DUPLO Deliverable D2.3.2

Transceiver circuits simulation, implementation and measurement report

Project Number:	**316369**
Project Title	**Full-Duplex Radios for Local Access – DUPLO**
Deliverable Type:	**PU**

Contractual Date of Delivery:	**April 30, 2015**
Actual Date of Delivery:	**April 30, 2015**
Editor(s):	**Dirk-Jan van den Broek (UT)**
Author(s):	**Cristina Lavín (TTI);** **Björn Debaillie, Barend van Liempd (IMEC);** **Eric Klumperink, Bram Nauta, Dirk-Jan van den Broek (UT).**
Work package:	**WP2**
Estimated person months:	**19**
Security:	**PU**
Nature:	**Report**
Version:	**1.0**

Keyword list: full-duplex, transmitter, receiver, transceiver, self-interference, CMOS IC, RF circuits

Abstract: This document briefly revisits the requirements, challenges and bottlenecks set on RF transceiver components by the full-duplex transmission paradigm (based on deliverable 2.3.1). Next, it describes a a new FD transceiver topology that meets these requirements for a short-range, low-power link scenario. It includes analysis and implementation details of a prototype transceiver and its measured performance. Finally, future integration and improvement opportunities are discussed.

Executive Summary

This deliverable describes the technical results obtained within the DUPLO project in work package 2, task 2.3. Within the frame of this task, an extensive system-level study resulted in a new self-interference cancelling transceiver architecture, introduced in D2.3.1 [1]. Whereas D2.3.1 put the emphasis on transceiver requirements in different full-duplex scenarios, this document only briefly revisits the *selected* scenario, before presenting in-depth design and measurement details of the implemented transceiver prototype.

The prototype transceiver aims to implement low-power, short-range full-duplex wireless with a moderate (~20dB) initial antenna isolation as a starting point. This allows it to operate in hand-held devices that require a compact antenna and suffer from a strongly varying antenna near-field. The transceiver uses a vector modulator downmixer to simultaneously apply phase shift, attenuation and downmixing to a copy of the transmit signal and subtracts it in the digital baseband, before amplification. This resuits in up to 27dB cancellation, which relaxes the requirements on the transmitter and downstream receiver components by the same 27dB. The vector modulator and main receiver are based on highly linear switched-resistor mixers to prevent the self-interference from inducing distortion in the receiver path that may mask the desired signal when performing digital cancellation.

Assuming only 20 dB isolation from the antenna, the SI-cancelling receiver offers nearly 90 dB link budget in 16.25MHz bandwidth around 2.5 GHz. Thanks to its good linearity performance, this link budget is sufficient for links up to 100m line-of-sight distance. The co-integrated transmitter offers nearly sufficient performance at 2.5 GHz to also meet this link budget, requiring only minor improvements by e.g. pre-distortion or non-linear modeling. The transceiver offers sufficient immunity to phase noise to operate with a commercial clock source, when clocking all mixers in the system from a shared clock. The functionality of the transceiver is not limited to the 2.5 GHz band only: both TX and RX offer a wide operating frequency range of 0.15 to 3.5 GHz. The work on the prototype transceiver was well received by the scientific community and has been presented at major events in the field, such as the IEEE ISSCC, RFIC and VTC conferences [2-4]. The work was also invited to appear in IEEE Journal of Solid State Circuits (JSSC).

This document concludes by proposing several ways to proceed with this transceiver architecture, both research and product-oriented. This concludes the activities of DUPLO task 2.3.

Authors

Partner	Name	Email
TTI Norte, SL (TTI)	Cristina Lavín	clavin@ttinorte.es
Interuniversitair Micro-Electronica Centrum vzw (IMEC)	Björn Debaillie	bjorn.debaillie@imec.be
	Barend van Liempd	barend.vanliempd@imec.be
Universiteit Twente (UT)	Eric Klumperink	e.a.m.klumperink@utwente.nl
	Dirk-Jan van den Broek	j.d.a.vandenbroek@utwente.nl
	Bram Nauta	b.nauta@utwente.nl

Table of Contents

Acronyms and Abbreviations

ADC	analog-to-digital converter
BB	baseband
BW	bandwidth
CMOS	complementary metal oxide semiconductor
DAC	digital-to-analog converter
dB	decibel
dBc	power ratio in decibels referenced to power at carrier frequency
dBm	power in decibels referenced to one milli-watt
DR	dynamic range
DUPLO	Full-duplex radios for local access project
EVM	error vector magnitude
f_c	carrier frequency in MHz
FDD	frequency division duplex
FD	full-duplex
GHz	giga (10^9) Herz
IC	integrated circuit
IEEE	Institute of Electrical and Electronics Engineers
IIPn	n-th order input-referred intercept point
IMn	n-th order intermodulation products
ISM	industrial, scientific and medical
LNA	low-noise amplifier
MHz	mega (10^6) Herz
NF	noise figure
OFDM	orthogonal frequency-division multiplexing
OIPn	n-th order output referred intercept point
PA	power amplifier
PN	phase noise
QAM	quadrature amplitude modulation
RF	radio frequency
RX	receiver
SI	self-interference
SIC	self-interference cancellation
SINDR	self-interference-to-noise-and-distortion ratio
SNR	signal to noise ratio
TDD	time division duplex
TX	Transmitter
Wi-Fi	WLAN products that are based on the IEEE 802.11 standards
WLAN	wireless local area network
WP	work package

1. INTRODUCTION

As stated in DUPLO deliverable 1.1 [5] and elaborated in deliverable 2.1 [6], the DUPLO project aims to increase the spectral efficiency of wireless communication devices by investigating the in-band full-duplex concept and by proposing effective techniques to enable concurrent transmission and reception using the same spectral resources.

The main physical layer challenge in full-duplex wireless is self-interference (SI): since transmission and reception occur on the same frequency simultaneously, the transmitted signal easily leaks into the receiver path, creating very strong in-band interference. Depending on the link scenario, the transmit power level may range from roughly 80 to 120 dB stronger than the receiver noise floor, as discussed in D2.3.1 [1] and [7]. Filtering this self-interference is obviously not an option, therefore it has to be rejected by other mechanisms:

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 self-interference from the RX path using knowledge of the TX signal.

Ideally, the combined effect of isolation and cancellation should yield the 80-120 dB self-interference rejection, to ensure the noise floor is not degraded compared to an equivalent half-duplex system, so as to obtain the full advantage and competitiveness of full-duplex.

Task 2.3 investigates *self-interference cancellation techniques at integrated transceiver level*. Until recently, most FD state-of-the-art concentrated on enabling full-duplex using commercial off-the-shelf transceivers. Only recently, DUPLO task 2.3 as well as other research groups worldwide have presented work on modifying the integrated radio itself, specifically for in-band FD. Modifying the transceiver allows us to make use of signals that are normally internal to the radio and to easily take these signals across the boundaries of the different domains involved: digital baseband, analog baseband and RF, while maintaining the benefits of an integrated solution. This allows for unique and innovative self-interference cancellation topologies. DUPLO task 2.3 was among the first to present an integrated self-interference cancelling transceiver aimed specifically at FD, and this work was well received by the scientific community [2-4].

D2.3.1 [1] has described possible application scenarios for FD, and the resulting requirements on transceiver components. Section 2 will briefly re-visit these requirements, but only for the scenario that was selected for the implemented prototype. Section 3 provides implementation details, design choices and analysis of the implemented transceiver prototype. Section 4 presents the measured performance of the prototype receiver and transmitter, as well as the phase noise properties of the overall transceiver. Section 5 discusses further integration and validation opportunities of the current design. Section 6 and 6 provide details on how the design could be further improved and extended for application in full-duplex products. Finally, section 7 concludes the work of task 2.3.

2. SYSTEM CONSIDERATIONS AND PROPOSED ARCHITECURE

As stated in the introduction and detailed in D2.3.1 [1], the main issue in achieving in-band full-duplex (FD) wireless is strong in-band (same-channel) crosstalk from transmitter to receiver, referred to as self-interference (SI), see Figure 1a. Recovering the (much weaker) desired signal from a remote transmitter necessitates SI-isolation and cancellation. Cancellation uses knowledge of the transmit signal from various points in the TX chain to subtract SI in the RX chain (Figure 1b).

