CMOS Front-End Techniques for In-Band Full-Duplex Radio

Dirk-Jan van den Broek

CMOS FRONT-END TECHNIQUES FOR IN-BAND FULL-DUPLEX RADIO

Dirk-Jan van den Broek

Samenstelling promotiecommissie:

Voorzitter en secretaris:

Prof.dr. P.M.G. Apers

Universiteit Twente

Promotor:

Prof.dr.ir. B. Nauta

Universiteit Twente

Co-promotor:

Dr.ing. E.A.M. Klumperink

Universiteit Twente

Leden:

Prof.dr.ir. F.E. van Vliet

Universiteit Twente

Dr.ir. M.J. Bentum

Universiteit Twente

Prof.dr.ir. P.G.M. Baltus

Technische Universiteit Eindhoven

Prof.dr. A. Pärssinen

University of Oulu

This work was supported by the European Union
Seventh Framework Programme (FP7/2007-2013)
under Grant 316369 - Project DUPLO.

Centre for Telematics and Information Technology
P.O. Box 217, 7500 AE
Enschede, The Netherlands

ISSN: 1381-3617 (CTIT Ph.D. Thesis Series No. 17-428)

ISBN: 978-90-365-4323-1

DOI: <http://dx.doi.org/10.3990/1.9789036543231>

Copyright © 2017 by Dirk-Jan van den Broek, Enschede, The Netherlands
All rights reserved.

Typeset with L^AT_EX.

This thesis was printed by GVO drukkers & vormgevers B.V., Ede, The Netherlands.

CMOS FRONT-END TECHNIQUES FOR IN-BAND FULL-DUPLEX RADIO

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
prof. dr. T.T.M. Palstra,
volgens besluit van het College voor Promoties
in het openbaar te verdedigen
op vrijdag 12 mei 2017 om 16:45 uur

door

Johannes Dingeman Antonius van den Broek
geboren op 13 april 1986
te Roosendaal

Dit proefschrift is goedgekeurd door:

de promotor prof.dr.ir. B. Nauta

de assistent-promotor dr.ing. E.A.M. Klumperink

Abstract

In-band full-duplex wireless communication (FD), i.e. transmission and reception at the same time at the same frequency, is an emerging research topic, driven by the ever increasing demand for mobile data traffic in the crowded radio spectrum. Besides a theoretical doubling of the spectral efficiency, the inherent channel reciprocity is an attractive physical layer aspect. In higher network layers, additional advantages are being explored such as collision prevention, low latency, security and simplified frequency planning.

The main issue in FD is strong self-interference (SI) from the transmitter (TX) into the local receiver (RX). In typical links, the transmitted signal is in excess of 90dB above the system noise floor, necessitating over 90dB total SI-rejection to fully compete with half-duplex links. This rejection is achieved by combining isolation at the antenna interface with SI-cancellation in different domains.

Impairments in the radio components limit the amount of achievable SI-cancellation in the digital domain. As a result, a system-level study shows that the SI should be rejected by at least 40dB in the analog and RF domain. In most scenarios, this cannot be achieved by antenna isolation alone, and requires analog SI-cancellation paths that adapt to a varying antenna environment. In conclusion, SI-cancellation techniques across multiple domains have to be combined, and achieving competitive link budgets with CMOS integration potential, small form factor, limited complexity and low power consumption remains challenging.

This work studies the feasibility of FD using a custom designed CMOS front-end, as opposed to using commercially available components. Full implementation details and analysis are presented of a 65-nanometer mixer-first front-end with a vector modulator (VM) downmixer for SI-cancellation. Using the implemented front-end as a research vehicle, several transceiver impairments that may limit its full-duplex operation were experimentally investigated, such as distortion, phase noise, image rejection and transmitter impairments.

The receiver was found to have over 90dB linear link budget potential in a 16.25 MHz bandwidth, when combined with only 20dB worst-case antenna isolation, thanks to its 21.5dBm effective IIP3 for SI present at the receiver input. It offers up to 27dB

cancellation of up to -16.4dBm SI at the RX input without generating distortion above the noise floor. The co-integrated transmitter was found to have almost sufficient performance to support this link budget, requiring minor linearization by e.g. pre-distortion. Transmitter and SI-cancelling receiver operate jointly over a wide range of center frequencies from 0.15 to 3.5GHz. Implementing a complete point-to-point link with digital cancellation is beyond the scope of this thesis.

An improved second front-end was developed in 65nm CMOS to investigate the limits of the proposed architecture and to support future research at system level. It targets over 100dB linear link budget, supporting a significantly higher transmit power (from 3.6 to 16dBm average output) and reduced noise floor (from 12.3 to 8dB worst-case noise figure). This required re-design of the VM and RX mixer for an effective IIP3 in the range of 30-40dBm. A two-stage class-AB power amplifier with partial on-chip matching was implemented to support the higher transmit power over a wide bandwidth of 1.5-4GHz. Also, the baseband section was revised to provide higher gain and lower noise.

Transistor-level simulations of the full system showed the reduced NF, increased linearity and expected cancellation behavior. Preliminary measurements on the stand-alone power amplifier are also presented.

Overall, this thesis demonstrates that vector-modulator downmixers are a promising building block for SI-cancelling full-duplex front-ends. Owing to its high linearity, the proposed architecture can improve upon a low and varying isolation at antenna-level without generating SI-induced distortion above the noise floor. This enables highly integrated full-duplex radios that can compete with relaxed half-duplex links.

Samenvatting

Als gevolg van de alsmear toenemende vraag naar mobiel dataverkeer in een drukbezet radiospectrum, wordt er momenteel onderzoek gedaan naar tegelijkertijd zenden en ontvangen op dezelfde frequentie. Deze ‘in-band full-duplex’ draadloze communicatie (FD) kan theoretisch de spectrale efficiëntie verdubbelen. Daarnaast is het radiokanaal inherent reciproom, wat een interessant aspect is van de fysieke laag. In hogere netwerklagen worden nog andere voordelen onderzocht, zoals het voorkomen van conflicten bij radioverkeer, lage vertragingen, beter beveiligde verbindingen en eenvoudigere verdeling van de beschikbare frequenties.

De belangrijkste uitdaging in full-duplex is sterke zelf-interferentie (SI) van de zender (TX) naar de lokale ontvanger (RX). In typische radioverbindingen ligt het uitgezonden signaal meer dan 90dB boven de ruisvloer van de ontvanger, waardoor meer dan 90dB SI-onderdrukking nodig is om in alle gevallen te kunnen concurreren met half-duplex verbindingen. Deze onderdrukking wordt behaald door isolatie op antenne-niveau te combineren met het verder onderdrukken van SI in verschillende domeinen.

De hoeveelheid SI-onderdrukking die te behalen is in het digitale domein is beperkt door imperfecties in de radiocomponenten: systeem-onderzoek laat zien dat er daarom al minstens 40dB onderdrukking behaald moet worden in het analoge en RF domein. In de meeste gevallen kan dit niet met enkel antenne-isolatie worden bereikt, en zijn analoge SI-onderdrukkingspaden nodig die zich aanpassen aan de antenne-omgeving. Kortom, SI-onderdrukkingstechnieken in meerdere domeinen moeten worden gecombineerd, en het blijft een uitdaging om concurrerende link-budgetten te combineren met CMOS-integratie, beperkte omvang, eenvoud en een laag energieverbruik.

Dit proefschrift onderzoekt de haalbaarheid van FD door middel van een specifiek ontworpen CMOS front-end, waar voorheen meestal gebruik gemaakt werd van commercieel verkrijgbare componenten. Het beschrijft een volledige implementatie en analyse van een mixer-first front-end in 65-nanometer CMOS met een vector modulator (VM) downmixer voor SI-onderdrukking. Met dit front-end wordt experimenteel de invloed onderzocht van verschillende radio-imperfecties die voor full-duplex beperkend kunnen zijn, zoals vervorming, faseruis, spiegelfrequentie-onderdrukking en

zenderimperfecties.

De ontvanger bleek in staat tot een 90dB lineair link budget in een 16.25MHz bandbreedte, met slechts 20dB minimale antenne-isolatie, dankzij de 21.5dBm effectieve IIP3 voor SI aan de ingang van de ontvanger. De ontvanger biedt tot 27dB onderdrukking van maximaal -16.4dBm SI aan de ontvanger zonder vervorming boven de ruisvloer te produceren. De mee-geïntegreerde zender bleek nagenoeg voldoende te presteren om dit linkbudget te ondersteunen, maar moet beperkt worden gelineariseerd, bijvoorbeeld door middel van pre-distorsie. De zender en SI-onderdrukkende ontvanger werken samen over een groot bereik aan klokfrequenties van 0.15 tot 3.5GHz. Het implementeren van een volledige link tussen twee van deze radio's, met digitale SI-onderdrukking, wordt niet beschreven in dit proefschrift.

Er werd een tweede, verbeterde radio ontwikkeld in 65nm CMOS om de grenzen van deze architectuur op te zoeken en om verder onderzoek op systeem-niveau mogelijk te maken. Dit ontwerp richt zich op 100dB lineair linkbudget, door middel van significant hoger zendvermogen (van 3.6 naar 16dBm gemiddeld vermogen) en een verlaagde ruisvloer (van 12.3 naar 8dB in het slechtste geval). Dit vereiste een herontwerp van de vector modulator en ontvangstmixer voor effectieve IIP3's in de orde van 30-40dBm. Voor het hogere zendvermogen is een twee-traps klasse-AB versterker op chip aanwezig, met een gedeelte van de impedantieaanpassing op-chip voor een grote bandbreedte van 1.5-4GHz. Ook werd de basisband-sectie aangepast voor een hogere versterking en lagere ruis.

Transistor-niveau simulaties van het volledige systeem bevestigen de afgenomen ruisvloer, toegenomen lineariteit en de verwachte SI-onderdrukking. Voorlopige metingen aan de los gefabriceerde versterker zijn ook inbegrepen.

Al met al laat dit proefschrift zien dat vector-modulator downmixers een veelbelovend bouwblok zijn voor SI-onderdrukkende full-duplex radio's. Dankzij de hoge lineariteit kan de beschreven architectuur een beperkte isolatie op antenne-niveau verbeteren zonder daarbij vervorming boven de ruisvloer te produceren. Daardoor worden vergaand geïntegreerde full-duplex radios mogelijk die kunnen concurreren met half-duplex-scenario's waarin de gewenste prestaties gematigd zijn.

Contents

Abstract	i
Samenvatting	iii
1 Introduction	1
1.1 Motivation	2
1.2 Challenges	4
1.3 Research goals and thesis outline	5
2 Full-Duplex Radio System Considerations	7
2.1 Introduction	7
2.2 Transmitter impairments	11
2.2.1 Transmitter nonlinearity	12
2.2.2 DAC dynamic range	12
2.3 Receiver impairments	13
2.3.1 Receiver nonlinearity	13
2.3.2 ADC dynamic range	13
2.4 System level impairments	14
2.4.1 System clock phase noise	14
2.4.2 Multi-path reflection	15
2.5 Conclusions	15
3 Survey of existing full-duplex front-end techniques	17
3.1 Antenna interfaces with isolation	17
3.1.1 Single-antenna interfaces	18
3.1.2 Dual-antenna interfaces	18
3.1.3 Three-antenna interfaces	20
3.2 Adding cancellation	20
3.3 Survey of full-duplex front-ends with integration potential	22
3.3.1 Dual-polarized patch antenna	22

3.3.2	Dual-polarized patch antenna pair	23
3.3.3	Replica transmit chain canceller	23
3.3.4	Electrical balance duplexing	25
3.3.5	RF multi-tap delay	26
3.3.6	N-path filter based canceller	27
3.3.7	Mixer-first with baseband duplexer	28
3.3.8	Integrated N-path filter based circulator	28
3.4	Conclusions	29
4	An Integrated Front-End for In-Band Full-Duplex Radio	33
4.1	System Considerations and Proposed Architecture	33
4.2	Implementation of a Full-Duplex Front-End	36
4.2.1	SI-Cancelling Receiver	36
4.2.2	Transmitter	44
4.3	Measurement results	44
4.3.1	Cancellation	44
4.3.2	Noise	46
4.3.3	Linearity	47
4.3.4	Broadband performance	50
4.3.5	Transmitter	50
4.3.6	Phase noise	52
4.3.7	Image rejection	58
4.3.8	Comparison	58
4.3.9	Antenna experiments	61
4.4	Conclusions	62
5	An improved front-end for in-band full-duplex radio	65
5.1	Design	66
5.1.1	SI-cancelling receiver	67
5.1.2	Transmitter	76
5.1.3	Clock divider and LO tree	84
5.2	Simulation results and measurements	86
5.2.1	Stand-alone amplifier measurements	86
5.2.2	Complete system simulations	90
5.2.3	Results overview	92
5.3	Future work	92
5.4	Conclusion	93

6	Conclusions	95
6.1	Summary and conclusions	95
6.2	Original contributions	97
6.3	Recommendations	97
A	Errata to published material	99
	Dankwoord	101
	Bibliography	104
	List of publications	113

Chapter 1

Introduction

Human beings are not built to speak and listen at the same time. The same holds for the wireless communication devices that have become an important part of our everyday lives.

Regarding humans, it is probably related to their biologically limited capabilities of multitasking, along with the fact that one's own voice will often resound louder than that of the person we are trying to understand. Both factors cannot be easily changed in humans: We cannot simply add more multitasking capabilities to our brain, and we cannot easily tap into our nerves to do some kind of 'echo cancellation'. Luckily, for humans, there is no real need to condense our conversations in half the time.

This is different for wireless communication devices: We have full control over their 'brains', which nowadays consist of millions of transistors and are more capable than ever, thanks to the evolution of integrated circuit technology. In addition, we have access to all of their internal signals, since these systems are designed from the ground up. Furthermore, in wireless devices, there is a real need to condense the exchange of information, since the radio frequencies that can be used for communication are becoming crowded, as illustrated in figure 1.1.

So why has it become commonplace for wireless devices to alternate between transmitting and receiving, or to use different frequencies for transmission and reception? For this, we should consider that in typical radio communication, the transmitted signal can be 100 decibels louder than the weakest signal to be received, or ten billion times! In human terms, imagine trying to listen to a whisper while speaking at the sound pressure of a jet engine. Understandably, it is very challenging to build wireless receivers that can withstand these enormous powers right on top of the weak signal they try to detect, followed by 'echo cancellation' with accuracies of one-billionth. All to condense communication by at most a factor of two in time or frequency.

1. INTRODUCTION

Figure 1.1: Allocation of the radio spectrum in the US from 700MHz to 6GHz as of January 2016 (Source: NTIA).

This seems almost impossible, and at least not worth the added effort and complexity. Yet, trying to break this paradigm has become a popular research topic in recent years, and this thesis describes system considerations and circuit techniques that contribute to this field of ‘in-band full-duplex wireless communication’.

1.1 Motivation

Today’s evolution in wireless communication is characterized by a tremendous growth and dynamism in data traffic and user access [1, 2]. Figure 1.2 illustrates this growth for the recent history and the near future.

As the fourth generation of mobile communications (4G) is being deployed, plans are made for a unified, ubiquitous and fast wireless infrastructure that will connect consumers, industries, sensors, medical and automotive applications under the umbrella of ‘5G’ [3].

From a consumer perspective, 5G is projected to support up to 1Gbps data throughput, driven by e.g. the demand for mobile high-definition video streaming [1]. This should come along with broader coverage and ultra-high reliability. For industrial

Figure 1.2: Global mobile data traffic is expected to grow to 30.6 exabytes per month by 2020, an eightfold increase over 2015 (from [1]).

and automotive applications, and for a real-time experience in consumer applications, ultra-low network latency in the order of 1ms is required. In addition, the internet-of-everything (IoE) trend will result in a steep increase in the number of connected devices [1].

However, the amount of available radio spectrum to sustain this evolution is limited, especially in the lower frequency bands (<6 GHz) that exhibit good propagation characteristics (figure 1.1). Current ideas to relax these spectrum issues for 5G are to increase the localization and directivity of traffic through e.g. beamforming and massive MIMO, moving to millimeter-wave frequencies despite the challenging propagation characteristics, using unlicensed spectrum, and reducing the size of cells [1].

In addition to this, improved air interface techniques are required to increase the spectral efficiency. As traditional spectral efficiency improvement techniques, such as more complex modulation schemes and traditional MIMO get exhausted, there is renewed interest in full-duplex wireless communication [4].

In-band full-duplex (FD) wireless communication¹ is an emerging, unconventional scheme for radio communication links: Transmission and reception occur simultaneously at the same frequency, thus utilizing the same spectral resources in two directions at once. Traditionally, wireless links separate the transmit and receive signals on the physical level, most commonly in time (i.e. time-division-duplexing (TDD))

¹In-band full-duplex wireless is also referred to in literature as ‘Same-channel full-duplex’, ‘Division-free duplexing’, ‘True full-duplex’, ‘Simultaneous transmit and receive’, ‘Two-way wireless’, ‘Co-Channel Co-Frequency Full-Duplex’ et cetera.

or in frequency (i.e. frequency-division-duplexing (FDD)).²

In the physical layer, full-duplex obviously promises up to 2x spectral efficiency. Another interesting physical layer characteristic is channel reciprocity: since the same channel is used in both directions simultaneously, its fading and propagation characteristics will be equal.

In higher network layers, different attractive network concepts have been developed [5–8] which exploit FD capabilities in wireless communication radios to improve the capacity and user access. Especially in cellular, for access points and mesh networks, FD has the potential to mitigate some fundamental problems like hidden terminals, bandwidth degradation and network latency [9–11]. Further advantages are being explored such as collision prevention, low latency and security [4]. Additionally, FD can simplify frequency planning for mobile operators. The benefits at network level however rely on the availability of full-duplex radios, which is the topic of this thesis.

1.2 Challenges

The main issue in achieving FD wireless is strong in-band (same-channel) crosstalk from transmitter to receiver, referred to as self-interference (SI), see figure 1.3a. Recovering the (much weaker) desired signal from a remote transmitter results in two key challenges in full-duplex radios:

1. *Isolation*: to prevent the RF-signal generated by the local transmitter (TX) from leaking onto its own receiver (RX), where it causes self-interference.
2. *Cancellation*: to subtract any remaining self-interference from the RX path using knowledge of the TX signal.

Cancellation uses knowledge of the transmit signal from various points in the TX chain to subtract SI at various points in the RX chain (figure 1.3b). From this generic view, many types of SI-cancellation can be conceived, and to some extent freely combined, ranging from RF to analog BB, to digital BB and even cross-domain cancellation.

As the receiver is capturing a signal coming from a distant source, the self-interference is much stronger in power in the absence of isolation/cancellation. As it occupies the same frequency band, it interferes with the reception and may hinder the receiver sensitivity and therefore the link throughput. The transmitted signal in a typical link is in excess of 100 dB above its receiver noise floor, requiring isolation and cancellation to provide roughly 100 dB rejection of the self-interferer if no

²Note that FD is only a duplexing scheme, and does not imply a multiple-access scheme to handle multiple users in e.g. a cellular system: A base-station can still serve the users multiplexed in time (TDMA), using orthogonal frequencies (OFDMA) or orthogonal codes (CDMA) et cetera.

Figure 1.3: a) Generic view of a full-duplex link between a local and a remote node, subject to 3 types of self-interference: A) Electrical crosstalk between TX and RX, B) RF coupling due to limited antenna isolation and a varying antenna near-field, C) SI reflected by the environment. b) Generic view of SI-cancellation in a single FD node, from various points in the TX chain to various points in the RX chain.

performance compromise can be accepted. These numbers and associated hardware bottlenecks will be refined for different application scenarios in chapter 2.

1.3 Research goals and thesis outline

Prior to the start of this work, FD was a recently emerging research topic and mainly researched using commercially available radio IC's that were not optimized for FD operation [12,13]. Additionally, promising architectures had impractical aspects such as large form factor antenna or circuit solutions, unfeasible requirements on radio performance or incompatibility with CMOS integration.

As part of the EU FP7 project 'Full-duplex Radios for Local Access' (DUPLO, [14]), the primary goal of this work was to make a fundamental step in full-duplex research by designing a dedicated radio IC. This allows full control over the radio structure and access to its internal signals, to explore new self-interference cancellation concepts with unique benefits.

To this end, chapter 2 of this thesis investigates how various full-duplex system scenarios translate into requirements on various transceiver blocks. Given typical impairments that are present in these blocks, the chapter assesses which link budgets are achievable in a CMOS integrated full-duplex radio and derives a set of building block specifications.

Chapter 3 presents a survey of recent integrated full-duplex radios, classifies their architectures and lists and quantifies their merits and drawbacks.

Next, chapter 4 describes in detail how a new self-interference cancelling architecture was derived and how it relates to the proposed link budget and building block specifications of chapter 2. The chapter proceeds with transistor-level implementation of a full-duplex radio according to the new architecture, and its characterization.

Chapter 5 explores the practical limits of the proposed architecture. It gives a preliminary description of an improved full-duplex radio, based on the results of chapter 4.

Chapter 6 summarizes and concludes this thesis, with recommendations on directions for future research.

Chapter 2

Full-Duplex Radio System Considerations

This chapter investigates the requirements that the full-duplex communication paradigm puts on various building blocks of a wireless transceiver. When commencing this thesis work, a comprehensive study of these aspects of full-duplex was not readily available.

As transmitted signals leak strongly into the receiver, the impact of any noise and distortion induced by these signals throughout the transmit and receive chains is more severe than in half-duplex scenarios, resulting in new challenges and requirements, specific to FD. While some of these impairments, such as deterministic distortion, can be estimated and compensated in digital processing, the premise of this chapter is to realize a competitive noise- and distortion-free FD link budget.

This chapter consists of material previously published in [15], with minor updates to reflect the errata of appendix A.

2.1 Introduction

Figure 2.1 illustrates a full-duplex link between two wireless radio nodes, the ‘local’ and ‘remote’ one. Assuming a symmetrical link, the discussions that will follow apply to either radio node. Hence only the ‘local’ node with its downlink is fully depicted, as well as the associated self- interference via various cross-talk and reflection paths.

In practice, the self-interference consists of multiple components as the transmit signal is corrupted by different impairments, such as nonlinearity, phase- and quantization noise [13]. Some of these by-products are noisy, others are deterministic. This transmit signal, including its by-products, is coupled into the receiver through various

Figure 2.1: FD-link between a ‘Local’ and ‘Remote’ radio node. Self-interference enters the receiver through various paths: direct crosstalk (A), limited antenna isolation (B) and reflections through the environment (C).

paths indicated in Figure 1, e.g. direct crosstalk (A), TX-RX antenna leakage due to limited isolation (B), and reflections on nearby objects in the environment (C). To achieve a receiver sensitivity similar to conventional half- duplex radios is very challenging, as all self-interference components should be suppressed to below the receiver noise floor. This likely requires isolation in the antenna solution combined with cancellation in the transceiver.

Figure 2.2 shows the key ‘local node’ signals limiting the FD link budget when receiving a remote transmitter signal, assuming both nodes operate at equal average transmit power. The locally transmitted signal (Local TX) consists of a clean signal and its by-products due to transmitter impairments (half-circle). Isolation at RF (e.g. antenna isolation) will attenuate the self-inferences coupled to the receiver (Local RX), along with its transmitter impairments. Additional by-products will arise on this large signal due to receiver impairments (circle). Cancellation techniques are required to further reduce the self- interference and its by-products towards the receiver digital baseband (Local BB), ideally to below the noise floor.

