of the output signal. The performance metrics of interest are the amount of deviation from the ideal transfer function (in the case of mismatch) and the output SNR (in the case of jitter).

We start studying the specifications of the CT-DSP delay cells by measuring the effects of mismatch in their values. Any deviation from the standard value results in a slight shift of the respective transfer function pole in the complex plane, which changes the overall transfer function of the filter. In the end, deviations from the ideal transfer function (which has been computed based on a filtering mask) shift the frequency response of the real filter outside of the previously specified mask, as shown in Figure 4.4.

The changes in the frequency response can be measured by extracting the following three parameters:

- δF_{pass} , the deviation from the ideal passband: the frequency difference between the ideal passband frequency (in this case $F_{pass} = 1\text{MHz}$) and the one obtained from a simulation with mismatch, defined by -3dB of rejection.
- δF_{stop} , the deviation from the ideal stopband: the frequency difference between the ideal stopband frequency (in this case $F_{stop} = 2\text{MHz}$) and the one obtained from a simulation with mismatch, defined by -30dB of rejection.

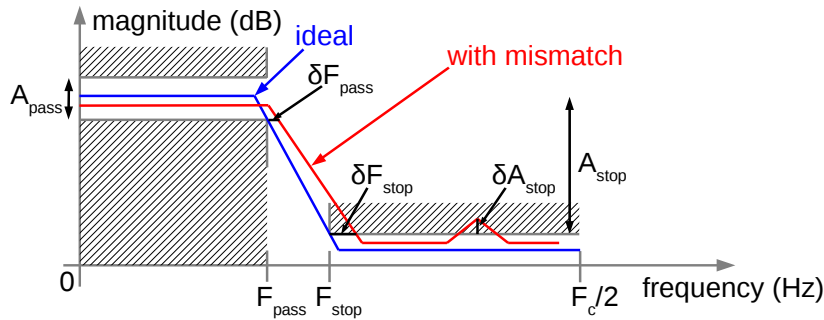



FIGURE 4.4: Changes observed by adding mismatch to an ideal CT-FIR transfer function.

- δA_{stop} , the deviation from the ideal rejection over the stopband: the difference in rejection between the ideal case (30dB) and the one obtained from a simulation with mismatch, over a bandwidth starting from the previously extracted F_{stop} .


To quantify these parameters, a series of simulations is done in which the values of the delay cells are obtained by sampling a uniform random variable, described by its average value (100ns in our case) and its standard deviation, which we call σ_τ . Next, behavioral level MC simulations of the 9th order FIR are conducted, which enable us to extract F_{pass} , F_{stop} and A_{stop} for different values of σ_τ , ranging from 0s to 1200ps. Results are plotted in Figure 4.5 on the facing page.

Note that the filter transfer function is studied *around* 50MHz, meaning that the absolute frequency of F_{pass} is shifted to 49MHz (corresponding to a difference of 1MHz with respect to the center frequency) and F_{stop} is shifted to 48MHz.


The uncertainty of the passband frequency, portrayed in Figure 4.5a on the next page, has a small effect on the filtering performance of the CT-DSP: since the signal data-rate is only 100kbps, most of its energy is located in several hundreds of kHz of bandwidth, much lower than the initial target of a total of 2MHz of bandwidth corresponding to an F_{pass} of 49MHz. Changes in F_{stop} , given in Figure 4.5b on the facing page, degrade the interferer rejection performance of the CT-DSP. Without any IF filtering, the effective bandwidth of the system is defined by the RF front-end filter, and is equal to 10MHz; consequently, the deviation of the stopband frequency from its initial target of 2MHz ($F_{stop} = 48\text{MHz}$) must be minimized, otherwise the CT-DSP is rendered useless, achieving a bandwidth similar to the one of the RF front-end filter. We settle for a stopband of 4.5MHz, equivalent to a minimum F_{stop} of 45.5MHz. Finally, A_{stop} (plotted in Figure 4.5c) directly influences the total interferer rejection. The initial target for the entire CT-DSP (CT-FIR and DF-CT-ADC) is 40dB. Estimating the contribution of the DF-CT-ADC to about 20dB of rejection, we conclude that A_{stop} for the CT-FIR must be above 20dB. Consequently, the maximum delay standard deviation (σ_τ) which can be tolerated is around 700ps.



(a) Degradation of F_{pass} for a center frequency of 50MHz; error bars indicate the performance of 90% of runs.



(b) Degradation of F_{stop} for a center frequency of 50MHz; error bars indicate the performance of 90% of runs.



(c) Degradation of A_{stop} ; error bars indicate the performance of 90% of runs.

FIGURE 4.5: Effects of delay cell mismatch on the transfer function of a 9th order FIR filter.

This value will serve as a specification for the delay cell calibration scheme, described later on in this chapter.

The second specification of the delay cells is related to the noise (jitter) they generate. This jitter is represented as a normal random variable which changes the value of the time delays affected to each event. These deviations occur randomly from the standard delay value and are different for the same delay cell from one event to another. Since the jitter has an average of 0 it does not affect the filter transfer function; on the other hand, it raises the noise floor at the FIR output by an amount proportional to the number of delay cells and to the standard deviation of their jitter, which we call σ_{τ}^{jitter} . Consequently, the delay cells must be designed so that the output of the FIR filter does not experience a significant rise in the noise floor compared to that observed at the CT-ADC output.

The noise introduced by the delay cells is not signal-dependent. Each pulse generated by the CT-ADC can be thought of as a signal element which is injected into the CT-FIR where it experiences jitter. Since all pulses are identical, regardless of being issued by a low frequency or a high frequency signal, it follows that they are affected by noise in the time domain in the same way.

To quantify the effects of jitter, we propose a series of simulations in which results issued from an “ideal” CT-ADC with a single tone 50MHz input are injected in the passband of the 9th order FIR filter designed previously. The delay cells have a nominal value of 100ns on top of which we add a random variable of average 0 and standard deviation σ_{τ}^{jitter} swept from 10ps to 400ps. The FIR output in-band SNR is plotted against the value of the σ_{τ}^{jitter} and compared to the measured CT-ADC SNR (which is between 32dB and 42dB). Results are plotted in Figure 4.6; we thus conclude that the delay cell jitter must be kept under 325ps in order to avoid adding supplementary noise to the FIR output.




FIGURE 4.6: Effects of delay cell jitter on the noise floor of the CT-DSP output.

Weighted Adder

The CT adder can deviate from its ideal behavior by two mechanisms:

- *coefficient quantization* – the previous study was done using MATLAB which employs 64 bit floating point numbers, the silicon implementation of the CT-adder will have its coefficients quantized over a finite number of bits.
- *coefficient mismatch* – for an analog implementation of the CT adder, likely to be used given the relatively high input frequency and low power budget of our application, the coefficients will inevitably be affected by mismatch, making them deviate from their nominal value.

Coefficient quantization imposes certain limits on the design of the CT-DSP transfer function: reducing the set of numbers from which coefficients can be chosen limits the best-case transfer function achieved by the filter. Ideally, we should minimize the number of bits used for the representation of the filter coefficients while guaranteeing sufficient stopband attenuation thereby minimizing the power consumption of the adder. To determine the number of bits required for the coefficient representation, we apply the minimum order equiripple FIR design method on the 9th order CT-FIR and plot the resulting transfer functions for different types of coefficient representations: 64 bit double precision floating point or fixed point over a various number of bits (between 2 to 4 bits). The resulting transfer functions are plotted in Figure 4.7.




FIGURE 4.7: CT-FIR transfer functions obtained using different representations of its coefficients.

Four bit coefficients achieve a transfer function very close to the original one; however, it can be seen that using only two bit coefficients allows us to achieve a rejection above 27dB. Moreover, for the design of highpass transfer functions, we also require *signed*

coefficients since the only way of implementing a transfer function zero near 0Hz is by using FIR filters with coefficients which have a sum of 0. Consequently, for the coefficient representation of both the IIR and the FIR filters, we use two bit signed numbers.

Finally, mismatch will occur in the value of the CT-DSP coefficients since an analog implementation of the adder is expected to be less power consuming. For a CT-ADC generating 10 events during a period of 20ns (corresponding to 1/50MHz) and injecting them into a 9th order CT-FIR, the CT adder would have to solve additions at an average frequency of $9 \cdot 10 \cdot 50\text{MHz} = 4.5\text{GHz}$, thereby requiring a large power budget if implemented digitally. Mismatch in the value of the CT-DSP coefficients affects the operation of the CT-DSP in a manner similar to the delay cell mismatch: it does not add any noise but it alters its transfer function. To measure this, the procedure used for the study of delay cell mismatch is repeated: the coefficient values are modeled as normal random variables around their nominal value. F_{pass} , F_{stop} and A_{stop} are used as performance metrics (transfer function is studied around 50MHz). Their values, observed over 100 MC simulations for different values of the mismatch standard deviation, are plotted in Figure 4.8.


By setting a limit of 20dB to A_{stop} , we conclude that the maximal allowed mismatch in the value of the 9th order FIR coefficients is around 9%, higher than the matching achieved by analog design kit components which are usually used to implement analog CT adders.

4.1.4 CT-DSP Specifications


The CT-DSP specifications derived previously are summarized in Table 4.1.

TABLE 4.1: CT-DSP specifications.


IIR	order	3
	configuration	lowpass/highpass
	minimum τ	66ns
	maximum τ	100ns
	rejection	20dB
FIR	order	9
	configuration	lowpass/highpass
	minimum τ	66ns
	maximum τ	100ns
	rejection	20dB
delay cells (100ns)	matching (st. dev.)	700ps
	jitter (st. dev.)	325ps
coefficients	# of bits	2 + 1 (sign bit)
	matching (st. dev.)	9%



(a) Degradation of F_{pass} for a center frequency of 50MHz; error bars indicate the performance of 90% of runs.



(b) Degradation of F_{stop} for a center frequency of 50MHz; error bars indicate the performance of 90% of runs.



(c) Degradation of A_{stop} ; error bars indicate the performance of 90% of runs.

FIGURE 4.8: Effects of coefficient mismatch on the transfer function of a 9th order FIR filter.

4.2 CT Delay Cell

This section discusses design choices related to the implementation of the delay cells used for the CT-DSP.

4.2.1 State of the Art for Asynchronous Delay Cells

There exist a number of ways of implementing asynchronous delay cells. The first possibility consists of using allpass filters, which have a constant gain in the frequency domain over their entire passband and a phase delay such that analog input signals come out delayed in the time domain. This solution is more adapted for the delay of analog signals rather than asynchronous digital signals. The implementation of allpass filters requires a set of poles usually constructed around OTAs which require a static power even when the input is “silent”, an undesired behavior for CT-DSPs.

Since we are dealing with digital signals, a purely digital possibility would be to use a chain of inverters to achieve the required delay. Despite drawing no static power (except leakage), this solution is not optimal due to the high number of switching events triggered by each input transition, thereby resulting in a large dynamic energy spent per input event. Furthermore, this solution offers a poor programmability: the only way to change the value of the delay is to bypass a certain number of its inverters.

Opting for a mixed-signal solution relieves us from the previously presented issues. One could reduce the number of switching events in the chain of inverters by “slowing down” the signal using lowpass RC filters, as seen in Figure 4.9. Depending on the time constant of the RC filter, the number of inverters of the previous digital solution can be reduced to only 1, which converts the signal on top of the capacitor, which is analog, into a digital signal. The delay value of can be tuned by either changing the value of the resistor or that of the capacitor. The main issue of this implementation is that, depending on the delay value, the voltage on top of the capacitor can spend a long time around the trip point of the inverter, making it draw a significant crowbar current from the supply rails. Furthermore, the lack of a reset of the capacitor makes this solution have “memory”: consecutive events may not be delayed by the same amount depending on the initial charge on top of the capacitor (Figure 4.9).

Lastly, mixed-signal delay cells can be implemented based on the operation of CMOS thyristors. A current mirror discharges (or charges) a capacitor until a certain threshold is reached, at that point, a positive feedback mechanism is triggered which pulls the capacitor voltage to V_{ss} (or V_{dd}) as seen Figure 4.10. The threshold is usually set as the trip point of an inverter while the feedback mechanism solves the problems related to the




FIGURE 4.9: RC based, mixed-signal delay cell.

previous implementation (based on the RC time constant). The reset of the capacitor voltage (not shown in Figure 4.10) ensures a memory-less operation, while the positive feedback mechanism minimizes the time the capacitor voltage spends around the trip point of the inverter. The duration of the delay can be programmed by tuning the reference current of the current mirror.




FIGURE 4.10: CMOS thyristor based, mixed signal delay cell.

4.2.2 Delay Cell Design

The design and layout of the proposed delay cell architecture has been done in collaboration with Columbia University. The chosen delay cell architecture is based on the implementation presented in [63], on top of which several simplifications have been made, with the goal of reducing the power consumption. The resulting transistor level view of the proposed elementary delay cell is presented in Figure 4.11.

The signal EN (active on low) is used to completely disable the delay cell in case a lower order filter is required. RST provides an external mechanism to reset all the internal nodes of the circuit; in normal operation, this action is accomplished by the trigger of the output signal V_{OUT} which ensures that the delay cell returns to its original state only once the input pulse has been completely and successfully delayed. In normal operation, an active on low pulse on V_{IN} triggers a transition to V_{dd} on V_{trip} which switches ON M_6 . V_C starts discharging through the current mirror $M_5 - M_{18}$, from the initial value of V_{dd}

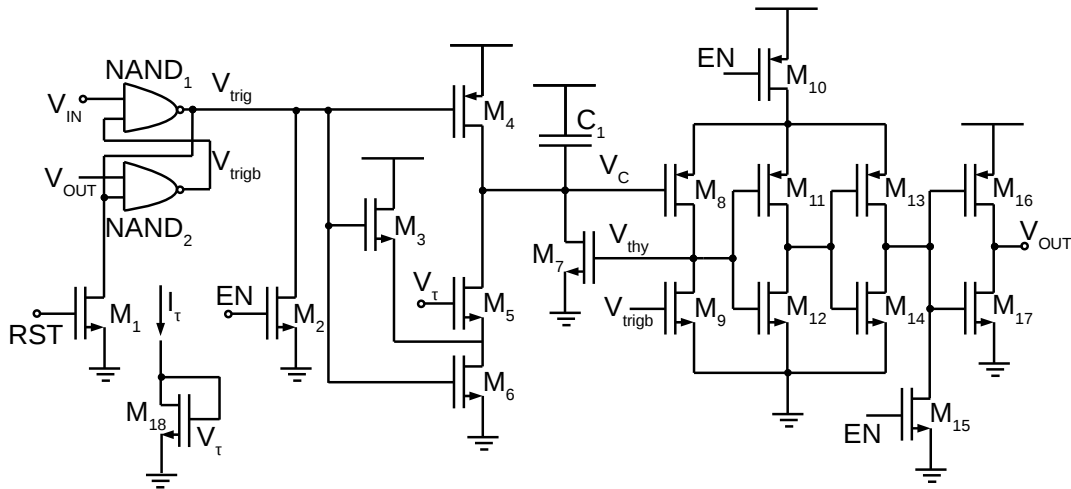


FIGURE 4.11: Transistor level implementation of the proposed delay cell.

to a trip point defined by the CMOS thyristor formed by M_7 and M_8 . The three output inverters are used to invert V_{thy} and also to make its transition sharper. Transistor M_9 pulls down V_{thy} immediately after the trip point has been reached, this recharges node V_C thereby reducing its swing from $V_{dd}-V_{ss}$ to $V_{dd}-V_{trig}$ and saving power.

The evolution of the input, output and some key internal nodes (V_{trig} , V_C and V_{thy}) is plotted in Figure 4.12.




FIGURE 4.12: Time domain evolution of different delay cell signals during a delay event.

The sizing of the transistors used in Figure 4.11 is detailed in Table 4.2. Transistor M_5 is used to mirror very small currents, hence a large value has been chosen for its length. Some transistors' widths have been optimized to minimize the crowbar currents observed

near switching events (M_{12} , M_{13} , M_{14} , M_{16} and M_{17}); otherwise, minimum size has been chosen to minimize dynamic power consumption of the delay cell.

TABLE 4.2: Sizing of the elementary delay cell components.

component	type	value	back-biasing
M_1	NMOS	80n/30n	0V
M_2	NMOS	80n/30n	0V
M_3	NMOS	80n/30n	0V
M_4	PMOS	80n/30n	0V
M_5	NMOS	200n/2u	0V
M_6	NMOS	80n/30n	0V
M_7	NMOS	80n/30n	0V
M_8	PMOS	320n/30n	0V
M_9	NMOS	80n/30n	0V
M_{10}	PMOS	80n/30n	0V
M_{11}	PMOS	80n/30n	0V
M_{12}	NMOS	480n/30n	0V
M_{13}	PMOS	164n/30n	0V
M_{14}	NMOS	108n/30n	0V
M_{15}	NMOS	80n/30n	0V
M_{16}	PMOS	164n/30n	0V
M_{17}	NMOS	108n/30n	0V
M_{18}	NMOS	200n/2u	0V
$NAND_1$	NAND	X2	0V
$NAND_2$	NAND	X2	0V
C_1	capacitance	2.11fF	n/a

The delay value versus the control current, I_τ (controlled through V_τ) is plotted in Figure 4.13. It is interesting to note that, regardless of the delay value τ programmed through V_τ , the elementary delay cell requires the same amount of energy to perform a single delay operation. This energy can be divided into a digital part, required to switch various digital circuits, and an analog part required to charge and discharge the capacitor, as shown in equation 4.2. Since the positive feedback mechanism is triggered at a voltage, V_{trip} , defined by the M_7 M_8 CMOS thyristor, it follows that for every delay event, the capacitor discharges from $0.65V$ to V_{trip} , making the analog energy constant regardless of I_τ .

$$E_{token}^{el} = E_{token}^{digital} + E_{token}^{analog} = E_{token}^{digital} + CV_{dd} \cdot (V_{dd} - V_{trip}) \quad (4.2)$$

In Table 4.3 we compare post-layout simulation results of our elementary delay cell with state of the art implementations from literature.