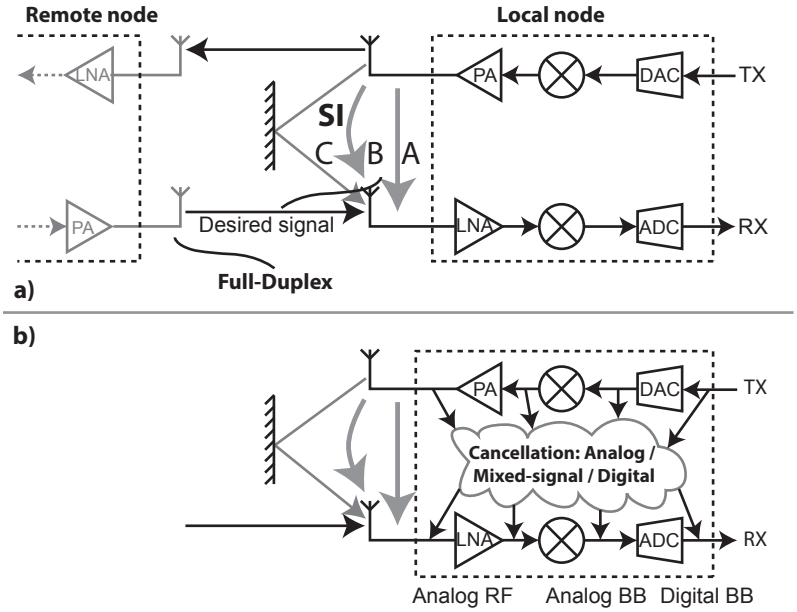

Figure 1: a) Generic view of a full-duplex link between a local and a remote node, subject to 3 types of self-interference: A) Electrical crosstalk between TX and RX, B) RF coupling due to limited antenna isolation and a varying antenna near-field, C) SI reflected by the environment. b) Generic view of SI-cancellation in a single FD node, from various points in the TX chain to various points in the RX chain.

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 detailed in D2.3.1. Figure 2 shows four recent approaches to SI-cancellation:

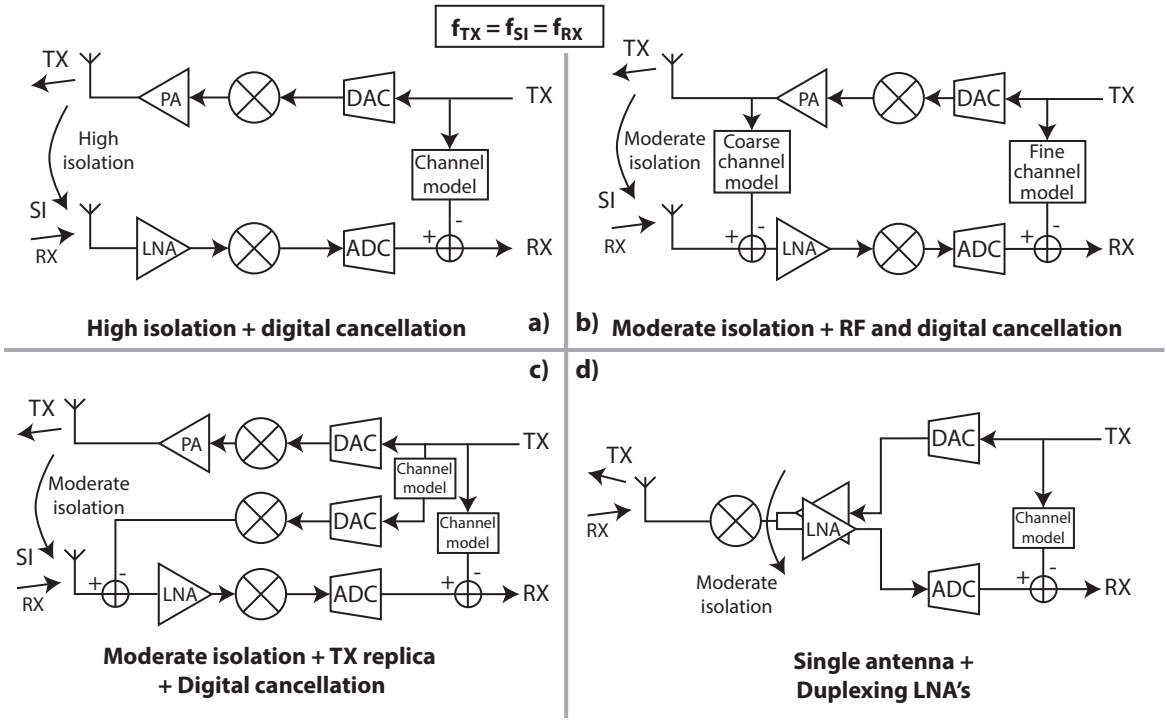

Figure 2: Four recent self-interference cancelling transceiver topologies for integrated full-duplex radios.

a) High isolation can be obtained at the antenna by sophisticated design, e.g. using cross-polarization [4], which is actively researched in DUPLO task 2.1. While successful in obtaining high isolation and very suitable in some application scenarios, it is difficult to combine robustness to a varying antenna near-field, compactness and frequency-agility. Another approach to turn a single-port antenna into a two-port circuit with high isolation, as investigated in task 2.2, is a tuneable electrical-balance-based duplexer using a hybrid transformer [4]. This can actively compensate for a varying near-field, and is to some extent frequency-agile.

b) Direct crosstalk as well as part of the reflected SI can be cancelled using an analog multi-tap filter at RF, combined with digital cancellation [8, 9]. This requires nanosecond-scale analog delays in its analog filter [8], which have only recently been integrated in the form of N-path filters [9]. This approach has potential to compete with high-end (802.11-style) half-duplex links [8], however, silicon / PCB area and power consumption remain high.

c) A replica TX chain can be used to regenerate the SI in the digital BB and cancel it at RF, combined with digital cancellation [10]. However, its ultimate cancellation performance is limited by uncorrelated (phase) noise and distortion sources between the two TX chains [11].

d) A mixer-first transceiver with baseband noise-cancelling, duplexing LNA's can be used that intrinsically copy a transmit signal to their antenna port, while rejecting it in their output [12]. Placing the LNA's in the baseband allows complex signal processing to tune their SI-rejection. Although very suitable for integration and capable of operating with a single-port antenna, the duplexing LNA's have limited capability to work with high TX powers [12].

As an alternative method, in D2.3.1 we developed an SI-cancelling transceiver for frequency-agile, low-power, short-range full-duplex. In this section, we briefly revisit the system considerations. The next sections proceed with more complete background information, implementation details, performance analysis and modelling of the design as presented by the end of D2.3.1. We then describe the measured performance of the implemented transceiver and relate this to the available full-duplex link budget.

2.1. System considerations

The SI-cancelling receiver developed in this task, aims to bring full-duplex to low-power, short-range communication devices. For this purpose, a TX power of 0dBm, a bandwidth of 16.25MHz (the active bandwidth of WLAN) and a 10dB RX noise figure are assumed. This results in an RX noise floor of roughly -90 dBm. Thus, in order not to degrade the noise floor, a combination of isolation and cancellation mechanisms 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, which requires a remaining 90-20 = 70dB from the cancellation.

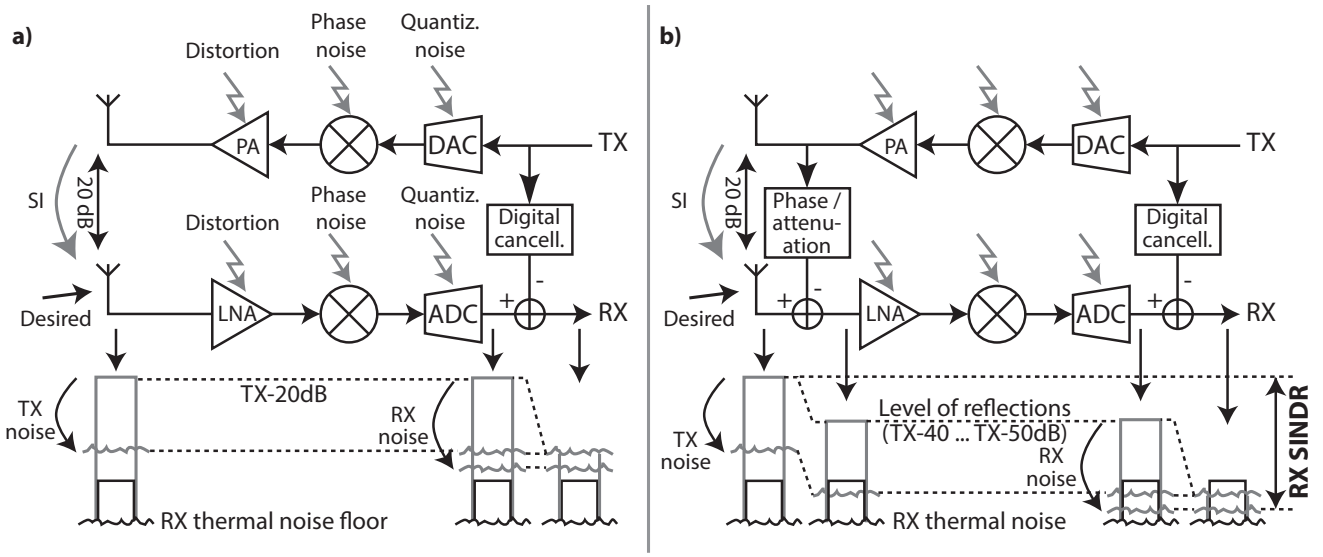

Figure 3: Full-duplex nodes assuming a moderate 20dB antenna isolation: a) The SI is affected by transmitter and receiver imperfections, which limit cancellation in the digital domain. b) A phase / amplitude based canceller can enhance a moderately isolating antenna in the analog domain, relaxing TX EVM, RX dynamic range and digital cancellation requirements.