It is well known from literature that high isolation is desired for FD-radios and some promising results have been achieved. What has not been explored much are the consequences of limited robustly achievable isolation, e.g. if compact low cost radios have to work under varying environmental conditions. The question then arises how transceiver requirements change as a function of the RF-isolation and link budget parameters, and whether a viable user scenario is still feasible. Table 2.1 shows some results of an analysis that will be detailed below. It analyzes several important transceiver requirements (**bold** fonts) as a function of several assumption (*italics*). Comparing the outcomes to typically feasible transceiver specifications taken from [16] as shown in the right side column, bottlenecks are identified and marked. These bot-

Figure 2.2: Relation between various in-band power levels in a full-duplex link budget. The combination of isolation and cancellation techniques suppress all self-interference components, preferably to below the receiver noise floor.

tlenecks can be resolved by increasing the amount of RF-isolation or by improvements in transceiver design. In the following, the equations needed for transceiver design are derived.

Starting from the link budget parameters bandwidth (BW), transmission power (P_{TX}) and receiver noise figure (NF), the in-band receiver noise floor (P_{noise}) is calculated as:

$$P_{noise}[\text{dBm}] = -174[\text{dBm/Hz}] + 10 * \log(BW[\text{Hz}]) + NF[\text{dB}] \quad (2.1)$$

In Table 2.1, typical numbers for a low-end, mid-end and high-end wireless link scenario are given, where the high-end specifications correspond to a commercial 54-Mbps WLAN link with 64-QAM OFDM [16]. The mid- and low-end scenarios have significantly relaxed transmit power and noise figure and a factor 2 lower bandwidth, which is deemed still viable for some shorter range links, e.g. for sensor networks.

Table 2.1 lists the outcome of equation 2.1 and also the resulting difference between the transmit power and the noise floor, indicating that 79 to 116 dB of isolation/cancellation is required to prevent sensitivity losses compared to a half-duplex link. In practical antenna solutions [9, 17], the effective isolation is limited to approx-

2. FULL-DUPLEX RADIO SYSTEM CONSIDERATIONS

Scenario difficulty RF-isolation (dB)	Eqn.	<i>Low-end</i>		<i>Mid-end</i>		<i>High-end</i>		<i>Feasible</i>
		<i>20</i>	<i>40</i>	<i>40</i>	<i>60</i>	<i>60</i>	<i>80</i>	
BW(MHz)		10	10	10	10	20	20	20
P_{TX} (dBm)		0	0	10	10	20	20	20
NF (dB)		25	25	15	15	5	5	5
P_{noise} (dBm)	2.1	-79	-79	-89	-89	-96	-96	-96
$P_{TX}-P_{noise}$ (dB)		79	79	99	99	116	116	116
Isolation (dB)		20	40	40	60	60	80	
$P_{SI} = P_{TX}-\text{Isolation}$ (dBm)		-20	-40	-30	-50	-40	-60	
EVM_{TX} (dB)	2.2	-59	-39	-59	-39	-56	-36	-40
$P_{IM3,TX}$ (dBc)	2.3	-59	-39	-59	-39	-56	-36	-40
$OIP3_{TX}$ (dBm)	2.4	30	20	40	30	48	38	40
$CP1dB_{TX} \approx OIP3_{TX}-10$ (dBm)		20	10	30	20	38	28	30
OBO_{TX} (dB)		20	10	20	10	18	8	10
DAC margin (dB)		15	15	20	20	25	25	25
DAC DR = $-EVM_{TX} + \text{margin}$ (dB)		74	54	79	59	81	61	65
DAC bits = DAC DR / 6.02		12	9	13	10	13	10	11
$P_{IM3,RX}$ (dBm)	=2.1	-79	-79	-89	-89	-96	-96	
$IIP3_{RX}$ (dBm)	2.5	10	-21	-1	-31	-12	-42	0
$CP1dB_{RX} \approx IIP3_{RX}-10$ (dBm)		0	-31	-11	-41	-22	-52	-10
ADC margin (dB)		20	20	25	25	30	30	30
ADC DR = $P_{SI}-P_{noise} + \text{margin}$ (dB)		79	59	84	64	86	66	70
ADC bits = ADC DR / 6.02		13	10	14	10	14	11	11
PN (dBc)	2.6	-62	-42	-62	-42	-59	-39	-40
PN (deg)	2.7	0.05	0.46	0.05	0.46	0.06	0.64	0.57

Table 2.1: Three FD link budget scenarios analyzed for variable RF-isolation. The right side column indicates typically achievable specifications (based on e.g. [16]), and requirements in excess of this are marked grey as ‘feasibility bottlenecks’.

imately 40 dB, also due to reflections from the environment. Therefore, an additional cancellation of 39 to 76 dB would be required to exceed the noise floor.

If the antenna solution is pushed to achieve more isolation, the self-interference path likely becomes more dominated by reflections from the environment, which can make the self-interference channel very frequency-selective [18]. Further cancellation of this frequency-selective self-interference can be addressed leveraging OFDM modulation and digital processing techniques to estimate the self-interference channel using pilot sequences and tones. Combining isolation and OFDM-based cancellation in the digital domain could theoretically form a full-duplex solution, as depicted in Figure 3. The transmit signal is fed through a digital estimate of the self-interference channel, and subtracted from the received signal in digital baseband. However, this solution puts stringent requirements on the transmitter and receiver, which will be illustrated in the next sections.

Figure 2.3: Self-interference suppression techniques are required to prevent the reduction of receiver sensitivity for FD e.g. isolation at the antenna(s) and cancellation in the digital baseband.

2.2 Transmitter impairments

Transmitter Error Vector Magnitude (EVM) is a commonly used metric to quantify the transmitter performance, which covers the main in-band impairments, albeit in a lumped fashion. In conventional half-duplex radios, the EVM toughest requirement results from the most complex modulation scheme to be used. E.g. for the high-end scenario, to demodulate 64-QAM OFDM for a 54-Mbps WLAN link, better than 5.6% (-25 dB) EVM is required [19].

In all three full-duplex scenarios mentioned before, an isolation of 40 dB and an EVM of -25 dB would result in a self-interference due to TX impairments well above the receiver noise floor, limiting the receiver sensitivity. To solve this problem, the TX EVM requirement should be better than:

$$\text{EVM}_{\text{TX}}[\text{dB}] \leq -(P_{\text{TX}}[\text{dBm}] - P_{\text{noise}}[\text{dBm}] - \text{Isolation}[\text{dB}]) \quad (2.2)$$

The resulting values are shown in Table 2.1, where each scenario is extended with an extra 20 dB of isolation to relax the EVM to a likely feasible value. Note that this EVM is no longer dictated by the modulation scheme, but by the FD constraint. Alternatively, extra analog cancellation can relax the EVM requirement by including the effects of the transmitter impairments in analog cancellation path [13].

EVM is a lumped term for errors that actually results from several causes, e.g. distortion, DAC dynamic range and phase noise. The related requirement will be modelled in the next paragraphs. Quadrature imbalance (I/Q amplitude and phase mismatch) is not modelled here as effective techniques exist to calibrate and suppress it [20].

2.2.1 Transmitter nonlinearity

In-band EVM due to transmitter nonlinearity is often mainly due to 3rd order intermodulation distortion. For low EVM a weakly nonlinear model with a 3rd-order output referred intercept point ($OIP3_{TX}$) can be adequate [21], while the 1 dB compression point provides an estimate of the upper limit of output power range.

Without distortion, the ‘clean self-interferer’ enters the receiver at a power level of P_{SI} . Ideally its distortion content should stay below P_{noise} . Therefore, the transmitted distortion products ($P_{IM3,TX}$) should satisfy (see also Table 2.1).

$$P_{IM3,TX}[dBc] \leq P_{SI}[dBm] - P_{noise}[dBm] \quad (2.3)$$

With a transmitted power P_{TX} , the required $OIP3$ at the transmitter thus equals

$$OIP3_{TX}[dBm] \geq P_{TX}[dBm] + \frac{P_{IM3,TX}[dBc]}{2} \quad (2.4)$$

Assuming a simple weakly nonlinear memory-less transmitter model, the 1 dB compression point will be approximately 10 dB below $OIP3_{TX}$ and the output back-off (OBO_{TX}) of the transmitter can be calculated (Table 2.1). In all scenarios, the transmitter has to be operated at a larger back-off than normally required for the corresponding link (e.g. in the high-end scenario, more than the normally required 6-8 dB for 802.11g WLAN [22]). This causes power-inefficient operation. A potential solution direction to reduce the required back-off is linearization by pre-distortion [23]. Alternatively, analog cancellation performed at RF includes the transmitter nonlinearities in the cancellation signal and therefore makes stronger distortion products acceptable [13].

2.2.2 DAC dynamic range

The main DAC requirement in a half-duplex transceiver is the dynamic range to transmit the most complex modulated signal with sufficient fidelity. E.g. for the high-end scenario, to transmit 64-QAM OFDM for a 54-Mbps WLAN link, about 8 bits are required in the DAC [22] resulting in about 50 dB dynamic range. Since the EVM requirement for this link is -25 dB, about 50 dB-25 dB=25 dB margin is taken in the DAC to make its EVM contribution (quantization noise and clipping noise due to high peak-to-average ratios) non-dominant. Table 2.1 lists typical DAC margins for this and the other scenarios.

In the full-duplex examples, more stringent EVM values are required in order to sufficiently cancel the self-interferer based on its digital representation, resulting in tougher DR requirements for the DAC. Assuming the same margins apply as in half-duplex, the resulting DAC dynamic range requirements are listed in Table 2.1. The

required resolution seems feasible [24], certainly for the ‘20 dB extra’ RF-isolation cases.

2.3 Receiver impairments

In a conventional half-duplex system, the receiver needs to capture the desired signal with sufficient fidelity to perform demodulation. In a full-duplex receiver, the self-interferer present at the receive port will usually be stronger than the desired receive signal (as illustrated in Figure 2). Hence, any by-products of capturing the self-interferer should not mask the underlying desired signal. The expected issues are nonlinearity in the receiver and limited ADC dynamic range.

2.3.1 Receiver nonlinearity

In the presence of a strong self-interferer, the receiver has to be sufficiently linear to prevent masking the desired signal with the receiver intermodulation products. If no analog cancellation is applied the required in-band input-referred 3rd-order intercept point (IIP3) [21] can be calculated, assuming a weak third-order nonlinearity. With a self-interferer power P_{SI} and a maximum strength of the 3rd-order distortion components $P_{IM3,RX}$, the required IIP3 equals:

$$IIP3_{RX}[\text{dBm}] \geq P_{SI}[\text{dBm}] + \frac{P_{SI}[\text{dBm}] - P_{IM3,RX}[\text{dBm}]}{2} \quad (2.5)$$

In Table 2.1 we see the resulting value for different scenarios. Now, for the low- and mid-end scenarios the IIP3 requirement is tougher but feasible given recent improvements achieved in in-band linearity [25]. However, care must be taken that the receiver can achieve the required IIP3 at the power level P_{SI} , which may require reduction of the front-end gain to avoid compression. Applying extra analog cancellation at RF may be useful, provided this does not add any (random) components that cannot be suppressed further in the digital domain [26].

2.3.2 ADC dynamic range

In order to perform self-interference cancellation in the digital domain, the ADC dynamic range has to cover the strong self-interferer, without masking the underlying desired signal with its quantization noise. Therefore, the demands on the ADC are tougher than in half-duplex systems.

In a half-duplex link budget, the ADC has to capture the signal at the most complex modulation scheme under fading conditions, plus several margins for gain control, quantization noise and peak-to-average ratio. Typical values for 64-QAM are about 30 dB for SNR and another 30 dB for various margins [16], resulting in 60 dB

ADC dynamic range. It is assumed here that the same margin applies to a full-duplex link.

The resulting ADC DR requirements and corresponding number of bits is listed in Table 2.1, and can be very tough. Again, the though requirement can be relaxed by means of analog cancellation, where a cancellation signal is subtracted before the ADC. Unlike analog cancellation used to relax RX linearity requirement, this need not necessarily be done at RF.

2.4 System level impairments

Two system level impairments are relevant to full-duplex radios: the system clock phase noise, and the multi-path components in the self-interference path.

2.4.1 System clock phase noise

Phase noise (PN), which is caused by the system clock generation and distribution system, degrades the SNR of the transmitted and the received signal. In case the transmitter and receiver operate on different system clocks, their PN will be uncorrelated. Then, TX and RX phase noise powers add and the combined noise limits the suppression that can be achieved by further cancellation at analog and digital baseband [26]. Assuming equal phase noise in the transmitter and receiver, the requirement of any single clock can be calculated:

$$PN_{RX}[dBc] = PN_{TX}[dBc] \leq EVM_{TX}[dB] - 3dB \quad (2.6)$$

For small values, this can be converted to degrees using:

$$PN[deg] \leq \arcsin(10^{\frac{PN[dBc]}{20}}) \quad (2.7)$$

Note that phase noise is integrated over the bandwidth BW here, and not presented as phase noise density. The calculated values in Table 2.1 are clearly very tough without sufficient isolation.

Integrated full-duplex radios could however share a common clock as the transmitter and receiver operate simultaneously at the same frequency. Ideally, their phase noise is fully correlated and the cancellation is no longer limited by phase noise. However, in practice there will be some delay between transmission and reception of the self-interferer due to reflections [27]. This delay reduces the correlation between the transmitted and received self-interference signal and hence degrades cancellation. As a result the phase noise cancellation degrades especially at higher offset frequencies. The amount of cancellation depends on several factors: the bandwidth of the wireless link, the phase noise profile of the PLL that is used and the reflection characteristics of the self-interference channel. This is an important topic for further research.

2.4.2 Multi-path reflection

In a realistic environment, transmitted signals may be reflected back to their own receiver through different paths which are each characterized by an attenuation and a delay. When leveraging OFDM modulation, the net effect is a subcarrier-dependent attenuation and phase shift, i.e., a frequency-selective channel. Multi-path reflections of the self-interferer can be cancelled by virtue of OFDM, but this requires a clean (noise-free) transmit signal. Realistically, the transmitter impairments add noise and distortion by-products which are also reflected via multiple paths. Distortion by-products can be reduced in the digital domain as was recently demonstrated in [13]. Noise by-products require, however, that an exact analog copy of the transmit signal is fed through circuitry that mimics the self-interference channel including the time delays. Implementing these delays at RF may provide robust cancellation [13], but leads to a bulky solution. Basically physical delay lines are needed with equal length as the delay path they model, divided by the ratio of the propagation velocities of associated propagation media (typically 2:1 for air compared to cables). Such a direct delay line implementation is not suitable for full CMOS integration, so alternative compact solutions are wanted.

2.5 Conclusions

The formulas derived in this chapter are very useful for spreadsheet calculation to assess overall feasibility and also as starting point for deriving sub-block specifications. The results in Table 2.1 indicate that even for a low-end very relaxed scenario 40 dB of RF-isolation is very much wanted. For the high-end case this increases to 60 or even 80 dB, which is extremely challenging. Table 2.1 indicates that not only phase noise but also transmitter linearity are very critical aspects, while receiver linearity becomes a bottleneck at low isolation values. Design innovations will hopefully move or remove some of the indicated bottlenecks, and the analysis above is believed to be very useful for future work on full-duplex transceiver design.

Chapter 3

Survey of existing full-duplex front-end techniques

As discussed in chapter 1 and further refined in chapter 2, self-interference is the main bottleneck when pursuing full-duplex wireless. Achieving competitive FD link budgets therefore necessitates isolation and cancellation. Chapter 2 showed that to achieve a competitive FD link budget with feasible component specifications under the assumption of only RF isolation + digital baseband cancellation, 40 to 60dB isolation would be required.

As any full-duplex radio starts with an antenna interface that provides some RF isolation, the next section briefly discusses common ways to obtain initial isolation at RF. It will intuitively explain that in most cases *adaptive* isolation or cancellation is a necessity to achieve the required 40-60dB SI-rejection in analog to relax the downstream radio components. Section 3.2 then intuitively shows that several cancellation paths and methods can be explored. Section 3.3 highlights some of the existing work that uses these cancellation methods to obtain (integrated) full-duplex radios. This chapter concludes with a high-level comparison of the treated full-duplex radios. Systematically studying the cancellation opportunities, combined with the known strengths of mixer-first receivers, led to an interesting new topology, which will be presented in chapter 4.

3.1 Antenna interfaces with isolation

This section briefly discusses some commonly used passive antenna interfaces that provide initial isolation for FD applications, their merits and drawbacks. The solutions are categorized by number of antennas. Limiting the number of antennas in FD

is beneficial for its competitiveness to half-duplex. Also, fewer antennas increases the potential compatibility of FD with MIMO systems.

3.1.1 Single-antenna interfaces

Passively interfacing a TX and RX simultaneously to a single antenna with low insertion losses on either side requires a non-reciprocal network. The microwave circulator is a commonly used component for this application.

Circulator

The circulator is a non-reciprocal microwave component with multiple ports, that couples energy from one port to the next in a circular fashion, while isolating in the reverse direction. Non-reciprocity can fundamentally only be achieved in a limited number of ways; in a circulator this is achieved by means of Faraday rotation using magnetic materials.

Circulator-based TX-RX isolation has the advantage of a single antenna port for transmit and receive, along with high linearity and low insertion losses from TX to antenna and from antenna to RX. Unfortunately, the presence of magnetic materials and the required physical size (in the order of the desired wavelength) make circulators bulky and expensive components.

Circulators are typically limited in their isolation by the matching of the antenna: TX energy reflected from the antenna due to imperfect matching cannot be distinguished from desired signals, and is coupled to the RX input. This typically limits the TX-RX isolation to 15-30dB [13, 28].

3.1.2 Dual-antenna interfaces

Using two antennas, the competitiveness of FD is reduced with respect to a single antenna, but this expense may be justified if FD is applied as an alternative mode of operation in devices that already possess dual antennas for e.g. MIMO or diversity applications. Commonly used methods to isolate two antennas are spatial separation, directivity, polarization and combinations of the three.

Antenna separation

Simple TX-RX antenna spacing of two omnidirectional antennas is used as a rudimentary way to obtain some isolation in full-duplex proof-of-concepts [29]. However, achieving the desired 40 to 60dB isolation requires a large antenna spacing [9]. Assuming a free space path loss model with 0dBi antenna gain at both ends and 2.45GHz operating frequency, figure 3.1 shows that a spacing of 1 metre would be required to

achieve 40dB isolation, already leading to an impractically large form factor for most applications.

Figure 3.1: Isolation between two omnidirectional antennas as a function of antenna separation, assuming a free-space path loss model at 2.45GHz operating frequency and 0dBi antenna gain at both ends.

Antenna directivity

In [30], a pair of directional antennas with 90-degree beam width is studied. Depending on the angle of placement and the spacing between the antennas, up to 45dB isolation was obtained, but in a reflective environment this quickly deteriorates to 36-37dB irrespective of the configuration and spacing. Similar results were obtained in [31]. It was found that two patch antennas mounted on two opposite sides of laptop computer frame with at least one plane orthogonal to each other, showed 35-40dB isolation.

Both sources agree that substantial isolation can be obtained by combining directional antennas with spacing, however this still results in a considerably large form factor and directional antennas are not desirable in many applications.

Antenna polarization

Transmitting and receiving in orthogonal polarizations can yield high isolation at antenna level. The orthogonal polarizations can be linear (horizontal - vertical) [32] or circular (left-handed - right-handed) [33]. Cross-polarization can be achieved in

pairs of antennas [30, 34] or in a single radiating aperture, such as a patch antenna with multiple feed points [32, 33].

Sophisticated antenna design showed 50-60dB isolation in [32], with good immunity to nearby reflective objects, making it promising for a full-duplex transceiver with only antenna isolation and digital baseband cancellation. However, cross-polarization isolation usually comes with directional antennas such as patches, and requires sophisticated antenna design for a fixed operating frequency, limiting its applications.

Combining two-antenna techniques

An interesting study of the effectiveness of antenna separation, combined with directivity and cross-polarization, was performed in [30]. Combining techniques, over 70dB isolation was obtained in an anechoic chamber, but reflections of a typical indoor environment again deteriorated the isolation to around 37-46dB. In [12], dipole antennas are placed orthogonally around a laptop computer frame, which reportedly achieves up to 60-70dB of isolation in office environments, however the environment and robustness are not extensively discussed.

3.1.3 Three-antenna interfaces

Some early full-duplex research used three-antenna solutions, mentioned here for the sake of completeness.

TX nulling at the RX antenna

Antenna isolation can also be achieved using two TX antennas, and placing an RX antenna in a near-field null [9]. However, this requires three antennas, making it less competitive in terms of hardware compared to employing these antennas for MIMO. Additionally, such structures were found to have poor far-field radiation patterns with unwanted nulls [32]. Furthermore, the TX nulling is only perfect at a specific operating frequency, and deteriorates quickly when moving away from this frequency.

3.2 Adding cancellation

The previous section showed that, while the desired 40-60dB isolation can be achieved in specific cases by careful antenna interface design, this is mainly feasible in environments with a static near-field environment. In cases where the near-field varies, such as hand-held devices, isolation will be limited and adaptive cancellation paths across multiple domains are a necessity.

Figure 3.2 illustrates that, for a simplified direct conversion transceiver, SI-cancelling front-end architectures can be conceptually classified according to the following criteria:

- Location of the TX tap:
 - RF at PA input
 - RF at PA output
 - Analog baseband
 - Digital baseband
- Location of the RX injection point:
 - RF at LNA input
 - RF at LNA output
 - Analog baseband
 - Digital baseband
- Complexity of channel model:
 - Phase shift + attenuation
 - Group delay
 - True time delay

Figure 3.2: Generic view of SI-cancellation indicating the trade-off in moving TX tap and RX injection point towards the antenna.

This leads to at least $4 \times 4 \times 3 = 48$ conceptual cancellation opportunities in a simplified direct conversion receiver. Additionally, several of these cancellation paths can be freely combined, further multiplying the possibilities, although not every combination leads to a sensible architecture.

The general trend (figure 3.2) is that tapping the TX signal closer to the antenna includes the impairments of more TX blocks in the cancellation signal, relaxing the accuracy requirements on these TX components given a desired link budget scenario. Similarly, injecting the cancellation signal closer to the antenna shields more RX blocks from the large SI, relaxing their dynamic range requirements.

Furthermore, the cancellation path can mimic not only the phase shift and attenuation of the SI-path, but also its delay, in order to cancel ambient reflections and other frequency-selective effects. Such delay can be implemented as true time delay [13], but since wireless communication systems typically use narrowband modulated carriers, group delay can also be sufficient to significantly extend the cancellation bandwidth compared to a phase/attenuation based canceller [35].

In conclusion, further relaxing the dynamic range requirements of the FD transceiver components generally requires more analog complexity in the cancellation path(s). This insight is useful to assess the full-duplex techniques presented next. Chapter 4 will present a new full-duplex architecture based on these insights. The remainder of this chapter discusses several state-of-the-art architectures in full-duplex research.

3.3 Survey of full-duplex front-ends with integration potential

This section lists several state-of-the-art architectures that demonstrate FD wireless capabilities and assesses their performance and integration potential. Schematic overviews help to classify the architectures along the lines of figure 3.2. Finally, a high-level comparison is made between the architectures in table 3.1.

3.3.1 Dual-polarized patch antenna

For specific applications, high isolation can be obtained by sophisticated antenna design. One way uses orthogonal linear polarization modes of a patch antenna for TX and RX, also referred to as cross-polarization isolation [36], schematically depicted in figure 3.3. This approach has a single radiating aperture with reasonably small form factor (60x60mm for 2.5GHz operation in [36]). It achieves better than 49dB isolation over a specific 20MHz channel in the ISM band, and 42dB isolation over the full 80MHz BW. However, in a varying antenna near-field, the isolation of the antenna can be degraded by over 10dB, which is why the patch antenna was complemented by a phase / attenuation based canceller at RF in [32].

A demonstrator of this system achieved over 60dB analog cancellation across a 20MHz channel at 2.4GHz, and up to 85-90dB including digital cancellation [37]. Functional full-duplex links of up to 16m distance were demonstrated [32].

Figure 3.3: Schematic representation of a full-duplex transceiver using a dual-polarized patch antenna, RF phase / attenuation based cancellation and digital baseband cancellation.