The previously presented delay cell allows us to implement an asynchronous digital delay with a programmable value, from 2ns to 15ns. In the next section we discuss how to




FIGURE 4.13: Value of the delay cell versus the tail current.

TABLE 4.3: Comparison of the proposed elementary delay cell with previous works.

metric	Schell [68]	Kurchuk [37]	Vezyrtzis [61]	this work
technology	90nm	65nm	130nm	28nm FDSOI
delay range	22ns–280ns	100ps–300ps	15ns–500ns	2ns–15ns
energy per token	35fJ	11fJ–13fJ	50fJ	3.8fJ–5fJ
power with inactive input	n/a	n/a	n/a	4.9nW
power when disabled	14.2nW	n/a	2.3nW	0.6nW
matching (st. dev.)	n/a	n/a	n/a	7.4%
V_{dd}	1V	1.2V	1V	0.65V

assemble instances of the proposed elementary delay cell in order to achieve the required 66ns to 100ns CT-FIR tap delays.

4.2.3 Delay Cell Architecture

Architecture level simulations of the DF-CT-ADC show that the minimum time between two consecutive events issued by the CT-ADC is 2ns (which we call T_{gran}) while the average time between such tokens is 5ns (T_{avg}). Consequently, a 100ns delay tap, required by the CT-FIR, must be constituted of at least 50 instances of the previously presented elementary delay cell programmed to a value of 2ns, as shown in Figure 4.14. Thus, the delay of every token generated by the DF-CT-ADC requires a total energy of 2.19pJ, as given by equation 4.3 with N_S – the number of elementary delay cells in a delay tap, N_{FIR} – the FIR order and E_{token}^{el} – the energy required by each elementary delay cell. We can now compute the average power required by the delay cells, given $T_{avg} = 5ns$, as being equal to $438\mu W$, more than 4 times above the power budget of the WU-RX.

$$E_{token}^{total} = N_S \cdot N_{FIR} \cdot E_{token}^{el} \quad (4.3)$$

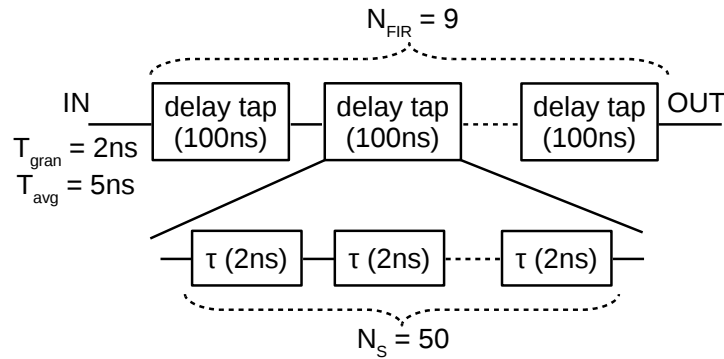


FIGURE 4.14: Classic delay cell architecture in the context of our FIR filter.

As it has been seen previously, the energy required by an elementary delay cell for a single operation does not depend on the actual delay value, given the capacitance is kept constant. This means that the energy requirements of a single tap can be reduced by lowering the number of series elementary delay cells it consists of. However, this requires using longer delay values for the elementary cells (to maintain a tap delay of 100ns) which in turn, requires an increased granularity at the ADC level. This can be achieved by adding parallel delay paths, as shown in Figure 4.15: the first event issued by the CT-ADC is sent to the first (upper-most) path, the second to the second path and so on, until the $N_P + 1$ -th event is reached, which is again sent to the first path.

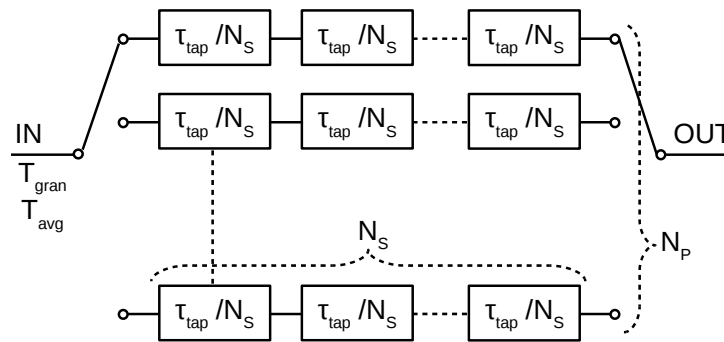


FIGURE 4.15: Delay tap architecture based on a parallel elementary delay arrangement.

Only 1 out of every N_P events is sent to any given delay path, effectively multiplying the CT-ADC granularity by a number equal to the number of parallel paths. This allows us to increase the delay value of each elementary delay, while maintaining a constant tap delay (τ), thus reducing the number of elementary delays in a delay tap and the power consumption by a factor of N_P . However, this reduction also has an impact on other important performance parameters of the delay tap such as jitter, size and matching.

In the next part of this chapter we study these trade-offs in the general case of a delay tap of value τ_{tap} , which consists of N_P parallel paths of N_S series elementary delay cells. The

input signal is an asynchronous train of events described by a minimum time between two consecutive events of T_{gran} and an average time between events of T_{avg} .

Energy

As seen earlier, the energy required by each token going into the delay tap is equal to the sum of energies of elementary delay cells triggered, as given by equation 4.4. Choosing a parallel delay architecture also imposes the use of a switch to dispatch each input token to the different delay paths; however, this energy is expected to be small compared to the total energy required to delay the token; we therefore ignore it.

$$E_{token}^{tap} = N_S \cdot E_{token}^{el} \quad (4.4)$$

The number of series delay cells is defined by the input granularity and the total tap delay value. Each parallel path sees the input granularity multiplied by N_P therefore each elementary delay can be programmed to a value of $N_P \cdot T_{gran}$. Finally, since the total tap delay value is τ_{tap} we conclude that the number of series delay cells is given by equation 4.5. Note that capacitance is kept constant regardless of the values chosen for N_S and N_P .

$$N_S = \frac{\tau_{tap}}{T_{gran} N_P} \quad (4.5)$$

This gives as a total of energy required per token given by equation 4.6.

$$E_{token}^{tap} = \frac{\tau_{tap}}{T_{gran} N_P} \cdot E_{token}^{el} \quad (4.6)$$

Jitter

Results from [76] show that the jitter of a CMOS thyristor based delay cell depends on its value (τ_{el}), on the steering current (I_τ), on a constant determined by different technology parameters (k_1), and on the value of the capacitance which is being charged (C_1), as shown in equation 4.7. Since the capacitance is kept constant throughout the design, it can be integrated in the technology constant k_1 , yielding k_2 .

$$Var[\tau_{el}] = \frac{C_1}{I_\tau^2} \cdot k_1 = \tau_{el} \cdot \frac{k_2}{I_\tau} \quad (4.7)$$

The delay value, τ_{el} , is inversely proportional to the steering current, I_τ , thus we can express the variance of an elementary delay cell as given in equation 4.8.

$$Var[\tau_{el}] = \tau_{el}^2 \cdot k_3 \quad (4.8)$$

The proposed delay tap is composed of N_S series elementary delay cells; its total delay can thus be expressed as a sum of Gaussian random variables with a variance given by equation 4.8. The variance of the total tap delay becomes equation 4.9.

$$Var[\tau_{tap}] = N_S \cdot Var[\tau_{el}] = N_S \cdot \tau_{el}^2 \cdot k_3 \quad (4.9)$$

Replacing the elementary delay value by the total tap delay divided by the number of series elements, we conclude that the jitter of the tap delay is described by equation 4.10 as a function of N_S or by equation 4.11 as a function of N_P .

$$Var[\tau_{tap}] = \frac{\tau_{tap}^2}{N_S} \cdot k \quad (4.10)$$

$$Var[\tau_{tap}] = N_P T_{gran} \tau_{tap} \cdot k \quad (4.11)$$

It is interesting to note that for a design where we choose to tune the current while keeping the capacitance value constant, the tap jitter increases proportionally to the number of parallel paths. On the other hand, it can be proven that if we choose to tune the capacitance while conserving the value of the current, the tap jitter will remain constant at the expense of increased energy requirements for the parallel path solution.

Size

Since we have opted to change the delay value of each elementary cell by tuning its steering current rather than changing its design, it follows that the total area of a tap is proportional to the number of delay cells it consists of: equation 4.12 with S^{el} – the area of an elementary delay cell.

$$S^{tap} = N_S \cdot N_P \cdot S^{el} \quad (4.12)$$

Replacing N_S with its value derived in equation 4.5 gives us a total area which is constant, given by equation 4.13.

$$S^{tap} = \frac{\tau_{tap}}{T_{gran}} S^{el} \quad (4.13)$$

Matching

Despite being an important parameter in the design of CT-DSPs, the matching achieved by individual delay taps can be compensated and minimized using supplementary circuits; it does not play a direct role in the true figure of merit of the delay tap. Mismatch compensation for the proposed delay cell architecture is discussed at a further point in this chapter: Section 4.2.4; the circuits used for this purpose only need to be periodically activated, thereby having a minimal impact on the power consumption.

Overall Figure of Merit

A figure of merit for the delay tap can be defined by combining the previously derived performance parameters as given in equation 4.14. It can be seen that this FoM is constant and does not depend neither on the number of series elements nor on the number of parallel paths. Essentially, by adding parallel paths we trade jitter for energy.

$$FoM_{tap} = E_{token}^{tap} \cdot Var[\tau_{tap}] \cdot S^{tap} \neq f(N_S, N_P) \quad (4.14)$$

If we admit that our system has a finite jitter specification for its delay taps, in our case 325ps as shown in Section 4.1.4 on page 96, it follows that this limitation defines the maximum number of parallel delay paths we can use. A set of post-layout simulations is done for different configurations of the delay tap (N_S and N_P), results are summarized in Table 4.4. The number of parallel paths can be increased up to 5, while maintaining an RMS jitter which does not affect the CT-ADC performance. The energy required by the delay taps is thus also divided by a factor of 5.

TABLE 4.4: RMS jitter of a delay tap with different configurations (N_S and N_P).

N_P	N_S	τ_{el}	$std(\tau_{el})$	$std(\tau_{tap})$
1	50	2ns	10ps	70ps
2	25	4ns	29ps	145ps
3	17	6ns	43ps	176ps
4	13	8ns	73ps	263ps
5	10	10ns	100ps	316ps
6	9	12ns	138ps	415ps

4.2.4 Calibration and Matching

The matching achieved naturally by the delay taps is not sufficient to preserve the performance of the 9th order transfer function of the CT-FIR. Consequently, a calibration

mechanism must be used to periodically fine tune the values of the delay taps such that the standard deviation of the total mismatch remains under 700ps, which represents the matching specification derived in Section 4.1.4 on page 96.

The elementary delays can be tuned by changing the tail current which controls the current mirror. Given the total number of elementary delay cells, it is impractical to control all of them individually through an analog control knob. Consequently, we propose a tap-level modification of the previous architecture which consists of introducing extra delay elements (with a nominal delay value of 10ns or 1ns) in each delay path which can be bypassed using a multiplexer, as shown in Figure 4.16. Tuning of the delay tap is achieved by adding / removing supplementary delay cells from the main delay path.




FIGURE 4.16: Programmable parallel delay tap architecture.

At this point, we want to minimize the number of extra elementary delays which need to be added and active in order to minimize the supplementary power consumption incurred by the proposed tuning mechanism. To compute N_{10ns} and N_{1ns} we develop a statistic model for the elementary delay which enables us to predict the best case matching performance obtained using different configurations of the delay tap.

The delay model is based on a set of MC simulations in which we study the evolution of an elementary delay versus the injected, reference current. A representation of the results over 10 runs is given in Figure 4.17. The value of each elementary delay can be written as a second order equation with respect to the reference current, as seen in equation 4.15.

$$\tau = aI_{\tau}^2 + bI_{\tau} + c \quad (4.15)$$

Observing the variations of the three parameters of equation 4.15 (a , b and c) by plotting their probability density functions (Figure 4.18), we can conclude that they can be




FIGURE 4.17: Value of a delay tap versus the control current for different MC runs.

described by three normal Gaussian random variables. Furthermore, simulations show a strong correlation between these three parameters, as it can be seen in Figure 4.19 where we plot b vs a and c vs a .




FIGURE 4.18: PDFs of the three parameters describing the statistical link between the value of each elementary delay and the control current.

We can now randomly draw a value for a , based on its PDF, compute b and c according to the previously extracted correlation functions and thus model the value of an elementary delay based on its steering current. This process is repeated for each delay cell in the architecture. We suppose that a single I_τ is used to control all 10ns delay elements of the 9th order FIR and another current, I'_τ , is used to control each 1ns elementary delay. An iterative search algorithm is then used to compute I_τ and I'_τ which yield the best-case matching performance for different values of N_{10ns} and N_{1ns} . Results are summarized in Table 4.5. The *average* number of active extra calibration delay cells per delay tap is also given; this allows us to estimate the energy cost of the proposed calibration method.


FIGURE 4.19: Correlation between b and a (left); c and a (right).

TABLE 4.5: Performance of the proposed delay tap calibration method.

	$N_{10ns} = 2$ $N_{1ns} = 3$	$N_{10ns} = 2$ $N_{1ns} = 5$	$N_{10ns} = 2$ $N_{1ns} = 7$	$N_{10ns} = 3$ $N_{1ns} = 3$	$N_{10ns} = 3$ $N_{1ns} = 5$	$N_{10ns} = 3$ $N_{1ns} = 7$
average τ_{tap}	99.95ns	99.95ns	99.98ns	99.93ns	99.96ns	99.98ns
matching	1.1ns	640ps	310ps	1.01ns	520ps	260ps
I_τ	90.75nA	91.25nA	90.25nA	99.75nA	98.25nA	94.75nA
I'_τ	40nA	40nA	40nA	40nA	40nA	45.25nA
# 10ns delays / tap	0.63	0.6	0.53	0.73	0.67	0.51
# 1ns delays / tap	0.46	0.52	0.5	0.52	0.47	0.49
power increase	+9.7%	+9.9%	+9%	+11.2%	+10.2%	+8.7%

We can see that using 3 delay cells of 10ns offers us minimal improvement with respect to a solution consisting of $N_{10ns} = 2$. Furthermore, the 700ps matching limit is reached when at least five 1ns delay elements are used. In average, each tap delay is expected to consume 9.9% more energy per CT-ADC token due to the added, active calibration delay cells. Note that, even though it might be tempting to increase the number of calibration cells, especially N_{1ns} , the area of the filter is also increased proportionally.

Ideally, the calibration is implemented on chip by injecting pulses in all parallel delay paths and changing their configuration based on a reference signal which defines the tap delay values. This calibration phase should be repeated periodically as a temperature drift may introduce a systematic shift in the delay values. Since the chip described in this chapter is intended to be an early stage prototype, the calibration is done off-chip and the effects of temperature drifts are not taken into account.

4.2.5 Delay Architecture Summary

A summary of the parameters and performance of the proposed delay tap architecture is given in Table 4.6.

TABLE 4.6: Summary of the delay cell architecture and its performance.

critereon	value
elementary delay	6.6ns–10ns
delay tap	66ns–100ns
series elements (N_S)	10
parallel paths (N_P)	5
calibration 10ns elements (N_{10ns})	2
calibration 1ns elements (N_{1ns})	5
energy per tap per token (E_{tap}^{token})	42fJ
mismatch ($std(\tau_{tap})$)	640ps
jitter (RMS)	316ps

4.3 CT Adder

The role of the CT adder is to sum the asynchronous pulses issued by the CT-ADC and to output them in the form of a digital bus or an analog signal.

4.3.1 Previous Work

Depending on the length of the FIR filter and on the average frequency of incoming pulses, previous works have either opted for a digital or for an analog CT adder implementation.

Digital CT-Adder

For digital implementations of CT adders, an important requirement is that the propagation delay through the adder is kept constant regardless which of its inputs is triggered. If this condition is not maintained, the asynchronous output signal risks undergoing distortion which would negatively affect its linearity. One digital CT architecture which guarantees constant propagation times is the Carry Save Adder (CSaA), which consists of several full-adders, each computing a single sum and carry bit. The schematic of an implementation which adds three 4-bit numbers is presented in Figure 4.20.

The speed requirement of the adder is defined by the average token frequency issued by the CT-ADC multiplied by the number of filter taps; in our case, this adds up to an *average* frequency of 1.8GHz. On the other hand, the *maximum* addition frequency is undefined as events can be issued by the CT-ADC at arbitrary intervals: consider a second order FIR with a 100ns delay cell along with a CT-ADC generating an event at $t = 0s$ and another one at $t = 100ns + \delta t$. Due to the asynchronous nature of the CT-ADC, δt can be arbitrarily small, making the maximum frequency of the adder arbitrarily high ($1/\delta t$). Consequently, to avoid metastability and assure that the CT




FIGURE 4.20: Architecture of a 3 input (4 bits) Carry Save Adder.

adder input does not change *while* an addition is being resolved, a pseudo-synchronization block is required at the input of the adder. This block has the role of latching all inputs which arrive while an addition is being resolved and passing them to the adder once the addition operation has finished.

The previously described metastability concerns are dependent on the average CT-ADC token frequency. Existing works achieved successful implementations of digital adders for CT-ADCs with output token frequencies up to 5MHz, [68]. However, our application described in the previous chapters employs an ADC operating at much higher frequencies, making the implementation of digital CT adders less attractive.