Figure 3a visualizes an attempt to cancel the remaining SI after antenna isolation all in the digital domain. Assuming digital cancellation can only cancel the deterministic, linear part of the self-interference, TX noise and distortion and SI-induced RX noise and distortion may still mask the desired signal [7]. To prevent this, roughly 70dB TX EVM[1] and 70dB RX dynamic range 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 robustly improve RF SI-rejection to a level where reflections from the environment become the limiting factor (figure Figure 3b). To also conquer reflections in this architecture, the canceller would need to incorporate time delays (i.e. a non-flat phase response), which easily becomes costly in silicon area. However, in an indoor scenario, the reflections are typically present at -40 to -50 dB [13]. As such, a frequency-flat canceller at RF can already reduce requirements on TX EVM, RX dynamic range and digital cancellation to 90 – (40 to 50) = 40 to 50dB, which is much more feasible than 70dB.

Therefore, 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 (level of the reflections). 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

[1] EVM, error vector magnitude, is commonly used to group TX noise and distortion into a single metric.

configuration. So, the canceller should privide 20dB fixed attenuation, about 30dB variable attenuation, and a full 360° phase shift.

For the RX, it is important that it has a reasonably low noise floor, and that this noise floor is not increased by noise and distortion products induced by strong SI. In other words, its SI-to-noise-and-distortion-ratio (SINDR) should be high to maximize the full-duplex link budget. The concept of SINDR is illustrated in Figure 3b.

For the TX, its EVM should ideally be at -40 to -50 dB, such that the digital cancellation only has to cancel its linear components in order to push the SI down to the noise floor. Although -40 to -50 dB EVM is feasible, it is quite a stringent value to achieve in a low-power integrated transceiver. However, if the distortion mechanisms of the TX are well understood, its EVM can be lowered by means of pre-distortion, or its EVM can be accounted for in the digital cancellation.

2.2. Proposed architecture

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 4a). 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 4b).

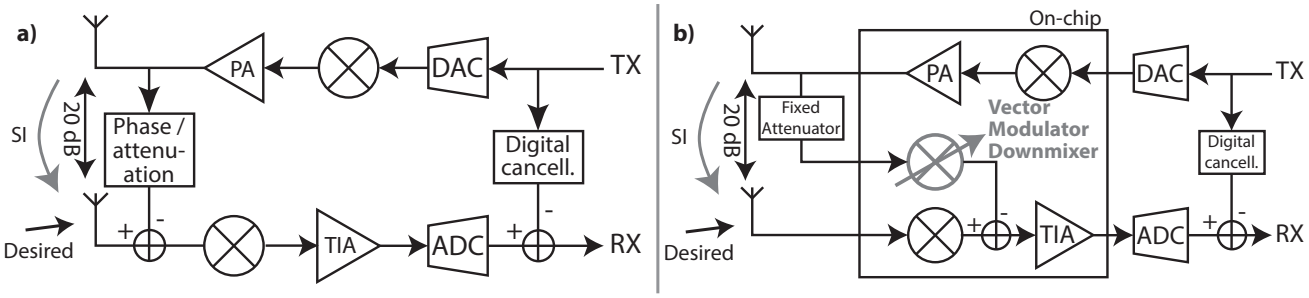

Figure 4: A generic phase shift/attenuation based canceller preceding a mixer-first receiver; b) The cancellation node is moved to the analog baseband and the phase shift, attenuation and down-mixing are combined into a vector modulator (VM) downmixer.

This topology taps the TX signal at the TX RF output, thus including TX impairments in the cancellation path, relaxing TX EVM requirements by the amount of cancellation achieved. It cancels SI before any active baseband amplification and before the 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 is kept external for versatility.

The topology in Figure 4b 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 transceiver prototype.

3. IMPLEMENTATION OF PROPOSED ARCHITECTURE

This section describes the implementation of an SI-cancelling transceiver in 65nm CMOS according to the topology of Figure 4b.

3.1. Receiver

As explained in section 2.2, 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 very 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 conventional FDD systems, there is no TX-RX frequency separation, so filtering cannot be used.

Hence, both the main RX and the VM are based on a highly linear passive mixer with series resistors into virtual ground nodes provided by transimpedance amplifiers (TIA's) [14]. Figure 5 shows an overview of the implemented receiver. The VM is a sliced version of the main RX, followed by static phase rotator switches that can rotate the current of each slice through the four virtual grounds. This way, the SI currents are diverted through highly linear passive networks and only the residual signal is amplified. The number of slices and other design details are motivated next.

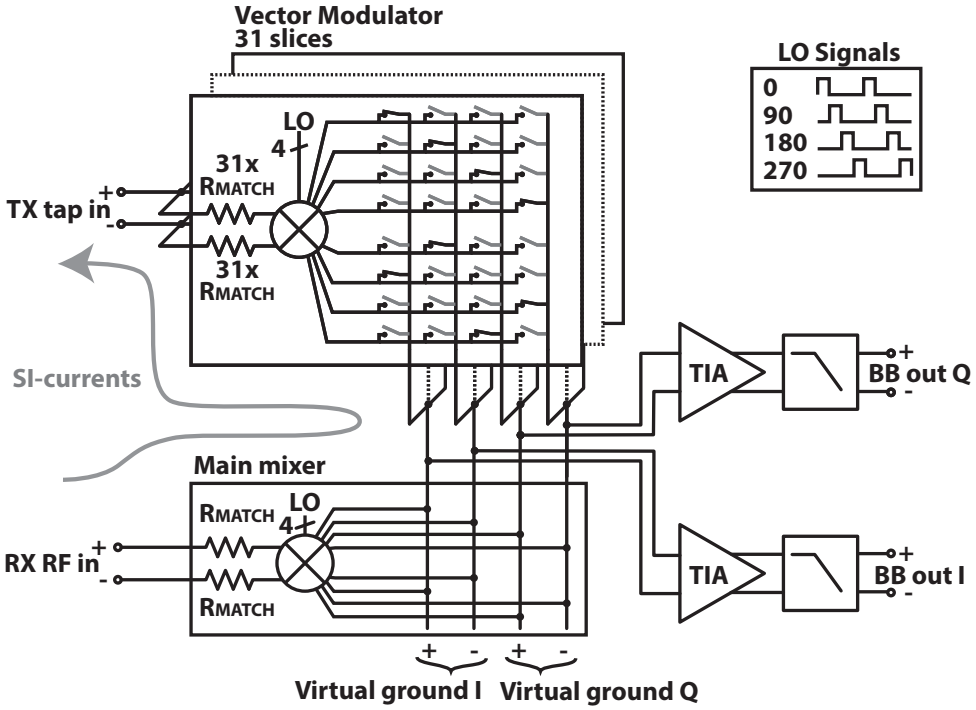

Figure 5: Implementation details of the SI-cancelling receiver. The vector modulator (top) is a 31-slice version of the main receiver (bottom), each slice followed by static phase rotator switches. The VM diverts self-interference currents through linear passive networks before amplification.

3.1.1. Resolution

The sliced VM principle is similar to the constant-GM vector modulator presented in [15], but implemented with *resistors to a virtual ground* rather than *active transconductors*. The amount of slices of the VM determines the number of phase / amplitude constellation points that can be covered and thus the amount of cancellation that can be achieved due to quantization effects. This is illustrated in Figure 6.