While the high directivity of the antenna may fit certain applications, such as static backhaul links, this solution requires alignment between antennas for optimum performance. Along with its dimensions, this solution is not suitable for hand-held applications. Also, it is highly tailored to a specific operating frequency.

3.3.2 Dual-polarized patch antenna pair

The effect of near-field variations on antenna isolation can be adaptively compensated using tuneable parasitic coupling paths between antennas [38].

An interesting practical implementation of this principle was presented in [39], for 60GHz full-duplex applications. The authors use a pair of patch antennas with orthogonal polarizations for TX and RX, als illustrated in figure 3.4. A parasitic coupling path inevitably exists from the TX antenna to the TX polarization in the RX antenna, and ultimately into the RX. By connecting a tuneable reflective load to the TX polarization in the RX antenna, this path can be effectively used to cancel self-interference. In [39], the reflective load consists of a tuneable parallel RLC-network that can approximate the frequency-selectivity of the self-interference path to further increase the bandwidth.

In [39], the authors demonstrate 50dB cancellation at 60 GHz, across a wide 13.5% fractional bandwidth.

3.3.3 Replica transmit chain canceller

As shown in figure 3.5, a replica TX chain can be used to regenerate the SI in the digital BB and cancel it at RF, combined with additional digital cancellation [40]. This cross-domain cancellation path has the advantage of a digital channel model, making it easy to implement time delay, for wideband cancellation of frequency-selective SI

Figure 3.4: Schematic representation of a full-duplex transceiver using a pair of polarized patch antennas with a tuneable parasitic coupling path, RF phase / attenuation based cancellation and digital baseband cancellation.

right at the antenna. This protects all RX blocks from the strong SI and part of its reflections, which made it a popular architecture in early FD research [40] and in demonstrators [29].

However, its ultimate cancellation performance is limited by uncorrelated noise and distortion sources between the two TX chains, and by phase noise if separate LO signals are used for the TX chains [26]. For the commercial radios used in [40], phase noise limits the combined effect of mixed-signal and digital cancellation to 35dB. Combined with a limited antenna isolation, this results in a limited link budget. The authors of [40] claim 85dB total cancellation, but this is achieved by virtue of a large (60-70dB) antenna isolation, which is questionable in light of other works and not extensively discussed.

Figure 3.5: Schematic representation of a full-duplex transceiver using a replica TX chain with digital channel model and digital baseband cancellation.

3.3.4 Electrical balance duplexing

Another approach to obtain adaptive TX-RX isolation is electrical balance duplexing [41, 42], as depicted in figure 3.6. The power of the PA is fed to the common-mode input of a 180-degree hybrid coupler. The antenna is attached to one common-mode port of the hybrid coupler and a tuneable load is attached to the other common-mode port. The load is tuned to mimic the antenna impedance in the band of interest, such that ideally the same signal is present on the antenna and the tuneable load. As a result, the difference port of the hybrid remains quiet unless a desired signal is presented to the antenna.

Figure 3.6: Schematic representation of a full-duplex transceiver using an electrical-balance duplexer and digital baseband cancellation.

Like the microwave circulator, this has the benefit of a single-antenna interface. Unlike the circulator, this approach lends itself to CMOS integration if the hybrid coupler is implemented as a transformer [41]. However, since the full transmit power is presented to the tuneable load, and any distortion generated by the load is presented as a differential signal on the output, this requires extreme linearity of the tuneable load, which was not readily achieved in [41]. In that work, the IIP3 was limited to 20dBm, allowing only -16dBm TX power to prevent running into distortion products during digital self-interference cancellation. This resulted in a 90dB theoretical link budget, with 70dB link budget demonstrated in practice at 0dBm TX power [37]. Remarkably, +70dBm IIP3 was later demonstrated for a tunable load in SOI CMOS [43], albeit for frequency division duplexing applications.

As shown in figure 3.7, electrical balance duplexing can also be combined with replica TX chain cancellation [42]. In their demonstrator, the authors used a laboratory-type hybrid coupler and impedance tuner for the duplexer, and instrument-grade transceivers with a low EVM and shared clocks for all up- and downconverters. This allowed 44dB rejection of SI by the electrical balance circuit and an additional 40dB of SI-cancellation by the replica TX, for a total link budget of 85dB over a 20MHz BW.

Figure 3.7: Schematic representation of a full-duplex transceiver using an electrical-balance duplexer and a replica TX chain for additional cancellation.

The drawback of any form of electrical balance duplexing is that unlike the circulator, a hybrid coupler is a reciprocal component. As a result, a strong antenna-RX coupling, required for low noise figure, results in a weak TX-antenna coupling and thus insertion losses in the transmitter. Similarly, strong TX-antenna coupling results in RX signal being coupled into the PA instead of the LNA. This can be partly justified by the absence of further duplexing filters which would also contribute insertion losses [41].

3.3.5 RF multi-tap delay

Direct crosstalk as well as part of the reflected SI can be cancelled in the RF domain using an analog multi-tap filter at RF, combined with digital cancellation [13]. Such a multi-tap filter takes power from the TX, splits it into many paths of varying electrical length to create many delayed versions of the TX signals, applies weights to these delayed signals and combines them for subtraction at the LNA input. Effectively this yields a FIR filter in the analog RF domain.

The authors of [13] combine this analog multi-tap filter with a circulator that provides about 15dB of initial isolation. To avoid tapping too much power from the TX for cancellation, the number of delayed TX copies is limited and the delays need to be chosen carefully to cancel the most relevant reflected SI components. In this case, the authors chose 8 delay taps to span the reflections of imperfect antenna matching, and 8 additional taps to span the nearby reflections.

Combined with non-linear digital baseband cancellation this architecture achieves the largest amount of total SI-cancellation reported of 110dB across an 80MHz bandwidth, making it competitive with high-end (802.11-style) half-duplex links [13].

A significant practical drawback of this architecture is the nano-second scale analog time delays required in the analog multi-tap filter, which physically require large area: The system consists of a 10x10cm PCB which does not permit CMOS integration.

Figure 3.8: Schematic representation of a full-duplex transceiver using a circulator, an analog multi-tap delay filter at RF and digital baseband cancellation.

3.3.6 N-path filter based canceller

Recently, the observation that time delay can be approximated locally in the band of interest as group delay, has led to a group-delay based canceller for wideband SI-cancellation: In [35], group delay was implemented in the form of two 2-port N-path filters with a tuneable gain, phase, quality factor and center frequency. Since an N-path filter can exhibit a high quality factor resonance at the center frequency, its transfer function shows a steep slope in phase around that frequency, which implies a large group delay.

This approach has the potential for full CMOS integration, and can add significant cancellation on top of antenna interfaces that already show high isolation and frequency selectivity. The authors add 20dB cancellation across 25MHz bandwidth to an antenna pair that itself has already 34dB isolation. Drawbacks are significant silicon area, power consumption and (tuning) complexity.

Figure 3.9: Schematic representation of a full-duplex transceiver using an n-path filter based canceller at RF and digital baseband cancellation.

3.3.7 Mixer-first with baseband duplexer

An interesting integrated approach is to use the bidirectional transparency of a mixer in a mixer-first front-end, and perform duplexing in the baseband [44], as depicted in figure 3.10.

To achieve baseband duplexing, the authors use modified differential noise-cancelling LNA's that intrinsically copy a transmit signal to their antenna port, while rejecting it in their output [44]. Placing the LNA's in the baseband allows complex signal processing in the form of I/Q cross-coupling paths to tune their SI-rejection.

Unlike a passive circulator, this architecture provides a single-port antenna solution suitable for integration and tuneable across a wide range of operating frequencies. However, the duplexing LNA's have limited capability to work with high TX powers and the TX performance will be limited by the loss of the mixers.

Across a 2MHz bandwidth, the authors demonstrate 30dB analog isolation, and assume 50dB digital cancellation, for a total link budget of 80dB [44].

Figure 3.10: Schematic representation of a full-duplex transceiver using a mixer-first front-end with tuneable baseband duplexing LNA's and digital baseband cancellation.

3.3.8 Integrated N-path filter based circulator

A recent approach breaks reciprocity for full-duplex applications using time-variance, in the form of an N-path-filter-based circulator [45,46]. A two-port N-path filter with a 90-degree phase shift between its two clocks provides a +90-degree phase shift in one propagation direction and -90-degree in the other, which is non-reciprocal behavior. When a 3/4-wavelength transmission line is wrapped around such a filter, and ports are attached to the transmission line with 1/4-wavelength spacing, circulator behavior can be obtained. Furthermore, if the transmission line is constructed out of lumped elements, almost the entire circulator can be integrated on-chip, unlike a microwave circulator. The structure is represented in figure 3.11. Interestingly, the authors

combine the circulator with phase / attenuation based cancellation in the analog baseband before proceeding to digital cancellation.

The authors demonstrate 20dB isolation from the integrated circulator across a 12MHz baseband bandwidth. By strategically locating the transmit, antenna and receive ports along the transmission line, the signal swings across the N-path filter can be limited, however linearity remains limited in this architecture, which will be detailed later in 4.3.8. In [46], the circulator is combined with analog baseband cancellation and rudimentary non-linear digital cancellation, 85dB total SI-rejection is demonstrated.

Figure 3.11: Schematic representation of a full-duplex transceiver using an N-path-filter-based circulator, analog baseband cancellation and digital cancellation.

3.4 Conclusions

Table 3.1 lists the cancellation performance, integration potential and other high-level characteristics of the main full-duplex wireless demonstrators presented to date and discussed in this chapter. As predicted by chapter 2, high amounts of overall cancellation are typically achieved by pushing the isolation and cancellation in the analog domain, combined with digital cancellation that takes varying degrees of TX/RX distortion into account. Practical FD link budgets in excess of 80dB across competitive bandwidths of over 10MHz seem within reach. However, each architecture also has its drawbacks such as form factor, complexity, power consumption or complex nonlinear digital cancellation.

The remainder of this thesis presents a new FD front-end architecture that compares favourably to other works in a number of these aspects, such as its high integration potential, limited complexity, modest power consumption and large linear link budget potential. At the end of chapter 4, more detailed comparisons will be made with the works from table 3.1 that have also demonstrated CMOS integration

potential to date.

Table 3.1: High-level comparison of relevant full-duplex demonstrators published at the time of writing.

Ref.	[42]	[37, 41]	[13]	[37, 41]	[35]	[44]	[40]	[46]	[34]
Operating principle	Hybrid-based electrical balance duplexing + replica TX chain	Transformer-based electrical balance duplexing + digital SIC	Circulator + multi-tap delay + nonlinear digital SIC	Dual-polarized patch antenna + RF canceller + digital SIC	N-path filter based canceller	Mixer-first with BB noise-cancelling duplexing LNA's + digital SIC	Replica TX chain + digital SIC	N-path filter based circulator + analog BB SIC + digital SIC	Dual-polarized patch antennas + reflective coupling path + digital SIC
Operating frequency	0.89-1.89GHz	2-2.5GHz	2.4-2.5GHz	2.4-2.8GHz	0.8-1.4 GHz	0.1-1.5GHz	2.4-2.5GHz	0.6-0.8GHz	57-66GHz
Form factor potential	Chip + hybrid + single antenna	Chip + single antenna	Chip + Circulator + 10x10cm PCB	Chip + 6x6cm antenna	Chip + antenna pair	Chip + single antenna	Chip + antenna pair	Chip + 3 inductors + single antenna	Chip + patch antenna pair
Full CMOS integration: Shown / potential	No/No	Yes/Yes	No/No	No/Yes	Yes/Yes	Yes/Yes	No/Yes	Yes/Yes	Yes/Yes
Single-antenna	Yes	Yes	Yes	Yes, two-port	No	Yes	No	Yes	No
Antenna iso	0 (single)	0 (single)	0 (single)	45-55dB across 20MHz	34dB across 25MHz	0 (single)	60-70dB across 20MHz	0 (single)	40dB across 1GHz
Analog isolation + cancellation	44dB across 20MHz	50dB across 6MHz	60dB across 80MHz	60dB across 20MHz	56dB across 25MHz	33dB across 300kHz	N/A	42dB across 12MHz	70dB across 1GHz
Total cancellation (analog + digital)	80dB across 20MHz	67dB across 20MHz	110dB across 80MHz	85-90dB across 20MHz	N/R	N/R	85dB across 20MHz	85dB across 12MHz	N/R (but shown)

Chapter 4

An Integrated Front-End for In-Band Full-Duplex Radio

This chapter describes the design and implementation of a direct-conversion CMOS front-end for small form factor full-duplex systems. It consists of material previously published in [36, 47–49], which was re-structured for readability in thesis form. The text was updated to reflect the errata of appendix A and to consistently use the latest set of measurements from [49] unless otherwise stated.

The chapter is structured as follows: First, we briefly review system considerations for FD and show how the proposed front-end topology emerges. Next, section 4.2 describes the implementation of the prototype SI-cancelling front-end. Section 4.3 describes the measured performance and relates it to FD link capabilities. Section 4.4 concludes this chapter.

4.1 System Considerations and Proposed Architecture

The SI-cancelling direct-conversion front-end developed here aims to bring full-duplex to low-power, short-range communication devices. For this purpose, a TX power of 0dBm is pursued, a bandwidth of 16.25MHz (the active bandwidth of WLAN) and a 10dB RX noise figure. This results in an RX noise floor of roughly -90 dBm. Thus, in order not to degrade the noise floor, isolation and cancellation mechanisms combined should reliably reject the SI by at least 90dB. Furthermore, we assume that a compact antenna solution in a changing near-field can achieve a worst-case isolation of only 20dB, requiring $90 - 20 = 70$ dB from cancellation.

Figure 4.1a visualises an attempt to cancel the remaining SI after antenna iso-

Figure 4.1: Full-duplex nodes assuming a moderate 20dB antenna isolation:
 a) The SI is affected by transmitter and receiver imperfections, which limit cancellation in the digital domain. b) A phase / amplitude based canceller can enhance a moderately isolating antenna in the analog domain, relaxing TX EVM, RX dynamic range and digital cancellation requirements.

lation all in the digital domain. This work pursues a large link budget under the assumption that only linear, deterministic self-interference will be cancelled in the digital baseband. As such, TX EVM and SI-induced RX noise and distortion may still mask the desired signal [15]. To prevent this, roughly 70dB TX EVM and 70dB RX dynamic range (DR) would be required, which is not feasible in a low-power FD node.

Introducing a frequency-flat phase shift / attenuation based canceller at RF can improve RF SI-rejection to a level limited by the frequency-selectivity of the antenna interface, environment and the desired bandwidth (figure 4.1b). To further cope with frequency-selectivity in this architecture, the canceller would need to incorporate multi-nanosecond time or group delay (i.e. a non-flat phase response) [35], which may become costly in silicon area. However, in an indoor scenario, the reflections in the 2.4GHz ISM band are reportedly present at -40 to -50dB [30]. For an antenna interface with limited additional frequency-selectivity, a frequency-flat canceller at RF may therefore reduce requirements on TX EVM, RX DR and digital cancellation to $90 - \{40 \text{ to } 50\} = 40 \text{ to } 50\text{dB}$, which is much more feasible than 70dB.

As such, the useful attenuation range for the canceller in this system with respect to the TX power ranges from 20dB (worst-case SI from the antenna) to 50dB (best-case level of the reflections). As for the phase shift, a full 360° range is desirable since the absolute phase of the SI can assume any value depending on the antenna configuration. So the canceller may consist of 20dB fixed attenuation, about 30dB variable attenuation, and a full 360° phase shift.

Figure 4.2: a) A generic phase shift/attenuation based canceller preceding a mixer-first receiver; b) The cancellation node is moved to the analog baseband and the phase shift, attenuation and down-mixing are combined into a vector modulator (VM) downmixer.

For a frequency-flat canceller, the tolerable group delay δ of the antenna solution can be evaluated mathematically. Assuming optimum cancellation in the center of bandwidth BW, the phase error at the band edge equals $\phi_e = \delta \times 2\pi BW/2$, resulting in an SI-cancellation at the band edge of $SIC = -20 \log_{10}(2 \sin(\phi_e/2))$. Rewriting yields the tolerable group delay

$$\delta = \frac{2}{\pi BW} \sin^{-1}\left(\frac{10^{-\frac{SIC}{20}}}{2}\right) \approx \frac{10^{-\frac{SIC}{20}}}{\pi BW} \quad (4.1)$$

using a small-angle approximation. Here, SIC is the desired *worst-case* SI-cancellation (at the band edge). Similarly, it can be shown that when *band-integrated* cancellation is considered, the tolerable group delay increases by a factor $\sqrt{3}$. For the aforementioned 20 to 30dB SIC on top of 20dB isolation integrated over 16.25MHz BW, the tolerable group delay is 3.4 to 1.1ns. For the following system design considerations, such values are assumed feasible.

According to these considerations, the full-duplex RX should realize a reasonable compromise between noise and SI-induced distortion. In other words, its SI-to-noise-and-distortion-ratio (SINDR) should be high for an optimum full-duplex link budget. SINDR is depicted in figure 4.1b.

Maintaining high in-band linearity under strong SI is crucial to obtain a high SINDR, which motivates interchanging the LNA and mixer and moving to a mixer-first architecture (figure 4.2a). Subsequently, the cancellation node may be moved to the analog baseband and the phase shift, attenuation and down-mixing can be combined in a single component, i.e. a Vector Modulator (VM) downmixer (figure 4.2b).

This topology taps the TX signal at the TX RF output, thus including TX impairments in the cancellation, relaxing TX EVM requirements by the amount of cancellation achieved. It cancels SI before the baseband amplifiers and ADC, relaxing their dynamic range requirements by the same amount. A fixed attenuator is added

to match the VM range to the worst-case isolation of the chosen antenna solution and kept external for versatility.

As for the transmitter, to permit 40 to 50dB linear digital cancellation without encountering its noise or distortion, its total error vector magnitude (EVM) should be in the order of -40 to -50dB when transmitting the desired 0dBm output power.

Clocking all three mixers in the system from a common LO source results in partial phase noise immunity during full-duplex operation, as discussed later in this section.

The topology in figure 4.2b has high integration potential and as discussed, it is applicable to low-power, short-range full-duplex nodes. The following section discusses implementation details of the prototype front-end.

4.2 Implementation of a Full-Duplex Front-End

This section describes the implementation of an SI-cancelling front-end in 65nm CMOS according to the topology of figure 4.2b.

4.2.1 SI-Cancelling Receiver

As explained in section 4.1, to allow cancellation of residual SI, including delayed SI-components, in digital and uncover the desired signal, the RX should have very high SINDR, and thus high in-band linearity under cancellation of strong SI. This prevents the SI from inducing distortion that raises the RX noise floor and masks the desired signal. In the proposed topology, this puts very strict in-band linearity requirements on both downmixers, as they both have to process the maximum TX leakage at their inputs. Furthermore, to prevent RX clipping under strong SI, cancellation has to take place before amplification. Contrary to traditional systems, there is no TX-RX frequency separation, so filtering cannot be used.

Hence, both the main RX and the VM are based on highly linear passive mixers with series resistors into virtual ground nodes provided by transimpedance amplifiers (TIAs) [50]. Figure 4.3 shows an overview of the implemented receiver. The VM is a sliced version of the main RX, followed by static phase rotator switches that route the current of each slice into the four virtual grounds. This way, the SI currents are diverted through highly linear passive networks and only the residue is amplified. The number of slices and other design details are motivated next.

Resolution

The sliced VM principle is similar to the constant- g_m vector modulator presented in [51], but implemented with *resistors to a virtual ground* rather than *active transconductors*. The amount of slices of the VM determines the number of phase / ampli-

Figure 4.3: Implementation details of the SI-cancelling receiver. The vector modulator (top) is a 31-slice version of the main receiver (bottom), each slice followed by static phase rotator switches. The VM diverts self-interference currents through linear passive networks before amplification.

tude constellation points it can cover and thus the amount of cancellation that can be achieved due to quantization effects. This is illustrated in figure 4.4. For n slices, the constellation consists of $n + 1$ by $n + 1$ points. The maximum quantization error occurs when the actual SI phase and amplitude represents a point right in the center of four VM constellation points. Normalizing the constellation to a square of 1×1 , the quantization error has a magnitude of $q_e = \frac{\sqrt{2}}{2n}$. Since the VM has to cover a full circle of phase shifts, the cancellation range is the largest circle that can be drawn through the constellation with maximum error q_e , which has a radius of $1/2 + 1/(2n)$. Thus, the worst-case cancellation given a number of slices is given by

$$\text{SIC [dB]} = 20 * \log_{10} \left(\frac{\frac{1}{2} + \frac{1}{2n}}{\frac{\sqrt{2}}{2n}} \right) = 20 * \log_{10}(n + 1) - 3\text{dB} \quad (4.2)$$

As discussed in section 4.1, a cancellation up to 30dB allows reducing the direct crosstalk to levels where frequency-selective components dominate the SI. Combined

Figure 4.4: The cancellation principle explained using a 3-slice VM (i.e. a 4-by-4 constellation): a fixed attenuation maps the VM cancellation range on the worst-case expected SI, the VM selects the phase / amplitude point closest to the incoming SI and some residual SI remains due to quantization effects.

with practical constraints, a resolution of $n = 31$ slices was chosen, allowing 27.1dB cancellation¹. 31 slices can be conveniently segmented and controlled with 5 bits.

Noise

Designs based on 50Ω resistive termination and 4-phase, 25% duty cycle mixing have a noise figure (NF) that is fundamentally limited to 3.9dB [52]. However, in the proposed design, the VM injects considerable noise current into the virtual ground nodes without contributing desired signal. Its noise contribution could be lowered by designing a weak TX coupler and scaling the VM impedance up from the 50Ω standard (i.e. weaker coupling of the SI into the RX path [35]), but in order to use standard external equipment, 50Ω matching was maintained also for the VM.

The VM noise depends on its setting. Analyzing this for all possible VM settings is mathematically involved, since each setting is a complex mapping of resistors and switches into each of the virtual ground nodes. However, three extremes can be analyzed to obtain upper and lower bounds for the NF:

1. The VM is disabled: the system acts as a conventional mixer-first receiver;

¹Slightly less than the 28.5dB mentioned in [47] as a result of more accurate calculation.

2. The VM is set to an I/Q corner of the constellation, i.e. all slices are configured equally and the VM essentially behaves like a regular mixer;
3. The VM is set to minimum amplitude, i.e. the center of the constellation, where half of the slices is set 180° out of phase with the other half.

The latter point cannot be reached in practice, due to the odd number of slices, but given sufficient VM resolution it can be well approximated. Similarly, the second point (maximum amplitude) is not used in practice, since the VM will only use the highest amplitude it can achieve over the desired full phase circle (section 4.2.1). However, both points provide useful bounds for the NF. Figure 4.5 depicts single-ended equivalent circuits in these three configurations, and their equivalent in-band LTI models for noise analysis according to [52]. For this analysis to be valid in-band, the time constants $(R_s + R_m + R_{sw})(1 + A)C_f$ and $R_f C_f$ are assumed much larger than $1/f_{LO}$, which is typically the case in this design. Out-of-band, the C_b shield the TIAs from high frequency IF components. For simplicity, the source impedances are considered resistive and frequency independent. Only thermal noise is considered.