Analog CT-Adder

There exist several analog CT adder designs which do not suffer from the metastability concerns of digital adders. One example, presented in [60], which achieves operation frequencies in the GHz range, uses charge pumps to charge and discharge a capacitor holding the result of the addition. The CT adder architecture is presented in Figure 4.21: a series of charge pumps are connected in parallel to a summing capacitor; the charge pumps are triggered by a digital control signal in the following manner: rising edges trigger a slight charge of the capacitance while falling edges trigger a symmetric discharge.

Since each charge pump operates independently, there are no minimum timing constraints between tokens arriving to the CT adder. Moreover, the FIR coefficients can be directly integrated with the CT adder by changing the weights affected to the charging (discharging) currents in the NMOS and PMOS current mirrors. The control blocks can also be reconfigured to react on either rising or falling edges thereby enabling the use of signed




FIGURE 4.21: Charge pump based analog adder.

coefficients. In the example given in Figure 4.21 all coefficients have the same absolute value, however the $i + 1$ coefficient is negative.

Analog CT adder implementations are low power and can operate at much higher speeds compared to their digital counterparts. The only drawback is related to the output signal which is in the analog domain rather than digital. However, this is not a concern for our application, since the CT-DSP output is supposed to be injected into the envelope detector which does not specifically require a digital signal. Consequently, we opt for an analog implementation of the CT adder; this is detailed in the next part of this chapter.

4.3.2 Proposed Weighted-CT-Adder

The proposed CT adder has been designed to function correctly with the encoding used for the data coming from the CT-ADC which presents itself as a train of narrow, asynchronous pulses. Furthermore, the multiplication with the tunable FIR coefficients is also embedded in the adder operation.

Elementary Cell

The schematic of a single elementary adder cell is given in Figure 4.22. Asynchronous pulses, coming either directly from the CT-ADC or from any delay tap, are injected from nodes V_+ and V_- to nodes a_{0+} , a_{1+} , a_{0-} and a_{1-} through a set of transmission gates and multiplexers which are controlled by the on-chip memory used to store the coefficients of the CT-DSP. Please note that the adder cell input (V_+ and V_-) does not represent a *differential* signal, it represents the two bit encoded data issued by the CT-ADC as shown in the previous chapter. We use 2 bits to program the value of the respective coefficients

(b_0 and b_1) and another bit (b_s) to set the sign of the coefficient. The ensuing pulses are then used to trigger a set of switches which control the total capacitance connected above the node V_{out} and directly influence its voltage. A capacitive ladder which controls V_{out} is thus created. A more detailed explanation of the adder operation and of the state of node V_{out} is given in the next section, where we study how the elementary adder cells are connected together to construct adders of arbitrary lengths.

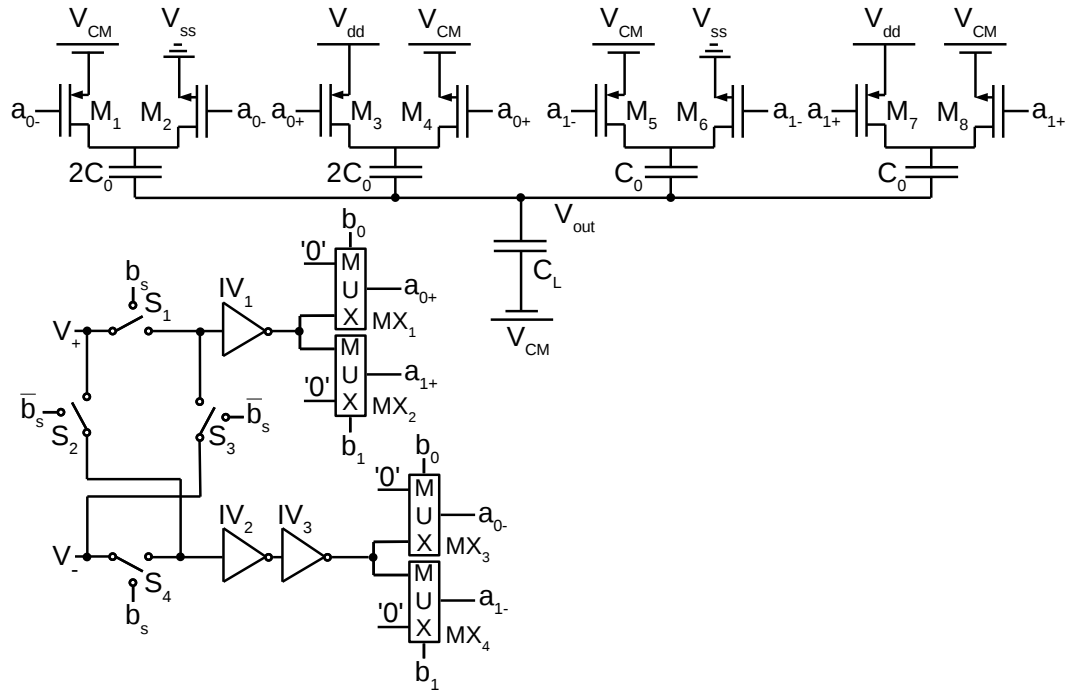


FIGURE 4.22: Proposed elementary adder cell.

Note that, for a single elementary adder cell, the output voltage V_{out} can be expressed shown in equation 4.16 with b_s – the sign bit taking values of 1 or -1 ; $\overline{b_0 b_1}$ – the binary value of the adder cell coefficient; $V_+ - V_-$ – the differential input voltage; k – a constant related to the implementation.

$$V_{out} = k \cdot b_s \cdot \overline{b_0 b_1} \cdot (V_+ - V_-) \quad (4.16)$$

The transistor sizes, their back-biasing, as well as the values of the capacitors are detailed in Table 4.7. Switches M_1 to M_8 have strong back-biases to minimize their resistive losses when conducting. Despite C_L having a value of 0.53fF no actual capacitance of that size has been used; when constructing a full adder composed of several elementary cells, several instances of C_L are placed in parallel and their values are lumped into a single component. We will also show that there is no matching requirement between C_L (or its lumped equivalent) and C_0 , thus the value of C_L can be chosen independently of C_0 .

TABLE 4.7: Component sizes of an elementary adder cell.

component	type	value	back-biasing
M_1	PMOS	$1\mu\text{m}/30\text{nm}$	-2V
M_2	NMOS	$1\mu\text{m}/30\text{nm}$	2V
M_3	PMOS	$1\mu\text{m}/30\text{nm}$	-2V
M_4	NMOS	$1\mu\text{m}/30\text{nm}$	2V
M_5	PMOS	$1\mu\text{m}/30\text{nm}$	-2V
M_6	NMOS	$1\mu\text{m}/30\text{nm}$	2V
M_7	PMOS	$1\mu\text{m}/30\text{nm}$	-2V
M_8	NMOS	$1\mu\text{m}/30\text{nm}$	2V
S_1-S_4	CMOS switch	$400\text{nm}/30\text{nm}$	0V
MX_1-MX_4	multiplexer	X5	0V
IV_1-IV_4	inverter	X10	0V
C_0	capacitance	5.1fF	n/a
C_L	capacitance	0.53fF	n/a

Full Adder

The previously presented elementary adder has a single differential input (consisting of V_+ and V_-) and therefore cannot perform any addition operation. An N_{tap} input full adder is obtained by connecting N_{tap} elementary adder cells in parallel, having a single net in common, V_{out} , as shown in Figure 4.23. A DC block composed of two $500\text{k}\Omega$ resistors is used to bias the DC point of the resulting structure and to avoid leaving net V_{out} floating.

A set of transmission gates controlled by local memory is used for each instance of the elementary delay cell. The full adder thus becomes a capacitive network in which asynchronous pulses control the amount of capacitance located between the output node V_{out} and the three reference nodes: V_{dd} , V_{CM} ($= V_{dd}/2$) and V_{ss} . To get a better understanding of the operation mode of the N_{tap} full adder we are going to suppose it is configured in the following manner: the coefficient of n branches is set to the value 10, the coefficient of p branches is set to 01 while the rest $N_{tap} - n - p$ branches are set to 00. To simplify the analysis, we also suppose that we compute V_{out} at a time instant when *all* elementary delay cell inputs V_+ are at state 1 (V_{dd}) while all V_- inputs are at state 0 (V_{ss}). A schematic view of the system in this configuration is given in Figure 4.24. The addition is successfully accomplished if V_{out} can be written as given in equation 4.17.

$$V_{out} = k \cdot (2n + p) \quad (4.17)$$

To compute V_{out} , we can simplify the full adder circuit in the following way:

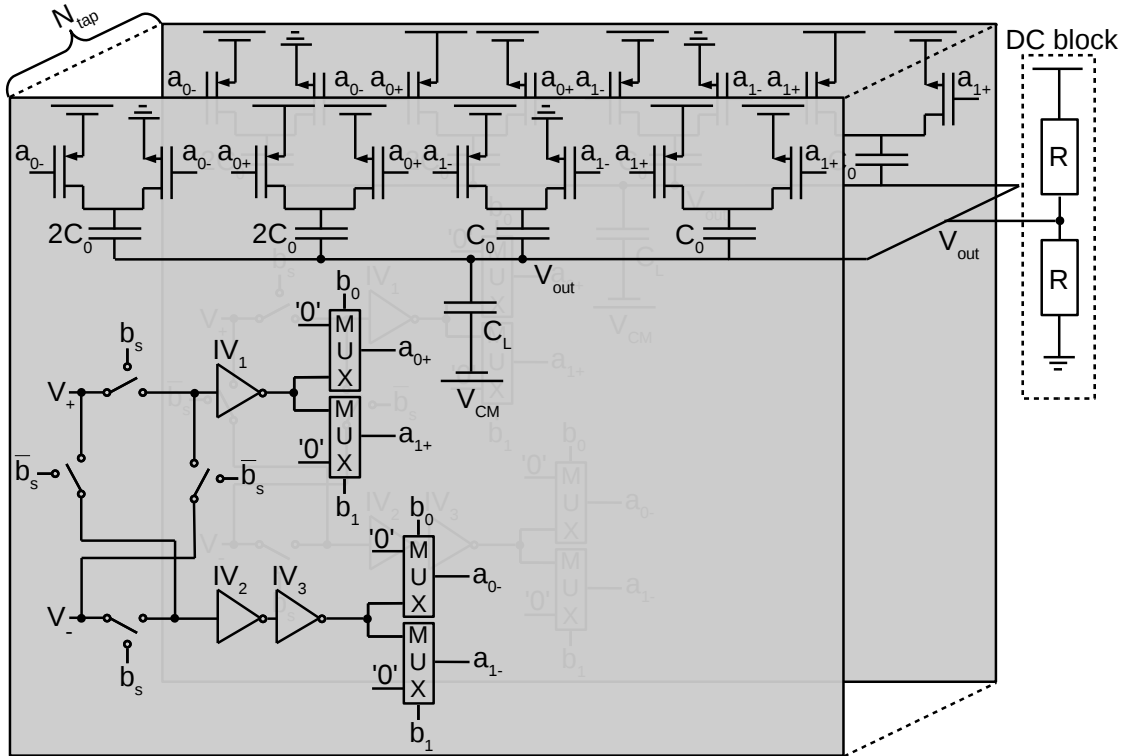


FIGURE 4.23: Proposed full adder cell.




FIGURE 4.24: Scenario used to illustrate the operation principle of the proposed weighted adder.

- branches with coefficient 10: n branches consist of a $2C_0$ capacitance connected between V_{out} and V_{dd} in parallel with a $4C_0$ capacitance connected between V_{out} and V_{CM} as well as C_L connected between V_{out} and V_{CM} , as shown in Figure 4.25.
- branches with coefficient 01: p branches consist of a C_0 capacitance connected between V_{out} and V_{dd} in parallel with a $5C_0$ capacitance connected between V_{out} and V_{CM} as well as C_L connected between V_{out} and V_{CM} , as shown in Figure 4.25.




FIGURE 4.25: Elementary adder with its coefficient equal to 10 and with input V_+ active (at V_{dd}).




FIGURE 4.26: Elementary adder with its coefficient equal to 01 and with input V_+ active (at V_{dd}).

- branches with coefficient 00: $N_{tap} - n - p$ branches consist of a $6C_0$ capacitance connected between V_{out} and V_{CM} in parallel with a C_L connected between V_{out} and V_{CM} , as shown in Figure 4.27.




FIGURE 4.27: Elementary adder with its coefficient equal to 00 and with input V_+ active (at V_{dd}).

If we now lump together the previously described circuit, we end up with the circuit presented in Figure 4.28, with C_{dd} , C_{ss} and C_{CM} given in equation 4.18–equation 4.20.

$$C_{dd} = n \cdot 2C_0 + p \cdot C_0 \quad (4.18)$$

$$C_{ss} = 0 \quad (4.19)$$

$$C_{CM} = N_{tap}C_L + n \cdot 4C_0 + p \cdot 5C_0 + (N_{tap} - n - p) \cdot 6C_0 \quad (4.20)$$




FIGURE 4.28: Lumped representation of the proposed CT adder

The initial condition (or the DC value) of the node V_{out} is set by the DC block to $V_{dd}/2$; moreover, we choose $V_{CM} = V_{dd}/2$. We can thus conclude that in the previously described configuration, V_{out} can be written as given in equation 4.21. Identifying the terms in the equation with those from equation 4.17 on page 114, it follows that the weighted addition has been successfully accomplished.

$$V_{out} = \frac{V_{dd}}{2} \frac{2nC_0 + pC_0}{N_{tap}(6C_0 + C_L)} + \frac{V_{dd}}{2} \quad (4.21)$$

4.3.3 Adder Performance

Among the critical performance parameters for the proposed CT adder we distinguish its power consumption as well as its conversion gain.

Power

The energy requirements of the proposed adder can be decomposed into two parts: a “digital” part linked to the energy required to trigger the inverters in the direct signal path (IV_1-IV_3) and an “analog” part required to charge the capacitive network. We can thus compute the energy per elementary adder cell (per FIR tap) depending on the value of the coefficients; results are summarized in Table 4.8.

TABLE 4.8: Energy required by an elementary adder cell to process a single CT-ADC token.

coefficient value	E_{ADDdig}^{tap}	E_{ADDan}^{tap}	$E_{ADDtotal}^{tap}$
00	3fJ	0fJ	3fJ
01	3fJ	3.9fJ	6.9fJ
10	3fJ	7.8fJ	10.8fJ
11	3fJ	11.7fJ	14.7fJ

We can see that, compared to a delay tap which requires 42fJ per CT-ADC token, the power requirements of the adder are significantly smaller.

Conversion Gain

It is difficult to define the conversion gain for a component which has a digital-to-analog conversion embedded in its operation mode. However, it will be seen that the “voltage gain” of the adder plays an important role in the design of the DF-CT-ADC feedback loop.

To compute the conversion gain of an N_{tap} input adder we suppose the following scenario: all coefficients are programmed to 11 (maximum value) and all N_{tap} positive inputs (V_+) are toggled. Ideally, the output voltage should be N_{tap} times V_{dd} , as shown by equation 4.22. This is clearly impossible since the passive implementation of our adder does not allow us to achieve output voltages above V_{dd} . The real adder output voltage, computed following the same reasoning as described in Section 4.3.2, is given in equation 4.23. The second term in the previous equation represents the common mode of the adder output and does not have any contribution to an AC analysis.

$$V_{ideal}^{add} = N_{tap}V_{dd} \quad (4.22)$$

$$V_{real}^{add} = \frac{V_{dd}}{2} \frac{3N_{tap}C_0}{N_{tap}(6C_0 + C_L)} + \frac{V_{dd}}{2} = \frac{V_{dd}}{2} \frac{3C_0}{6C_0 + C_L} + \frac{V_{dd}}{2} \quad (4.23)$$

We now express the adder conversion gain as the *real* output swing divided by the *ideal* output swing, as in given in equation 4.24.

$$G_{conv} = 20\log_{10} \frac{V_{real}^{add}}{V_{ideal}^{add}} = 20\log_{10} \left(\frac{1}{2N_{tap}} \frac{3C_0}{6C_0 + C_L} \right) \quad (4.24)$$

Replacing C_0 and C_L with values given previously results in a conversion gain for the 9th order output FIR of -32dB . This is an attenuation applied to all input frequencies and therefore does not affect the filter transfer function. On the other hand, the 3rd order FIR embedded in the DF-CT-ADC operates at a conversion gain of -24dB , which has a direct impact on the DF-CT-ADC transfer function, as shown in Figure 4.2 on page 89. Consequently, an active voltage gain stage is required to compensate for the attenuation of the CT adder.

4.4 DF-CT-ADC

The previously presented delay taps and adder enable us to construct a CT-FIR of an arbitrary order. However, for the implementation of the DF-CT-ADC, several elements

are missing, as the CT-FIR analog output needs to be interfaced with the CT-ADC input.

4.4.1 CT-ADC – CT-FIR Integration

The classic solution for implementing the addition at the input of the CT-ADC is through the use of an analog adder, constructed around an OTA. However, this implies the use of a DC power even when the DF-CT-ADC input and the CT-FIR output are “quiet“. Alternatively, the CT-ADC ”front-end“ also features a $G_m - C$ cell, which enables us to implement an addition in the current domain, by simply connecting a separate transconductance to the same capacitance. The second transconductance is driven by the CT-FIR output and, in terms of specifications, is similar to the one used for the CT-ADC, thus it doesn't need to be too power consuming. A detailed view of the resulting DF-CT-ADC architecture is presented in Figure 4.29. Note that the circuit which splits the CT-ADC output into the 5 parallel data streams of the CT delay cells is described at a later point in this chapter; the pulses are recombined at the output of each delay tap using a 5-input OR gate.