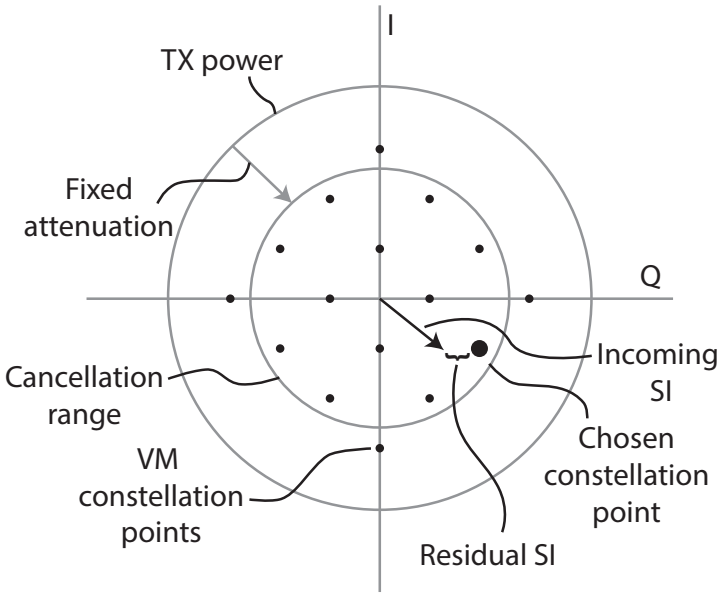

Figure 6: The cancellation principle explained using a 3-slice VM (i.e. a 4-by-4 constellation): a fixed attenuation maps the VM cancellation range on the worst-case expected SI, the VM selects the phase / amplitude point closest to the incoming SI and some residual SI remains due to quantization effects.

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 1x1, 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 $\frac{1}{2} + \frac{1}{2n}$. Thus, the worst-case cancellation given a number of slices is given by

$$SIC[dB] = 20\log_{10}\left(\frac{\frac{1}{2}+\frac{1}{2n}}{\frac{\sqrt{2}}{2n}}\right)$$

As discussed in section 2.1, a cancellation around 30dB allows reducing the direct crosstalk to levels where ambient reflections dominate the SI. Combined with practical constraints, a resolution of n=31 slices was chosen, allowing 27.1dB cancellation[2]. 31 slices can be conveniently segmented and controlled with 5 bits.

[2] Slightly less than the 28.5dB mentioned in [2] as a result of more accurate calculation.

3.1.2. 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 [16]. 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 scaling its matching impedance up from the 50Ω standard (i.e. weaker coupling of the SI into the RX path [9]), but in order to use standard external attenuators, 50Ω matching was maintained also for the VM.

The VM noise depends on its setting. Analyzing this for all possible VM settings is mathematically intensive, since each setting is a specific, time-variant 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;
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 approximated rather accurately. 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 3.1.1). However, both points provide useful bounds for the NF. Figure 7 depicts single-ended equivalent circuits in these three configurations, and their equivalent in-band LTI models for noise analysis according to [16]. 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 (LTV) circuit are represented by a shunt resistance $R_{sh} = \frac{4\gamma}{1-4\gamma}(R_m + R_{sw})$ in the linear time-invariant (LTI) equivalent [16]. Here, $\gamma = \frac{2}{\pi^2}$. The feedback amplifier is modeled by a noiseless amplifier preceded by its input impedance $R_b = \frac{R_f}{1+A}$ and two correlated noise voltages $v_{n,amp}$and $i_{n,amp}R_b$, where $i_{n,amp}^2 R_b^2 = \frac{4kTR_f}{(A+1)^2} + \frac{v_{n,amp}^2}{(A+1)^2}$ [16]. The noise factor is then given by [16]:

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

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 the noise factor equation above by an equivalent resistor

$$R_{eq} = R_{sh} \| R_{sh2} \| (R_{sw2} + R_{m2} + R_{s2})$$

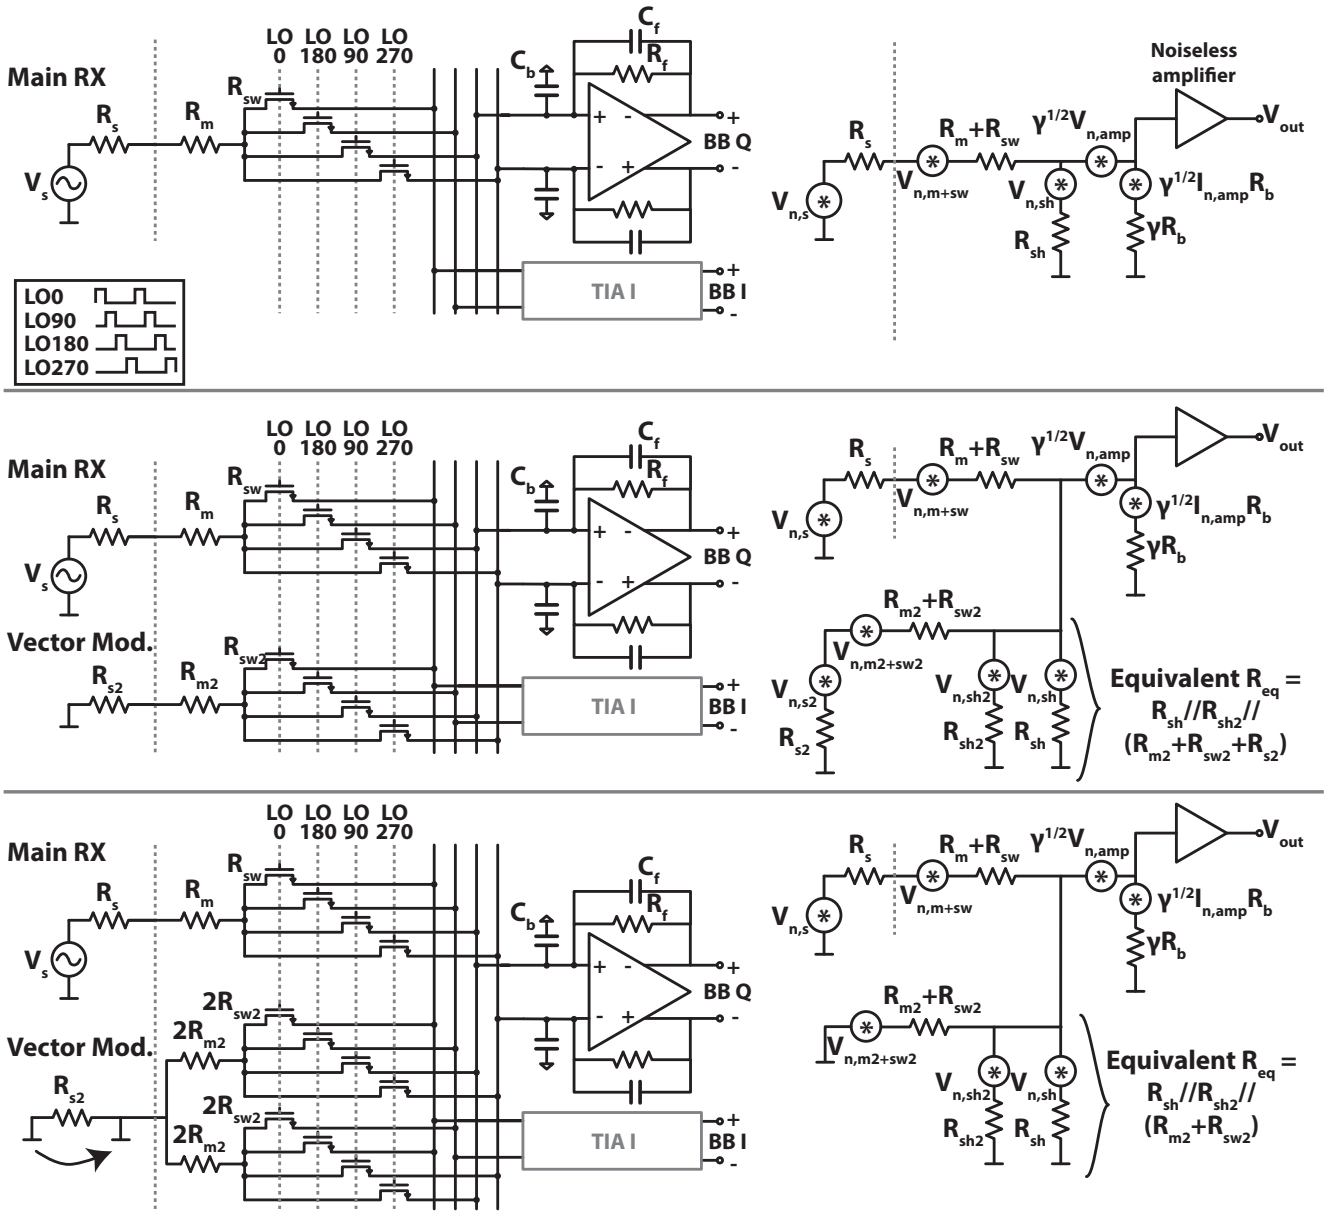

Figure 7: Left: single-ended representations of the receiver in three configurations, Right: their equivalent noise models following the approach of [16]. Top: cancellation path disabled; Center: VM set to maximum amplitude; Bottom: VM set to minimum amplitude.