In situation 1), the mixer can be represented by a resistor $R_m + R_{sw}$, due to the non-overlapping nature of the LO signals. The noise and impedance folding effects of the linear time-variant circuit are represented by a shunt resistance $R_{sh} = \frac{4\gamma}{1-4\gamma}(R_m + R_{sw})$ in the LTI equivalent [52]. Here, $\gamma = 2/\pi^2$. The feedback amplifier is modeled by a noiseless amplifier preceded by its input impedance $R_b = R_f/(1+A)$ and two correlated noise voltages $v_{n,amp}$ and $i_{n,amp}R_b$ Where $i_{n,amp}^2 R_b^2 = \frac{4kTR_f}{(A+1)^2} + \frac{v_{n,amp}^2}{(A+1)^2}$ [52]. The noise factor is then given by [52]:

$$F = 1 + \frac{R_m + R_{sw}}{R_s} + \frac{R_{sh}}{R_s} \left(\frac{R_s + R_m + R_{sw}}{R_{sh}} \right)^2 + \gamma \frac{R_f}{R_s} \left(\frac{R_s + R_m + R_{sw}}{\gamma R_f} \right)^2 + \gamma \frac{v_{n,amp}^2}{4kTR_s} \left(\frac{R_s + R_m + R_{sw}}{\gamma R_f} + \frac{R_s + R_m + R_{sw} + R_{sh}}{R_{sh}} \right)^2 \quad (4.3)$$

In situation 2), the VM can be represented like the main mixer by a source resistance R_{s2} , a switch and matching resistance $R_{m2} + R_{sw2}$ and a shunt resistance $R_{sh2} = \frac{4\gamma}{1-4\gamma}(R_{s2} + R_{m2} + R_{sw2})$ accounting for the time variant effects. This network is effectively in parallel with the original shunt resistance, so we can replace R_{sh} in equation 4.3 by an equivalent resistor

$$R_{eq} = R_{sh} // R_{sh2} // (R_{sw2} + R_{m2} + R_{s2}) \quad (4.4)$$

In situation 3), the input of the VM can be considered a differential ground: the source resistance R_{s2} does not contribute any noise in this case, but the VM itself

Figure 4.5: Left: single-ended representations of the receiver in three configurations, Right: their equivalent noise models following the approach of [52]. Top: cancellation path disabled; Center: VM set to maximum amplitude; Bottom: VM set to minimum amplitude.

directly acts as a shunt resistor with value $R_{m2} + R_{sw2}$, which can be modeled in the LTI circuit as $R_{sh2} = \frac{4\gamma}{1-4\gamma}(R_{m2} + R_{sw2})$. The equivalent total shunt resistance now equals:

$$R_{eq} = R_{sh} // R_{sh2} // (R_{sw2} + R_{m2}) \quad (4.5)$$

The noise figure can be evaluated for the three scenarios by introducing practical values. R_s was kept at 50Ω for both inputs. R_f is chosen $1.5k\Omega$ for 24dB overall receiver gain. A two-stage, telescopic op-amp was used with $A = 1000\times$ open loop gain. The main noise contributors of the op-amp are the input pair ($g_{m1} = 2 \times 23.4\text{mS}$) and the active loads of the input stage ($g_{m2} = 2 \times 12.8\text{mS}$). Assuming a noise excess

	Analysis	Simulation
VM disabled	6.4 dB	6.2 dB
VM maximum	9.8 dB	9.9 dB
VM minimum	12.3 dB	12.5 dB

Table 4.1: Calculated and simulated RX noise figure.

factor of 1, the input-referred op-amp noise can be calculated as $v_{n,amp}^2 = 4kT(g_{m1} + g_{m2})/(g_{m1}^2)$. Taking into account a non-zero baseband impedance due to finite op-amp gain, matching is achieved by setting $R_{sw} + R_m = 48\Omega$ and $R_{sw2} + R_{m2} = 48\Omega$. The results are listed in table 4.1.

Beside the analysis, simulations were performed at 2.5GHz LO frequency, with the real baseband amplifier, but ideal mixers, resistors and sources. C_f was chosen 8pF for 13MHz BW and $C_b = 10$ pF capacitors were put on the virtual grounds to filter higher harmonics. Table 4.1 lists the simulated NF at 10MHz offset, to minimize the influence of flicker noise. Analysis and simulation are in close agreement. In conclusion, the VM contributes the largest amount of noise at small amplitude settings, and enabling the cancellation path degrades the system NF by up to roughly 6dB.

Linearity

This work considers SI-induced RX distortion as limiting for digital cancellation, since cancelling this in digital requires precise models of the TX, the channel and the RX distortion behavior, as well as added signal processing. Hence we target minimizing the SI-induced distortion. In this FD mixer-first design, both SI-induced second-order non-linearity (IM2) and third-order intermodulation (IM3) fall directly in the band of interest and deteriorate the system noise+distortion floor for desired signals. Thus, we aim for sufficient in-band IIP2 and IIP3 by design. Given the targeted 16.25MHz BW, 20dB worst-case isolation, and 12.3dB NF, figure 4.6a plots the required in-band IIP2 and IIP3 to keep the SI-induced IM2 and IM3 equal to the system noise floor, as a function of transmit power. For illustrative purposes, the case for a 6dB NF is also drawn. As motivated in section 4.1, this work targets at least 0dBm TX power, resulting in in-band IIP2 and IIP3 requirements of roughly 20dBm and 50dBm, respectively. Note that we aim for sufficiently low distortion to achieve analog cancellation while preserving the noise floor. We do not pursue sufficiently low distortion to further increase the TX power, as this would again put unfeasible requirements on the TX EVM and TIA / ADC DR (see section 4.1).

Ideally, for 0Ω switches and a perfectly linear 50Ω matching resistor, there is no signal swing across the switches and therefore no IM3-currents are induced by the SI before cancellation. The only source of IM3 are the TIAs that process residual SI and

Figure 4.6: a) Required IIP2 and IIP3 to keep the SI-induced IM2 and IM3 equal to the system thermal noise floor, as a function of TX power, assuming 16.25MHz BW, 20dB isolation and 12.3dB or 6dB NF; b) Theoretical IIP3 as a function of switch on-resistance.

the (usually weaker) desired signal. Therefore, every 1dB of cancellation of the SI would result in a 3dB reduction of the SI-induced IM3, boosting the effective IIP3 by 1.5dB. However, low-ohmic mixer switches are power-hungry to drive, resulting in a trade-off between power consumption and IIP3 for switched-resistor mixers. Assuming simple square-law behavior of the switch devices and ideal virtual grounds, the in-band linear and third order components can be computed to be $a_1 = (R_s + R_m + R_{sw})^{-1}$ and $a_3 = -(R_s + R_m)R_{sw}^2 / (2V_{OD}^2(R_{sw} + (R_s + R_m))^5)$, where V_{OD} is the overdrive voltage of the switches [53]. Then $IIP3 = \sqrt{(3/4)|a_1/a_3|}$. Using $V_{OD} = 800\text{mV}$ and taking $R_s = 50\Omega$ and $(R_m + R_{sw}) = 50\Omega$, the IIP3 is plotted as a function of R_{sw} in figure 4.6b. For $>20\text{dBm}$ IIP3, the design was implemented with 25Ω resistors, with the remaining 25Ω distributed over the switch resistance, virtual ground impedance and routing parasitics. The bulk of the mixer switches was tied to the baseband side for reduced on-resistance and better linearity. The multiplexer switches of the VM were sized wide and low-ohmic, since parasitics are absorbed in the baseband capacitance and since they are driven by static control signals. This allows negligible increase of the virtual ground impedance.

For low IM2, a fully differential structure was adopted for both mixers with carefully balanced parasitics, and a common centroid layout scheme was used for the VM slices.

The TIAs were not specifically designed for linearity, and therefore will dominate the system IM3 performance when cancellation is deactivated. However, they perform such that under 27dB cancellation, the mixers will dominate the IM3 performance by a large margin. In addition, the TIAs were further linearized by a differential negative

Figure 4.7: Implementation details of one VM slice for one LO phase, and the TIA linearized by negative conductance.

conductance at their inputs [50]. While not strictly necessary for this application, it allows us to eliminate the TIA as a linearity bottleneck in measurements and study the raw linearity achieved by the mixers, even with cancellation disabled. Figure 4.7 shows an implementation detail of one fully differential VM slice for one LO phase, and one of the negative-conductance-assisted TIA's. The TIA's are implemented as high-gain, two-stage OTA's with a telescopic input stage and a push-pull output stage [50].

LO generation and input matching

The 25% duty cycle LO is generated by an on-chip divide-by-two and logic operations on the four resulting phases. The final stages of LO drivers are AC-coupled to the mixer switches to allow level shifting the LO signals for reduced switch on-resistance. Figure 4.8 shows the level shifting circuit for two clock phases and two switches. The AC coupling capacitors are slowly charged by small switches during the intervals where the LO is low. The level shift voltage is set between 0V and mid-supply by a 5-bit R-2R DAC, allowing digital control of input matching. This allows good input matching over process spread. Independent DACs are used for the VM and the main mixer, to overcome any differences in e.g. layout parasitics. In measurements, the RX and VM were tuned for matching once and the resulting DAC values were used throughout.

Figure 4.9: Transmitter topology: Four-phase sampling mixer and common source class-A PA with thick oxide cascodes. Dashed line = chip boundary. The clocks are shared with the RX.

Figure 4.10: Die photo with relevant blocks indicated. The 65nm design measures 1.4×1.4 mm.

Figure 4.11: a) Over 100 arbitrary, spiral-shaped phase/amplitude points emulate the SI channel; b) The on-chip VM finds the corresponding setting for best cancellation; c) The residual SI is always at least 27dB below the VM range (circle).

4.11a/b. The results, plotted in figure 4.11c, show better than 27dB cancellation which is very close to the calculated 27.1dB from section 4.2.1. This is expected, since despite the minimal practical sizing of the VM slices, matching was found to be much better than strictly required for the 31-slice VM.

4.3.2 Noise

In the thermal noise limited region, a noise figure was measured of 6.3dB without cancellation enabled; 10.3dB with cancellation set for maximum SI (i.e. the VM is set to a point on the maximum circle it can cover) and 12.3dB when set for small SI (i.e. the VM is set to a minimum amplitude). These values correspond very well with analysis and simulation as listed in table 4.1. The 1/f noise corner of the RX

was measured to reside at roughly 2MHz.

4.3.3 Linearity

For a symmetrical point-to-point link based on this design, the available link budget² will at first increase linearly with increasing transmit power (i.e. an increasing SINDR, see section 4.1). However, at some point the increasing SI will induce distortion in the RX that raises the noise floor, limits digital cancellation, and thus decreases the link budget again. This also holds under cancellation, due to the finite linearity of the RX and VM mixers. In other words, there is an optimum SI power for which the system achieves the highest SINDR and thus the largest link budget. To find this optimum, a two-tone self-interferer was applied and its power was swept under cancellation.

First, the IM3 products were observed. Under cancellation, an effective in-band IIP3 can be defined with respect to the SI³. The peak SINDR can then be calculated as:

$$\text{SINDR [dB]} = \frac{2}{3}(\text{Effective IIP3 [dBm]} - \text{Noise Floor [dBm]}) - 3 \text{ dB} \quad (4.6)$$

where the 3dB is due to the RX noise floor and SI-induced IM3 products adding as powers.

Due to the discrete nature of the VM, it is difficult to guarantee exactly 27dB cancellation, therefore the measurement was performed under 26dB cancellation, in order not to be optimistic. Figure 4.12a shows the results without cancellation. Drawing a noise floor in 16.25MHz allows deriving the SINDR. Figure 4.12b shows how the results change under 26dB cancellation. Again, the RX noise floor can be included to derive the SINDR (figure 4.12c). Both SINDRs are shown in figure 4.12d. Note that the performance has improved slightly with respect to [47], to reflect the most recent measurements. The peak SINDR of the system increases from 66.5dB without cancellation, to 71.5dB under cancellation, indicating a 5dB increase in link budget when cancellation is enabled. The point of maximum link budget has moved from -27.6dBm to -16.4dBm SI at the RX input. Also, if the system operates slightly above the optimum amount of SI (e.g. the external attenuator is chosen conservatively or the TX power is slightly larger than expected), the link budget degrades smoothly, whereas the original RX would suffer from output stage clipping (figure 4.12d).

The measurements show that the IIP3 increases from 9dBm to an effective 21.5dBm when cancellation is enabled: an increase of 12.5dB. The fact that the IIP3 does not increase by the full $1.5 \times 27\text{dB}$ (section 4.2.1) indicates that the linearity bottleneck

²‘Link budget’ in this work assumes 0dB SNR at the receiver and does not include any fading and AGC margins, to obtain a standard-independent metric.

³Effective in-band IIP3 is similar to effective out-of-band IIP3, as used in interference-cancelling FDD systems, e.g. [55]

Figure 4.12: Results of a two-tone linearity test: a) Cancellation disabled, including 16.25MHz noise floor and defining SINDR; b) Cancellation enabled; c) Cancellation enabled, including noise floor and SINDR; d) SINDR with and without cancellation for various SI powers. Note that the performance has improved slightly with respect to [47], to reflect the most recent measurements.

	Without cancell.	With cancell.
Maximum link budget (SINDR + Isolation)	86.5 dB	91.5 dB
Digital cancellation requirement (SINDR - Cancellation)	66.5 dB	44.5 dB
TX power @ max. link budget (SI + Isolation)	-7.6 dBm	3.6 dBm

Table 4.2: Summary of cancellation, noise and linearity effects on overall full-duplex link performance, assuming 20 dB antenna isolation.

has moved from the TIA to the nonlinear RX and VM switches. Since enabling the cancellation increases the effective IIP3 by 12.5dB but also increases the noise floor by 6dB, equation 4.6 shows why the 27dB cancellation only yields a 5dB link budget increase.

However, the main intention of the canceller was *not to improve link budget, but to relax TX EVM, TIA / ADC dynamic range and digital cancellation requirements*, and all of these are still relaxed by the full 27dB of cancellation, minus the 5dB link budget increase. Table 4.2 summarizes the effect of the cancellation on the link budget of the system, under the assumption of 20dB antenna isolation. Its main merit is bringing the digital cancellation, TX EVM and TIA / ADC dynamic range requirements down from an unfeasible 66.5dB to a realistic 44.5dB.

Given the optimum TX power based on IM3, the IM2 was evaluated. Referring to figure 4.6 at 3.6dBm TX power, 56dBm IIP2 would be required for IM2 equal to the noise floor. Measuring the beat component of two in-band tones, +60dBm IIP2 was measured, which is sufficient by some margin and similar to that achieved in other mixer-first designs. Note that in this mixer-first design, IM2 is dominated by the mixers and therefore is not reduced by cancellation. As such, defining an effective IIP2 is not useful. Since the design required a post-production routing fix, a limited number of functional samples was available and the IIP2 was not characterized over multiple samples.

In figure 4.12c, to find the SINDR, the fundamentals were extrapolated from the case without cancellation. This assumes that under cancellation, the SI does not compress the RX for the SI power range of interest. This can be validated by applying a third tone, representing the desired signal, and monitoring its conversion gain. Figure 4.13 shows the result: under cancellation, the RX can handle in excess of 1.5dBm of SI before the desired signal is compressed; at this point, the residual SI is strong enough to saturate the TIA, despite the cancellation. This is 24dB higher SI than without cancellation and justifies the extrapolation made in figure 4.12b/c.

Figure 4.13: Conversion gain for desired signal with increasing SI, without and with cancellation.

As mentioned in section 4.2.1, the TIAs can also be eliminated as linearity bottleneck by enabling the differential negative conductance present at their inputs. With the cancellation disabled, this allows us to observe the raw linearity of the main RX mixer, which results in an IIP3 of 19dBm. The fact that the effective IIP3 under cancellation is even 2.5dB higher, can be explained by two phenomena: 1) distortion cancellation mechanisms occurring between the RX and VM; 2) the fact that cancelled SI does not cause signal swing on the virtual grounds, whereas received signal does. Note that the measurements in figure 4.12 and 4.13 were performed without negative conductance.

4.3.4 Broadband performance

Although the aforementioned results were obtained at 2.5GHz LO frequency, the receiver employs frequency-agile operation and cancellation principles. Figure 4.14 shows several performance characteristics over a broad range of LO frequencies. NF and RX gain are reasonably flat over the entire operating range from 0.15 to 3.5GHz. Due to the discrete nature of the VM, the cancellation performance varies, as expected, but always exceeds 27dB. Power consumption increases linearly with frequency with a static component, as expected.

4.3.5 Transmitter

As demonstrated, the SI-cancelling RX achieves up to 71.5dB SI-to-Noise-and-Distortion-Ratio (SINDR) at 27dB SI-cancellation in a 16.25MHz BW. Hence the digital SI-cancellation potential is $71.5 - 27 = 44.5$ dB without RX noise and distortion limitation. Combined with a modest 20 dB antenna isolation, a link budget of $71.5 + 20 = 91.5$ dB

Figure 4.14: Performance of the receiver over a wide range of LO frequencies.

Specification	Value
Output IP3	20.1 dBm
Image rejection ratio	38 dB (uncalibrated)
LO radiation	-49 dBm (uncalibrated)
Power consumption	108 mW from 2.0V (PA) 21 mW from 1.2V (LO)
Maximum single-tone output	12.4 dBm
Efficiency at max. output	13%
Operation frequency	0.05 - 3.5 GHz
EVM @ 0dBm 802.11a output	-40 dB

Table 4.3: Transmitter specifications @ 2.5 GHz carrier frequency.

is available at 3.6 dBm TX power, enough for short-range links. However, to achieve 44.5dB digital cancellation without non-linear TX modelling, the TX should ideally have better than -44.5dB error vector magnitude (EVM) at 3.6dBm output power.

Table 4.3 lists some key measured parameters of the TX. It covers the RX operating frequency range of 0.15 to 3.5 GHz. For 0 dBm TX power, its EVM is -40 dBc, allowing almost the full amount of linear digital cancellation mentioned above. Bridging the gap to -44.5dB EVM at 3.6dBm output power by e.g. pre-distortion is left for future research.

Figure 4.15: FD radio topology implemented in [40], using an additional up-converter for SI-cancellation at RF. All 3 up- and downconverters are separate transceivers with their own PLL.

4.3.6 Phase noise

Phase noise (PN) can be troublesome for FD [15, 26], specifically in architectures where cancellation occurs across frequency domains (e.g. RF and IF / BB). In this section, the phase noise benefits of clocking all three up- and down-conversion mixers from the same LO source are explored.

Early work on FD generally used off-the-shelf radios, specifically the WARP research platform [13, 40]. One particular topology, shown in figure 4.15, uses an additional upconverter chain for SI cancellation at the receiver input [40]. In [26], the authors observe that in this design, the combined amount of mixed-signal and digital cancellation never exceeds 35dB, leaving remaining SI far above the thermal noise floor. This is a result of using separate off-the-shelf radio IC's for both upconverters, each with their own PLL, generating -38dBc uncorrelated phase noise in the band of interest, inducing a combined noise floor of -35dBc that limits cancellation.

In contrast, the integrated front-end in our work uses the same LO for the upconverter and both downconverters, which, ideally, would make the system insensitive to phase noise in the SI. We will now show that for short SI-paths, the RX noise floor is indeed hardly deteriorated by PN when using a PLL of the WARP platform. However, for long SI-paths via a reflective environment, the PN of the delayed SI becomes decorrelated from the RX clock [26]; we will evaluate this effect and examine its implications.

To exploit the maximum $\text{SINDR}=71.5\text{dB}$ of our RX and thus its maximum digital cancellation potential, the receiver should operate under its full 27 dB analog cancellation while receiving -16.4dBm SI (section 4.3.3). For higher SI powers, digital cancellation is impaired by SI-induced RX-distortion; for lower SI powers, dynamic

Figure 4.16: Simplified set-up used to measure the effect of correlated and non-correlated phase noise between TX and RX/VM, for short and long self-interference paths. The power levels are all referred to the RX input.

range is limited by the RX noise floor. Hence, targeting the maximum link budget, the PN impact should be evaluated under 27dB analog cancellation of -16.4dBm SI at the receiver input. The remaining input-referred SI after analog cancellation will be $-16.4\text{dBm} - 27\text{dB} = -43.4\text{dBm}$. The RX noise figure (NFDSB) in its full bandwidth (50kHz - 12 MHz) is 11.7 dB (slightly more than reported before since flicker noise is included down to 50kHz). Therefore, its input-referred DSB noise floor is $-174\text{dBm/Hz} + 10 \cdot \log_{10}(24\text{MHz}) + 11.7\text{dB} = -88.5\text{dBm}$. This leaves room for $-43.4 - -88.5 = 45.1\text{dB}$ of digital cancellation towards the noise floor. Therefore, the presence of e.g. -35 dBc uncorrelated phase noise between TX and RX can easily raise the noise floor and deteriorate the amount of achievable digital cancellation. The following experiments investigate if a common clock helps to prevent noise floor degradation due to phase noise present in the residual SI.

Since more preliminary link budget measurements showed a peak SINDR of 69dB at -18dBm SI, the system was actually evaluated at 27dB cancellation of -18dBm SI instead of -16.4dBm, however, the conclusions of the following experiments remain applicable. The measurement setup is shown in figure 4.16. Experiments were done at 2.5 GHz carrier frequency, requiring a 5 GHz LO. This LO is either generated by a generator or by an unmodulated carrier from the actual radio used in the WARP platform. The generator has better than -50 dBc PN, hardly degrading the RX noise floor, which lies at 43.5 dB below the residual SI. However, the WARP radio with its -38 dBc in-band PN is expected to significantly degrade the noise floor.

Figure 4.17: Measured input-referred noise spectra for the four clocking configurations with a short SI-path: A shared WARP clock shows comparable PN performance to a low-PN generator.

A very low offset-frequency single tone of 1 kHz is used from the TX. This permits filtering out TX imperfections like LO-radiation, image and distortion by AC-coupling toward the spectrum analyzer, allowing sensitive measurements of only the phase noise impact on the noise floor. In reality, such TX imperfections will also deteriorate cancellation, but since they are largely deterministic, they can be treated in the digital domain by e.g. calibration, pre-distortion and nonlinearity estimation [13], in contrast to phase noise.

Table 4.4 shows for short SI channels how the PN deteriorates the RX noise floor under cancellation, and Figure 4.17 shows the corresponding measured spectra. With un-correlated sources, the noise floor degrades by 9.4 dB, and clearly shows the presence of the WARP phase noise profile. However, by virtue of correlated clocks, the noise floor is hardly affected by PN if TX and RX/VM share a single WARP clock source, but stays comparable to using a shared, low-PN generator.

If the SI travels long distances, its phase noise may become decorrelated from that at the RX for larger offset frequencies (i.e. quick variations in LO phase), introducing a new, fundamental noise floor in the system that may limit digital cancellation [26]. An experiment was set up with an available delay line, providing again -18 dBm of SI at the RX input and 27 dB analog cancellation, but with 115ns of delay in the SI path. The results are shown in Figure 8 and Table 4.5. The effect of PN

TX clock	RX / VM clock	DSB noise figure	Noise floor degradation
Disabled	WARP	11.7 dB	-
Shared low-PN generator		13.0 dB	1.1 dB
Low-PN gen.	WARP	26.6 dB	9.3 dB
WARP	Low-PN gen.	20.2 dB	9.4 dB
	Shared WARP	13.0 dB	1.2 dB

Table 4.4: Effect of phase noise on RX noise floor, short SI path.

Figure 4.18: Measured input-referred noise spectra for the four clocking configurations with 115ns delay in the SI-path, resulting in phase noise decorrelation at high carrier offsets.

decorrelation is observed: clocking all upconverters from a shared WARP source does not bring the noise level down to that of a shared generator, but instead only offers 6.2 dB improvement compared to clocking one side with WARP, consistent with further analysis.

However, this is a very pessimistic scenario: in a practical environment, a 115ns reflection will not return in the RX at such high signal strengths: Assuming a path-loss exponent of 2 in the SI path (corresponding to a perfect reflection against a flat metal surface), such a reflection would encounter 73.3dB path loss. For a TX power of 0dBm, the reflected SI would come in at -73.3dBm, which is only 15.1dB above the RX noise floor. Therefore, even if the phase noise were fully decorrelated, it would

TX clock	RX / VM clock	DSB noise figure	Noise floor degradation
Disabled	WARP	11.8 dB	-
Shared low-PN generator		13.2 dB	1.2 dB
Low-PN gen.	WARP	19.5 dB	9.8 dB
WARP	Low-PN gen.	18.5 dB	9.5 dB
Shared WARP		15.6 dB	3.3 dB

Table 4.5: Effect of phase noise on RX noise floor, long SI path.

still be well below the noise floor.