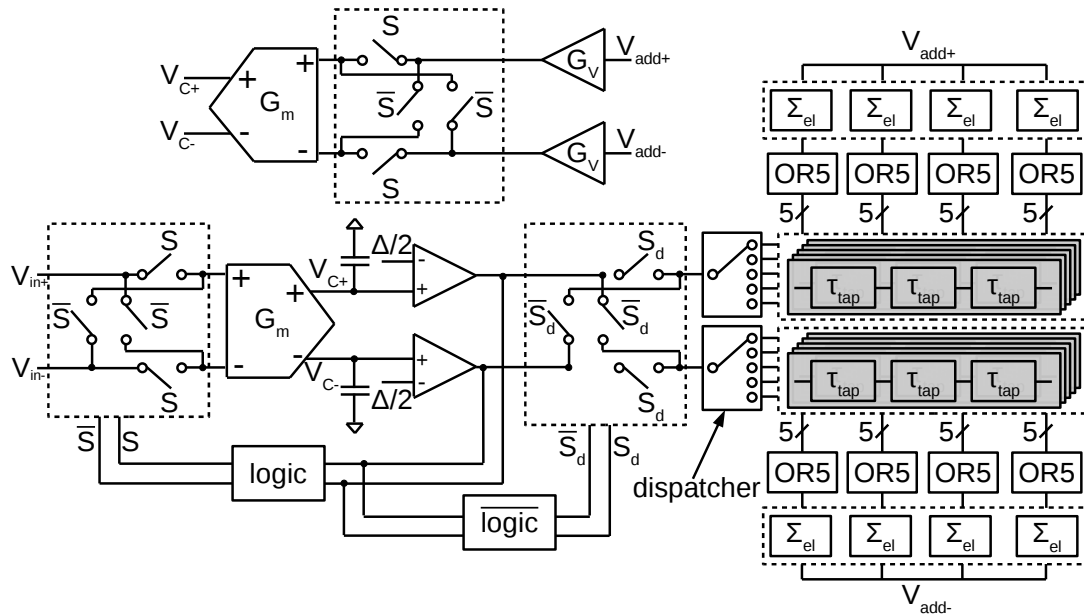


FIGURE 4.29: View of the complete DF-CT-ADC implementation.

We have opted for a fully differential implementation of the feedback path; due to the low CT adder conversion gain, a voltage amplifier has been used at its output. Moreover, since the current injected on the CT-ADC capacitance is scrambled by the flipping switches present at its input, it is clear that, if we want to apply the feedback at this point, a similar operation is required for the feedback signal. Consequently, a set of

flipping switches has been introduced in the feedback path, after the amplification of the adder signals, which are controlled by the *same* signals as the flipping switches located at the CT-ADC input. We now discuss the implementation of the extra circuitry required by the DF-CT-ADC.

4.4.2 Dispatcher

The dispatcher is used to split the CT-ADC output stream in five parallel streams thereby increasing each stream's granularity (T_{gran}) by a proportional factor and enabling the use of the chosen delay tap architecture. A schematic view of this circuit is given in Figure 4.30; a summary of the sizes of the components used is also given in Table 4.9.

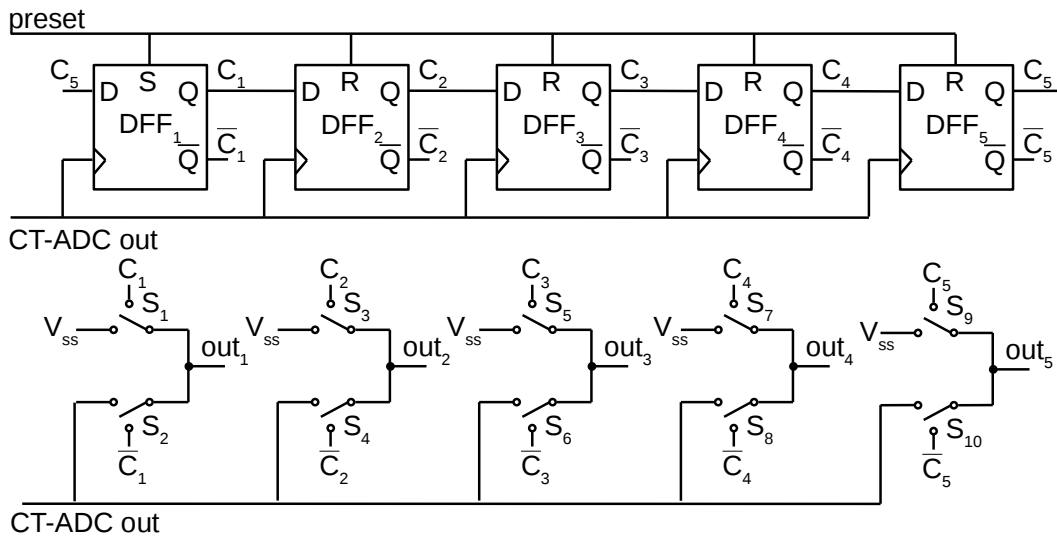


FIGURE 4.30: Schematic of the dispatching circuit used to split the CT-ADC output in 5 parallel streams.

TABLE 4.9: Sizing of the components used for the dispatcher.

component	type	value	back-biasing
DFF_1	DFF with set	X10	0V
DFF_2 - DFF_5	DFF with reset	X10	0V
S_1 - S_{10}	CMOS switch	180nm/30nm	0V

The 5 D-flip-flops are set to an initial state in which only one of their outputs is equal to V_{dd} while the rest are pulled down to V_{ss} . Each CT-ADC token toggles the clock inputs of the D-flip-flops thereby shifting the position of the logic 1 in a cyclic manner. Finally, each DFF controls a transmission gate which is used to select the input to be connected to each stream: choosing between either the CT-ADC output or ground (V_{ss}). Compared to the adder or tap delay, the energy required by the dispatcher does not scale linearly

with the number of taps, making its power consumption relatively small. Post-layout simulations have shown an energy requirement of 17fJ per CT-ADC token.

4.4.3 Voltage Gain and Filtering

As seen in Section 4.3.3 on page 118, the -24dB of CT adder conversion gain must be compensated by the use of an active gain stage. Furthermore, the CT-ADC generates strong high power components which must be attenuated in the feedback path before being injected into the input of the CT-ADC. These two functions are incorporated in an active, filtering voltage gain stage G_V (Figure 4.29 on page 119); its transistor-level implementation is given in Figure 4.31. The sizing of the components used is detailed in Table 4.10. Note that a current mode implementation of this stage is rendered difficult by its filtering requirement.


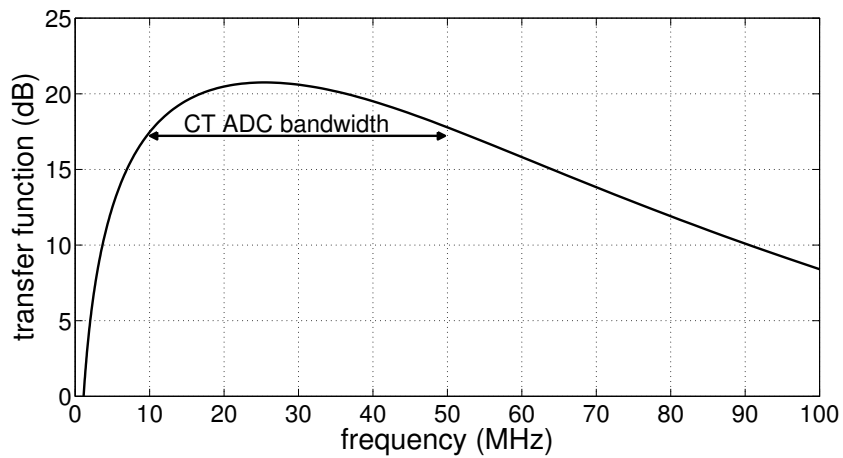


FIGURE 4.31: Schematics of the active voltage gain stage used in the feedback loop of the DF-CT-ADC.

TABLE 4.10: Sizing of the components used in the active voltage gain stage.


component	type	value	back-biasing
M_1	PMOS	$4.5\mu\text{m}/40\text{nm}$	-0.75V
M_2	NMOS	$1.5\mu\text{m}/150\text{nm}$	0V
R_1	resistance	$400\text{k}\Omega$	n/a
R_2	resistance	$400\text{k}\Omega$	n/a
R_3	resistance	$200\text{k}\Omega$	n/a
C_1	capacitance	118fF	n/a
C_2	capacitance	118fF	n/a
C_3	capacitance	250fF	n/a
C_4	capacitance	20fF	n/a

A second order low-pass filter is used to remove high power components from the CT-ADC output signal. The series capacitance C_3 is used to decouple the CT adder output from the second G_m cell input, by adding a zero in its transfer function at 0Hz . The push-pull stage is self biased using the R_3 resistance and its output current is injected into a capacitance to obtain the required voltage gain. The simulated transfer function of the voltage gain stage is given in Figure 4.32.

FIGURE 4.32: Voltage transfer function of the G_V block.

4.4.4 Feedback G_m

For the feedback G_m we have used a design very similar to the one utilized in the CT-ADC; its detailed schematic is given in Figure 4.33. Since both G_m cells have their differential outputs connected together, it follows that only one common mode feedback block is required. We thus use the common mode feedback control generated in the CT-ADC G_m cell to control the gates of the active load transistors in the feedback G_m cell. The sizing of the various components used in this design is given in Table 4.11.

FIGURE 4.33: Schematics of the feedback G_m cell.

Depending on the configurations of the switches controlling the degeneration resistance, R_{degFB} , the resistance can be varied from 0Ω to $50k\Omega$. The transconductance is thus decreased from $43\mu S$ to $21\mu S$, for a standard bias current of $2\mu A$.

TABLE 4.11: Sizing of various components used in the design of the feedback G_m cell.

component	type	value	back-biasing
M_1	PMOS	$1.28\mu\text{m}/400\text{nm}$	-0.75V
M_2	PMOS	$1.28\mu\text{m}/400\text{nm}$	-0.75V
M_3	NMOS	$2\mu\text{m}/60\text{nm}$	0.75V
M_4	NMOS	$2\mu\text{m}/60\text{nm}$	0.75V
M_5	NMOS	$2\mu\text{m}/200\text{nm}$	0.75V
M_6	NMOS	$2\mu\text{m}/200\text{nm}$	0.75V
M_7	NMOS	$2\mu\text{m}/200\text{nm}$	0.75V
$R_{deg1}-R_{deg5}$	resistance	$10\text{k}\Omega$	n/a
S_1-S_5	CMOS switch	$1\mu\text{m}/30\text{nm}$	0V

4.5 Simulation/Measurements Results

The goal of this section is to highlight the performance of the proposed CT-DSP system. Unfortunately, due to unexpected delays in the fabrication (the delivery date was postponed by more than 4 months), the chip is not back from the foundry; results presented in this section are based on transistor level simulations only. The rest of this section is organized as follows: first, the achievable DF-CT-ADC transfer functions are presented. The same analysis is then conducted for the output FIR filter. A CT-DSP system configuration is chosen and tested with a two interferer input scenario; a breakdown of the system power consumption is thus given. Finally, we show the scalability of the proposed system's power consumption versus the input signal characteristics as well as its single tone noise performance. The proposed CT-DSP is also compared to other existing implementations from literature.

4.5.1 DF-CT-ADC Performance

The DF-CT-ADC has four tunable parameters which influence its transfer function:

- CT adder coefficients: a_0 , a_1 , a_2 and a_3 – determine the type of transfer function;
- delay reference current I_τ – sets the transfer function repetition frequency;
- feedback loop degeneration resistance R_{degGM} – contributes to the feedback loop direct gain, k_{ad} , thereby setting the quality factor of the DF-CT-ADC IIR transfer function;
- feedback loop reference current I_{refGM} – contributes to the feedback loop direct gain, k_{ad} , thereby setting the quality factor of the DF-CT-ADC IIR transfer function.

- direct path CT-ADC G_m and comparator threshold Δ – influence the overall conversion gain of the system, but do not contribute to the filter transfer function.

These tuning options are demonstrated in the following paragraphs.

Coefficients Tuning

Here, we set I_τ such that the tap delays are configured to 100ns and we use an R_{degGM} of 10k Ω along with an I_{refGM} of 2 μ A. Three configurations for the CT adder coefficients are chosen as showed in Table 4.12, corresponding to a highpass, lowpass and bandpass implementation.

TABLE 4.12: Coefficients which demonstrate the transfer function reconfigurability of the DF-CT-ADC.

coefficient	highpass	lowpass	bandpass
a_0	0.75	0.75	0.75
a_1	0.5	-0.5	-0.5
a_2	-0.5	-0.5	0.75
a_3	0.25	-0.25	0.25

The resulting, simulated transfer functions are plotted in Figure 4.34. Due to the 100ns configuration of the delay cells, only 10MHz of spectrum is required to define the DF-CT-ADC transfer function over the entire frequency range; consequently, we choose to plot the 4th repetition of its transfer function, located between 40MHz and 50MHz.




FIGURE 4.34: Highpass, lowpass and bandpass configurations for the DF-CT-ADC transfer function.

We conclude that tuning the DF-CT-ADC coefficients allows us to achieve any type of transfer function.

Delay Tap Tuning

In this section we configure the DF-CT-ADC in the previously presented highpass configuration with a feedback loop gain determined by choosing $R_{degGM} = 10k\Omega$ and $I_{refGM} = 2\mu A$. The delay reference current is then varied from 117nA to 147nA such that the repetition frequency is increased from 12MHz to 14MHz. The resulting DF-CT-ADC transfer functions are plotted in Figure 4.35.




FIGURE 4.35: Tuning the center frequency of the DF-CT-ADC signal transfer function.

Tuning the delay tap values as well as reconfiguring the IIR filter from a highpass to a lowpass configuration allows us to cover the entire 40MHz bandwidth of the CT-ADC.

Feedback Gain Tuning

Tuning the feedback gain of the DF-CT-ADC loop has a direct impact on the “sharpness” of the its signal transfer function. We choose a highpass configuration of the DF-CT-ADC along with a tap delay of 100ns. Three feedback loop configurations are studied, the amount of rejection achieved in these configurations is summarized in Table 4.13.

TABLE 4.13: Simulated IIR feedback configurations.

configuration #	I_{refGM}	R_{degGM}	simulated rejection
1	$3\mu A$	$10k\Omega$	23dB
2	$2\mu A$	$10k\Omega$	14dB
3	$3\mu A$	$50k\Omega$	11dB

The resulting DF-CT-ADC transfer functions are plotted in the Figure 4.36. Increasing the feedback loop gain by 9.5dB (equivalent from switching from configuration #3 to configuration #1) enables us to increase the rejection of the filter from 11dB to 23dB

which translates into an increase of A_{att} from 4.5dB to 8.5dB along with an increase of A_{gain} from 6.5dB to 14.5dB (Figure 2.19 on page 47). These results are summarized in Table 4.14.




FIGURE 4.36: DF-CT-ADC transfer functions for different configurations of the feedback loop voltage gain.

TABLE 4.14: Rejection performance achieved by the proposed DF-CT-ADC.

configuration #	A_{att}	A_{gain}
1	8.5dB	14.5dB
2	6.6dB	7.4dB
3	4.5dB	6.5dB

4.5.2 CT-FIR Performance

In this section we study the performance of the 9th order output FIR. As previously, the parameters of interest are the amount of rejection achieved along with a tunability of the central frequency and of the type of transfer function. The CT-FIR transfer function can be changed through the following parameters:

- CT adder coefficients: b_0, b_1, \dots, b_9 – determine the type of transfer function;
- delay reference current: I_τ – sets the transfer function repetition frequency;

Coefficient Tuning

The search algorithm described in Section 1.3 on page 9 requires a reconfigurable FIR filter capable of changing its transfer function from lowpass to highpass and bandpass.

To showcase this, the FIR filter is studied in three configurations, corresponding to the coefficients given in Table 4.15, with all tap delays configured to 100ns. The spectrum repetition frequency thus becomes 10MHz, meaning that it is sufficient to simulate the filter with an input signal which has its frequency swept between 40MHz and 50MHz to completely determine the CT-FIR frequency response.

TABLE 4.15: Highpass, lowpass and bandpass CT-FIR configurations.

coefficient	lowpass	highpass	bandpass
b_0	0.25	0.25	0
b_1	0.5	-0.5	0.25
b_2	0.5	0.5	-0.25
b_3	0.75	-0.75	-0.5
b_4	0.75	0.75	0.5
b_5	0.75	-0.75	0.5
b_6	0.75	0.75	-0.5
b_7	0.5	-0.5	-0.25
b_8	0.5	0.5	0.25
b_9	0.25	-0.25	0

The resulting transfer functions are plotted in Figure 4.37. We conclude that by tuning the coefficients of the CT-FIR adder, the resulting frequency response can be changed from lowpass to highpass and bandpass.




FIGURE 4.37: Highpass, lowpass and bandpass configurations for the CT-FIR transfer function.

Delay Tap Tuning

We set the FIR in highpass configuration and tune the value of its delay taps such that its spectrum repetition frequency is shifted from 12MHz to 14MHz. The resulting FIR

transfer functions, simulated over the 40MHz to 50MHz frequency range, are plotted in Figure 4.38.




FIGURE 4.38: Tuning the central frequency of the CT-FIR filter transfer function.

We have thus shown that by tuning the FIR coefficients and its delay taps, we can generate a transfer function with a pass band which can be shifted over the entire CT-ADC frequency range, from 10MHz to 50MHz.

4.5.3 Interferer Rejection

We now study the CT-DSP with a worst case input scenario consisting of two interferers and a useful signal, which have a total input swing equal to the full scale of the CT-ADC. The position of the useful signal is fixed inside the passband of the DF-CT-ADC and CT-FIR at 45MHz assuming a highpass configuration of both filters. Concerning the interferers, one is kept at 40MHz in the stopband of the CT-DSP, while the other is injected at 42.25MHz, such that it generates a third order intermodulation term inside the CT-DSP passband, at 44.5MHz. A schematic representation of the input is given in Figure 4.39.