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 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:

$$R_{eq} = R_{sh} \| R_{sh2} \| (R_{sw2} + R_{m2})$$

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 1.5kΩ for 24dB overall receiver gain. A two-stage, telescopic op-amp was used with A=1000x open loop gain. The main noise contributors of the op-amp are the input pair (gm_1 = 2 x 23.4mS) and the active loads of the input stage (gm_2 = 2 x 12.8mS). Assuming a noise excess factor of 1, the input-referred op-amp noise can be calculated as $v_{n,amp}^2 = \frac{4kT(gm_1+gm_2)}{gm_1^2}$. 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 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 = 10pF capacitors were put on the virtual grounds to filter higher harmonics. Table 1 lists the simulated NF at 10MHz offset, to minimize the influence of flicker noise. Analysis and simulation are in close agreement. From the results in Table 1, we conclude that the VM contributes the largest amount of noise at small amplitude settings. The results also show that enabling the cancellation path degrades the system NF by up to roughly 6dB.

Table 1: Calculated and simulated RX noise figure.

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

3.1.3. Linearity

Ideally, for 0Ω switches and a perfectly linear 50Ω matching resistor, there is no signal swing across the switches and therefore no non-linear currents are induced by the SI before cancellation. The only source of non-linearity are the TIAs that process residual SI and the (usually weaker) desired signal. Therefore, every 1dB of cancellation of the SI would result in a 2dB reduction of the SI-induced IM3 products, boosting the effective IIP3 by 1dB. However, low-ohmic mixer switches are power-hungry to drive, resulting in a trade-off between power consumption and linearity for switched-resistor mixers. The design was implemented with 25Ω resistors, with the remaining 25Ω distributed over the switch resistance, virtual ground impedance and routing parasitics to achieve input matching. 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.

The TIAs were further linearized by a differential negative conductance at their inputs [14]. 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. Figure 8 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 [14].

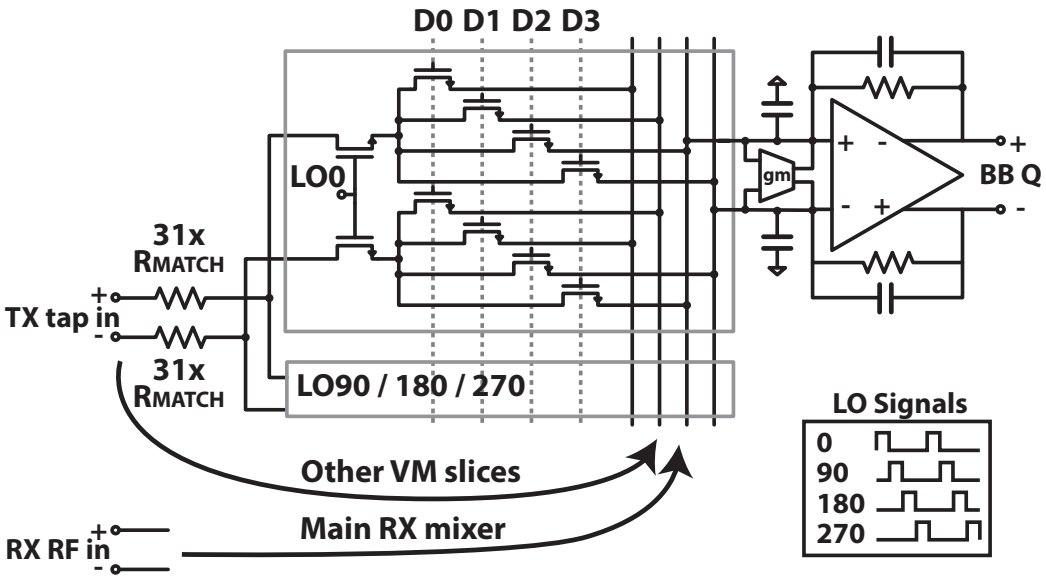

Figure 8: Implementation details of one VM slice for one LO phase, and the TIA linearized by negative conductance.

3.1.4. 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 9 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.

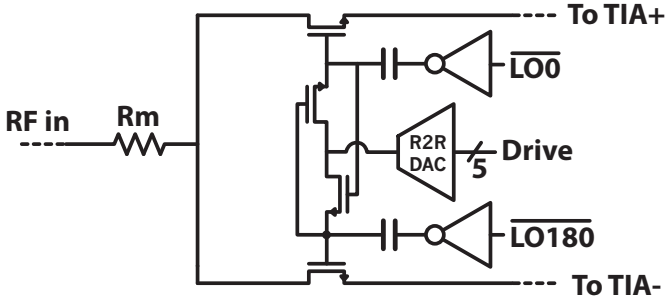

Figure 9: 5-bit tuneable level shifting of the LO for tuneable input matching.

3.2. Transmitter

The implemented 65nm front-end is depicted in figure 3 and its differential transmitter structure is shown in Figure 10. It consists of a four-phase sampling mixer, sampling its output on the gates of a class-A common source PA [17]. Whereas in [17], the design is promoted for FDD applications due to its low out-of-band noise potential, we chose it for its high in-band accuracy (low 1/f noise and predictable distortion). This is important for full-duplex, as highly deterministic self-interference is more easily cancelled digitally. Choosing this transmitter architecture also has the benefit that it can be clocked with the same four-phase, non-overlapping full-swing clock signals readily available in the receiver. AC-coupling between the sampling mixer and PA allows low drain / source voltages in the mixer, for reduced on-resistance. The PA is biased using a 5-bit R-2R DAC for

tunable transconductance (a similar DAC used in the levelshifting of the LO for the RX and VM driving). Thick oxide cascodes enable biasing with RF choke inductors from a 2V supply (Vdd_PA).

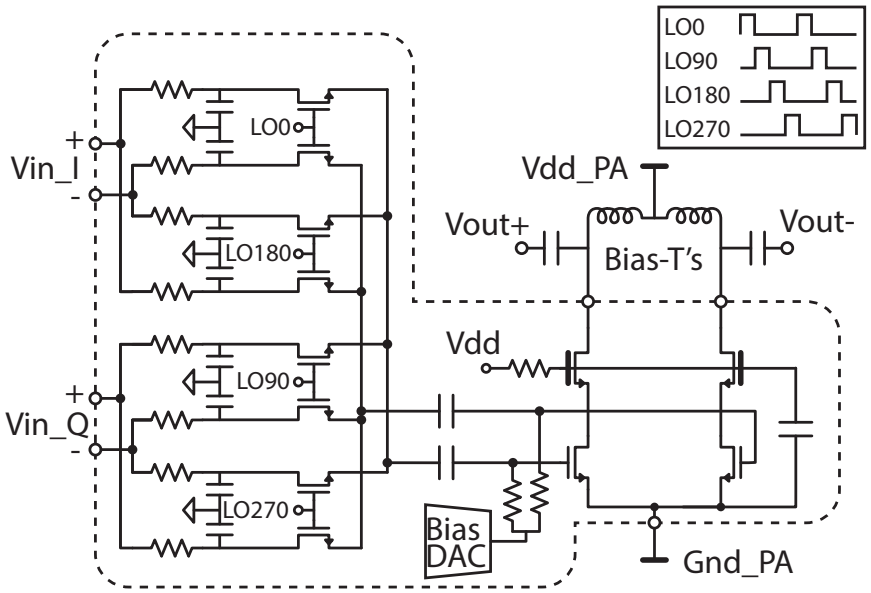

Figure 10: Transmitter topology: Four-phase sampling mixer and common source class-A PA with thick oxide cascodes. Dashed line = chip boundary. The clocks are shared with the RX.