Since this is a discrete measurement point at 115ns delay, further analysis is required to show that there are no reflection scenarios for which decorrelated PN will raise the RX noise floor and deteriorate the link budget. For this purpose, the PN profile of WARP was captured and further processed in MATLAB. After applying decorrelation, path loss and ideal phase / amplitude based cancellation for several distances of the hypothetical reflector, the residual PN was integrated from 50kHz to 12MHz offset. Figure 4.19 shows the results for path losses $n=1.4$, 1.7 and 2.0 , and in the hypothetical case of no loss ($n=0$). For reference, typical path loss models for point-to-point links in the same room use $n=1.6$ to 1.8 [56]. For values of $n<2$, a distance can be identified where reflections are essentially most harmful to the link. However, for practical values of n , the link budget is not deteriorated by decorrelated PN.

The vertical lines in figure 4.19 allow us to relate these simulations to literature and the experimental results. In figure 4 of [27], it is shown that for a 42ns delay, using a single WARP instead of two uncorrelated clocks improves the cancellation potential by roughly 10dB. The same situation offers roughly 14dB improvement in our simulations (first vertical line). Differences can arise from e.g. the exact bandwidth taken into account. Relating figure 4.19 to the 115ns-measurements in table 4.5 can be done by observing the second vertical line and applying the 20dB isolation and 27dB cancellation to the plotted lines for $n=0$. The decorrelated noise will appear at roughly $-48-(20+27) = -95\text{dBc}$, or 3.5dB below the noise floor. Therefore, a minor noise floor degradation of about 1.3dB is expected. In practice (table 4.5), a 1.9dB degradation is observed compared to the situation with a shared generator, and a 3.3dB degradation compared to the situation with decorrelation of WARP PN. This confirms that the decorrelated PN floor is indeed close to the thermal noise floor as predicted by this analysis.

In conclusion, by virtue of a common clock for all mixers in the system, the phase noise arising from operation off a commercially available PLL with -38 dBc phase noise does not degrade the 91.5dB link budget potential calculated previously.

Figure 4.19: Integrated PN after ideal phase/amplitude-based cancellation with respect to transmitted power, for varying reflector distances and path losses (modeled by a loss coefficient n). Dashed lines indicate both mixers clocked with uncorrelated WARP PN profiles, solid lines indicate one shared WARP clock suffering from decorrelation. For comparison with literature and experiments, the vertical lines indicate reflections of 42ns and 115ns respectively as described in [27].

4.3.7 Image rejection

A concern of the proposed topology is image rejection: The RX and the VM process the full SI power, but ideally, the received image of the SI should be below the noise floor. As such, about 71dB image rejection is required from the mixers, which is not a feasible value without on-chip calibration mechanisms. However, if the image rejection is over 27dB, it does not limit analog cancellation, and the residual image can be dealt with in digital cancellation [57]. The prototype achieves 37dB image rejection, sufficient for analog cancellation by a margin of 10dB, but the image must be accounted for in digital to reach the full 44.5dB digital cancellation potential.

4.3.8 Comparison

Table 4.6 compares this work to three previously published integrated FD receivers. For fair comparison, no antenna isolation is assumed for all designs. The work of [44] features a single antenna, and uses the bidirectional transparency of a mixer to achieve TX-RX duplexing in the baseband. TX-RX isolation is achieved in baseband by injecting the TX signal into the common gate node of a noise-cancelling BB LNA. In [35], wideband cancellation is achieved by means of an RF canceller path comprising two-port N-path filters. Such filters have four degrees of freedom (phase, gain, center frequency, and quality factor) and can therefore emulate group delay of the self-interference path to some extent. Finally, in [46], a two-port N-path filter provides a non-reciprocal 90-degree phase shift, which is combined with a $3/4$ -wavelength transmission line to construct a highly integrated circulator. This topology uses a single antenna.

The peak SINDR of the other works was calculated using equation 4.6. The SI power at which the peak SINDR occurs is given by:

$$\text{SI [dBm]} = \text{Effective IIP3 [dBm]} - \frac{1}{3} * (\text{Effective IIP3 [dBm]} - \text{Noise Floor [dBm]}) \quad (4.7)$$

where the noise floor depends on the NF and RX BW. Although this work features the highest peak SINDR, and thus the highest link budget potential given a fixed amount of antenna isolation, it should be noted that the architecture of [35] can theoretically achieve significant cancellation over a wide bandwidth even when the initial antenna isolation is high, thanks to its ability to address delayed SI components. When combined with 34dB antenna isolation, the total linear link budget of [35] can slightly exceed that of the proposed design, despite its lower SI-handling capability at the RX input. However, whether 34dB antenna isolation can be maintained under a changing antenna near-field is questionable.

Although the gain of this design is relatively low due to limited range of the BB feedback network, experiments using an external $10\text{k}\Omega$ feedback network resulted in 39.3dB gain at the cost of a reduced compression level and baseband bandwidth, but without compromising the effective in-band IIP3 and thus the linear link budget potential.

In conclusion, the proposed design enables FD links at practical TX powers in excess of 0 dBm with sufficient dynamic range for reception, compared to the compressed -17.3dBm of [44], while maintaining low complexity and power consumption when compared to [35].

Table 4.6: Comparison with other integrated FD transceivers, assuming no antenna isolation.

	[35]	[44]	[46]	This work
Topology	Dual-port N-path filter based canceller + noise-cancelling receiver	Mixer-first architecture + Noise-cancelling duplexer LNA's	N-path filter based circulator + noise-cancelling receiver + analog BB canceller	Mixer-first receiver + SI-cancelling VM-downmixer
Technology	65 nm CMOS	65 nm CMOS	65nm CMOS	65 nm CMOS
Supply	N/R	1.2V (LO) / 2.5V (BB)	1.2V (RX) / 1.3/2.2V (BB)	1.2V
Operating freq.	0.8-1.4 GHz	0.1-1.5 GHz	0.6-0.8 GHz	0.15-3.5 GHz
Max. gain	42 dB	51-55 dB	42 dB	24 dB
NF in HD mode	4.8 dB	5-8 dB	5 dB	6.3dB
NF degradation in FD mode	0.9-1.3 dB	N/R	6.5 dB	4-6 dB
Power consumption	63 - 69 mW (RX) + 44 - 182mW (Canc.)	43 - 56 mW (incl. TX)	129mW (RX/circ.) + 30mW (Canc.)	22 - 46 mW (RX) + 1 - 10 mW (Canc.) ¹⁾
Baseband BW	>30 MHz (-15 to +15)	6-192 MHz	12 MHz	24 MHz (-12 to +12)
In-band IIP3	-20 dBm	-32.7 dBm	-33 dBm	+9 / +19 dBm (Neg. conductance off / on) ²⁾
Effective IB-IIP3 with respect to SI	2 dBm	-0.7 dBm ³⁾	-18 dBm	21.5 dBm
SINDR in 16.25 MHz BW	62.5 dB peak @ -30.7 dBm SI	60.8 dB peak @ -32.6 dBm SI	45.9 dB peak @ -42.1 dBm SI	71.5 dB peak @ -16.4 dBm SI
Typical RF isolation	34 dB (antenna pair)	0 dB (single antenna)	20 dB (integrated circulator)	20 dB (antenna pair)
Linear link budget in 16.25MHz BW	96.5 dB	60.8 dB	65.9 dB	91.5 dB
Out-of-band IIP3	17 dBm	22.5 dBm	19 dBm	22.0 dBm
Resolution-limited SI Cancellation	N/R	N/A	42dB	27 dB
SI power @ 1dB RX compression	$\gg$ -8 dBm	-17.3 dBm	-5dBm ⁶⁾	$>$ +1.5 dBm ⁴⁾
In-band IIP2	+10 dBm	+7 dBm ⁵⁾	N/R	+60 dBm
Effective IB-IIP2 with respect to SI	+68 dBm	+24 dBm ⁵⁾	N/R	+60 dBm
1/f Noise corner	N/R	N/A	1 MHz ⁶⁾	2 MHz
Practical cancellation details	20 dB worst-case in 25 MHz BW, 34 dB initial iso. from 1.4 GHz dipole pair, 8 ns peak group delay ⁷⁾	33.5 dB in $\sim$ 1 MHz BW ⁸⁾ , with single-port antenna	42 dB across 12 MHz BW, with single-port antenna, 85dB shown incl. digital SIC	15.6 dB worst-case, 21 dB integrated in $\sim$ 16 MHz BW, 25 dB initial iso. from crossed 2.5 GHz dipoles, 4 ns peak group delay
Area	4.8 mm ²	1.5 mm ²	1.4 mm ²	2 mm ²

Notes: Several values of [44] and this work were updated with respect to [47] to reflect the most recent data sets. 1) The transmitter adds 129mW at 2.5GHz, as detailed in [48], 2) Negative conductance gives about 1.5dB NF penalty [50], 3) From -38.7 dBm IIP3 and 38 dB IIP3 improvement @33.5 dB isolation, 4) 135 kHz spacing [44], under 27 dB cancellation, 5) Estimated from [44], figure 31, 6) Estimated from [46], figure 9.8.5, 7) Cancellation was optimized for wide bandwidth, 8) From [44], figure 25

Figure 4.20: Test setup for cancellation measurements using the built-in transmitter and crossed WLAN dipoles as FD antenna.

4.3.9 Antenna experiments

To verify the claims of 20dB as a representative worst-case antenna isolation and -40 to -50dB as the level where frequency-selective components dominate the SI in 16.25 MHz BW, some experiments were performed with the front-end, using a crossed pair of commercial WLAN dipoles as a simple FD TX/RX antenna pair. The setup was located in an RF lab environment without taking any specific measures to induce or reduce reflections. The crossed antennas were characterized first using a vector network analyzer. In this environment without special precautions, this antenna solution provides typically 25dB isolation with 4ns peak group delay and 2.5dB amplitude variation.

Next, the crossed dipoles were connected to the front-end as illustrated in Figure 4.20. Connections were kept short to avoid introducing unnecessary propagation delay. A test signal of 20 tones with random phases over a 20MHz bandwidth was transmitted at 2.5GHz while the RX baseband bandwidth was set to 24MHz. The external attenuator is set to 20dB to configure the canceller for a worst-case antenna isolation of 23dB (including the 3dB attenuation of the VM to describe a circle in a square constellation).

Figure 4.21 shows typical measured spectra in the system: at the transmitter output, at the receiver input and after cancellation (referred to the receiver input). At a transmit power level of about 0dBm, the antenna isolation integrated across the band equals 28.9dB. Added with a band-integrated VM cancellation of 16.8dB, the combined SI rejection equals 45.7dB. The residual SI becomes strongly frequency-selective: combined SI-rejection degrades to 40.6dB at the worst band edge. This indicates the presence of reflected SI in the residue, since this type of SI is delayed and therefore not well approximated by a phase shift & attenuation. For narrow bandwidths, the VM phase shifting and attenuation provides a better fit, resulting in

Figure 4.21: Spectra and integrated power levels at various points in the FD front-end, all referred back to the antenna ports for comparison. The shaded area indicates the 25 MHz integration bandwidth.

increased SI cancellation. Delay can be implemented in an RF canceller with long delay lines [13], or approximated as group delay [35], but the former is incompatible with CMOS integration and the latter adds complexity, non-linearity and power consumption. Therefore, in the proposed compact radio, reflected, delayed self-interference should be dealt with in the digital domain, requiring sufficient TX EVM and ADC DR as discussed.

Furthermore, heavily influencing the antenna near-field with a hand showed that 20dB is a reasonable worst-case isolation for this FD antenna.

4.4 Conclusions

This chapter presented an integrated self-interference (SI) cancelling front-end, aiming to bring in-band full-duplex wireless communication to compact low-power devices. Starting from full-duplex system considerations, we found that a phase / amplitude based SI-canceller in the analog domain is useful to improve upon low and varying antenna isolation.

The proposed receiver takes an attenuated copy of the transmit signal, and pro-

vides simultaneous tuneable phase shift, amplitude scaling and downmixing using a vector modulator (VM) downmixer, for SI-cancellation in the RX analog baseband. The main RX and VM are based on a highly linear switched-resistor mixer-first architecture, to cancel SI with highly linear passive circuits, prior to amplification of the residue. This keeps SI-induced distortion low and thus maximizes the digital cancellation and link budget potentials.

For the sliced vector modulator, the cancellation performance was derived as a function of the number of slices. We also showed how to analytically obtain upper and lower bounds for the setting-dependent noise performance of the receiver including VM. Other design choices, such as the VM resolution, were also motivated. The SI-to-noise-and-distortion ratio (SINDR) of the system was defined as a crucial figure for link budget performance.

The prototype with 31-slice VM achieves up to 27 dB cancellation. When configured for only 20 dB worst-case isolation from the antenna, the maximum link budget is achieved at 3.6 dBm TX power, where SI-introduced distortion equals the RX noise floor. Given its 12.3 dB worst-case noise figure with cancellation enabled, this results in up to 91.5dB link budget in a 16.25MHz bandwidth, enough for short-range links. Since the TX is inside the cancellation loop, and cancellation occurs before amplification, the 27dB cancellation reduces the requirements on TX EVM, baseband amplifiers and ADC to feasible levels.

The co-integrated TX has -40dB EVM performance at 0dBm TX power, requiring only minor improvement by e.g. pre-distortion to produce the desired 3.6dBm output power at -44.5dB EVM and thus allow the full 91.5dB link budget potential of the system.

The ability to clock the TX and RX from a shared clock offers increased immunity to phase noise compared to previous demonstrators based on commercial integrated transceivers, allowing operation from a typical commercial PLL with -38dBc phase noise. Decorrelation of the phase noise due to delays in the self-interference path was found not to deteriorate the system's overall link budget potential.

The entire system was mainly demonstrated at 2.5GHz, but offers frequency-agile operation and cancellation from 0.15 to 3.5GHz LO frequency.

Chapter 5

An improved front-end for in-band full-duplex radio

Chapter 4 presented an integrated front-end for in-band full-duplex wireless communication. Many of its analog impairments and their impact on a full-duplex link budget were analyzed and experimentally investigated. However, because of a required silicon fix, too few functional samples were obtained to set up a fully functional point-to-point link and experiment with digital cancellation. Therefore, a second front-end was designed aiming to maximize the full-duplex link budget, based on the same concept and assumptions. Note that *the work presented in this chapter is preliminary*, since measurements on the complete system were not performed within the time span of this thesis work.

The following constraints and assumptions of the previous design remained unchanged:

- 20dB worst-case isolation is assumed from the antenna interface in a changing near-field environment,
- 27dB analog cancellation is targeted to push self-interference down to the level of ambient reflections,
- 20MHz baseband bandwidth is targeted,
- frequency-selective, delayed SI will be addressed in digital,
- only cancellation of linear SI is preferred in the digital domain; SI-induced RX distortion should remain below the noise floor.

The simplified definition of link budget used in this thesis is the distance from the TX power of a node to its noise floor (ignoring any fixed margins for simplicity). Since the targeted baseband bandwidth remains unchanged, there are only two degrees of freedom to increase the link budget: Reducing the RX noise figure and increasing the

	Previous design	New design (target)
TX power	3.6 dBm	16 dBm
Noise figure	12.3 dB	8 dB
Baseband BW	16.25 MHz	16.25 MHz
Antenna ISO	20 dB	20 dB
Noise floor	-90 dBm	-94 dBm
Link budget	91.5 dB	107 dB
SI at RX in	-16.4 dBm	-4 dBm
IIP3 w.r.t. SI	21.5 dBm	41 dBm
VM resolution	27 dB	27 dB
1/f noise corner	2 MHz	200 KHz
Linear digital SIC potential	44.5 dB	60 dB
TX OIP3	20.1 dBm	46 dBm
Power consumption	47mW (RX/Canc) + 129mW (TX)	100mW (RX/Canc.) + 1.2W (TX)

Table 5.1: Achieved performance of the previous design and targeted performance of the new design.

TX power. As detailed in the next section, the noise figure of the current architecture can be substantially reduced, but not indefinitely. Targeting a NF reduction from 12.3 to 8dB, and an increase in TX power from 3.6 to 16 dBm, the impact on link budget and the main TX and RX specifications is derived in table 5.1.

We observe that to accommodate the increased SI power without generating distortion above the slightly lowered noise floor, the effective RX IIP3 needs to be substantially increased. If this is satisfied, the main burden is now put on the ADC and DAC dynamic range, which need to exceed 60dB by a significant margin (e.g. for fading, signal peaking and quantization noise) to enable 60dB digital cancellation. This is feasible for demonstrator purposes using measurement-grade baseband equipment, but for practical applications, such requirements are quite tough.

The increased power budget is spent on obtaining a lower noise figure, driving mixers with increased linearity and generating higher TX power with typically low CMOS PA efficiencies.

5.1 Design

Figure 5.1 compares top-level views of the previous architecture and the new design. The most notable changes on architectural level include:

Figure 5.1: a) top-level view of the previous architecture, b) new architecture. Compared to the previous design, a two-stage power amplifier will be integrated, the transmit taps are brought on-chip, and a second vector modulator is added for optional cartesian feedback.

- An integrated two-stage power amplifier,
- A second vector modulator for optional Cartesian feedback linearization of the amplifier,
- Both transmit taps are brought on-chip.

The next sections adress the implementation details of all modified blocks with respect to the previous design. Starting from the noise figure reduction, it becomes apparent why the transmit taps are brought on-chip. After discussing the required modifications to the receive mixers and baseband amplifiers, implementation of the transmitter is addressed.

Whereas the previous design was implemented in 65nm CMOS from ST micro-electronics, this design uses 65nm CMOS from TSMC.

5.1.1 SI-cancelling receiver

This section discusses the improvements made to the front-end at the SI-cancelling receiver side. First, steps taken to improve the noise figure are discussed. Next, the linearity is reconsidered given the new link budget. Subsequently, the baseband

amplifiers are improved and the impact on the common-mode level and clocking strategy of the receiver is treated.

Noise

The fundamental noise figure of a 50Ω 4-phase mixer into an ideal virtual ground is 3.9dB (section 4.2.1). However, the additional noise contribution of the vector modulator and the baseband amplifiers should be minimized. Since the VM is allowed to attenuate 20dB compared to the main mixer (i.e. the isolation of the antenna interface, see figure 5.1), better noise performance can be obtained by scaling it up from the 50Ω impedance standard. Doing so prevents us from interfacing the VM to a 50Ω environment, which is why the TX tap is brought on-chip in the new design, by sensing the output voltage of the TX using a capacitive divider. This will be detailed later.

Next, the available 20dB ‘attenuation budget’ must be carefully spent: Firstly, since the VM outputs a rotated square constellation (upper part of figure 5.3), 3dB is lost to make sure the VM can cancel the maximum SI across a full 360-degree circle. Secondly, the capacitive TX tap will be designed to attenuate by a factor of 2 to prevent the large voltage swings from damaging the VM devices. This causes another 6dB of attenuation. Finally, the TX tap will be in a 50Ω differential environment, owing to the design of its output matching network (detailed later). Cancellation will take place in a 100Ω differential environment. This subtracts another 3dB from the attenuation budget. As a result, $(20-3-3-6)=8\text{dB}$ attenuation is left for scaling up the VM impedance, allowing a VM impedance of $50 \times 10^{(8/20)} = 125\Omega$. Unfortunately, since the capacitive TX tap is a short for high frequencies and the VM essentially senses the TX output *voltage*, unlike the previous design there is no 50Ω source impedance in series with the noisy VM, resulting in limited NF improvement.

In an effort to further reduce the NF, we observe that the previous VM architecture contributed most of its noise at its minimum amplitude setting, as detailed in section 4.2.1. This occurs since by design, all slices are always on, contributing noise even at minimum signal contribution (effectively acting as a noisy differential shunt resistance at minimum amplitude setting). In the new design, additional switches were added to disable slices, shorting their outputs differentially to the common mode voltage, in order to keep TX tap loading by the VM impedance constant. The performance of this common-mode voltage is not critical in this differential architecture, but simply mixing to a floating short will result in an ill-defined voltage if a large number of slices is disabled.

One issue of the previous design was its requirement of 31 slices x 2 (double-balanced) x 4 phases x 4 multiplexer switches = 992 switches (top part of figure 5.2). This made the VM a physically very wide structure, with the associated difficulties

	Old meas.	New sim.
VM disabled	6.3 dB	5.2 dB
VM maximum	10.3 dB	8.9 dB
VM minimum	12.3 dB	5.2 dB

Table 5.2: Measured RX NF of the previous design v.s. simulated RX NF of the new design at 2.5 GHz operating frequency.

in interconnecting the cancellation node with a low series resistance. The new design would even require 1240 multiplexer switches, making it more impractical. However, adding the off state to the slices adds new points to the constellation, resulting in 3dB better resolution, or $\sqrt{2}$ times less slices for the same resolution (bottom part of figure 5.3). As a result, the number of slices was reduced from 31 to 24. To further reduce the number of multiplexer switches, a new VM scheme was adopted, with only three switches per slice: 0, 90 degrees or off (i.e. a differential short). The bottom part of figure 5.2 shows this new architecture. The result is a VM covering only a quarter of the constellation; the other three quarters are covered by LO rotation. The LO rotator will be covered in section 5.1.3. As a result of these modifications, the VM now requires only $24 \times 2 \times 4 \times 3 = 576$ switches.

In contrast to the previous design, the worst-case noise figure now occurs at the highest VM output amplitude. At the lowest amplitude setting, all VM slices mix into the differential short and the noise contribution of the VM is negligible.

One final way to reduce the NF is to decrease the input-referred noise of the BB amplifiers by slightly increasing the bias in their input stages: their transconductance was increased by a factor 1.5 in this design.

Table 5.2 compares measurement results of the previous design to simulation results of the new design, at 2.5GHz operating frequency. In simulation, all blocks were transistor implementations on schematic level.

In other words, the half-duplex NF of the design has improved from 6.3dB in measurements to 5.2dB in simulation, and the NF degradation in FD operation has improved from 4-6dB in measurements to 0-3.7dB in simulation. The worst-case FD NF has improved from 12.3dB to 8.9dB. For lower operating frequencies, the performance further improves to 4.8dB half-duplex NF with 0-2.9dB degradation in FD mode. The deterioration at higher clock frequencies is presumably caused by the more pronounced dynamic effects of clock overlap compared to the previous design, as the switch resistance has been considerably lowered for increased linearity, which is discussed next.

Figure 5.2: a) previous architecture of the vector modulator, b) new architecture, with reduced number of slices, reduced number of switches and LO rotation and slice disable functionality added.

Figure 5.3: Constellation and cancellation mechanism of the two vector modulator architectures for an equal number of 3 slices: a) the previous VM architecture, b) the new VM architecture with slice disable functionality, showing the increased resolution.

Figure 5.4: a) linearity requirements for the previously achieved and targeted link budget, b) associated switch resistance requirement.

Linearity

As seen in table 5.1, to prevent intermodulation products above the noise floor, the linearity has to be substantially improved with respect to the previous design, to 41dBm effective IIP3. This is further illustrated in the left plot of figure 5.4, which was taken from section 4.2.1.

As a consequence, the distribution of the 50Ω over the switch on-resistance and its preceding matching resistor needs to be modified. The right plot of figure 5.4 indicates that on-resistances below 10Ω are required, preceded by roughly 40Ω of matching resistance.

As such, the NMOS switches were designed for 3Ω on-resistance, preceded by 40Ω matching resistors (116Ω for the VM). This leaves some margin for nonzero interconnect resistance at RF.