FIGURE 4.39: Input scenario used to test the performance of the proposed CT-DSP.

Simulations are done for an input signal-to-interferer ratio of -23dB , close to the limit of robust reception according to the results from Section 1.6 on page 28. Figure 4.40 shows the spectrum of the input, DF-CT-ADC and CT-FIR output.




FIGURE 4.40: Spectrum of the signal at various points in the proposed system (DF-CT-ADC & CT-FIR).

We note that, going through the DF-CT-ADC, the SIR improves from -23dB to 4.5dB , meaning that the DF-CT-ADC achieves an out-of-band rejection of 27.5dB . Moreover, the DF-CT-ADC operates at 12.8dB of SFDR, meaning that the ratio between the useful signal and third order intermodulation term is 17.3dB , sufficient not to compromise robust reception of input signals (as specified in Section 1.5 on page 20, the baseband demodulator only requires 12dB of SNR). Finally, at the output of the FIR, the SIR further improves, from 4.5dB to 27.8dB , meaning that the CT-FIR adds an extra 23.3dB of rejection. Note that the CT-FIR improves the signal to intermodulation ratio only marginally, by around 4dB , as both the useful signal and the third order intermodulation term from the interferers are located inside the CT-FIR bandwidth.

To quantify the improvements offered by the DF-CT-ADC, the same input signal is used, this time into a system consisting only of the CT-ADC followed by the CT-FIR. In this scenario, there is no feedback path around the CT-ADC, hence there is no predictive

interferer cancellation. The spectra of the signal at the input of the system, CT-ADC output and FIR output are showed in Figure 4.41.




FIGURE 4.41: Spectrum of the signal at various points in the proposed system with no feedback loop around the ADC (CT-ADC & CT-FIR).

Besides observing less interferer attenuation, there is another fundamental, more important difference between the results of the two previous simulations. With DF-CT-ADC feedback, the average activity observed at the output of the CT-ADC was around 130MHz (one event every 7.5ns), while the simulation without the DF-CT-ADC, yielded an average activity of 256MHz (one event every 3.9ns). As a consequence, the power of the output CT-FIR drastically increases in the second scenario in comparison with the first. This difference is highlighted in Table 4.16.

A power breakdown for a full-scale input of the complete DF-CT-ADC-DSP system is given in Figure 4.42. Note that the power consumption of most presented blocks (IIR delays, IIR adder, FIR delays and FIR adder) scales linearly with the amplitude of the input swing, making the results a worst case estimation of the system power consumption.

TABLE 4.16: Power breakdown of the proposed CT-ADC-DSP system with and without DF-CT-ADC feedback.

component	DF-CT-ADC & CT-FIR	CT-ADC & CF-FIR
<i>token frequency</i>	130MHz	256MHz
CT-ADC core	17.5 μ W	26.2 μ W
IIR delays	13.5 μ W	n/a
IIR adder	3.2 μ W	n/a
IIR analog FB	6.4 μ W	n/a
FIR delays	40.8 μ W	78.5 μ W
FIR adder	8.8 μ W	16.9 μ W
total	90.2 μ W	121.6 μ W

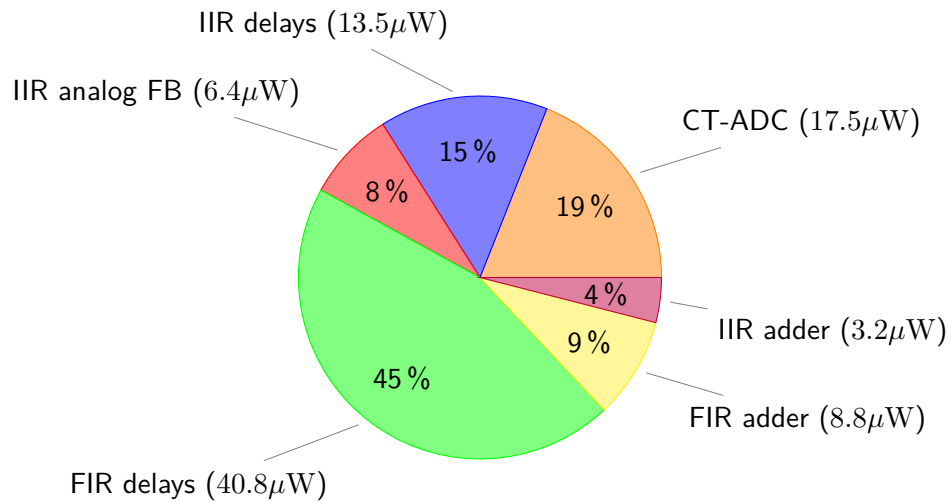


FIGURE 4.42: Breakdown of the complete DF-CT-ADC-DSP power consumption.

4.5.4 Power Consumption Scaling

To demonstrate the scalability of the proposed system's power consumption with respect to the input signal characteristics, we perform a set of simulations in which the amplitude of the input is swept. We choose a single tone input signal, located at the high end of the bandwidth, 50MHz, inside the passband of both filters. It is important to note that the DF-CT-ADC performs in a very similar manner to the CT-ADC: the rate of events it generates at the output is *proportional to the amplitude* of the input signal and *independent of its frequency*. Thus, we predict that the CT-FIR as well as the DF-CT-ADC itself will have a power consumption proportional to this event rate. Simulation results are plotted in Figure 4.43.

We can now extract the linear relation linking the power consumption of the two parts of the system to the input amplitude. The resulting equations are given in equation 4.25, for the CT-FIR, and equation 4.26, for the DF-CT-ADC. The standby power of the proposed system is computed by activating all blocks without applying any input; simulations




FIGURE 4.43: CT-DF-ADC and CT-FIR power consumption versus the input amplitude (system event rate). Note that the power consumption of the CT-ADC is included in both previously defined systems.

show that the CT-FIR thus draws a static power of $15.7\mu\text{W}$ ($11.1\mu\text{W}$ for the CT-ADC and another $4.6\mu\text{W}$ for the other CT processing blocks used in the CT-FIR) while the DF-CT-ADC draws a static power of $20\mu\text{W}$ ($11.1\mu\text{W}$ in the CT-ADC and $9.6\mu\text{W}$ in the feedback loop). Most of the latter comes from the static power of the CT-ADC and from the continuous power required by the feedback loop: more specifically, by the transconductance and filtering, voltage gain stage.

$$P_{CT-FIR} = 0.93 \left(\frac{\mu\text{W}}{\text{mV}_{p-p}} \right) \cdot V_{in}(\text{mV}_{p-p}) + 13.8\mu\text{W} \quad (4.25)$$

$$P_{DF-CT-ADC} = 0.36 \left(\frac{\mu\text{W}}{\text{mV}_{p-p}} \right) \cdot V_{in}(\text{mV}_{p-p}) + 20.76\mu\text{W} \quad (4.26)$$

We thus confirm that the power of the CT-DSP (CT-IIR and CT-FIR, excluding the CT-ADC) scales linearly with the input amplitude requiring a static power consumption under $14.2\mu\text{W}$, which can be broken down into: $9.6\mu\text{W}$ for the DF-CT-ADC active feedback path and another $4.6\mu\text{W}$ consumed by the rest of the CT-DSP blocks.

4.5.5 Noise

To compute the noise in the proposed CT-ADC/DSP, we input a -3dBFS single tone signal at the input, at 50MHz , inside the passband of both filters. A transient noise simulation is done and the resulting noise is integrated over the entire ADC bandwidth: $[10\text{MHz } 50\text{MHz}]$. The extracted SNR is 31dB , close to the specification defined in Figure 4.6 on page 94: the proposed filter does not degrade the worst-case noise performance

of the lone CT-ADC. This corresponds to an average, in-band, input referred noise of $252\text{nV}/\sqrt{\text{Hz}}$.

The spectrum of the output signal is plotted in Figure 4.44. The low frequency noise spectral density is determined by the delay cell noise of the CT-FIR, and hence is white and not first order shaped. As the frequency increases, the CT-ADC noise becomes dominant and a 20dB/dec increase in the noise floor is observed. For very high frequencies, the noise power decreases, hinting at the presence of several high frequency poles. Throughout the spectrum, small peaks are observed in the noise power at multiples of 10MHz, which correspond to different order repetitions of the CT-DSP passband.




FIGURE 4.44: Output spectrum of a full-scale, 50MHz single tone input, for a transient noise simulation of the entire system.

4.5.6 Comparison with State of the Art

Previous CT-ADC-DSP implementations are either aimed to process low frequency voice signals ([68] and [61]) or GHz range ultra-wide band RF signals ([37]), therefore a direct comparison between these solutions is not evident. Nevertheless, a comparison of the main performance metrics is provided in Table 4.17.

We now compare the proposed, continuous time, method of achieving tunable filtering with existing analog or discrete time implementations. Consequently, we slice the proposed CT-ADC-DSP in two parts and draw a comparison between the proposed DF-CT-ADC and existing analog or discrete time IIR filter implementations along with another comparison between the proposed CT-FIR and existing analog or discrete time FIR implementations. The first comparison is highlighted in Table 4.18; as a figure of merit, we use the energy figure of merit usually employed to compare IIR filters, its formula is given in equation 4.27 with P_{DC} – the filter power consumption (in W), N –

TABLE 4.17: Comparison between the proposed DF-CT-ADC-DSP and other, state of the art, CT-DSP implementations.

performance	Kurchuk ([37])	Schell ([68])	this work
technology	65nm	90nm	28nm FDSOI
supply voltage	1.2V	1V	0.65V
max frequency	3.2GHz	20kHz	50MHz
bandwidth	0.8GHz – 3.2GHz	20Hz – 20kHz	10MHz – 50MHz
area (CT-ADC)	0.06mm ²	n/a	0.06mm ² *
area (CT-FIR)	0.55mm ²	0.073mm ²	0.09mm ²
DSP order	6 (FIR)	14 (FIR)	9 (FIR) & 3 (IIR)
power	1.1mW – 10mW	0.35mW – 1.71mW	25μW–215μW

(*) DF-CT-ADC: includes the continuous time digital feedback loop

the filter order (no unit), f_c – the cutoff frequency (in Hz) and $SFDR$ – a measure of linearity (in dB). It is important to note that the $SFDR$ used here is *not* the two tone $SFDR$ presented previously in this chapter but rather a measure of the circuit linearity – through its third order interception point, IIP_3 (measured in dBm) – combined with a measure of the circuit noise – P_n (integrated over the filter passband and measured in dBm) – as shown in equation 4.28.

TABLE 4.18: Comparison between the proposed DF-CT-ADC and existing analog or digital IIR filter implementations.

performance	Drost [77]	Zhao [78]	Lechevallier [79]	Tohidian [80]	D’Amico [81]	Liscidini [82]	Oskooei [83]	this work
implementation	ring oscillator integrator	switched capacitor*	$G_m - C$	DT-IIR*	source follower based	active RC	$G_m - C$	CT-IIR
technology	90nm	65nm	28nm FD-SOI	65nm	180nm	90nm	90nm	28nm FD-SOI
supply	0.55V	1.2V	0.7V	1.2V	1.8V	2.5V	1V	0.65V
order	4	4	3	7	4	4	6	3
power consumption	2.9mW	0.68mW	4mW	1.96mW	4.1mW	1.25mW	4.35mW	activity-dependent 20μW – 79μW
input-referred noise	62.4μVRMS	19.5 $\frac{nV_{RMS}}{\sqrt{Hz}}$	5.9 $\frac{nV_{RMS}}{\sqrt{Hz}}$	2.85 $\frac{nV_{RMS}}{\sqrt{Hz}}$	7.5 $\frac{nV_{RMS}}{\sqrt{Hz}}$	32μVRMS	75 $\frac{nV_{RMS}}{\sqrt{Hz}}$	252 $\frac{nV_{RMS}}{\sqrt{Hz}}$
TF type	lowpass	lowpass	lowpass	lowpass	lowpass	lowpass	lowpass	all**
tunability	cutoff	cutoff	no	cutoff	no	no	cutoff	full***
tuning range	7MHz–30MHz	0.5MHz–10MHz	n/a	0.4MHz–30MHz	n/a	n/a	8.1MHz–13.5MHz	10MHz–50MHz
IIP_3	28dBm	17.6dBm	5.2dBm	16dBm	12dBm	36dBm	22dBm	14.8dBm
area	0.29mm ²	0.02mm ²	0.04mm ²	0.42mm ²	0.26mm ²	n/a	0.239mm ²	0.06mm ²
FoM_{IIR} ****	25.1aJ	20.5aJ	61.33aJ	7.15aJ	83aJ	3.3aJ	287aJ	200aJ****

(*) requires clock generation

(**) lowpass, bandpass, highpass & notch

(***) computed according to the formula given in equation 4.27

(****) obtained for a –3dBFS single tone input signal which corresponds to a power consumption of 49.6μW; $f_c = 1MHz$

$$FoM_{IIR} = \frac{P_{DC}}{N \cdot f_c \cdot 10^{SFDR/10}} \quad (4.27)$$

$$SFDR = \frac{2}{3} \cdot (IIP_3 - P_n) \quad (4.28)$$

To provide a fair comparison with the state of the art, we have chosen an $f_c = 1MHz$ for the proposed DF-CT-ADC which corresponds to a bandwidth equal to half of that of a standard bandpass configuration of our filtering CT-ADC. We can see that despite having an inferior energy efficiency, the proposed DF-CT-ADC achieves by far the smallest

power consumption as well as the best tunability. Moreover, the proposed system's output coincides with that of the CT-ADC, meaning that it has an embedded CT analog to digital conversion, potentially easing the implementation of subsequent filtering / data processing blocks. Note that the input signal configuration used to extract the data required to compute FoM_{FIR} implies the use of a single tone, near-full scale signal (i.e. -3dBFS) and fails to take into account the scalability of the power consumption with respect to the signal characteristics.

We now compare the proposed CT-FIR with discrete time or analog FIR filter implementations. As a figure of merit we use the FoM_{FIR} which is the power per pole per hertz per number of input bits per number of coefficient bits as given in equation 4.29, with P_{DC} – the power consumption, N – the filter order, f_s – the sampling frequency, B_{in} – the number of input bits and B_{coeff} – the number of bits used for the filter coefficients. Note that this figure of merit cannot be computed if the filter input is analog, as the B_{in} would be undefined.

TABLE 4.19: Comparison between the proposed CT-FIR and existing analog or digital FIR filter implementations.

performance	Chang [84]	Ma [85]	Jin [86]	O'hAnnaidh [87]	Hwang [88]	this work
implementation	DT-FIR	DT-FIR	analog FIR	analog FIR	DT-FIR	CT-FIR
technology	90nm	130nm	500nm	45nm	130nm	28nm FD-SOI
order	8	14	8	16	8	9
input bits	8	8	n/a	n/a	8	5
coefficient bits	8	8	8	6	8	3
supply	300mV	360mV	3.3V	1.1V	85mV	650mV
power consumption	0.74 μ W	5.9mW	160mW	64mW	40nW	activity-dependent 15.7 μ W–163 μ W**
frequency	148kHz	187MHz	180MHz	3.2GHz	240Hz	50MHz
area	0.48mm ²	9.91mm ²	0.28mm ²	1.21mm ²	1.54mm ²	0.09mm ² ***
FoM_{FIR} **** (nW/MHz/order/in – bit/coeff – bit)	9.76	35.31	n/a*	n/a*	325	8.2****

(*) cannot be computed since the number of input bits is undefined (analog implementation)

(**) excluding the CT-ADC

(***) computed according to the formula given in equation 4.29

(****) obtained for a -3dBFS single tone input signal which corresponds to a CT-FIR power consumption of 88.5 μ W; $f_s = 2 \cdot f_{BW} = 80\text{MHz}$

$$FoM_{FIR} = \frac{P_{DC}}{N \cdot f_s \cdot B_{in} \cdot B_{coeff}} \quad (4.29)$$

We can see that the FoM_{FIR} achieved by the proposed implementation is competitive with respect to existing, state of the art designs of discrete time FIR filters. Moreover, the proposed system is capable of scaling its power consumption, achieving a standby (0 input) power consumption of only 15.7 μ W.

4.6 CT-DSP Conclusion

In this chapter we have presented the implementation of a CT-DSP designed to enable an ultra low power, tunable filtering transfer function suitable for interferer rejection and channel selection in low power radios. We have showed that the proposed system is capable of achieving highpass, lowpass and bandpass transfer functions, with rejection levels which can be traded off for an even lower power consumption by switching OFF a part of the filter. The central frequency of the filters is also continuously programmable through the tuning of an external current reference.

We have also proposed a filtering loop around the CT-ADC which enables us to cancel out unwanted frequency components *before* injecting them into the CT-ADC. In high interferer situations, the proposed scheme effectively reduces the total signal swing, as seen by the CT-ADC, thereby reducing its power consumption. Simulations have shown that a power reduction of a factor between 2 and 3 is achievable.

Chapter 5

Conclusion

THE conclusion of this manuscript is divided into three parts which cover the following topics: first, a summary of the contributions to the state of the art is presented, followed by a section covering possible improvements of the proposed circuits and architecture. Finally, the last part of this chapter discusses the possibilities opened by the proposed low power, programmable interferer rejection stage in terms of receiver and IoT network design.

5.1 Motivations and Contributions of this Work

The success of IoT networks depends on the implementation of easily deployable, cheap, low-cost, low-maintenance wireless nodes. These nodes are expected to have a lifetime of several years with very limited human interaction, meaning that they should be independent from a power perspective, either relying on batteries or on different forms of energy harvesting. This imposes a hard limit on the amount of power available for the different electronic systems used in each node. One of the most power consuming systems is the wireless receiver which can consume up to several milliwatts when turned ON.