4. MEASUREMENT RESULTS

The design was implemented in 65nm CMOS; a die photo is shown in Figure 11. This section describes the measured performance of this prototype.

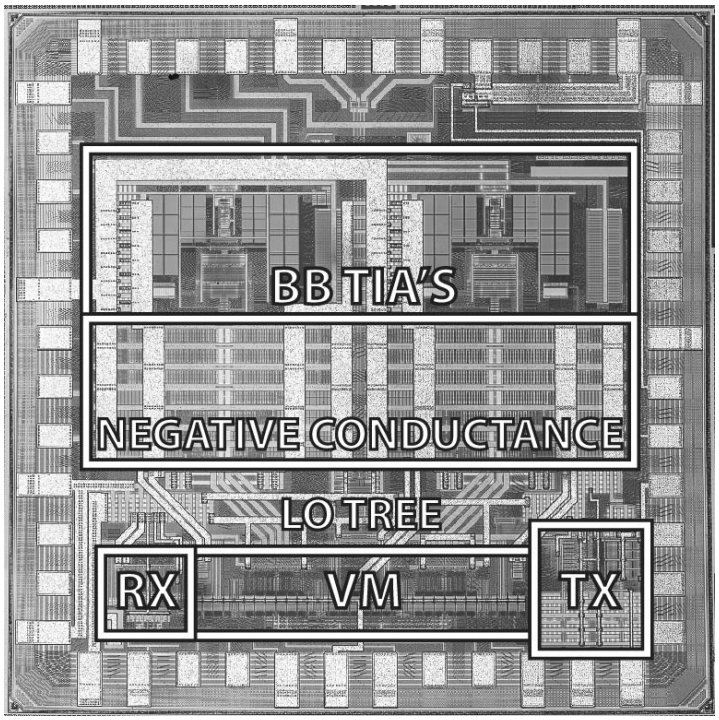

Figure 11: Die photo with relevant blocks indicated. The 65nm design measures 1.4x1.4mm.

4.1. Receiver

The receiver was characterized first for cancellation performance, to ensure it could reach the 27dB cancellation expected from the chosen resolution. Next, its noise figure and linearity performance were determined, to check if the system indeed showed the low noise and distortion floor under cancellation of strong self-interference.

All measurements in this section were performed in a controlled setup, using cables and 50-ohm measurement equipment, to verify the raw performance and cancellation potential of the receiver. A real-world experiment including actual antennas is described in section 4.4.

Simplified measurement setups for the cancellation and linearity measurements are shown in Figure 12. The setups are discussed in detail in the corresponding paragraphs.

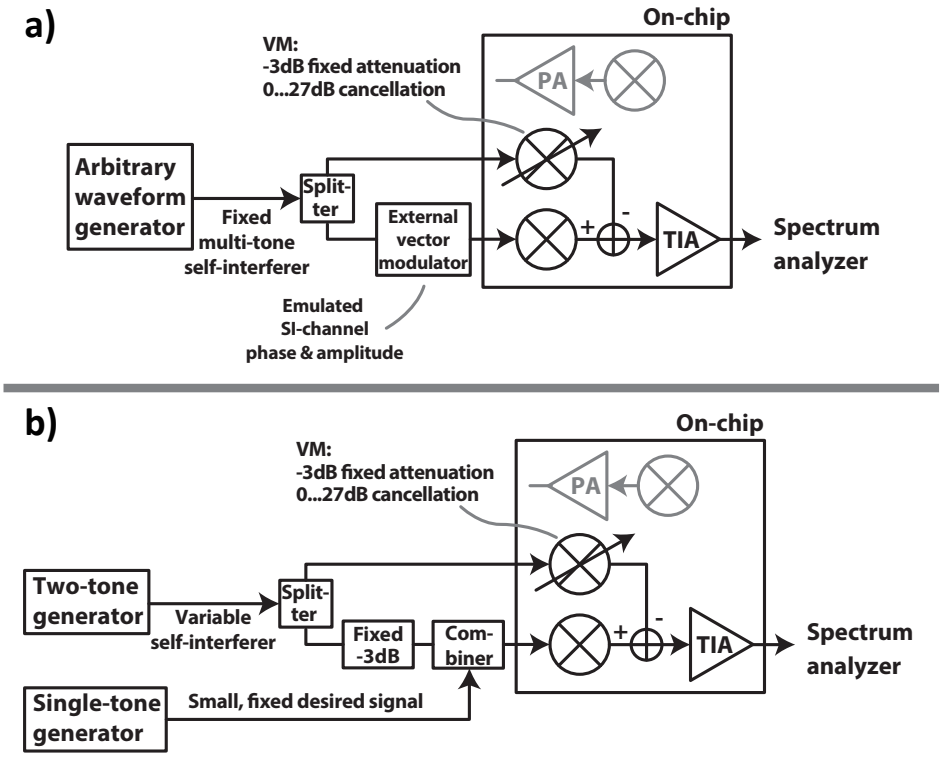

Figure 12: Simplified measurement setups used for a) cancellation performance and b) linearity.

4.1.1. Cancellation

The cancellation performance of the circuit was evaluated using an 802.11g-like TX signal of 52 tones with random phases in 16.25MHz bandwidth centered at 2.5GHz. This multi-tone signal was generated by an arbitrary waveform generator, and the SI channel was emulated by a commercial high-resolution vector modulator as shown in Figure 12a. More than 100 arbitrarily chosen phase / amplitude points were evaluated within the cancellation range of the VM, as shown in Figure 13a. An iterative search algorithm based on received power minimization was used to find the VM setting for best cancellation for each point, shown in Figure 13b. The residual SI power was measured for each point, relative to the maximum power the VM could cancel (i.e. the gray circle in figures Figure 13a/b. The results, plotted in Figure 13c, show better than 27dB cancellation which is very close to the calculated 27.1dB derived in section 3.1.1. This is expected, since the sizing of the components in the VM slices was based mostly on practical layout constraints. This resulted in better slice-to-slice component matching than strictly required for the 31-slice VM, and therefore close-to-theoretical quantization errors.

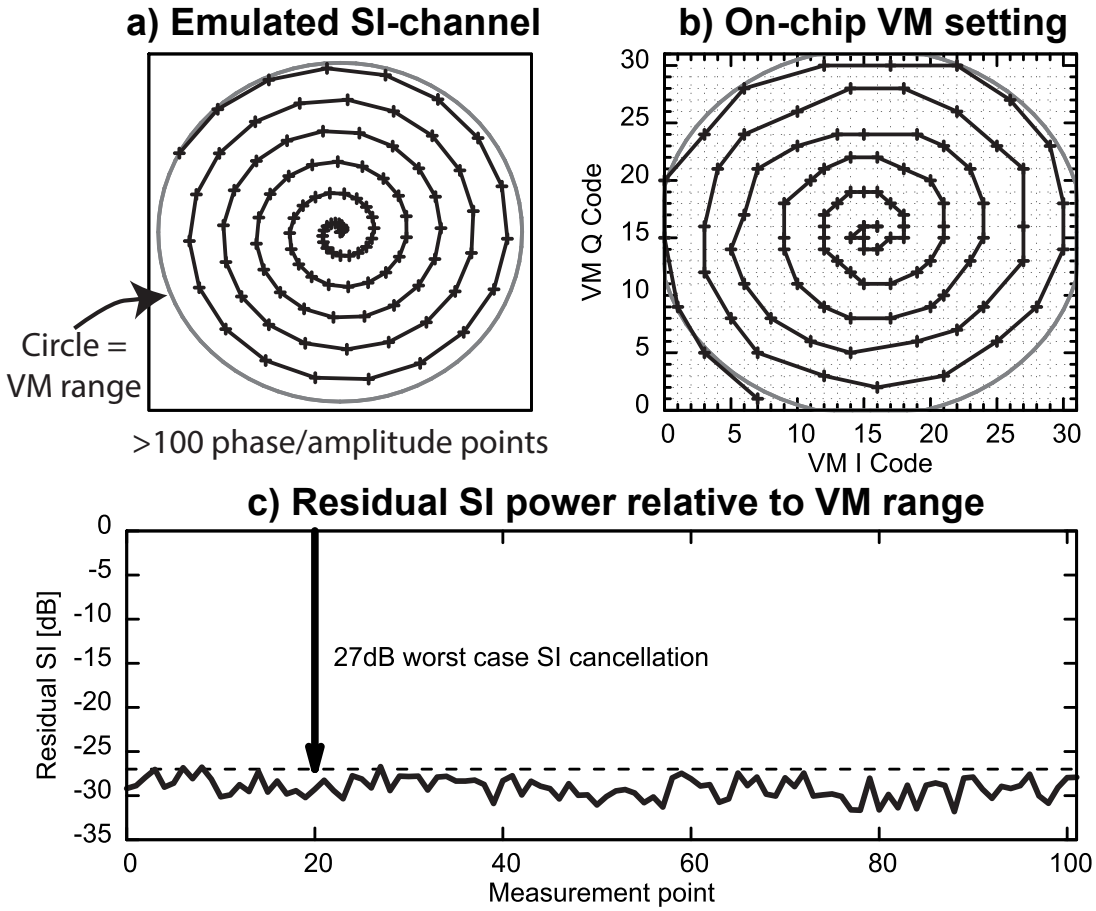

Figure 13: a) Over 100 arbitrary, spiral-shaped phase/amplitude points emulate the SI channel; b) The on-chip VM finds the corresponding setting for best cancellation; c) The residual SI is always at least 27dB below the VM range (circle)

4.1.2. 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 1. The 1/f noise corner of the RX was measured to reside at roughly 2MHz.