Whereas for desired signals, the virtual ground is allowed to have a non-zero impedance, for self-interference the baseband nodes of the two mixers should be interconnected with very low series resistance. This way, the SI-currents will not generate swing on the baseband nodes. This is where we benefit from the new, smaller VM layout and the ultra-thick metal layer of the TSMC 65nm process. Using these, the VM was connected to the RX mixer with under 0.5Ω series resistance.

Ideally, very high IM2 values in excess of 70dBm are also required (figure 5.4), necessitating IIP2 calibration to prevent envelope detection products from appearing above the noise floor. However, for lack of time, best effort was made during lay-out to establish high IIP2, but no tuning techniques were implemented. As long as the IIP2 is sufficiently high to allow 27dB cancellation, its effects can be investigated and possibly corrected while implementing digital cancellation.

Baseband amplifier

As shown in section 4.3.8, the previous RX design had only 24dB maximum gain setting by design. This raised comments on whether the reported NF could be maintained during further amplification of the output to drive an ADC. For this reason, the new receiver was designed with a higher maximum gain.

However, at 27dB cancellation of -4dBm SI at the RX input (table 5.1), a competitive 40dB overall conversion gain would already result in 9dBm SI at baseband¹, corresponding to 1.8Vpp signal swing across 50Ω , even without margins for high peak-to-average ratios and non-ideal cancellation.

Since such signal swings easily cause output clipping of the 1.2V-supply BB amplifiers of the previous design, the amplifiers were redesigned based on thick-oxide I/O transistors, operating from a 2.4V supply. Adopting the push-pull output stage from the previous design, a differential output swing close to 4.8V was obtained.

To address the high flicker noise corner of the previous design, the input stage was adapted to a PMOS input pair in moderate inversion. To reduce the overall NF, the input transconductance was increased by roughly 1.5x at a slight power penalty. Obviously, the 2x higher supply voltage will also result in a doubling of the power consumption.

In most other aspects, the op amp is an impedance-scaled version of the previous design. Its schematic is shown in figure 5.5. The cross-coupled conductance technique to linearize the baseband section was no longer implemented in this design, since under SI-cancellation, the mixers are expected to dominate the overall linearity.

The previous design suffered from a significant peaking in virtual ground impedance beyond the signal bandwidth, reducing the linearity for nearby blockers. In this design, the baseband capacitance was therefore increased to 100pF per branch. For reliability, thick-oxide MOS capacitors were used, with a fringe stack on top for slightly higher capacitance density. Although large in area, this amount of baseband capacitance is a direct consequence of low-pass filtering in a 50Ω environment at tens of MHz bandwidth. However, if a wider bandwidth op amp is constructed, or reduced robustness to nearby blockers can be tolerated, it can be reduced.

Next, the feedback network can be designed based on the desired gain and bandwidth range. A 4-bit tuneable range of 1...16kOhm was selected for the feedback resistor. Aiming for 20MHz BW across this gain range, a 4-bit tuneable feedback capacitor range of 0.5...7.5pF was chosen.

As the loop gain of the baseband TIA crosses unity far above the cutoff frequencies of the baseband capacitors and the feedback network, its stability is determined by the capacitive division of the feedback and baseband capacitors (figure 5.5). As such,

¹For simplicity, we ignore the fact that conversion gain of a direct-conversion receiver is poorly defined [21] and we assume the baseband swing to be present across a 50Ω load.

Figure 5.5: Schematic of the two-stage operational amplifier used in the new design. The input stage is telescopic, with a PMOS input pair for low flicker noise. The output stage is push-pull. All transistors are 2.5V thick oxide. Indirect compensation is used [58].

the compensation capacitor needs to be tuned along with the feedback capacitor. To obtain a 60-degree phase margin across all feedback capacitor settings, a 4-bit tuneable compensation capacitor of 0.25...3.75pF was implemented.

An important consequence of the new BB design is an increased common-mode voltage of 1.2V. This has an impact on the RX mixer and vector modulator design.

Common mode level and clocking

To accommodate the higher BB common mode level, some additional modifications to the RX and VM mixers were needed. The drain and source of the mixer switches now reside at the 1.2V V_{dd} , allowing the gates to be driven to $2 \times V_{dd}$ instead of $1.5 \times V_{dd}$, for minimum on-resistance. As such, for robustness, the cross-coupled LO level shifting switches (section 4.2.1) were replaced by high-voltage devices.

Since the targeted technology did not offer high-density MIM capacitors, native- V_t MOS capacitors were used for area-efficient AC coupling of the clock from the last driver stages to the mixer gates².

The multiplexer switches in the VM are statically controlled thin-oxide switches (for reduced size and on-resistance). They will suffer from reliability issues if driven

²Hugo Westerveld is acknowledged for proposing this method.

Figure 5.6: Detailed schematic of a VM slice in the new design. Notable changes include the clamps to protect the switches from high swings, the PMOS multiplexer devices, the reduced amount of switches and the decoding logic for 2-bit individual slice control.

at $2 \times V_{dd}$ if the mixers are still in their start-up phase. Therefore, they were implemented as PMOS switches driven to ground.

In contrast to the previous design, the 24 VM slices cannot be easily grouped in a binary fashion, since full flexibility is required in order to reach every coordinate in a quarter constellation: an arbitrary number of the 24 slices can be enabled, and of these enabled slices, an arbitrary number can be steered to 0 degrees and the remaining slices are steered to 90 degrees. Fortunately, 24 slices is a manageable amount for individual control. As shown in figure 5.6, some logic was added locally to each slice to reduce the amount of control bits from 3×24 to 2×24 . The logic uses 3 identical NAND gates for convenient layout, and each multiplexer switch is driven from a local logic gate, essentially re-using the gates as buffers.

A clamp to common mode was added after the resistors (figure 5.6), to protect the mixer devices from the high voltage swings present at the TX tap, in case there is no clock signal present at the vector modulator. These swings can reach as much as $2.5V_p$ -p single-ended as the PA output peaks at 24dBm.

Figure 5.7: Schematic of the complete TX chain: Upconversion mixer, driver amplifier, power stage and output matching network + off-chip balun. The two feedback taps are embedded in the matching network.

5.1.2 Transmitter

To set a target for the TX power, literature suggests that CMOS power amplifiers can achieve saturated powers up to 24dBm before complex techniques such as power combining become essential to prevent excessive losses in parasitic resistances [59]. As such, 24dBm output-referred 1dB compression point was targeted. Furthermore, in 802.11-like OFDM systems, 8dB back-off is required due to the high peak-to-average ratio of the OFDM signals involved. This prevents excessive clipping of the signal peaks, which causes EVM, deteriorating the bit error rate [21]. At 8dB backoff from 24dBm CP1dB, the average TX power will be 16dBm.

To achieve reasonable efficiency, without excessive complexity, a textbook-based pseudo-differential class-AB design was adopted [21]. The schematic of the transmit chain is shown in figure 5.7.

The power stage is a pseudo-differential cascoded common-source stage. The gain devices are thin-oxide, the cascodes are 3.3V thick-oxide devices.

The PA is driven by another pseudo-differential stage, in class A to ensure its linearity is not the bottleneck in the TX chain. This driver stage is inductively loaded, and DC-coupled to the PA using a low-frequency tracking loop, to avoid the added parasitics associated with AC coupling, which would easily double the capacitive loading of the driver [21]. The center frequency and quality factor of the loading tank are tuneable to compensate the frequency behavior of the entire TX chain for process

variations.

The driver amplifier is driven by a complementary passive mixer in the sampling mode. Like in the previous design, direct quadrature voltage modulation [54] is used for its low-EVM capabilities.

Power amplifier

Sizing the TX chain starts with the power stage. The thick oxide cascodes allow a safe differential output swing of roughly 6Vp-p. Expecting at most 70% passive efficiency from the output matching network and off-chip balun (discussed later), we target 25.5dBm CP1dB at the drains of the output stage. Achieving 25.5dBm at 6V swing requires a 12.5Ω differential input impedance of the matching network.

Along the lines of [21], we can observe that a single-ended output current of 480mA is required into 6.25Ω . The single-ended PA input swing is limited to 0.6Vpp, keeping some margin in case higher driving levels are required in practice. This results in a transconductance of 0.8S for each branch, or 1.6S in total. In practice, 2.4S was required to push the 1dB compression point of the PA to the desired 25.5dBm at the drain nodes. To achieve this, the PA was biased at 470mA from a 3V-supply, dissipating 1.4W, with an expected drain efficiency of roughly 20%. This corresponds to the 30% drain efficiency of the textbook design [21], multiplied by the 70% passive efficiency expected from a practical output network. The peak efficiency and peak output power should be quite comparable to typical CMOS Class-AB designs with on-chip matching networks [60].

The transconductance devices and cascodes are kept at minimum lengths, which are $0.06\mu\text{m}$ and $0.5\mu\text{m}$, respectively. Choosing the gate bias voltage at 650mV, the width of the transconductance devices was sized to obtain the aforementioned 470mA bias current and 2.4S of total transconductance. The cascode devices were then be sized to keep the stack in saturation across the widest output swing.

As a consequence of the design for large backoff, the PA operates in class A for most of its output power range; for peak output, it operates slightly in class AB.

Driver amplifier

With the PA sized, the required output swing of the driver amplifier is known ($2 \times 0.6 = 1.2\text{Vpp}$ differential). For reasons detailed in section 5.1.2, the PA operates across a very wide fractional bandwidth of 1.5GHz to 4GHz, and the driver amplifier should match this response. Therefore, the quality factor of the driver amplifier load inductor can be chosen deliberately low for small area ($Q=8.5$ in this design), and needs to be lowered even further, close to a value of $(\text{BW}/\text{Fc}) = (2.75/2.5) = 1.1$. To accomplish this over process variations, a 2-bit tuneable parallel resistor was added to the tank,

with a nominal value of 100Ω differential and about 30% tuning range in both directions. To tune the resonance frequency of the load to compensate for process spread, a 2-bit tuneable capacitor was also added to the tank with a range of 0 to 2.25pF.

The fixed part of the Q-lowering resistor is used for common-mode sensing to control the tracking loop. The tracking amplifier consists of a single differential pair. Its reference is a diode-connected transistor that is located close to the PA transconductors. This will ensure stable biasing across temperature, which is essential since the power stage will heat up significantly.

The driver was operated in class A, to prevent it from limiting the linearity of the TX chain. The driver was biased at 32mA from the 1.2V supply, for a transconductance of 83mS per side and 12dB voltage gain. This ensures sufficient gain to drive the PA inputs at safe driver input swings, even as the operating frequency moves away from the resonance frequency of the load.

Upconversion mixer

Like in the previous design, the upconversion mixer is a passive mixer operating in the sampling region. This is also referred to as direct quadrature voltage modulation [54]. The mixer was designed for half-supply (600mV) common-mode level and AC coupled to the PA driver to make the driver bias independently tuneable.

As described in the previous section, when moving away from resonance, the driver amplifier will have to be driven at quite high input swings to still obtain the saturated output power of the transmitter. Such swings will modulate the on-resistance of the sampling mixer switches, leading to incomplete settling, especially at higher frequencies. For a more constant on-resistance of the sampling mixer without the need for bootstrapping, a complementary (transmission gate) upconversion mixer was used.

To absorb charge injection peaks, each switch pair was preceded by a first-order RC lowpass filter [54]. To prevent signal loss, the resistance of this filter should be significantly lower than the switched-capacitor resistor seen when looking into the upconversion mixer, which is in the order of $R_{swC} = \frac{1}{C_{in,driver}f_{clock}}$. The input capacitance of the driver is approximately 140fF per branch, which is roughly doubled to 280fF by the 10% parasitic of the 10x larger AC coupling capacitor. As such, at 4GHz switching frequency, $R_{swC} \approx 893\Omega$. Since the capacitor is not reset to a fixed voltage in between samples, the effective value seen from e.g. the I branch may further decrease depending on the voltage present in the Q branch. Therefore, based on simulations, the lowpass filters were designed 19 Ω and 5pF, resulting in a cutoff of 1.7 GHz.

The sampling switches were sized to achieve settling within 25% of the minimum $1/(4\text{GHz})$ sampling time, or within 62.5ps. Targeting roughly 10 RC time constants

Figure 5.8: Principle of cartesian feedback linearization.

for settling (corresponding to roughly 14-bit settling), with $C = 280\text{fF}$, a switch resistance of $R_{\text{on}} = 26\Omega$ was selected for the NMOS switch. Preferring balanced charge injection over constant on-resistance, the PMOS was sized equally large.

According to [54], the thermal spot noise of the upconversion mixer (ignoring phase noise) can be calculated as $N_{\text{mixer}}[\text{dBc/Hz}] = 10 \times \log \frac{4kTR_{\text{on}}}{V_p^2} + \text{PAR}$, where V_p is the single-ended peak swing of the sampling mixer output, and PAR is the peak-to-average ratio in dB of the modulated signal. Simulations indicate that typically 200mV single-ended peak voltage from the mixer will drive the PA into compression. At 8dB PAR and 26Ω mixer resistance, this results in -166dBc/Hz of spot noise, or -93dBc in 20MHz. At 16dBm average output power, this equals -109dBc, which is more than sufficient in-band performance for full-duplex applications. Thermal noise is generally more troublesome out-of-band, if filterless co-existence of several standards is desired [54].

Cartesian feedback option

Assuming a weakly nonlinear model for the PA, as a rule of thumb, its OIP3 will be about 10dB above its CP1dB, or, considering 8dB backoff for 802.11, at 18dB above the average operating power. Therefore, on average, its 3rd-order intermodulation products lie at -36dBc. As shown in table 5.1, about 60dB of linear digital cancellation is desired, so preferably the TX intermodulation should be at -60dBc, requiring 60-36=24 dB linearization. This corresponds to increasing the OIP3 from 34 dBm to the 46dBm listed in table 5.1. Therefore, in the demonstrator, we aim to facilitate Cartesian feedback linearization of the PA.

The principle of Cartesian feedback linearization is shown in figure 5.8. A feedback loop is wrapped around the TX chain comprising a TX tap, downmixer, subtraction point and low-pass loop filter in baseband. The principle is called Cartesian since in a direct conversion receiver, the baseband components are in complex baseband (I and Q).

As the signal propagates through the TX chain, it experiences phase shift, causing

signals that originated in the I path to mix down into the Q path and vice versa. This can cause loop instability. Therefore, a prerequisite for Cartesian feedback linearization is careful phase alignment between upconversion and downconversion mixer [61]. In our design, this will be accomplished by introducing a second vector modulator downmixer, dedicated to the optional Cartesian feedback loop.

As detailed in [61], roughly 10 degrees of phase alignment resolution is sufficient for negligible stability degradation in a typical Cartesian feedback system, which relaxes the resolution requirement compared to our SI-cancelling VM. Therefore, this vector modulator was constructed of 12 slices, which gives approximately 6 degrees of phase resolution, reducing the physical width of the VM by roughly a factor of 2 compared to the SI-cancelling VM.

In this architecture, the subtraction point will be a virtual ground and subtraction takes place in the current domain. This means that in feedback configuration, the overall TX chain has a current input. This means that the feedback VM transconductance sets the closed loop transimpedance of the entire TX chain. Additionally, the feedback VM should be highly linear as its distortion will not be suppressed by the loop gain.

In practice, the baseband signal will be supplied by a 50Ω -signal source. Typical commercial PA's offer 30dB gain from such a source to the output load, therefore this is also targeted as closed-loop gain in our system. However, terminating 50Ω directly into the virtual ground results in a 20mS input transconductance. For 30dB closed-loop gain, this requires a large feedback transconductance, resulting in very low impedance levels for the feedback VM and the first loop filter stage, and an unnecessarily low thermal noise output from the PA.

Based on these insights, compared with explorative loop gain and noise simulations, the Cartesian feedback VM impedance was scaled. A slice of this VM is shown in figure 5.9. Since higher attenuation is permitted than in the SI-cancelling VM, the capacitive TX tap attenuates by a factor of 5 to keep the VM input swings within the supply, obviating the need for clamping diodes. The VM resistors were chosen $6.71k\Omega$ for each of the 12 slices, such that $R_{match} = 560\Omega$.

In this linearizing loop, the ultimate achievable OIP3 of the TX chain is determined by the IIP3 of the feedback path. Aiming for the 46dBm TX OIP3 from table 5.1, section 4.2.1 shows that the required switch resistance is very similar to that of the SI-cancelling VM. Therefore, the switch size was kept unchanged, allowing re-use of this part of the VM slice design.

The upconverter input is designed for half-supply (600mV) common-mode level, and since this CM level propagates through the loopfilter to the subtraction point, the VM will be designed for that level as well. This means that in contrast to the SI-canceller, this VM can have NMOS multiplexer switches.

Taking some margin over the required 24dB, 30dB linearization is targeted across

Figure 5.9: Detailed schematic of a slice of the cartesian feedback VM. Compared to the SI-cancelling VM, this VM consists of 12 slices. It has no input clamp, it has NMOS multiplexer switches and NOR-based 2-bit decoding logic instead of NAND.

the full bandwidth and its IM3 components. Since the targeted 16MHz RF signal bandwidth extends from DC to 8MHz in baseband, its IM3 products will extend to 24MHz. In summary, 30dB loop gain in a 24MHz bandwidth is desired from the cartesian loop.

A major factor in the stability of this loop is group delay of the PA chain. From a complex baseband perspective, group delay manifests itself as time delay τ_g , i.e. an added phase shift that increases linearly with increasing offset frequency. For loop stability, the best that can be achieved from the loop filter around the unity gain crossing is a first-order roll-off, or 90 degrees phase shift. Targeting e.g. a 45-degree phase margin, $180-90-45 = 45$ degrees are allowed for the time delay, putting the maximum unity gain frequency f_0 of the loop at $f_0 = \frac{1}{8} \times \frac{1}{\tau_g}$.

A 24MHz bandwidth at 30dB loop gain with first-order roll-off would cross unity loop gain after $(30/20)=1.5$ decade, or at $10^{1.5} \times 24\text{MHz} = 759\text{MHz}$. This would permit 160ps of group delay, which was found not to be a feasible value during PA design. Even with very wide bandwidth matching, group delays in the order of 300ps were found to be more realistic. Thus, a higher order loop filter would be required, with added zeros to return to first-order roll-off before the zero crossing.

Although the loop filter was not implemented in the scope of this thesis, this illustrates the importance of maintaining low group delay, to realize 30dB linearization with a loop filter of a limited order and complexity.

Matching network

A wideband output match has two advantages:

- It permits a relatively flat PA group delay over a wide band, allowing a low-order Cartesian feedback loop filter,
- It increases the overlap in operating frequencies of TX and SI-cancelling RX, allowing them to operate jointly over a wider range of center frequencies.

Drawbacks are increased complexity, loss and component count of the matching network, and less inherent filtering of the spurious PA output such as harmonics. This work favors a wideband match, for an operating frequency range that better complements the RX, as detailed below.

Design of the matching network started from the constraints:

- 12.5 ohm differential impedance seen from the PA output
- 1.6-2.0nH estimated bondwire impedance
- 50 ohm load

A center frequency of roughly 2.5 GHz is targeted. Since no transformers were available in the design kit, LC-type matching networks were implemented.

Using a 1.8nH bondwire directly with an off-chip shunt capacitor as part of a single LC-lowpass matching section would result in a second-order equiripple filter response. To optimize the match at $f=2.5\text{GHz}$, the 2nd-order Chebyshev polynomial $T_2(\frac{f}{f_c}) = 2(\frac{f}{f_c})^2 - 1$ can be set equal to zero, which would result in a cutoff frequency f_c around $\sqrt{2} \times 2.5\text{GHz} = 3.5\text{GHz}$ [62]. Applying the other constraints, i.e. a source resistance of 6.25Ω and a series inductance of 1.8nH, we can scale the filter back to its prototype and find that the normalized inductance is 3.5H [62]. A second-order equiripple filter with this normalized inductance has an impedance transformation ratio of roughly 16 [63], resulting in about 200 ohm differential output match. Using an off-chip balun with an impedance ratio of 4, this could be a suitable matching network. However, tolerating e.g. 1dB passband ripple of this matching network will result in a relatively narrowband match with only 30% relative bandwidth [63]. This work prefers a larger operating frequency range of RX and TX, to obtain a more flexible, more ‘software-defined’ FD front-end.

To widen the matching BW, a second, on-chip LC section was added. Tolerating again 1dB passband ripple, a second-order match achieves almost 100% relative bandwidth with an impedance ratio of about 8x. The second LC-section of this filter uses an inductor of $\frac{0.5 \times 8}{2\pi \times 2.5\text{GHz}} = 2.5\text{nH}$, close to the estimated value of the bond wire inductance. The final off-chip impedance will be $8 \times 12.5 = 100\Omega$ differential, allowing the use of standardized transmission lines and an off-chip balun with an impedance ratio of 2. As such, this matching network was selected for the implementation. The

impedance level before the bondwires will be approximately the geometric mean of the two impedances, or $\sqrt{12.5 \times 100} = 35\Omega$. Finally, two more inductors were used to feed the supply to the PA drain nodes. These inductors were tuned to resonate with the drain capacitances to provide a real output impedance.

Manually optimizing the matching network, it was found that adding some explicit capacitance to the drain node helped to reduce the impedance seen at the second harmonic to suitable values. The feed inductor was modified to resonate with the added capacitance and restore the desired match in the band of interest.

Since the series inductors and feed inductors need to be high-Q to obtain high passive efficiency (a Q of 14 was used), a center-tapped feed inductor would save area. Unfortunately, the center tap of the provided inductors had an electromigration issue when feeding high DC current, necessitating a split inductor.

The Minicircuits TCM1-43X+ was found to be a suitable balun³. Its 3-port S-parameters were obtained from the manufacturer and applied in the schematic. Package parasitics and PCB traces were also estimated and added to the schematic. This allowed co-optimization of the matching network and the off-chip balun. Among other effects, it was found that adding some series capacitance and a slight length of PCB transmission line towards the balun improved the bandwidth of the match when pessimistic bond wire inductances and realistic balun frequency response was taken into account⁴.

After layout, parasitics on the drain and bond pad nodes were extracted and estimated, and subtracted from the explicit capacitances implemented at these nodes.

Figure 5.10 shows simulated S-parameters of the matching network from the 12.5 Ω PA drain node to the 50 Ω antenna. The simulation setup includes the PA drain capacitances, extracted on-chip parasitics, real inductor models, estimated bondwire inductance, pad and pin capacitances, off-chip series capacitance and transmission lines and 3-port S-parameters of the balun.

The S₂₁ has a plateau at -2dB insertion loss, or 63% passive efficiency, indicating our initial estimate of 70% was reasonable.

At this point, before proceeding to integration in the larger system, the driver amplifier and PA were taped out as a stand-alone block, with a 50 Ω terminated, AC-coupled input. Measurement results are presented in section 5.2.1.

Next, for the full system, part of the matching capacitance was split up and dedicated to the required transmit taps.

³Despite its nominal impedance ratio of 1, the S-parameters provided by the manufacturer indicated a ratio closer to 2.

⁴Ali Ghahremani is gratefully acknowledged for his assistance in optimizing the matching network.

Figure 5.10: Simulated S-parameters of the matching network, from 12.5Ω differential PA drains to single-ended 50Ω antenna load.

Transmit taps for cancelling and linearizing

Two TX taps are required: one for the cancelling VM and one for the Cartesian FB VM. Both taps will be voltage-sensing capacitive dividers, absorbed in the pad capacitance after the first matching stage, as shown in figure 5.7. At this stage, the impedance is 50Ω differential, and therefore the differential swing will closely represent the single-ended signal swings at the 50Ω antenna.