Current low power wireless communication standards rely on receiver duty cycling (turning ON the receiver for only very short periods of time) to limit its *average* power consumption. This solution works well when the wireless node needs to exchange data relatively often, such that the synchronization overhead is small compared to the energy requirements for the data exchange. However, some applications can have an average time between two consecutive packets of several tens of minutes, greater than the network synchronization period. This means that for each useful data packet, the receiver needs to

be turned ON several times to guarantee it maintains its synchronization with the network, thereby greatly increasing its power requirements. Moreover, another disadvantage of dramatically decreasing the duty cycle is related to the communication delays which increase proportionally.

Recently, a new receiver paradigm has been proposed which consists of embedding a second, ultra low power, low performance receiver (WU-RX) next to the main receiver. This opens up the design space, as it is now possible to monitor the channel activity while consuming a power two orders of magnitude smaller than that of the initial data receiver. An analysis of WU-RX state of the art has shown that the main feature missing from existing implementations is the possibility of robustly receiving the desired signal in the presence of interferers.

Consequently, this thesis focuses on the implementation of an ultra low power IF filtering stage suitable to be used in the context of wake-up receivers. Its specifications have been derived from the WU-RX architecture presented in the introduction of this manuscript. The filter is implemented based on the operation principle of CT digital signal processing systems which enable a clock-less, low power operation which natively scales its energy consumption based on the characteristics of the input signal.

The CT-ADC-DSP chain uses a CT-ADC inspired from the operation of delta modulators, which, thanks to its looped structure has modest hardware requirements (number of comparators) regardless of the number of converted bits. This enables us to achieve a very low power consumption for a moderately precise, in terms of ENOB, ADC. The greatest strength of the proposed solution also happens to be its greatest weakness, as the inevitable delay of the feedback loop limits the maximum input frequency. To combat this, we have opted for an architecture with a very simple feedback path, consisting only of a comparator, two logic gates and a set of switches, which achieves loop delays, in simulation, of only 250ps. What results is an alias-free, error-shaping, asynchronous ADC with a Walden figure of merit between $3fJ/\text{conv} - \text{step}$ and $10fJ/\text{conv} - \text{step}$, among the best reported for the selected frequency range [10MHz 50MHz] which also represents an improvement of at least 3X compared to existing CT-ADCs.

In the last part of the manuscript we have presented the CT-DSP designed specifically to operate alongside the previously presented CT-ADC. The first issue encountered for this block was related to the increased CT-ADC activity in the presence of high power, out-of-band signals which also generates a proportionally increased power consumption in the digital part of the CT-DSP. To counter this, we have proposed a feedback structure in which the CT-ADC output is injected in a CT-DSP which is then injected back to the CT-ADC input; its operation resembles that of a CT-IIR filter. Thus, the CT-DSP can be programmed such that it cancels certain frequencies at the CT-ADC input while

boosting others. This structure has the advantage of attenuating interfering signals while at the same time increasing the linearity of the conversion with respect to these out-of-band components. Simulations show, that in high power interferer configurations, the proposed CT-DSP structure reduces CT-ADC output activity by factors between 2 and 3 corresponding to an equivalent power reduction in the digital part of the system.

The proposed DF-CT-ADC achieves an out-of-band rejection of up to 20dB; supplementary rejection is implemented by adding a CT-FIR at the output of the filtering CT-ADC structure. We have opted for a 9th order CT-FIR; simulations show that it adds around 25dB of rejection. Moreover, its power consumption also benefits from the reduced filtering CT-ADC output activity, as the number of tokens which need to be processed is reduced.

5.2 Improvements of the Proposed Design

In this section we discuss possible ways of improving the different blocks of the proposed CT-ADC-DSP system.

5.2.1 CT-ADC

By design, the CT-ADC scales the average time between two consecutive output events according to the swing of the input signal it is presented: high amplitude signals generate a high number of output events while low amplitude signals generate a low number of output events. In the case of a purely digital behavior, the CT-ADC would scale its power accordingly, i.e. reducing the average input swing by a factor of 4 should result in a 4 times lower power consumption. However, measurement results from Figure 3.37 on page 85 show that reducing the amplitude of the input signal from $200\text{mV}_{\text{p-p}}$ to $40\text{mV}_{\text{p-p}}$ (factor of 5) results in a power which decreases from $18.7\mu\text{W}$ to $11\mu\text{W}$ (factor of 1.7).

The main reason for this loss in scalability is the high power required by the following, mostly analog, components:

- *comparators*: according to results from Figure 3.31 on page 80 the two core comparators draw 50% of the CT-ADC power and the two overflow comparators draw 15% of the CT-ADC power. The offset compensated, inverter-based architecture chosen for their implementation requires a power which can be divided into two categories according to its scalability with respect to the characteristics of the input

signal. On the one hand, the first inverter stage draws a significant crowbar current, as it is behaving like a high gain stage; it requires a power which depends only on the difference between the instantaneous capacitor voltage and the comparator threshold. On the other hand, the last 3 inverter stages have an almost purely dynamic power consumption, proportional to the average number of CT-ADC tokens.

- $G_m - C$: the purely analog implementation of the transconductance imposes a constant power consumption regardless of the characteristics of the input signal. This means that even in the case of a 0 input (stuck to common mode), the $G_m - C$ will still require $2.6\mu\text{W}$ of power.

Optimizing the behavior of the previously presented items is not trivial. The $G_m - C$ cell requires a more power-efficient architecture, however it is difficult to imagine an implementation which can scale its power consumption according to the input. The comparator design, on the other hand, needs to be optimized to minimize the crowbar current of its first stage. This poses a problem because this also means decreasing the gain of this stage which results into an increased comparison time, thereby reducing the maximum frequency the CT-ADC is able to process. One solution which allows an escape from the previously presented trade-off consists in using two comparators: one designed for low crowbar current and slow operation and another using a high input gain resulting in a very fast operation but functioning at a higher threshold. This concept is described in Figure 5.1.

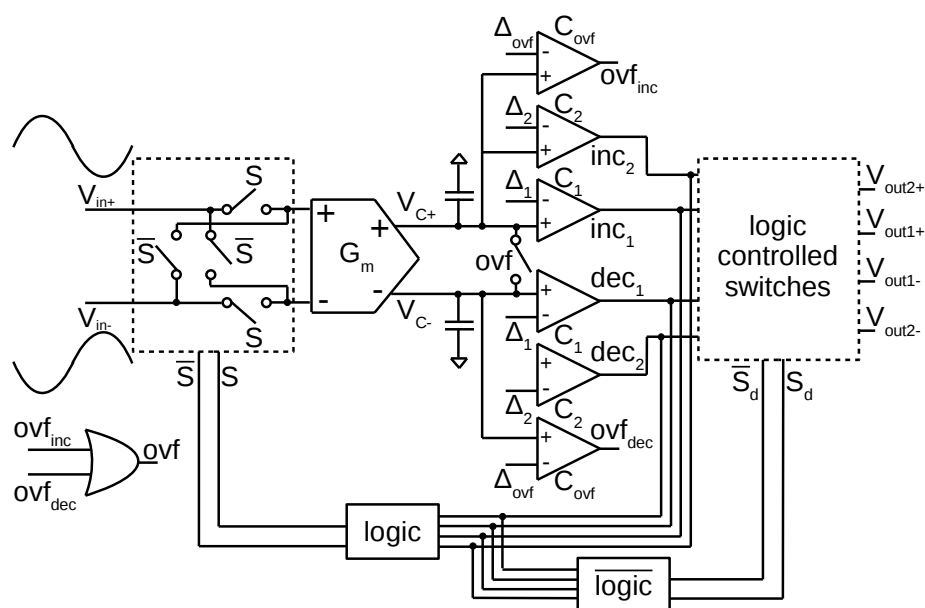


FIGURE 5.1: Concept of the improved CT-ADC design.

The CT comparators draw a high crowbar current when the signal is close to their threshold voltages; designing low crowbar current comparators is not an option as their speed inevitably degrades. We suggest an architecture level solution which consists in the addition of an extra set of comparators, labeled C_2 on top of the comparators from the original design, C_1 (core) and C_{ovf} (overflow). With the new design, comparators labeled C_1 can be designed to consume a low static power and have a poor delay performance; as a consequence, there will be an increased overflow event rate with respect to their threshold (Δ_1). A second set of comparators C_2 can be optimized for high speed, and hence draw a significant static current when the capacitor voltages approach their threshold (Δ_2). However, this event will not occur very often, as the capacitor voltages are most of the time bounded by Δ_1 . Finally, in the unlikely event that C_2 comparators are overflowed, an extra set of comparators is used, labeled C_{ovf} , which reset the capacitors when they are triggered.

The newly proposed system has the information encoded on 4 wires (inc_1 , inc_2 , dec_1 and dec_2) which need to be flipped before being sent to the output. This results in a more complex and power consuming logic controlling the input and the output switches, which has the advantage of scaling its power consumption according to the CT-ADC output activity.

5.2.2 DF-CT-ADC Feedback Path

For the implementation of the DF-CT-ADC feedback path we have opted for a voltage domain adder for the CT-FIR filter. Its output is amplified and then sent into a G_m cell, similar to the one located in the direct CT-ADC path. The reason behind this is the fact that, as seen in Chapter 3, the CT-ADC has high frequency components at its output which need to be filtered *before* being sent back to its input. This filtering is accomplished in the voltage domain, either by the amplifier which has a first order roll-off beyond 50MHz or by a passive second order $R - C$ filter located before the amplifier input, as seen in Figure 4.31 on page 121.

A more direct approach would be to inject the outputs of the digital delay cells into a series of switchable current sources which act as a current mode adder and are connected to the CT-ADC capacitors. An example of such a current source is given Figure 5.2.

A reference current, I_{ref} , is injected into the diode connected transistor M_1 and is used to tune the gain of the feedback path. A switchable current mirror ($M_{21} - M_{2N}$), whose configuration sets the value of the respective coefficient, copies I_{ref} and injects it into a second diode connected transistor M_3 . Finally, digital pulses coming from the continuous

The CT-FIR cannot be used for such filtering, as its transfer function is periodic with a period equal to τ_{tap} : if a signal of frequency f lies in the passband of the CT-FIR, then all of its harmonics will also lie in higher order repetitions of its transfer function. One solution, which has not been investigated, would be to apply pulse shaping to the current impulses generated by each current mode adder cell.

5.2.3 CT-DSP Delay Cells

The CT-DSP delay cells need to be calibrated so that the value of their standard deviation is small enough not to degrade the CT-DSP transfer function. In the current version of the CT-DSP, this calibration is done manually, meaning that the delay cell inputs can be disconnected from the CT-ADC enabling the injection of external pulses which are then sent through the delay taps and finally to a system output. The resulting delays of each individual tap are then measured and a value is computed for the control current, I_τ , as well as the number of programming 10ns elements and 1ns elements which need to be added.

In a more advanced version of the CT-ADC-DSP system, a delay calibration system should be embedded on chip. However, its design is not trivial, as the value of each delay cell is affected by two parameters: a local variable, corresponding to the value of the memory controlling the supplementary 10ns and 1ns elements, as well as a global variable, the current I_τ , which controls the values of all delay cells.

5.3 Future Work

The obvious follow-up of the work presented in this manuscript would be the implementation of the proposed CT-ADC-DSP system in a wake-up radio. However, in order to efficiently leverage the interferer filtering capabilities of the proposed circuits, several considerations should be kept in mind.

The main feature of the proposed system lies in its capability of scaling its power consumption according to the characteristics of the input signal (presence or lack of strong interferers). Therefore, the optimization of its power consumption requires a thorough analysis of the expected environment conditions *before* deciding on the best signal search strategy as well as on the required performance levels of the CT-DSP. Moreover, the proposed system could also benefit from being coupled with adaptive learning algorithms. An analysis of the *average* network-wide environment conditions would lead to a network-level optimized power consumption of the receivers but it would

fail to take into account the fact, once deployed, some nodes would find themselves in more favourable conditions than others. Consequently, the long term power consumption of each individual receiver could be optimized at a node level, by having every node sense its environment and change its filtering strategy accordingly.

Appendix A

Squarer Noise Analysis

In this appendix we analyse the expression and shape of the noise at the output of a squarer circuit, employed for the asynchronous demodulation of OOK signals. The input scenario is presented in the Figure A.1: the input signal is composed of a single tone signal, situated at frequency of f_{sig} , along with bandlimited white noise, bounded by BW_{noise} , determined by the filtering achieved in the receiver front-end. This signal is then inserted into a squarer, at which point its SNR is computed over a narrow bandwidth, BW_{BB} , determined by the communication data rate.




FIGURE A.1: Scenario used to study the noise performance of a squarer.

We start by remarking that the self-mixing (squaring) operation generates frequency components around DC and $2f_{sig}$. The shape as well as the bandwidth of the noise observed at the output is discussed and explained at a later point in this appendix.

We define the input signal-to-noise ratio (SNR_{in}) as the ratio between the input tone power and the input noise power integrated from 0 to infinity. Since we are considering bandlimited white noise, the previous relation can be rewritten as the ratio between the input signal power divided by the noise power spectral density (PSD) multiplied by the noise bandwidth, as shown in equation A.1. Note that the total noise input power (σ_{in}^2)

can be expressed as given in equation A.2.

$$SNR_{in} = 10 \log \frac{P_{sig}^{in}}{\int_0^\infty PSD_{noise}^{in}(f) \cdot df} = 10 \log \frac{P_{sig}^{in}}{PSD_{noise}^{in} \cdot BW_{noise}} \quad (A.1)$$

$$\sigma_{in}^2 = PSD_{noise}^{in} \cdot BW_{noise} \quad (A.2)$$

To compute the output signal-to-noise ratio, we consider the ratio between the power of the DC signal component and the noise integrated from DC to BW_{BB} (determined by the data rate of the communication). The corresponding equation is given in equation A.3. Note that for OOK modulation, if we are targeting a bit error rate of 1e-3, 12dB of SNR_{out} is required.

$$SNR_{out} = 10 \log \frac{P_{sig}^{out}}{\int_0^{BW_{BB}} PSD_{noise}^{out}(f) \cdot df} \quad (A.3)$$

We now express the mathematical equations describing the input and the output signal in order to determine the relation between PSD_{noise}^{in} and PSD_{noise}^{out} (equation A.4 and equation A.5). In the latter expression, k is the conversion gain of the squarer, measured in V^{-1} , necessary to adjust the unit of V_{out} .

$$V_{in} = s(t) + n(t) = A_{in} \sin(2\pi f_{in} t) + n(t) \quad (A.4)$$

$$\begin{aligned} V_{out} &= k \cdot (A_{in} \sin(2\pi f_{in} t) + n(t))^2 \\ &= k \cdot \left(A_{in}^2 \sin^2(2\pi f_{in} t) + n(t)^2 + 2 \cdot A_{in} \sin(2\pi f_{in} t) \cdot n(t) \right) \\ &= k \cdot \left(\frac{A_{in}^2}{2} - \frac{A_{in}^2}{2} \cos(2\pi 2f_{in} t) + n(t)^2 + 2 \cdot A_{in} \sin(2\pi f_{in} t) \cdot n(t) \right) \end{aligned} \quad (A.5)$$

Analyzing the expression of V_{out} (equation A.5) we can conclude that the noise is made up of two components: self mixed noise ($n^2(t)$) and noise mixed with the signal ($2 \cdot A_{in} \sin(2\pi f_{in} t) \cdot n(t)$). The DC signal component is given by $\frac{A_{in}^2}{2}$ while the component at twice the signal frequency, $\frac{A_{in}^2}{2} \cos(2\pi 2f_{in} t)$, can be ignored since it is not inside the BB_{BW} . To compute the output noise power we use the definition of the variance, given in equation A.6.

$$\sigma_{out}^2 = E[V_{out}^2] - (E[V_{out}])^2 \quad (A.6)$$

We replace V_{out} by $k \cdot (S_{in} + N_{in})^2$ in equation A.6; the result is given in equation A.7. Note that the odd moments of N_{in} have been removed since they are all equal to 0, i.e.

$E[N_{in}]$, $E[N_{in}^3]$ etc. This derivation has been developed in [38].

$$\begin{aligned}\sigma_{out}^2 &= k \cdot \left(E \left[(S_{in}^2 + 2S_{in}N_{in} + N_{in}^2)^2 \right] - \left(E[(S_{in} + N_{in})^2] \right)^2 \right) \\ &= k \cdot \left(S_{in}^4 + 6S_{in}^2 E(N_{in}^2) + E[N_{in}^4] - \left(E[S_{in}^2 + 2S_{in}N_{in} + N_{in}^2] \right)^2 \right)\end{aligned}\quad (\text{A.7})$$

We now replace the odd moments by 0 and the even moments by the values given by theory: $E[N_{in}^2] = \sigma_{in}^2$ and $E[N_{in}^4] = 3\sigma_{in}^4$. The output noise power thus becomes as given in equation A.8, which is divided into two equal parts, one corresponding to the noise around the DC point and another corresponding to the noise around $2f_{sig}$.

$$\sigma_{out}^2 = k \cdot \left(4S_{in}^2\sigma_{in}^2 + 2\sigma_{in}^4 \right)\quad (\text{A.8})$$

Analyzing equation A.8, we note that the first term corresponds to the noise resulting from mixing between the signal component and the bandlimited white noise at IF, while the second part corresponds to noise created by self-mixing of the IF noise. The power spectral densities of these two components are analyzed next.

Noise Mixed with the Signal

To determine the shape of the baseband noise issued by the cross-mixing of the signal with the noise component, we use the multiplication – convolution duality with respect to the Fourier transform. The spectrums of the signal, noise and baseband noise (issued from the convolution of the previous two components) are plotted in the Figure A.2.