4.1.3. Linearity

For a symmetrical point-to-point link based on this design, the available link budget[3] will at first increase linearly with increasing transmit power (i.e. an increasing SINDR, see section 2.2). 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; the set-up to oerform this measurement is shown in Figure 12b. The IM3 products were observed. Due to the discrete nature of the VM, it is difficult to guarantee exactly 27dB cancellation, therefore the measurement

[3] `Link budget' in this document assumes 0dB SNR at the receiver and does not include any fading and AGC margins, to obtain a standard-independent metric.

was performed under 28dB cancellation. Figure 14a shows the results without cancellation. Based on the noise floor in 16.25MHz, the SINDR can be derived. Figure 14b shows how the results change under 28dB cancellation. Again, the SINDR can be derived with respect to the noise floor (Figure 14c). Both SINDRs are shown in Figure 14d. The peak SINDR of the system increases from 66.5dB without cancellation, to 69.5dB under cancellation, indicating a 5dB increase in link budget when cancellation is enabled. The point of maximum link budget has moved from -28dBm to -18dBm 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 14d).

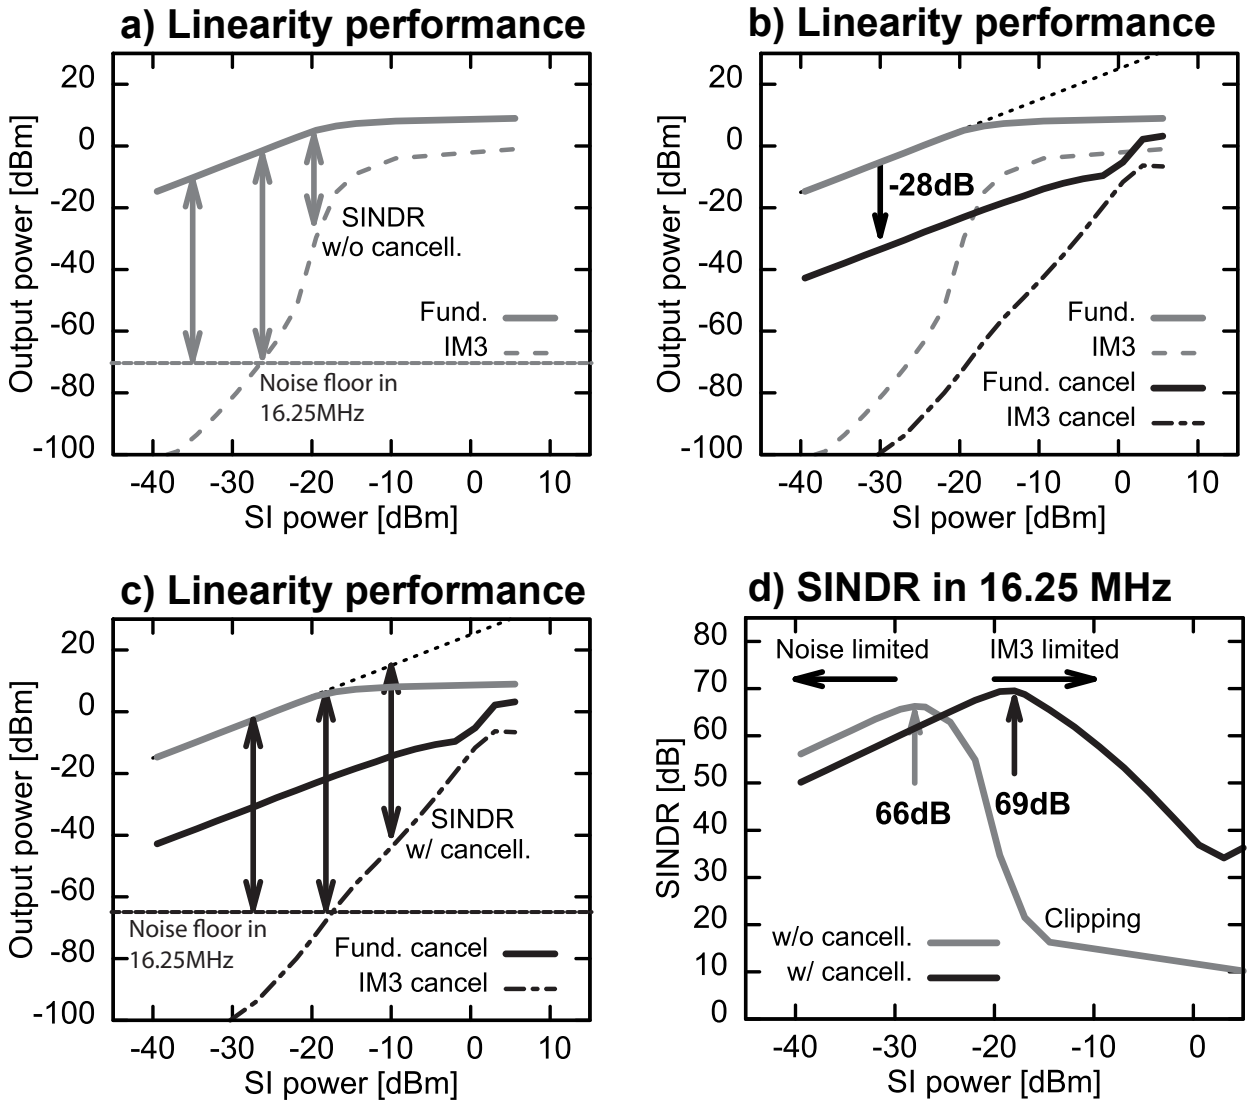

Figure 14: Results of a two-tone linearity test: a) Cancellation disabled, including 16.25MHz noise floor and defining SINDR; b) Cancellation enabled; c) Cancellation enabled, including noise floor and SINDR; d) SINDR with and without cancellation for various SI powers. Note that the performance has improved slightly with respect to \cite[2], to reflect the most recent measurements.

Under cancellation, an effective in-band IIP3 can be defined with respect to the SI. The measurements show that the IIP3 increases from 8dBm to an effective 19dBm when cancellation is enabled: an increase of 11dB. The fact that the IIP3 does not increase by the full 27dB indicates that the linearity bottleneck has moved from the TIA to the nonlinear RX and VM switches. The peak SINDR can also be calculated as:

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

where the 3dB is due to the RX noise floor and SI-induced IM3 products adding as powers. Since enabling the cancellation increases the effective IIP3 by 11dB but also increases the noise floor by 6dB, this equation shows why the 27dB cancellation only yields a 3dB 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. Table 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 42.5dB.

Table 2: Summary of cancellation, noise and linearity effects on overall full-duplex link performance, assuming 20 dB antenna isolation.

	Without cancellation	**With cancellation**
Maximum link budget (i.e. SINDR + isolation)	86.5 dB	89.5 dB
Digital cancellation requirement (SINDR – Cancellation)	66.5 dB	**42.5 dB**
TX power @ max link budget (i.e. SI power + isolation)	-8 dBm	2 dBm

In Figure 14c, 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 (see Figure 12b). Figure 15 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 14b/c.