As detailed in 5.1.1, the TX tap for the cancelling VM will attenuate capacitively by a factor 2. Dedicating C_{tap1} of the matching capacitance to the tap will require two capacitors of $2C_{\text{tap1}}$ in series as illustrated in figure 5.11. So to limit the area overhead, care has to be taken how much of the matching capacitance will be dedicated to this tap. The 126Ω VM will be driven by a source impedance of $4C_{\text{tap1}}$ in parallel. To keep the highpass corner well below the lowest PA operating frequency, C_{tap1} was chosen 1.5pF .

Similarly, the 560Ω Cartesian feedback VM will be driven by a $5\times$ attenuating tap. Dedicating C_{tap2} of the matching capacitance to this tap will require one cap of $\frac{5}{4}C_{\text{tap2}}$ in series with four caps of $\frac{5}{4}C_{\text{tap2}}$, and will present $\frac{25}{4}C_{\text{tap2}}$ as driving impedance to the VM. In this case, C_{tap2} was chosen to be 160fF .

5.1.3 Clock divider and LO tree

To generate the four-phase clock, the on-chip divide-by-two structure of the previous design was re-used (not shown), with two main modifications.

Figure 5.11: Close-up of the matching network showing the distribution of the matching capacitance over the transmit taps.

First, since the RX and VM mixer switches have much lower on-resistance, guaranteed non-overlapping clocks with a low crossing point are essential to prevent shorting the baseband nodes briefly during transitions, deteriorating the noise performance. Since the original LO divider performed logic recombination of 4-phase, 50% duty cycle clocks to obtain 25% duty cycle versions, guaranteed non-overlapping 25% duty cycle signals were easily obtained by inserting one additional gate delay before logic recombination⁵.

Second, since the TX uses a complementary upconversion mixer (or transmission gate mixer), complementary clocks were required, i.e. both 75% and 25% duty cycle. This was achieved by combining the same set of 50% duty cycle signals not only by NAND gates, but also by NOR gates.

The LO rotator used for the coarse 4-quadrant rotation of both VM's is shown in figure 5.12. The static inverters controlling the transmission gates were deliberately combined into centralized inverters, to allow cancellation of charge injections of the LO edges onto the control lines by the complementary LO edges, reducing the required inverter size. The rotator includes functionality to disable the clock to the corresponding VM.

⁵Hugo Westerveld is acknowledged for this modification.

Figure 5.13: Complete overview of the stand-alone PA+driver with 50Ω-input.

Figure 5.14: Die photo of the stand-alone PA+driver in TSMC 65nm CMOS.

Figure 5.15: Amplifier bonded in its QFN package and soldered to the PCB. Special care was taken to minimize the bond wire inductance using two bond wires for each output and placing the output pads on the edge of the multi-project die (right side).

Figure 5.16: Complete PCB containing amplifier, matching network and off-chip balun, and headers for supply, bias and control.

onto the PCB. The PCB was designed in ADS, and EM-simulations of the transmission lines were co-simulated with models of the output matching network and the balun to verify the matching network from PA drains up to the single-ended 50Ω output. The complete PCB is shown in figure 5.16.

Figure 5.17 shows preliminary measurements of gain and group delay, compared to simulation results. Port extension was used on the network analyzer to de-embed the loss and delay of the PCB transmission lines connecting to the PA inputs and

Figure 5.17: Measured and simulated S21 (gain) and group delay of the stand-alone power amplifier.

balun outputs (the balun itself was not de-embedded). The driver capacitance and Q-lowering resistance were set to the same values in simulation and measurement. For most of the range, the gain of the PA falls approximately 4dB behind the simulated value. Group delay shows a good fit for most of the operating range thanks to the port extension, with a slight discrepancy in the center of the operating band.

Figure 5.18 presents the measured 1dB-compression point (CP1dB) versus frequency. Similar to the gain, the measured CP1dB was roughly 4dB lower than the 24dBm expected from simulations. Suspecting the off-chip balun losses to deteriorate the CP1dB, figure 5.18 includes some rough transient-based CP1dB simulations that include the S-parameters of the off-chip balun, provided by the manufacturer and valid from 1MHz to 7GHz. The frequency range of these simulations was intentionally limited to the range of interest. In the region around 3GHz, where the measured gain of the PA best matches the simulated values, we also observe a good match in CP1dB, indicating the off-chip balun, while meeting its manufacturer specifications, is indeed responsible for part of the output power loss. Unfortunately, a 4dB power loss dramatically reduces the already limited PA efficiency at the P1dB (roughly 20%) by another factor of 2.5.

In conclusion, despite careful co-simulations of the IC, PCB and off-chip balun, the PA performance is promising, but not completely up to par with simulation results. Many effects can be of influence here, such as input and supply bond wires, transistor degradation due to excessive voltage stress or soldering heat, high operating temperature due to poor thermal conductance to the package and underestimated parasitics. However, the measured performance should be sufficient to realize the significantly increased linear link budget in the complete system.

Figure 5.18: Measured and simulated 1dB-compression point of the stand-alone power amplifier.

Figure 5.19: Completed layout of the improved SI-cancelling front-end in TSMC 65nm CMOS.

5.2.2 Complete system simulations

Figure 5.19 shows the completed lay-out of the system shown in 5.1, measuring 1×1.7 mm. Significant blocks are indicated. This section presents some relevant simulated results. First, the noise performance is discussed, followed by the achievable RX linearity. Finally the cancellation is given some attention. No measurement results are presented within the scope of this thesis, because of the complexity of such measurements and time constraints.

Figure 5.20: Simulated gain and noise figure of the complete receiver v.s. base-band frequency.

Noise

Figure 5.20 shows the simulated noise performance and gain of the SI-cancelling receiver. All blocks are schematic-level implementations without layout parasitics and the center frequency is 2.5GHz. An RX-path gain setting of 35dB was chosen and is maintained over the 20MHz signal bandwidth. The flicker noise corner was reduced to 100kHz compared to the 2MHz of the previous design, owing to the large-sized PMOS input stage of the BB amplifier. The VM no longer contributes most of its noise in the center of the constellation, but instead is completely off in that case, reducing the NF from roughly 9dB to 5dB.

Linearity

As a result of the new baseband amplifier design and its increased supply voltage, the in-band IIP3 for desired signals was simulated to be 13dBm, a 4dB improvement with respect to the previous design without negative-conductance linearization.

Simulating the high effective IIP3 of the front-end under cancellation conditions is more troublesome, since it depends on many modeling degrees of freedom in the mixers, such as:

- The resistor model used for the switched resistor (ideal, TSMC, corrected TSMC et cetera),
- The transistor model used (TSMC BSIM, with or without correcting for its discontinuities at $V_{ds} = 0V$),

- The amount of capacitance included on the RF node in between resistor and switches,
- Vector modulator setting and associated partial distortion cancellation.

Depending on these and other factors, 31 to 37dBm effective IIP3 was observed in transistor-level simulation, with a combination of the most elaborate models resulting in 34dBm. Earlier experiences with this distribution of resistor and switch on-resistance showed effective IIP3's of 41dBm using this configuration [64]. However, even the worst-case 31dBm IIP3 should result in a significant link budget improvement with respect to the previous design, as detailed in section 5.2.3.

Cancellation

The cancellation mechanism of the new VM architecture was verified in simulation. To prevent the erroneous interchange of LO phases which was present in the previous design, the LO tree downstream from the fork to the RX and VM was kept untouched afterwards.

5.2.3 Results overview

Table 5.3 summarizes the simulated and preliminary measurement results and compares them to the design targets. Despite the reduced PA output power and possibly pessimistic linearity results from simulations, a linear link budget in excess of 100dB should still be feasible.

5.3 Future work

Due to time constraints in the design process, some aspects of the design were left for further investigation by experimentation. Since measurements were not yet performed at the time of writing, these aspects are mentioned here.

First, the impact of finite mixer IIP2 and image rejection due to mismatch in the mixers was not studied or simulated. These effects are expected not to limit analog cancellation, but may appear during digital cancellation. Although calibration and cancellation techniques for these impairments are well-known, none were implemented in this design.

Second, the lack of harmonic rejection in the RX and VM may result in the front-end receiving harmonic output of the power amplifier. This may limit the operating frequency range of the system to TX frequencies where harmonics are sufficiently filtered, or may require modifications to the output matching network.

Third, the main performance metric of the design is effective IIP3 under cancellation. However, without TX linearization, the TX IM3 will mask the RX/VM IM3.

	New design (target)	New design (sim./PA meas.)
TX power	16 dBm (24dBm P1dB)	12 dBm (20dBm P1dB)
Noise figure	8 dB	8.9 dB
Baseband BW	16.25 MHz	16.25 MHz
Antenna ISO	20 dB	20 dB
Noise floor	-94 dBm	-93 dBm
Link budget	107 dB	102 dB
SI at RX in	-4 dBm	-8 dBm
IIP3 w.r.t. SI	41 dBm	31-37 dBm
VM resolution	27 dB	27 dB
1/f noise corner	200 KHz	100 KHz
Linear digital SIC potential	60 dB	55 dB
TX OIP3	46 dBm	N/A
Power consumption	100mW (RX/Canc) + 1.2W (TX)	120mW (RX/Canc.) + 1.4W (TX)

Table 5.3: Targeted performance of the new design and achieved performance based on simulation results and preliminary measurements of the stand-alone PA.

In practice, the PA will be linearized by pre-distortion or Cartesian feedback, but in simulations this was not implemented, and the RX effective IIP3 was obtained with an idealized TX. Therefore, TX-RX interactions may be present in practice which were not accounted for in simulations.

5.4 Conclusion

This chapter presented the design procedure and simulation results and preliminary measurements of an improved full-duplex front-end based on the concept presented in chapter 4. Measurements on the complete system were not performed within the time span of this thesis work.

Targeting a robust demonstrator front-end with an increased linear link budget in excess of 100dB, the new front-end was designed for reduced noise figure and increased transmit power.

To lower the noise figure, the VM and transmit tap were taken on-chip and scaled away from the 50-ohm external standard. Also, the baseband amplifiers were designed for 1.5x higher input transconductance to lower their noise at the cost of power con-

sumption. The noise figure was reduced from from 10-12.3dB in measurement to 5.2-8.9dB in simulation.

To increase the transmit power to over 20dBm, a two-stage class-AB power amplifier was co-integrated. For wideband operation, a two-stage matching network including the bond wire was designed, with the capacitive transmit taps absorbed into the shunt capacitors of the matching network.

An additional vector modulator was implemented to permit Cartesian feedback linearization of the on-chip transmitter for reduced EVM, targeting 46dBm TX OIP3, which should help to achieve the proposed link budget using mostly linear digital cancellation in a later stage.

To accommodate the higher transmit power, the distribution of switch resistance and matching resistance in the mixers was reconsidered, resulting in increased effective IIP3 of the receiver from 21.5dBm in measurements to 31-37dBm in simulation.

The baseband amplifiers were modified for lower flicker noise, higher baseband gain and larger output swings, and the impact of their increased common mode level on the VM and RX mixer was addressed.

Preliminary measurements on the stand-alone power amplifier show full functionality, but measured efficiency and peak output power are not yet fully on par with simulations and under investigation.

Points that require further investigation in practice are the IIP2 performance, harmonic rejection and image rejection of the implemented front-end and their impact on the achievable link budget.

The overall system is expected to realize a linear link budget in excess of 100dB and allow investigation of full-duplex point-to-point links based on this architecture.

Chapter 6

Conclusions

This chapter concludes and summarizes this thesis, and lists the original contributions present in this work. It proceeds with recommendations for future work.

6.1 Summary and conclusions

In-band full-duplex wireless communication (FD), i.e. transmission and reception at the same time at the same frequency, is an emerging research topic, driven by the ever increasing demand for mobile data traffic in the crowded radio spectrum (chapter 1). Besides a theoretical doubling of the spectral efficiency, the inherent channel reciprocity is an attractive physical layer aspect. In higher network layers, additional advantages are being explored such as collision prevention, low latency, security and simplified frequency planning.

The main issue in FD is strong self-interference (SI) from the transmitter into the local receiver. In typical links, the transmitted signal is in excess of 90dB above the system noise floor, necessitating over 90dB total SI-rejection to fully compete with half-duplex links. Chapter 2 presents a system-level study of how the FD paradigm influences the requirements on different building blocks of a typical radio transceiver. These requirements were derived under the premise of a fully linear link budget, meaning that SI-induced distortion should not deteriorate the system noise floor. Along with SI-cancellation in the digital baseband, in excess of 40dB SI-rejection should be achieved in the analog domain to be competitive to half-duplex links with feasible component requirements.

Applying just passive isolation at antenna level is generally not sufficient to robustly achieve 40dB SI-rejection in analog (chapter 3), which is especially true for applications that suffer from a varying antenna near-field, such as mobile devices. For this reason, SI-cancellation techniques across multiple domains have to be combined.

Achieving competitive link budgets with high integration potential, compact form factor, limited complexity and low power consumption remains challenging (chapter 3).

This work was one of the first to study the feasibility of FD using a custom designed CMOS front-end, as opposed to using commercially available components. This allowed access to all of the internal signals and domains to explore new SI-cancelling concepts. In chapter 4, full implementation details and analysis are presented of a 65-nanometer mixer-first front-end with a vector modulator downmixer for SI-cancellation. Using the implemented front-end as a research vehicle, several transceiver impairments that may limit its full-duplex operation were experimentally investigated, such as distortion, phase noise, image rejection and transmitter impairments.

The receiver was found to have over 90dB linear link budget potential in a 16.25MHz bandwidth, when combined with only 20dB worst-case antenna isolation, thanks to its 21.5dBm effective IIP3 for SI present at the receiver input. It offers up to 27dB cancellation of up to -16.4dBm SI at the RX input without generating distortion above the noise floor. The co-integrated transmitter was found to have almost sufficient performance to support this link budget, requiring minor linearization by e.g. pre-distortion. Transmitter and SI-cancelling receiver operate jointly over a wide range of center frequencies from 0.15 to 3.5GHz.

Due to the limited availability of functional samples of the first front-end, this research did not proceed to setting up a point-to-point link with digital cancellation. Rather, an improved second front-end was developed in 65nm CMOS to investigate the limits of the proposed architecture and to support future research (chapter 5). It was designed for over 100dB linear link budget, supporting a significantly higher transmit power (from 3.6 to 16dBm average output) and reduced noise floor (from 12.3 to 8dB worst-case noise figure). This required re-design of the cancelling and RX mixers for an effective IIP3 in the range of 30-40dBm. A two-stage class-AB power amplifier with partial on-chip matching was implemented to support the higher transmit power over a wide bandwidth. Also, the baseband section was revised to provide higher gain and lower noise.

Preliminary measurements on the stand-alone power amplifier demonstrate its functionality, but the measured gain and compression point are not up to par with simulations across the full 1.5-4GHz bandwidth, reducing the PA efficiency. Characterization is ongoing, but the current performance should still allow a link budget exceeding 100dB in the integrated system, which is being manufactured at the time of writing. Transistor-level simulations of the full system showed the reduced NF, increased linearity and expected cancellation behavior.

Overall, this thesis has demonstrated that vector-modulator downmixers are a promising building block for SI-cancelling full-duplex front-ends. Owing to their high

linearity, the proposed architecture can improve upon a low and varying isolation at antenna-level without generating SI-induced distortion above the noise floor, to enable highly integrated full-duplex radios that can compete with relaxed half-duplex links.

6.2 Original contributions

This thesis contains several original contributions to the field:

- A system-level discussion of the impact of the full-duplex wireless paradigm on the sub-block specifications of a typical transceiver (Chapter 2 and [15]),
- A survey and high-level comparison of existing transceiver architectures for full-duplex wireless (Chapter 3),
- The introduction of switched-resistor vector modulator downmixers as a highly linear means of cancelling self-interference in integrated (full-duplex) front-ends (Chapter 4 and [47–49]),
- Full design and implementation details of a front-end with the proposed vector modulator downmixer as an SI-cancelling component (Chapter 4 and [49]),
- Systematic experimental studies of front-end impairments that may limit full-duplex link budgets and discussions of their impact (Chapter 4 and [47–49]),
- Full implementation details of an improved front-end with an increased linear link budget based on the switched-resistor vector modulator downmixer as a SI-cancelling component (Chapter 5).

6.3 Recommendations

Although increasing spectral scarcity has sparked renewed interest in full-duplex radio, this research direction is meeting rightful skepticism. Being competitive with half-duplex link budgets with limited additional resources and complexity remains challenging. Inter-compatibility of SI-cancellers with MIMO, a throughput improvement technique which is already widely deployed, is another open issue.

That said, this work has demonstrated a relevant new SI-cancelling front-end architecture and put it into perspective with other recent developments. It has theoretically and experimentally investigated many aspects of the implemented front-ends that may prevent digital SI-cancellation from reaching the noise floor. However, the ultimate cancellation potential can only be verified by actually implementing digital cancellation, which was not done in the scope of this thesis but should be the next step in this research. Another aspect that was insufficiently treated in this work is even-order nonlinearity and image rejection. The impact of these impairments are

expected to show when attempting digital cancellation. Although calibration techniques are widely known, they were not implemented in the front-ends designed in the scope of this thesis.

The presented phase/attenuation-based canceller is limited in its analog cancellation by ambient reflections that are frequency-selective. As a result, as the presented architecture is pushed towards larger link budgets, the analog cancellation will not improve and the burden of the larger link budget is put on components further down the transmitter and receiver paths. In chapter 5, this resulted in measurement-equipment-grade requirements for these blocks, reducing competitiveness with half-duplex. Therefore, in real-world applications, wideband cancellation techniques that implement (group) delay to at least partially suppress ambient reflections are a necessity. Several were briefly addressed in chapter 3.

Throughout this work, the importance of a linear link budget was emphasized, since especially SI-induced RX distortion is a complex interplay between SI and desired signal that cannot be compensated easily in the digital domain. However, recent developments have shown that using a non-linear channel estimate, even this kind of distortion may be compensated in digital [46]. While not convincingly demonstrated for modulated signals, this could shed a new light on the link budget potential of the presented architecture, as it would allow the RX to be driven almost into compression by the SI.

Finally, owing to the passive impedance matching and the presence of two noise-contributing paths into the virtual ground, the presented design is fundamentally limited in its noise figure. Other works apply noise-cancelling techniques in various paths for significantly reduced noise figures [35, 44, 46], however none of these achieve comparable linearity with respect to self-interference. Future research could address how to combine the best of both worlds.

Appendix A

Errata to published material

This appendix lists errata and updated results with respect to published material, that have been incorporated the body of this thesis, wherever parts of these publications have been re-used.

Errata to: B. Debaillie, **D.-J. van den Broek**, C. Lavin, B. van Liempd, E. Klumperink, C. Palacios, J. Craninckx, B. Nauta, and A. Parssinen, “Analog/RF solutions enabling compact full-duplex radios,” in Selected Areas in Communications (JSAC), IEEE Journal on, vol. 32, no. 9, pp. 1662-1673, 2014:

1. The $CP_{1dB,RX}$ in table 3 is correctly expressed as $IIP3_{RX}-10$ [dBm], but this expression was incorrectly evaluated in the table body. From left to right, the values -20, -40, -30, -50, -40, -60, -30 have been replaced by 0, -1, -11, -41, -22, -52, -10
2. The DAC DR in table 3 incorrectly refers to an equation, while it is derived from TX EVM and margin as $DAC\ DR = -EVM_{TX} + \text{margin}$. This is correctly evaluated in the table body.

Errata to: **D.-J. van den Broek**, E. Klumperink, and B. Nauta, “An In-Band Full-Duplex Radio Receiver with a Passive Vector Modulator Downmixer for Self-Interference Cancellation,” in IEEE Journal of Solid-State Circuits (JSSC), 2015, vol. 32, no. 12, pp. 3003–3014.

1. Section III-C:

“Therefore, every 1 dB of cancellation of the SI would result in a 2 dB reduction of the SI-induced IM3, boosting the effective IIP3 by 1 dB.”

This is incorrect, it should be:

*“Therefore, every 1 dB of cancellation of the SI would result in a **3 dB** reduction of the SI-induced IM3, boosting the effective IIP3 by **1.5 dB**.”*

2. Section IV-C:

“The fact that the IIP3 does not increase by the full 27 dB indicates that the linearity bottleneck has moved from the TIA to the nonlinear RX and VM switches.”

This should be, in correspondence with the previous correction:

*“The fact that the IIP3 does not increase by the full **1.5×27 dB (section III-C)** indicates that the linearity bottleneck has moved from the TIA to the nonlinear RX and VM switches.”*

3. Section III-C:

$$a_3 = (-R_s R_{sw}^2) / (2V_{OD}^2 (R_{sw} + R_s)^5)$$

This should be:

$$a_3 = -(R_s + R_m) R_{sw}^2 / (2V_{OD}^2 (R_{sw} + (R_s + R_m))^5)$$

The plotted curve is nonetheless valid.

4. A square was missing in equation 3. Original equation:

$$\begin{aligned} F = 1 + \frac{R_m + R_{sw}}{R_s} + \frac{R_{sh}}{R_s} \left(\frac{R_s + R_m + R_{sw}}{R_{sh}} \right)^2 \\ + \gamma \frac{R_f}{R_s} \left(\frac{R_s + R_m + R_{sw}}{\gamma R_f} \right) \\ + \gamma \frac{v_{n,amp}^2}{4kTR_s} \left(\frac{R_s + R_m + R_{sw}}{\gamma R_f} + \frac{R_s + R_m + R_{sw} + R_{sh}}{R_{sh}} \right)^2 \end{aligned} \quad (A.1)$$

Corrected equation:

$$\begin{aligned} F = 1 + \frac{R_m + R_{sw}}{R_s} + \frac{R_{sh}}{R_s} \left(\frac{R_s + R_m + R_{sw}}{R_{sh}} \right)^2 \\ + \gamma \frac{R_f}{R_s} \left(\frac{R_s + R_m + R_{sw}}{\gamma R_f} \right)^2 \\ + \gamma \frac{v_{n,amp}^2}{4kTR_s} \left(\frac{R_s + R_m + R_{sw}}{\gamma R_f} + \frac{R_s + R_m + R_{sw} + R_{sh}}{R_{sh}} \right)^2 \end{aligned} \quad (A.2)$$

Dankwoord

De afgelopen vier jaar heb ik in een hechte vakgroep mogen groeien van een verse masterstudent tot iemand die zich voorzichtig chipontwerper mag noemen, met dit proefschrift als eindresultaat. Zo'n promotie wordt geen succes zonder de hulp van vele mensen om je heen.

Allereerst gaat veel dank uit naar mijn dagelijkse begeleider Eric, een begeleider die niet snel vragen voor je zal beantwoorden, maar eerder helpt je vragen verder aan te scherpen. Dat was wennen voor mij, als praktisch ingesteld persoon op zoek naar antwoorden, maar uiteindelijk van onschatbare waarde. En natuurlijk bedank ik ook Bram, voor het in de eerste plaats mogelijk maken van deze opdracht en voor het bij elkaar brengen en houden van deze gepassioneerde chipontwerpgroep.

My gratitude also goes to the other members of my graduation committee: Frank, for his introduction to microwave engineering and his help when applying this to my designs. Aarno Pärssinen, for his initial role in the DUPLO project and for his suggestions regarding my ISSCC submission. Peter Baltus, for his cheerful appearance at, among other occasions, ICTOpen. And Mark Bantum, for also taking the time to scrutinize my thesis and his help with early antenna experiments.

The first 2.5 years I had the pleasure of collaborating in the fruitful European project DUPLO. Meeting up with colleagues from all over Europe really helped to define and shape this research, with the nice benefit of visiting Oulu, Santander, Paris, Surrey and Leuven (and sampling their local cuisines :)), and welcoming the entire consortium in Twente. Many thanks to our chairman Kari, work-package leader Björn, fellow “Roosendaler” Barend, and Visa, Wei-Li, Alok, Cristina, Carmen, Hirley, Ari, Mir, Pei and Jawad.