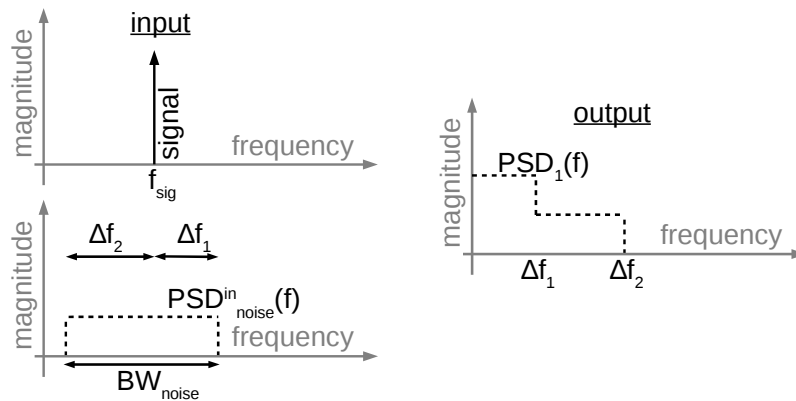


FIGURE A.2: Baseband noise generated by the cross-mixing of the IF noise with the signal component.

We note that the noise power spectral density until Δf_1 is twice that between Δf_1 and Δf_2 . This can be explained by the fact that, from a spectrum perspective, the noise is

symmetric with respect to the signal in the $[f_{sig} - \Delta f_1 \ f_{sig} + \Delta f_1]$, thus generating a double convolution product, while the noise located inside the $[f_{sig} - \Delta f_2 \ f_{sig} - \Delta f_1]$ is single side banded with respect to the signal, generating a single convolution product. The value of the PSD_1 can be computed by integrating this noise and equating it to the corresponding noise power derived in equation A.8, as done in equation A.9. We thus compute PSD_1 and express it in equation A.10.

$$k \cdot (2S_{in}^2 \sigma_{in}^2) = \int_0^{\Delta f_2} PSD_1(f) df = 2 \cdot PSD_1 \cdot \Delta f_1 + PSD_1 \cdot (\Delta f_2 - \Delta f_1) \quad (A.9)$$

$$PSD_1(f) = \begin{cases} \frac{2k \cdot (2S_{in}^2 \sigma_{in}^2)}{\Delta f_1 + \Delta f_2} & f < \Delta f_1 \\ \frac{k \cdot (2S_{in}^2 \sigma_{in}^2)}{\Delta f_1 + \Delta f_2} & \Delta f_1 < f < \Delta f_2 \\ 0 & \Delta f_2 < f \end{cases} \quad (A.10)$$

Noise Self-Mixing

The PSD of the noise generated by noise self mixing can be computed in the same way as done previously. A representation of the spectrums which mix at the input along with the resulting output is given in Figure A.3. The convolution of the two bandlimited white noise spectrums at the input yields a triangular shaped baseband PSD.




FIGURE A.3: Baseband noise generated by the self-mixing of the bandlimited IF noise.

The baseband noise PSD can thus be written as equation A.11. Integrating it from 0 to BW_{noise} and equating it with half of the corresponding noise power computed in equation A.8, we can deduce the expression of the $PSD_2(f)$, which is given in equation A.12.

$$PSD_2(f) = -cst \cdot f + cst \cdot BW_{noise} \quad (A.11)$$

$$PSD_2(f) = \begin{cases} \frac{2k\sigma_{in}^4}{BW_{noise}^2} \cdot (BW_{noise} - f) & f < BW_{noise} \\ 0 & BW_{noise} < f \end{cases} \quad (A.12)$$

Compact Expression of the Baseband Noise PSD

Combining results from equation A.10 and equation A.12 we can derive the compact expression of the baseband noise PSD, which is given in equation A.13.


$$PSD(f) = \begin{cases} \frac{2k\sigma_{in}^4}{BW_{noise}^2} \cdot (BW_{noise} - f) + 2 \cdot \frac{k \cdot (2S_{in}^2 \sigma_{in}^2)}{\Delta f_1 + \Delta f_2} & f < \Delta f_1 \\ \frac{2k\sigma_{in}^4}{BW_{noise}^2} \cdot (BW_{noise} - f) + \frac{k \cdot (2S_{in}^2 \sigma_{in}^2)}{\Delta f_1 + \Delta f_2} & \Delta f_1 < f < \Delta f_2 \\ 0 & f > \Delta f_2 \end{cases} \quad (A.13)$$

The last step in computing the noise performance of the squarer consists of replacing the previous expression in the output SNR definition, equation A.3 on page 146. Given that the noise PSD is a piecewise defined function, the requirement for the SNR_{in} for 12dB of SNR_{out} may depend on the frequency of the signal as well as the precise noise bandwidth. This result is further discussed in the main body of the manuscript as it is beyond the scope of this appendix.


Appendix B

The 28nm UTBB FDSOI CMOS Technology

The Ultra Thin Body and Buried oxide Fully Depleted Silicon On Insulator (UTBB FDSOI) CMOS technology is a planar process which consists of isolating a thin layer of silicon (the channel of the MOSFETs) between a top gate oxide and a thin layer of body oxide; the entire structure is created on top of a silicon base. A cross sectional view as well as a schematic view of an UTBB FDSOI NMOS transistor is given in Figure B.1.



(a) Photo of an actual device.



(b) Schematic view.

FIGURE B.1: Cross sectional view of an UTBB FDSOI transistor [89].

The main advantages of the 28nm UTBB FDSOI technology are all linked to the addition of the second transistor gate (sometimes referred to as backgate) which enhances the electrostatic control of the channel.

Threshold Voltage Control

The most important feature of the FDSOI technology consists in the possibility of controlling the threshold voltage of the transistors by applying a strong bias on their backgate. This enables a dynamic control of the state of the transistor from low speed, low leakage to high speed, high leakage depending on the requirements of the circuit at the specific time instant. This technique is well known and has been employed in bulk technologies, however, the FDSOI technology allows for an 85mV/V threshold voltage scaling, while a classic bulk node only allows for 25mV/V slope.

A cross sectional view of a PMOS transistor next to an NMOS transistor is given in Figure B.2 along with the allowed range of variation of their backbias voltages. Simulated threshold voltages of the FDSOI transistors (PMOS and NMOS) are compared against those of bulk transistors of the same node, Figure B.3.




FIGURE B.2: Sectional view of an FDSOI PMOS transistor next to an FDSOI NMOS transistor, along with the allowed backgate bias voltages.




FIGURE B.3: Threshold voltage variation of FDSOI and standard bulk transistors.

Channel Control

The presence of the second gate as well as the very thin channel thickness account for two important features which have a significant contribution to the continued reduction of channel length. First, the very precise control of the channel potential enabled by the presence of the backgate allows for the use of an undoped channel which significantly reduces mismatch as the channel length of the transistors is reduced. Bulk processes generally require the use of a doped channel, with the total number of dopants proportional to the surface of the channel. As the channel length is reduced, the *average* number of dopants per channel is also reduced, while the *deviation* of this number is constant. This results in a higher spread of the characteristics of a transistor as its dimensions are scaled down. Furthermore, the BOX reduces the cross section of the channel thereby also reducing the source to drain leakage.

Bibliography

- [1] Jean-Baptiste Waldner. Nanoinformatique et intelligence ambiante: Inventer l'ordinateur du xxi-ème siècle. Hermès science publications, 2007.
- [2] D Blaauw, D Sylvester, P Dutta, Y Lee, I Lee, S Bang, Y Kim, G Kim, P Pannuto, D Yoon, W Jung, Z Foo, S Oh, S Jeong, and M Choi. IoT Design Space Challenges: Circuits and Systems. In *Symposium on VLSI Circuits*, pages 1–2, 2014.
- [3] Yao-hong Liu, Christian Bachmann, Xiaoyan Wang, Yan Zhang, Ao Ba, Benjamin Busze, Ming Ding, Pieter Harpe, Gert-jan Van Schaik, Georgios Selimis, Hans Giesen, Jordy Gloudemans, Adnane Sbai, Li Huang, Hiromu Kato, Guido Dolmans, Kathleen Philips, and Harmke De Groot. A 3.7mW-RX 4.4mW-TX Fully Integrated Bluetooth Low-Energy/IEEE802.15.4/Proprietary SoC with an ADPLL-Based Fast Frequency Offset Compensation in 40nm CMOS. In *International Solid-State Circuits Conference*, pages 236–238, 2015.
- [4] Nathan M Pletcher, Simone Gambini, and Jan Rabaey. A $52\mu\text{W}$ Wake-Up Receiver With 72 dBm Sensitivity Using an Uncertain-IF Architecture. *Journal of Solid-State Circuits*, (1):269–280, 2009.
- [5] Salvatore Drago, Domine M W Leenaerts, Bram Nauta, Fabio Sebastiano, Kofi A A Makinwa, and Lucien J Breems. A $200\mu\text{A}$ Duty-Cycled PLL for Wireless Sensor Nodes in 65 nm CMOS. *Journal of Solid-State Circuits*, (7):1305–1315, 2010.
- [6] Robert Bogdan Staszewski, John L Wallberg, Sameh Rezek, Chih-ming Hung, Oren E Eliezer, Sudheer K Vemulapalli, Chan Fernando, Ken Maggio, Roman Staszewski, Nathen Barton, Meng-chang Lee, Patrick Cruise, Mitch Entezari, Khuram Muhammad, and Dirk Leipold. All-Digital PLL and Transmitter for Mobile Phones. *Journal of Solid-State Circuits*, (12):2469–2482, 2005.
- [7] Xiang Gao, Eric A M Klumperink, Mounir Bohsali, and Bram Nauta. A 2.2GHz 7.6mW Sub-Sampling PLL with -126dBc/Hz In-Band Noise and 0.15psRMS Jitter in $0.18\mu\text{m}$ CMOS. In *International Solid-State Circuits Conference*, pages 392–394, 2009.

- [8] Ben W Cook, Axel Berny, Alyosha Molnar, Steven Lanzisera, and Kristofer S J Pister. Low-Power 2.4-GHz Transceiver With Passive RX Front-End and 400-mV Supply. *Journal of Solid-State Circuits*, (12):2757–2766, 2006.
- [9] Xuan Zhang and Alyssa B Apsel. A Low-Power, Process and Temperature Compensated Ring Oscillator With Addition-Based Current Source. *Transactions on Circuits and Systems*, (5):868–878, 2011.
- [10] K R Lakshmikumar and V Mukundagiri. A Process and Temperature Compensated Two-Stage Ring Oscillator. In *Custom Integrated Circuits Conference*, pages 691–694, 2007.
- [11] R. Vijayaraghavan, S. K. Islam, M. R. Haider, and L. Zuo. Wideband Injection-Locked Frequency Divider Based on a Process and Temperature Compensated Ring Oscillator. *IET Circuits, Devices and Systems*, (5):259–267, 2009.
- [12] Somayeh Abdollahvand, Luis B Oliveira, Luis Gomes, and Nuno Paulino. Low Phase-Noise Temperature Compensated Self-Biased Ring Oscillator. In *International Conference on Circuits and Systems*, pages 2489–2492, 2012.
- [13] F Bruccoleri, E A M Klumperink, and B Nauta. Noise Cancelling in Wideband CMOS LNAs. In *International Solid-State Circuits Conference*, number July 2001, pages 406–407, 2002.
- [14] Christian Hambeck, Stefan Mahlknecht, and Thomas Herndl. A $2.4\mu W$ Wake-Up Receiver for Wireless Sensor Nodes with -71dBm Sensitivity. In *International Conference on Circuits and Systems*, pages 534–537, 2011.
- [15] Xiongchuan Huang, Ao Ba, Pieter Harpe, Guido Dolmans, Harmke De Groot, and John R Long. A 915 MHz, Ultra-Low Power 2-Tone Transceiver With Enhanced Interference Resilience. *Journal of Solid-State Circuits*, (12):3197–3207, 2012.
- [16] Carl Bryant and Henrik Sjolund. A 2.45GHz, $50\mu W$ Wake-up Receiver Front-end with -88dBm Sensitivity and 250kbps Data Rate. In *European Solid States Circuit Conference*, pages 235–238, 2014.
- [17] Maarten Lont, Dusan Milosevic, Arthur H M Van Roermund, and Guido Dolmans. Ultra-Low Power FSK Receiver for Body Area Networks with Automatic Frequency Control. In *European Solid States Circuit Conference*, pages 430–433, 2012.
- [18] Camilo Salazar, A Kaiser, Andreia Cathelin, and Jan Rabaey. A -97dBm -Sensitivity Interferer-Resilient 2.4GHz Wake-Up Receiver Using Dual-IF Multi-N-Path Architecture in 65nm CMOS. In *International Solid-State Circuits Conference*, pages 242–244, 2015.

- [19] Specification of the Bluetooth System, 2010. URL https://www.bluetooth.org/docman/handlers/downloaddoc.ashx?doc_id=229737.
- [20] IEEE Standard for Local and Metropolitan Area Networks, 2011. URL <https://standards.ieee.org/getieee802/download/802.15.4-2011.pdf>.
- [21] Simone Gambini, Nathan Pletcher, and Jan M Rabaey. Sensitivity Analysis for AM Detectors. Technical report, 2008.
- [22] Nathan E Roberts and David D Wentzloff. 915MHz Ultra Low Power Receiver using Sub-Vt Active Rectifiers. In *Subthreshold Microelectronics Conference*, pages 1–3, 2012.
- [23] Lee Jae-Seung, Kim Joo-Myoung, Lee Jae-Sup, Han Seok-Kyun, and Lee Sang-Gug. 227pJ/b -83dBm 2.4GHz Multi-Channel OOK Receiver Adopting Receiver-Based FLL. In *International Solid-State Circuits Conference*, pages 234–236, 2015.
- [24] S. Ramasamy, B. Venkataramani, S.M. Stalin, and K. Venkatachalam. Tunable Band pass Gm-C Filter with Switched Transconductance Cells. In *TENCON*, number 2, pages 1–5, 2008.
- [25] Lucia Acosta, Mariano Jimenez, Ramon G Carvajal, Antonio J Lopez-Martin, and Jaime Ramirez-Angulo. Highly Linear Tunable CMOS Gm-C Low-Pass Filter. *Transactions on Circuits and Systems*, (10):2145–2158, 2009.
- [26] Armin Tajalli and Yusuf Leblebici. Low-Power and Widely Tunable Linearized Biquadratic Low-Pass Transconductor-C Filter. *Transactions on Circuits and Systems*, (3):159–163, 2011.
- [27] Hussain A Alzaher. A CMOS Highly Linear Digitally Programmable Active-RC Design Approach. *Transactions on Circuits and Systems*, (11):2636–2646, 2011.
- [28] Shinichi Hori, Tadashi Maeda, Hitoshi Yano, Noriaki Matsuno, Keiichi Numata, Nobuhide Yoshida, Yuji Takahashi, Tomoyuki Yamase, Robert Walkington, and Hikaru Hida. A Widely Tunable CMOS Gm-C Filter with a Negative Source Degeneration Resistor Transconductor. In *European Signal Processing Conference*, pages 449–452, 2003.
- [29] Yannis Tsvividis. Integrated Continuous-Time Filter Design - An Overview. *Journal of Solid-State Circuits*, (3):166–176, 1994.
- [30] Jhin-fang Huang, Jiun-yu Wen, and Yan-cheng Lai. Chip Design of an 8 MHz CMOS Switched-Capacitor Low-Pass Filter for Signal Receiver Applications. In *International Congress on Image and Signal Processing*, number 1, pages 4037–4041, 2010.

- [31] Jacqueline S Pereira and Antonio Petraglia. Optimum Design and Implementation of IIR SC Filters Using Small-Order FIR Cells. *Transactions on Circuits and Systems*, (8):529–538, 2002.
- [32] Jhin-fang Huang, Jiun-yu Wen, and Yen-jung Lin. Chip Design of A 10-MHz Switched Capacitor Low-Pass Filter for Wireless Application. In *International Conference on Wireless Communications and Signal Processing*, pages 1–5, 2014.
- [33] W Aloisi, G Giustolisi, and G Palumbo. Exploiting the High-Frequency Performance of Low-Voltage Low-Power SC Filters. *Transactions on Circuits and Systems*, (2):77–84, 2004.
- [34] Un-Ku Moon. CMOS High-Frequency Switched-Capacitor Filters for Telecommunication Applications. *Journal of Solid-State Circuits*, (2):212–220, 2000.
- [35] Yannis Tsvividis. Event-Driven Data Acquisition and Digital Signal Processing—A Tutorial. *Transactions on Circuits and Systems*, (8):577–581, 2010.
- [36] Yannis Tsvividis. Event-Driven, Continuous-time ADCs and DSPs for Adapting Power Dissipation to Signal Activity. In *International Symposium on Circuits and Systems*, pages 3581–3584, 2010.
- [37] Mariya Kurchuk, Colin Weltin-wu, Dominique Morche, and Yannis Tsvividis. Event-Driven GHz-Range Continuous-Time Digital Signal Processor With Activity-Dependent Power Dissipation. *Journal of Solid-State Circuits*, (9):2164–2173, 2012.
- [38] Xiongchuan Huang, Guido Dolmans, Harmke De Groot, and John R Long. Noise and Sensitivity in RF Envelope Detection Receivers. *Transactions on Circuits and Systems*, (10):637–641, 2013.
- [39] Bob Schell and Yannis Tsvividis. Analysis of Continuous-Time Digital Signal Processors. In *International Symposium on Circuits and Systems*, number 5, pages 2232–2235, 2007.
- [40] Filipp Akopyan, Rajit Manohar, and Alyssa B Apsel. A Level-Crossing Flash Asynchronous Analog-to-Digital Converter. In *International Symposium on Asynchronous Circuits and Systems*, pages 11–22, 2006.
- [41] F Aeschlimann, E Allier, L Fesquet, and M Renaudin. Asynchronous FIR Filters: Towards a New Digital Processing Chain. In *International Symposium on Asynchronous Circuits and Systems*, pages 198–206, 2004.
- [42] H. Inose, T. Aoki, and K. Watanabe. Asynchronous Delta-Modulation System. *Electronics Letters*, (3):95–96, 1966.

- [43] Bob Schell and Yannis Tsividis. A Clockless ADC/DSP/DAC System with Activity-Dependent Power Dissipation and No Aliasing. In *International Solid-State Circuits Conference*, pages 550–552, 2008.
- [44] Wei Tang, Ahmad Osman, Dongsoo Kim, Brian Goldstein, Chenxi Huang, Berin Martini, Vincent A Pieribone, and Eugenio Culurciello. Continuous Time Level Crossing Sampling ADC for Bio-Potential Recording Systems. *Transactions on Circuits and Systems*, (6):1407–1418, 2013.
- [45] Colin Weltin-wu and Yannis Tsividis. An Event-driven Clockless Level-Crossing ADC With Signal-Dependent Adaptive Resolution. *Journal of Solid-State Circuits*, (9):2180–2190, 2013.
- [46] Yongjia Li and Wouter A Serdijn. A Continuous-Time Level-Crossing ADC with 1-Bit DAC and 3-Input Comparator. In *International Symposium on Circuits and Systems*, pages 1311–1314, 2012.
- [47] Yongjia Li, Duan Zhao, Marijn N Van Dongen, and Wouter A Serdijn. A 0.5V Signal-Specific Continuous-Time Level-Crossing ADC with Charge Sharing. In *Biomedical Circuits and Systems Conference*, pages 381–384, 2011.
- [48] Yongjia Li, Duan Zhao, and Wouter A Serdijn. A 0.8V 8-Bit Low-Power Asynchronous Level-Crossing ADC with Programmable Comparison Windows. In *Biomedical Circuits and Systems Conference*, pages 138–141, 2013.
- [49] Bob Schell and Yannis Tsividis. Analysis and simulation of Continuous-Time Digital Signal Processors. *Signal Processing*, (10):2013–2026, 2013. ISSN 0165-1684.
- [50] Y Tsividis. Continuous-Time Digital Signal Processing. *Electronics Letters*, (21): 1–2, 2003.
- [51] Karen Guan and Andrew C Singer. Opportunistic Sampling by Level-Crossing. In *International Conference on Acoustics, Speech and Signal Processing*, pages 1513–1516, 2007.
- [52] Marek Miskowicz. Analytical Approximation of the Uniform Magnitude-Driven Sampling Effectiveness. In *International Symposium on Industrial Electronics*, pages 407–410, 2004.
- [53] Marek Miskowicz. Efficiency of Level-Crossing Sampling for Bandlimited Gaussian Random Processes. In *International Workshop on Factory Communication Systems*, pages 137–142, 2006.

- [54] Saeed Mian Qaisar, Laurent Fesquet, and Marc Renaudin. Spectral Analysis of a Signal Driven Sampling Scheme. In *European Signal Processing Conference*, number Eusipco, pages 1–5, 2006.
- [55] Hadi Alasti. Non-Uniform-Level Crossing Sampling for Efficient Sensing of Temporally Sparse Signals. *IET Wireless Sensor Systems*, (1):27–34, 2014.
- [56] Mariya Kurchuk and Yannis Tsvividis. Signal-Dependent Variable-Resolution Quantization for Continuous-Time Digital Signal Processing. In *International Conference on Circuits and Systems*, pages 1109–1112, 2009.
- [57] Dieter Bruckmann, Tobias Feldengut, Bedrich Hosticka, Rainer Kokozinski, Karsten Konrad, and Nima Tavangaran. Optimization and Implementation of Continuous Time DSP-Systems by Using Granularity Reduction. In *International Symposium on Circuits and Systems*, pages 410–413, 2011.
- [58] Christos Vezyrtzis and Yannis Tsvividis. Processing of Signals Using Level-Crossing Sampling. In *International Symposium on Circuits and Systems*, number 1, pages 2293–2296, 2009.
- [59] E Allier, G Sicard, L Fesquet, and M Renaudin. A New Class of Asynchronous A/D Converters Based on Time Quantization. In *International Symposium on Asynchronous Circuits and Systems*, pages 196–205, 2003.
- [60] Mariya Kurchuk, Colin Weltin-wu, Dominique Morche, and Yannis Tsvividis. GHz-Range Continuous-Time Programmable Digital FIR with Power Dissipation that Automatically Adapts to Signal Activity. In *International Solid-State Circuits Conference*, number May 2010, pages 232–234, 2011.
- [61] Christos Vezyrtzis, Weiwei Jiang, Steven M Nowick, and Yannis Tsvividis. A Flexible, Event-Driven Digital Filter With Frequency Response Independent of Input Sample Rate. *Journal of Solid-State Circuits*, (10):2292–2304, 2014.
- [62] Bob Schell. *Continuous-Time Digital Signal Processors: Analysis and Implementation*. PhD thesis, 2008.
- [63] Bob Schell and Yannis Tsvividis. A Low Power Tunable Delay Element Suitable for Asynchronous Delays of Burst Information. *Journal of Solid-State Circuits*, (5):1227–1234, 2008.
- [64] Gyudong Kim, Min-kyu Kim, Byoung-soo Chang, and Wonchan Kim. A Low-Voltage, Low-Power CMOS Delay Element. *Journal of Solid-State Circuits*, (7):966–971, 1996.

- [65] Karsten Konrad, Dieter Bruckmann, Nima Tavangaran, Jidan Al-eryani, Rainer Kokozinski, and Thomas Werthwein. Delay Element Concept for Continuous Time Digital Signal Processing. In *International Symposium on Circuits and Systems*, pages 2775–2778, 2013.
- [66] Mariya Kurchuk and Yannis Tsvividis. Energy-Efficient Asynchronous Delay Element with Wide Controllability. In *International Symposium on Circuits and Systems*, pages 3837–3840, 2010.
- [67] A. Enis Cetin, Omer N. Gerek, and Yasemin Yardmici. Equiripple FIR Filter Design by the FFT Algorithm. *Singal Processing Magazine*, (2):60–64, 1997.
- [68] Bob Schell and Yannis Tsvividis. A Continuous-Time ADC/DSP/DAC System With No Clock and With Activity-Dependent Power Dissipation. *Journal of Solid-State Circuits*, (11):2472–2481, 2008.
- [69] W R Bennett, B M Oliver, J R Pierce, C E Shannon, S P Lloyd, Bell Laboratories, H Holzwarth, H Steinhaus, J Max, M D Paez, T H Glisson, P Fleischer, P F Panter, W Dite, B Smith, G M Roe, B Widrow, V R Algazi, H Gish, J N Pierce, T J Goblick, J L Holsinger, and Statistics Re. Oversampled, Linear Predictive and Noise-Shaping Coders of Order AM. *Transactions on Circuits and Systems*, (7):436–447, 1978.
- [70] V Majidzadeh, A Schmid, and Y Leblebici. Low-distortion Switched-Capacitor Event-Driven Analogue-to-Digital Converter. *Electronics Letters*, (20):1373 – 1374, 2010.
- [71] K Yoshioka, A Shikata, R Sekimoto, T Kuroda, and H Ishikuro. 0.0058mm^2 7.0 ENOB 24MS/s 17fJ/conv. Threshold Configuring SAR ADC with Source Voltage Shifting and Interpolation Technique. In *Symposium on VLSI Circuits*, pages 266–267, 2013.
- [72] J Tsai, Y Chen, M Shen, and P Huang. A 1-V, 8b, 40MS/s, $113\mu\text{W}$ Charge-Recycling SAR ADC with a $14\mu\text{W}$ Asynchronous Controller. In *Symposium on VLSI Circuits*, pages 264–265, 2011.
- [73] Geert Van Der Plas and Bob Verbruggen. A 150MS/s $133\mu\text{W}$ 7b ADC in 90nm digital CMOS Using a Comparator-Based Asynchronous Binary- Search sub-ADC. In *International Solid-State Circuits Conference*, pages 242–244, 2011.
- [74] Lane Brooks and Hae-seung Lee. A Zero-Crossing-Based 8b 200MS/s Pipelined ADC. In *International Solid-State Circuits Conference*, pages 460–462, 2007.
- [75] Colin Weltin-wu and Yannis Tsvividis. An Event-Driven, Alias-Free ADC with Signal-Dependent Resolution. In *Symposium on VLSI Circuits*, pages 28–29, 2012.

- [76] Mariya Kurchuk. *Signal Encoding and Digital Signal Processing in Continuous Time*. PhD thesis, 2011.
- [77] Brian Drost, Mrunmay Taleganokar, and Pavan Kumar Hanumolu. A 0.55V 61dB-SNR 67dB-SFDR 7MHz 4th-Order Butterworth Filter Using Ring-Oscillator-Based Integrators in 90nm CMOS. In *International Solid-State Circuits Conference*, number 9, pages 360–362, 2012.
- [78] Yaohua Zhao, Pui-in Mak, Rui P Martins, and Franco Maloberti. A 0.02mm² 59.2dB SFDR 4th-Order SC LPF With 0.5-to-10 MHz Bandwidth Scalability Exploiting a Recycling SC-Buffer Biquad. *Journal of Solid-State Circuits*, (9):1988–2001, 2015.
- [79] Joeri Lechevallier, Remko Struiksma, Hani Sherry, Andrea Cathelin, Eric Klumperink, and Bram Nauta. A Forward-Body-Bias Tuned 450MHz Gm-C 3rd-Order Low-Pass Filter in 28nm UTBB FD-SOI with > 1dBVp IIP3 over a 0.7-to-1V Supply. In *International Solid-State Circuits Conference*, pages 96–98, 2015.
- [80] Massoud Tohidian, Iman Madadi, and Robert Bogdan Staszewski. A 2mW 800MS/s 7th-Order Discrete-Time IIR Filter with 400kHz-to-30MHz BW and 100dB Stop-Band Rejection in 65nm CMOS. In *International Solid-State Circuits Conference*, pages 174–176, 2013.
- [81] Stefano D’Amico, Matteo Conta, and Andrea Baschiroto. A 4.1-mW 10-MHz Fourth-Order Source-Follower- Based Continuous-Time Filter With 79-dB DR. *Journal of Solid-State Circuits*, (12):2713–2719, 2006.
- [82] Antonio Liscidini, Alberto Pirola, and Rinaldo Castello. A 1.25mW 75dB-SFDR CT Filter with In-Band Noise Reduction. In *International Solid-State Circuits Conference*, pages 336–338, 2009.
- [83] Mostafa Savadi Oskooei, Nasser Masoumi, Mahmud Kamarei, and Henrik Sjolund. A CMOS 4.35-mW +22-dBm IIP3 Continuously Tunable Channel Select Filter for WLAN/WiMAX Receivers. *Journal of Solid-State Circuits*, (6):1382–1391, 2011.
- [84] Ik Joon Chang, Sang Phill Park, and Kaushik Roy. Exploring Asynchronous Design Techniques for Process-Tolerant and Energy-Efficient Subthreshold Operation. *Journal of Solid-State Circuits*, (2):401–410, 2010.
- [85] Wei-Hsiang Ma, Jerry C Kao, Visvesh S Sathe, and Marlos C Papaefthymiou. 187 MHz Subthreshold-Supply Charge-Recovery FIR. *Journal of Solid-State Circuits*, (4):793–803, 2010.
- [86] Huawen Jin and Edward K F Lee. A 350 MHz Programmable Analog FIR Filter Using Mixed-Signal Multiplier. In *European Solid States Circuit Conference*, pages 152–155, 2000.

- [87] Eoin O'hAinidh, Emmanuel Rouat, Sarah Verhaeren, Stephane Le Tual, and Christophe Garnier. A 3.2GHz-Sample-Rate 800MHz Bandwidth Highly Reconfigurable Analog FIR Filter in 45nm CMOS. In *International Solid-State Circuits Conference*, pages 90–92, 2010.
- [88] Myeong-Eun Hwang, Arijit Raychowdhury, Keejong Kim, and Kaushik Roy. 4OnW Process-Tolerant Subthreshold 8x8 FIR Filter in 13Onm. In *Symposium on VLSI Circuits*, pages 154–155, 2007.
- [89] Philippe Magarshack, Philippe Flatresse, and Giorgio Cesana. UTBB FD-SOI: a Process/Design symbiosis for breakthrough energy-efficiency. In *Design, Automation & Test in Europe Conference & Exhibition*, pages 952 – 957, 2013.

Personal Bibliography

- [PRMT15a] Sharvil Patil, Alin Ratiu, Dominique Morche, and Yannis Tsividis. A 3-10fJ/conv – step 0.0032mm² Error-Shaping Alias-Free Asynchronous ADC. In *Symposium on VLSI Circuits*, pages 160–161, 2015.
- [PRMT15b] Sharvil Patil, Alin Ratiu, Dominique Morche, and Yannis Tsividis. A 3 – 10fJ/conv – step Error-Shaping Alias-Free Continuous-Time ADC. *Journal of Solid-State Circuits (submitted for publication)*, 2015.
- [PTRM15] Sharvil Patil, Yannis Tsividis, Alin Ratiu, and Dominique Morche. An Asynchronous Flipping ADC, Patent submitted for US region in 2015.
- [RMA⁺13] Alin Ratiu, Dominique Morche, Bruno Allard, Xuefang Lin-Shi, and Jacques Verdier. Efficient Simulation of Continuous Time Digital Signal Processing RF Systems. In *International Conference on Sampling Theory and Applications*, pages 416–419, 2013.
- [RMA⁺14] Alin Ratiu, Dominique Morche, Bruno Allard, Xuefang Lin-Shi, and Jacques Verdier. Continuous Time ADCs in Interferer Resistant Wake-Up Radios. In *Conference on Ph.D. Research in Microelectronics and Electronics*, pages 1–4, 2014.
- [RMA⁺15] Alin Ratiu, Dominique Morche, Bruno Allard, Xuefang Lin-Shi, and Jacques Verdier. Continuous Time Digital Filtering Analog to Digital Converter. *Electronics Letters (submitted for publication)*, 2015.
- [RMPT15] Alin Ratiu, Dominique Morche, Sharvil Patil, and Yannis Tsividis. Continuous Time ADC and Filter, Patent submitted for the US and EU regions in 2015.
- [VRM⁺14a] Matthieu Verdy, Alin Ratiu, Dominique Morche, Emeric de Focauld, and Suzanne Lesecq. Cost-Driven Statistical Analysis for Selection of Alternative Measurements of Analog Circuits. In *International Conference on Electronics Circuits and Systems*, pages 104–107, 2014.

- [VRM⁺14b] Matthieu Verdy, Alin Ratiu, Dominique Morche, Emeric de Focauld, Suzanne Leseq, Jean-Pascal Mallet, and Cedric Mayor. Weight based Feature Selection Algorithm for Alternative Measurements. In *European Test Symposium Workshop*, 2014.

FOLIO ADMINISTRATIF

THÈSE SOUTENUE DEVANT L'INSTITUT NATIONAL
DES SCIENCES APPLIQUÉES DE LYON

NOM : Ratiu (avec précision du nom de jeune fille, le cas échéant)	DATE de SOUTENANCE : october 2, 2015
Prénoms : Alin	
TITRE : Continuous Time Signal Processing for Wake-Up Radios	
NATURE : Doctorat	Numéro d'ordre : 2015ISAL0078
Ecole doctorale : EEA	
Spécialité : Electronics	
RESUME :	
<p>Wake-Up Receivers (WU-RX) have been recently proposed as candidates to reduce the communication power budget of wireless networks. Existing architectures achieve very high sensitivities for power consumptions below $50\mu\text{W}$ but severely degrade their performance in the presence of out-of-band blockers. We attempt to tackle this problem by implementing an ultra low power, tunable, intermediate frequency filtering stage using Continuous Time Digital Signal Processing (CT-DSP). A CT-DSP chain can be divided into two parts: the CT-ADC and the CT-DSP itself; the specifications of these two blocks, given the context of this work, are also discussed.</p> <p>The CT-ADC is based on a novel, delta modulator-based architecture which achieves a very low power consumption; its maximum operation frequency was extended by the implementation of a very fast feedback loop. Moreover, the CT nature of the ADC means that it does not do any sampling in time, hence no anti-aliasing filter is required. The proposed ADC requires only $24\mu\text{W}$ to quantize signals in the [10MHz 50MHz] bandwidth for an SNR between 32dB and 42dB, resulting in a figure of merit of $3 - 10\text{fJ}/\text{conv} - \text{step}$, among the best reported for the selected frequency range.</p> <p>Finally, we present the architecture of the CT-DSP which is divided into two parts: a CT-IIR and a CT-FIR. The CT-IIR is implemented by placing a standard CT-FIR in feedback loop around the CT-ADC. If designed correctly, the feedback loop can now cancel out certain frequencies from the CT-ADC input (corresponding to those of out-of-band interferers) while boosting the power of the useful signal. The effective amplitude of the CT-ADC input is thus reduced, making it generate a smaller number of tokens, thereby reducing the power consumption of the subsequent CT-FIR by a proportional amount. The CT-DSP consumes around $100\mu\text{W}$ while achieving more than 40dB of out-of-band rejection; for a bandpass implementation, 2MHz passband can be shifted over the entire ADC bandwidth.</p>	
MOTS-CLÉS : continuous time digital signal processing, continuous time ADC, continuous time DSP, wake-up radio	
Laboratoire (s) de recherche : Ampère	
Directeur de thèse: Bruno Allard	
Président de jury : Dominique Dallet	
Composition du jury : Dominique Dallet, Pieter Harpe, Hassan Aboushady, Yannis Tsvividis, Bruno Allard, Stephane Le Tual, Dominique Morche	