The whole PhD would not be a success without a wonderful bunch of colleagues and former colleagues. Many thanks to Bram, Mark, Jasper, Johan, Erik, Hugo, Remko, Claudia, Harish, Y-C, Joeri, Ali, Inês, Labrinus, Chris, Bart, Maikel, Zhiliang, Vijay, Sajad, Daniël, Dawei, Milad, Lorenzo, Nori and Reda. Special thanks to former PhD's Dlovan, Saqib and Michiel, since I learned a great deal from measuring their existing chips and digging into their designs. And our neighbouring group, semiconductor components, only added to a pleasant time at the UT: thanks to Jurri-

aan, Cora, Ray, Sander, Tom, Dirk, Rob, Alexey, Remke, Satadal, Sourish, Oguzhan, Gaurav, Jiahui, Mengdi, Elmer, Arnoud, Svetlana, and former SC-colleagues TomvH, Boni and Anubha. And from TE I would like to highlight Arjan Meijerink.

Bedankt Gerdien, het stille brein achter alle drukke agenda's en communicatiestromen. Gerard, bedankt voor je ondersteuning richting de fabricage van de ontwerpen en Henk voor alle hulp als de chips eenmaal binnen waren. En Anne-Johan en Ronan, die me lang geleden het eerste stukje universitaire elektrotechniek bijbrachten: Het was erg leuk om jullie later weer bij te staan in onderwijs en voorlichting, en bedankt voor jullie input als ik advies zocht over op-amps en devices. Ed, bedankt voor de bijdrage in mijn stageopdracht, alweer wat langer geleden.

Master- en bacheloropdrachten begeleiden biedt de mogelijkheid om zijsporen van je onderzoek te laten verkennen en geeft je nieuwe inzichten: Bedankt Tobias, Arno en Alexander voor jullie bijdragen, en hopelijk heb ik wat van mijn aanpak en enthousiasme naar jullie overgebracht. Bedankt Jelmer, Thijs, Jippe en Olivier voor jullie B2-opdracht die een veel gebruikte beamforming-demo heeft opgeleverd. Robert van Veldhoven van NXP, bedankt voor je vertrouwen in de groep en in mij, van masterstudent tot promovendus tot begeleider van Alexander.

De afgelopen jaren is mijn generatie flatgenoten van Residentie 7-2 een hechte vriendengroep gebleven, en hopelijk mag dat nog lang zo blijven. Thomas, Floris, Dirk, Coen, Arjan, Ires, Robert, Bert, Tom, Hugo en Jules, mooi dat we nog altijd regelmatig samenkomen en dat er altijd wel iemand in is voor een borrel en een goed gesprek. Over de borrel gesproken mag Jarno ook niet ontbreken als vaste vrijdagmiddagcrew.

Mijn familie, in het bijzonder mijn ouders, ontzettend bedankt dat ik destijds dankzij jullie op deze universiteit kon beginnen, dat jullie nog altijd de afstand naar Twente voor lief nemen en dat jullie mij, Jeroen en Marloes steunen in alle keuzes die we maken in het leven.

Lieve Inge, bedankt voor het doorstaan van de pieken en dalen die kwamen kijken bij mijn promotie, van weekenden doorwerken tegen een deadline tot periodes van vast zitten in een vervelend probleem. Fijn dat ik je af en toe eens wat uit mocht leggen over interferers en Smith charts, en dat jij mij er op jouw beurt uitlegt hoe de wereld wel meer moet zijn dan moleculen, reacties en stroompjes. En bedankt voor de leuke manier waarop je met mijn collega's en vrienden kunt opschieten.

Upon receiving the IEEE SSCS Benelux Student Chip Design award 2015, with the runners-up Niels van Thienen and Hui Jiang and Patrick Reynaert and Filip tavernier (SSCS).

Bibliography

- [1] Cisco report, “Cisco visual networking index: Global mobile data traffic forecast update, 20152020,” Cisco, Tech. Rep., February 2016.
- [2] ITU report, “Measuring the information society: 2012,” ITU, Tech. Rep., 2012.
- [3] 5G PPP Architecture Working Group, “View on 5g architecture,” 5G PPP, White paper, July 2016, version 1.0.
- [4] A. Sabharwal, P. Schniter, D. Guo, D. Bliss, S. Rangarajan, and R. Wichman, “In-band full-duplex wireless: Challenges and opportunities,” *IEEE J. Sel. Areas Commun.*, vol. 32, no. 9, pp. 1637–1652, 2014. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6832464>
- [5] D. Bliss, P. Parker, and A. Margetts, “Simultaneous transmission and reception for improved wireless network performance,” in *Proc. IEEE/SP 14th Workshop Statistical Signal Processing*, Aug. 2007, pp. 478–482.
- [6] X. Fang, D. Yang, and G. Xue, “Distributed algorithms for multipath routing in full-duplex wireless networks,” in *Proc. IEEE Eighth Int. Conf. Mobile Ad-Hoc and Sensor Systems*, Oct. 2011, pp. 102–111.
- [7] A. Sahai, G. Patel, and A. Sabharwal, “Asynchronous full-duplex wireless,” in *Proc. Fourth Int. Conf. Communication Systems and Networks (COMSNETS 2012)*, Jan. 2012, pp. 1–9.
- [8] P. Weeraddana, M. Codreanu, M. Latva-aho, and A. Ephremides, “The benefits from simultaneous transmission and reception in wireless networks,” in *Proc. IEEE Information Theory Workshop (ITW)*, Aug. 2010, pp. 1–5.
- [9] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, “Achieving single channel, full duplex wireless communication,” in *Proceedings of the sixteenth annual international conference on Mobile computing and networking*. ACM, 2010, pp. 1–12.

- [10] A. Goldsmith, *Wireless communications*. Cambridge university press, 2005.
- [11] P. C. Weeraddana, M. Codreanu, M. Latva-aho, and A. Ephremides, “On the effect of self-interference cancelation in multihop wireless networks,” *EURASIP Journal on Wireless Communications and Networking*, vol. 2010, no. 1, p. 1, 2010.
- [12] M. Duarte, “Full-duplex wireless: Design, implementation and characterization,” Ph.D. dissertation, Rice University, 2012.
- [13] D. Bharadia, E. McMillin, and S. Katti, “Full duplex radios,” *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, pp. 375–386, Aug. 2013. [Online]. Available: <http://doi.acm.org/10.1145/2534169.2486033>
- [14] “Eu fp7 project duplo, full-duplex radios for local access,” 2012-2015. [Online]. Available: <http://www.fp7-duplo.eu/>
- [15] B. Debaillie, D.-J. van den Broek, C. Lavin, B. van Liempd, E. Klumperink, C. Palacios, J. Craninckx, B. Nauta, and A. Parssinen, “Analog/RF solutions enabling compact full-duplex radios,” *IEEE J. Sel. Areas Commun.*, vol. 32, no. 9, pp. 1662–1673, 2014. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6832471>
- [16] I. Vassiliou, “Wireless transceiver system design for modern communication standards,” Tutorial in IEEE Int. Solid-State Circuits Conf., 2013.
- [17] D. Korpi, M. Valkama, T. Riihonen, and R. Wichman, “Implementation challenges in full-duplex radio transceivers,” in *Finnish URSI Convention on Radio Science*, 2013.
- [18] M. Duarte, A. Sabharwal, V. Aggarwal, R. Jana, K. Ramakrishnan, C. Rice, and N. Shankaranarayanan, “Design and characterization of a full-duplex multi-antenna system for WiFi networks,” *IEEE Transactions on Vehicular Technology*, vol. 63, no. 3, pp. 1160–1177, Mar. 2014.
- [19] IEEE 802.11 Standardization Board, “Ieee standard for information technology–telecommunications and information exchange between systems–local and metropolitan area networks–specific requirements–part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications,” *IEEE Std*, p. 91, 1999.
- [20] B. Debaillie, P. V. Wesemael, G. Vandersteen, and J. Craninckx, “Calibration of direct-conversion transceivers,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, no. 3, pp. 488–498, Jun. 2009.

-
- [21] B. Razavi, *RF microelectronics*. Prentice Hall New Jersey, 1998, vol. 1.
- [22] B. Come, R. Ness, S. Donnay, L. V. der Perre, W. Eberle, P. Wambacq, M. Engels, and I. Bolsens, “Impact of front-end non-idealities on bit error rate performance of WLAN-OFDM transceivers,” in *Proc. IEEE Radio and Wireless Conf. RAWCON 2000*, 2000, pp. 91–94.
- [23] S. Bensmida, K. Morris, J. Lees, P. Wright, J. Benedikt, P. Tasker, M. Beach, and J. McGeehan, “Power amplifier memory-less pre-distortion for 3GPP LTE application,” in *Proc. European Microwave Conf. EuMC 2009*, Sep. 2009, pp. 1433–1436.
- [24] J. Tsai, Y.-J. Chen, Y.-F. Lai, M. Shen, and P. Huang, “A 14-bit 200MS/s current-steering DAC achieving over 82dB SFDR with digitally-assisted calibration and dynamic matching techniques,” in *Proc. Int VLSI Design, Automation, and Test (VLSI-DAT) Symp*, Apr. 2012, pp. 1–4.
- [25] D. Mahrof, E. Klumperink, M. O. Alink, and B. Nauta, “A receiver with in-band iip3 >20dBm, exploiting cancelling of opamp finite-gain-induced distortion via negative conductance,” in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Jun. 2013, pp. 85–88.
- [26] A. Sahai, G. Patel, C. Dick, and A. Sabharwal, “On the impact of phase noise on active cancelation in wireless full-duplex,” *IEEE Trans. Veh. Technol.*, vol. 62, no. 9, pp. 4494–4510, 2013. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6523998>
- [27] —, “Understanding the impact of phase noise on active cancellation in wireless full-duplex,” in *Proc. Systems and Computers (ASILOMAR) 2012 Conf. Record of the Forty Sixth Asilomar Conf. Signals*, Nov. 2012, pp. 29–33.
- [28] L. Belov, S. Smolskiy, and V. Kochemasov, *Handbook of RF, Microwave, and Millimeter-wave Components*, ser. Artech House microwave library. Artech House, 2012.
- [29] WARP platform: Wireless Open Access Research Project, “Rice university full duplex communication example,” Online, 2013, visited October 7, 2016. [Online]. Available: https://warpproject.org/trac/wiki/Projects/Rice_FullDuplex
- [30] E. Everett, A. Sahai, and A. Sabharwal, “Passive self-interference suppression for full-duplex infrastructure nodes,” *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, pp. 680–694, 2014. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6702851>

- [31] M. Ghoraiishi, V. Tapio, A. Sethi, K. Rikkinen, D. Liang, G. Chen, and W. Li, “D3.1: Digital baseband solutions for full-duplex transceiver,” DUPLO, Tech. Rep., 2014.
- [32] B. Debaillie, C. Lavin, C. Palacios, B. Hershberg, B. van Liempd, E. Klumperink, and D.-J. van den Broek, “D2.2: Integration report of rf and antenna self-interference cancellation techniques,” DUPLO, Tech. Rep., 2014.
- [33] M. Rousstia and M. Herben, “High performance 60-GHz dielectric rod antenna with dual circular polarization,” in *Proc. European Radar Conf. (EuRAD)*, Oct. 2013, pp. 359–362.
- [34] T. Dinc, A. Chakrabarti, and H. Krishnaswamy, “A 60 GHz CMOS full-duplex transceiver and link with polarization-based antenna and RF cancellation,” *IEEE Journal of Solid-State Circuits*, vol. 51, no. 5, pp. 1125–1140, May 2016.
- [35] J. Zhou, T.-H. Chuang, T. Dinc, and H. Krishnaswamy, “Receiver with >20MHz bandwidth self-interference cancellation suitable for fdd, co-existence and full-duplex applications,” in *Solid- State Circuits Conference - (ISSCC), 2015 IEEE International*, 2015, pp. 1–3. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7063066>
- [36] B. Debaillie, D.-J. van den Broek, C. Lavin, B. van Liempd, E. Klumperink, C. Palacios, J. Craninckx, and B. Nauta, “RF self-interference reduction techniques for compact full duplex radios,” in *Proc. IEEE 81st Vehicular Technology Conf. (VTC Spring)*, May 2015, pp. 1–6.
- [37] C. Lavn, K. Rikkinen, A. Sethi, V. Tapio, L. Gonzlez, C. Lavn, R. Marante, B. Debaillie, M. Mikhael, H. Suys, B. van Liempd, M. Ghoraiishi, and H. Malik, “D5.2: Final proof-of-concept validation, results and analysis,” DUPLO, Tech. Rep., 2014.
- [38] A. Wegener, “Broadband near-field filters for simultaneous transmit and receive in a small two-dimensional array,” in *Microwave Symposium (IMS), 2014 IEEE MTT-S International*, 2014, pp. 1–3. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6848672>
- [39] T. Dinc and H. Krishnaswamy, “A t/r antenna pair with polarization-based reconfigurable wideband self-interference cancellation for simultaneous transmit and receive,” in *Microwave Symposium (IMS), 2015 IEEE MTT-S International*, 2015, pp. 1–4. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7167135>

-
- [40] M. Duarte, C. Dick, and A. Sabharwal, "Experiment-driven characterization of full-duplex wireless systems," *IEEE Trans. Wireless Commun.*, vol. 11, no. 12, pp. 4296–4307, 2012. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6353396>
- [41] B. van Liempd, B. Debaillie, J. Craninckx, C. Lavin, C. Palacios, S. Malotiaux, J. Long, D. van den Broek, and E. Klumperink, "RF self-interference cancellation for full-duplex," in *Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM), 2014 9th International Conference on*, 2014, pp. 526–531. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6849740>
- [42] L. Laughlin, C. Zhang, M. Beach, K. Morris, and J. Haine, "A widely tunable full duplex transceiver combining electrical balance isolation and active analog cancellation," in *Proc. IEEE 81st Vehicular Technology Conf. (VTC Spring)*, May 2015, pp. 1–5.
- [43] B. van Liempd, B. Hershberg, K. Raczkowski, S. Ariumi, U. Karthaus, K.-F. Bink, and J. Craninckx, "A +70dBm iip3 single-ended electrical-balance duplexer in 0.18um SOI CMOS," in *Solid- State Circuits Conference - (ISSCC), 2015 IEEE International*, 2015, pp. 1–3. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7062851>
- [44] D. Yang, H. Yüksel, and A. Molnar, "A wideband highly integrated and widely tunable transceiver for in-band full-duplex communication," *IEEE J. Solid-State Circuits*, vol. 50, no. 5, pp. 1189–1202, May 2015.
- [45] N. Reiskarimian and H. Krishnaswamy, "Magnetic-free non-reciprocity based on staggered commutation," *Nature communications*, vol. 7, 2016.
- [46] J. Zhou, N. Reiskarimian, and H. Krishnaswamy, "Receiver with integrated magnetic-free n-path-filter-based non-reciprocal circulator and baseband self-interference cancellation for full-duplex wireless," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Jan. 2016, pp. 178–180.
- [47] D.-J. van den Broek, E. Klumperink, and B. Nauta, "A self-interference-cancelling receiver for in-band full-duplex wireless with low distortion under cancellation of strong TX leakage," in *Solid- State Circuits Conference - (ISSCC), 2015 IEEE International*, 2015, pp. 1–3. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7063067>
- [48] —, "A self-interference cancelling front-end for in-band full-duplex wireless and its phase noise performance," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, May 2015, pp. 75–78.

- [49] D. van den Broek, E. Klumperink, and B. Nauta, "An in-band full-duplex radio receiver with a passive vector modulator downmixer for self-interference cancellation," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3003–3014, Dec. 2015.
- [50] D. Mahrof, E. Klumperink, Z. Ru, M. Oude Alink, and B. Nauta, "Cancellation of opamp virtual ground imperfections by a negative conductance applied to improve RF receiver linearity," *IEEE J. Solid-State Circuits*, vol. 49, no. 5, pp. 1112–1124, 2014. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6708481>
- [51] M. Soer, E. Klumperink, B. Nauta, and F. van Vliet, "A 1.0-to-2.5 GHz beamforming receiver with constant-gm vector modulator consuming < 9mW per antenna element in 65nm CMOS," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, 2014, pp. 66–67. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6757340>
- [52] C. Andrews and A. Molnar, "Implications of passive mixer transparency for impedance matching and noise figure in passive mixer-first receivers," *IEEE Trans. Circuits Syst. I*, vol. 57, no. 12, pp. 3092–3103, 2010. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=5518350>
- [53] A. Molnar and C. Andrews, "Impedance, filtering and noise in n-phase passive CMOS mixers," in *Custom Integrated Circuits Conference (CICC), 2012 IEEE*, 2012, pp. 1–8. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6330616>
- [54] X. He and J. van Sinderen, "A low-power, low-EVM, SAW-less WCDMA transmitter using direct quadrature voltage modulation," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 12, pp. 3448–3458, Dec. 2009.
- [55] J. Zhou, A. Chakrabarti, P. Kinget, and H. Krishnaswamy, "Low-noise active cancellation of transmitter leakage and transmitter noise in broadband wireless receivers for fdd/co-existence," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 3046–3062, 2014. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6922161>
- [56] T. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall PTR, 2001.
- [57] D. Korpi, L. Anttila, V. Syrjala, and M. Valkama, "Widely linear digital self-interference cancellation in direct-conversion full-duplex transceiver," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 9, pp. 1674–1687, 2014. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6832439>

- [58] B. Ahuja, “An improved frequency compensation technique for CMOS operational amplifiers,” *IEEE Journal of Solid-State Circuits*, vol. 18, no. 6, pp. 629–633, Dec. 1983.
- [59] A. Afsahi and L. Larson, “Monolithic power-combining techniques for Watt-level 2.4-GHz CMOS power amplifiers for WLAN applications,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 61, no. 3, pp. 1247–1260, Mar. 2013.
- [60] H. Wang, C. Sideris, and A. Hajimiri, “A 5.2-to-13GHz class-ab CMOS power amplifier with a 25.2dBm peak output power at 21.6% pae,” in *Proc. IEEE Int. Solid-State Circuits Conf. - (ISSCC)*, Feb. 2010, pp. 44–45.
- [61] A. Geurts, “Transmitter linearization by cartesian feedback in a full-duplex context,” Master’s thesis, University of Twente, Jul. 2015.
- [62] D. M. Pozar, *Microwave engineering*. John Wiley & Sons, 2009.
- [63] G. Matthaei, “Tables of Chebyshev impedance –transforming networks of low-pass filter form,” *Proceedings of the IEEE*, vol. 52, no. 8, pp. 939–963, Aug. 1964.
- [64] H. Westerveld, E. Klumperink, and B. Nauta, “A cross-coupled switch-RC mixer-first technique achieving +41dBm out-of-band IIP3,” in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, May 2016, pp. 246–249.

List of publications

Peer-reviewed

- B. Debaillie, **D.-J. van den Broek**, C. Lavin, B. van Liempd, E. Klumperink, C. Palacios, J. Craninckx, B. Nauta, and A. Parssinen, “Analog/RF solutions enabling compact full-duplex radios,” in *Selected Areas in Communications (JSAC)*, IEEE Journal on, vol. 32, no. 9, pp. 1662-1673, 2014.
- B. van Liempd, B. Debaillie, J. Craninckx, C. Lavin, C. Palacios, S. Malotau, J. Long, **D.-J. van den Broek**, and E. Klumperink, “RF self-interference cancellation for full-duplex,” in *Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM)*, 2014 9th International Conference on, 2014.
- **D.-J. van den Broek**, E. Klumperink, and B. Nauta, “A self-interference-cancelling receiver for in-band full-duplex wireless with low distortion under cancellation of strong TX leakage,” in *Solid-State Circuits Conference - (ISSCC)*, 2015 IEEE International, 2015, pp. 1-3.
- **D.-J. van den Broek**, E. Klumperink, and B. Nauta, “A self-interference cancelling front-end for in-band full-duplex wireless and its phase noise performance,” in *Radio Frequency Integrated Circuits Symposium (RFIC)*, 2015 IEEE, 2015.
- B. Debaillie, **D.-J. van den Broek**, C. Lavin, B. van Liempd, E. Klumperink, C. Palacios, J. Craninckx, and B. Nauta, “RF self-interference reduction techniques for compact full-duplex radios,” in *Vehicular Technology Conference (VTC) Spring*, 2015 IEEE, 2015.
- **D.-J. van den Broek**, E. Klumperink, and B. Nauta, “An In-Band Full-Duplex Radio Receiver with a Passive Vector Modulator Downmixer for Self-Interference Cancellation,” in *IEEE Journal of Solid-State Circuits (JSSC)*, 2015, vol. 32, no. 12, pp. 3003–3014, 2015.

- B. Debaillie, B. van Liempd, B. Hershberg, J. Craninckx, K. Rikkinen, **D.-J. van den Broek**, E.A.M. Klumperink, B. Nauta, “In-Band Full-Duplex Transceiver Technology for 5G Mobile Networks,” in European Solid State Circuits Conference (ESSCIRC), 2015 IEEE, 2015.
- M.C.M. Soer, E.A.M. Klumperink, **D.-J. van den Broek**, B. Nauta, and F.E. van Vliet, “A Beamformer with Constant-Gm Vector Modulators and its Spatial Intermodulation Distortion,” accepted for publication in IEEE Journal of Solid-State Circuits (JSSC).

Other

EU FP7 project DUPLO output:

- V. Tapio, W. Li, K. Rikkinen, P. Pirinen, C. Lavin, L. Gonzalez, B. Debaillie, B. van Liempd, E. Klumperink, **D.-J. van den Broek**, M. Ghorraishi, Y. Ko, H. Khalife, “DUPLO deliverable D1.1: System Scenarios and Technical Requirements for Full-Duplex Concept,” April 2013.
- B. Debaillie, C. Lavin, C. Palacios, B. van Liempd, E. Klumperink, **D.-J. van den Broek**, “DUPLO deliverable D2.1: Design and Measurement report for RF and antenna solutions for self-interference cancellation,” April 2014.
- B. Debaillie, C. Lavin, C. Palacios, B. Hershberg, B. van Liempd, E. Klumperink, **D.-J. van den Broek**, “DUPLO deliverable D2.2: Integration report of RF and antenna self-interference cancellation techniques,” October 2014.
- **D.-J. van den Broek**, C. Lavin, C. Palacios, B. Debaillie, B. van Liempd, E. Klumperink, “DUPLO deliverable D2.3.1: Transceiver Circuits Simulation, Implementation and Measurement report,” April 2014.
- **D.-J. van den Broek**, C. Lavin, B. Debaillie, B. van Liempd, E. Klumperink, B. Nauta, “DUPLO deliverable D2.3.2: Transceiver Circuits Simulation, Implementation and Measurement report,” April 2015.
- K. Rikkinen, A. Pouttu, V. Tapio, H. Alves, B. Debaillie, C. Lavin, M. Ghorraishi, J. Seddar, **D.-J. van den Broek**, “DUPLO deliverable D6.3: Full-Duplex Radios for Local Access, final report,” June 2015.
- **D.-J. van den Broek**, B. Nauta, K. Rikkinen, E. Klumperink, “DUPLO deliverable D6.4: Dissemination report,” May 2015.

And contributions to:

- E.A.M. Klumperink, “Interference Rejection exploiting Switch-R-C Techniques Compatible with CMOS,” invited talk at the 2015 IMS/RFIC Workshop on RF Interference Mitigation Techniques, Phoenix, USA, 17 May, 2015.
- B. Nauta, “Multiphase RF techniques in CMOS: Applied to beamforming and full duplex receivers: CICC 2015 educational session,” Custom Integrated Circuits Conference (CICC), 2015 IEEE, pp. 1-147, September 2015.
- B. Nauta, “In-Band Full Duplex: System Considerations and Transceiver Design,” Forum Presentation in Forum 3: Radio Architectures and Circuits Towards 5G, at Solid- State Circuits Conference - (ISSCC), 2016 IEEE International, 2016.