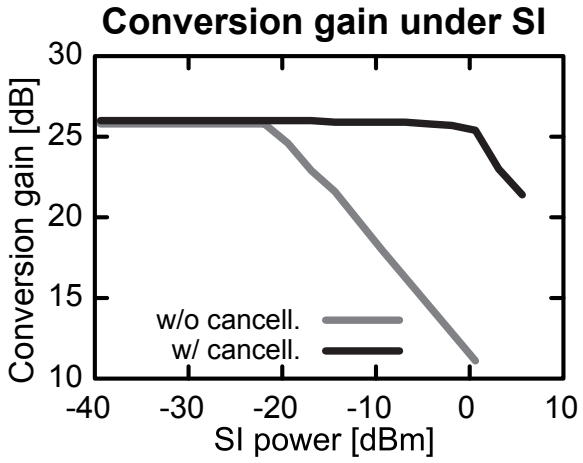

Figure 15: Conversion gain for desired signal with increasing SI, without and with cancellation.

As mentioned in section 3.1.3, 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 16.2dB. The fact that the effective IIP3 under cancellation is even 2.8dB higher, can be explained by two phenomena: 1) distortion cancellation mechanisms occurring between the RX and VM, and 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 14 and Figure 15 were performed without negative conductance.

4.1.4. Broadband performance

Although the aforementioned results were obtained at 2.5GHz LO frequency, the receiver employs frequency-agile operation and cancellation principles. Figure 16 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. The power consumption increases linearly with frequency, as expected.

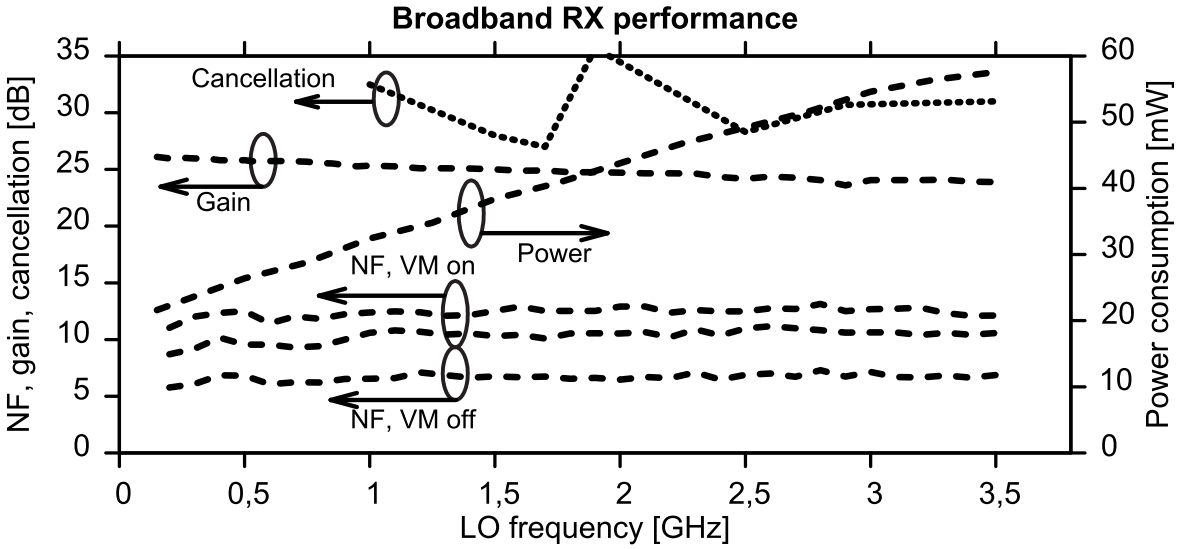

Figure 16: Performance of the receiver over a wide range of LO frequencies.

4.1.5. Image rejection

A drawback of the proposed topology is the required 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 69dB image rejection is required from the mixers, which is not a feasible value. However, if the image rejection is better than 27dB, it does not limit the analog cancellation, and the residual image can be dealt with in digital cancellation [18]. The prototype achieves 37dB image rejection, which is sufficient not to limit analog cancellation by a margin of 10dB, but the image must be accounted for in digital to reach the full 42.5dB digital cancellation potential.

4.1.6. Comparison

Table 3 compares the presented receiver to two previously published integrated FD receivers. For fair comparison, no antenna isolation is assumed for all designs. The peak SINDR of the other works was calculated using the SINDR equation from section 4.1.3. The SI power at which the peak SINDR occurs is given by:

SI [dBm] = Effective IIP3 [dBm] - 1/3*(Effective IIP3 [dBm] - Noise Floor [dBm])

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 [9] can theoretically achieve significant cancellation over a wide bandwidth even when the antenna isolation is already high, thanks to its ability to address delayed SI components.

Table 3: Comparison with other integrated FD transceivers, assuming no antenna isolation.

	[9]	[12]	This work
Topology	Dual-port N-path filter based canceller + noise-cancelling receiver	Mixer-first architecture + Noise-cancelling duplexer LNA's	Mixer-first receiver + SI-cancelling VM-downmixer
Technology	65 nm CMOS	65 nm CMOS	65 nm CMOS
Supply	N/R	1.2V (LO) / 2.5V (BB)	1.2V
Operating freq.	0.8-1.4 GHz	0.1-1.5 GHz	0.15-3.5 GHz
Max. gain	42 dB	51-55 dB	24 dB
NF	5.7 - 6.3 dB (4.8 in HD)	5.5 dB	10.3-12.3 dB (6.3 in HD)
Power consumption	63 - 250 mW	43 - 56 mW (incl. TX)	23 - 56 mW[4]
Baseband BW	>30 MHz (-15 to +15)	6-192 MHz	24 MHz (-12 to +12)
In-band IIP3	-20 dBm	-32.7 dBm	+8 / +16.2 dBm (Neg. conductance off / on)[5]
Effective in-band IIP3 with respect to SI	2 dBm	-0.7 dBm[6]	19 dBm
SINDR in 16.25 MHz BW	62.5 dB peak @ -30.7 dBm SI	60.8 dB peak @-32.6 dBm SI	**69.5 dB peak @ -18 dBm SI**
Out-of-band IIP3	17 dBm	22.5 dBm	22.0 dBm
SI Cancellation	20 dB	33.5 dB	27 dB
SI power @ 1dB RX compression	>>-8dBm	-17.3 dBm	>+1.5 dBm[7]
1/f Noise corner	N/R	N/R	2 MHz
Cancellation details	20 dB in 25 MHz BW, 34 dB initial isolation from 1.4 GHz dipole pair	33.5 dB in ~1 MHz BW[8], with single-port antenna	27 dB in ~16 MHz BW, 20 dB initial isolation from crossed 2.5 GHz dipoles
Area	4.8 mm^2	1.5 mm^2	2 mm^2

[4] The transmitter adds 129mW at 2.5GHz, as detailed in [3]
[5] Negative conductance gives about 1.5dB NF penalty [14]
[6] From -38.7 dBm IIP3 and 38 dB IIP3 improvement @33.5 dB isolation
[7] 135 kHz spacing [12], under 27 dB cancellation
[8] From [12], figure 25

4.2. Transmitter

As discussed, the SI-cancelling receiver achieves up to 69.5dB SI-to-Noise-and-Distortion-Ratio (SINDR) at 27dB SI-cancellation in 16.25MHz. Hence the digital SI-cancellation potential is 69.5-27 = 42.5dB without RX noise and distortion limitation. Given 20dB worst-case antenna-isolation, a link budget of 69.5+20 = 89.5dB is available. The peak link budget occurs at -18dBm self-interference at the RX input, so the transmit power required to achieve the maximum link budget is -18+20 = 2dBm.

To achieve this maximum link budget puts direct requirements on the transmitter. Ideally, it should output the required 2dBm of power at 42.5dB EVM performance, in order not to hamper digital cancellation. Table 4 lists some key measured parameters of the transmitter at 2.5 GHz carrier frequency. For 0 dBm TX power, its EVM is -40 dBc. While not immediately sufficient to provide 2dBm at 42.5dB EVM performance, the small performance gap may be bridged by either pre-distortion of the TX, or by modeling the PA nonlinearity in the digital cancellation. This is left for future research.

Note that like the receiver, the chosen transmitter structure is frequency agile, allowing operation of the entire SI-cancelling transceiver at carrier frequencies from 0.15 to 3.5 GHz.

Table 4: Transmitter specifications @ 2.5 GHz carrier frequency

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

4.3. Phase noise

Phase noise (PN) can be troublesome for FD. Previous work on FD generally used off-the-shelf radios, specifically the WARP research platform [8, 10]. One particular topology, shown in figure 5, uses an additional upconverter chain for SI cancellation at the receiver input [10]. In [11], 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.

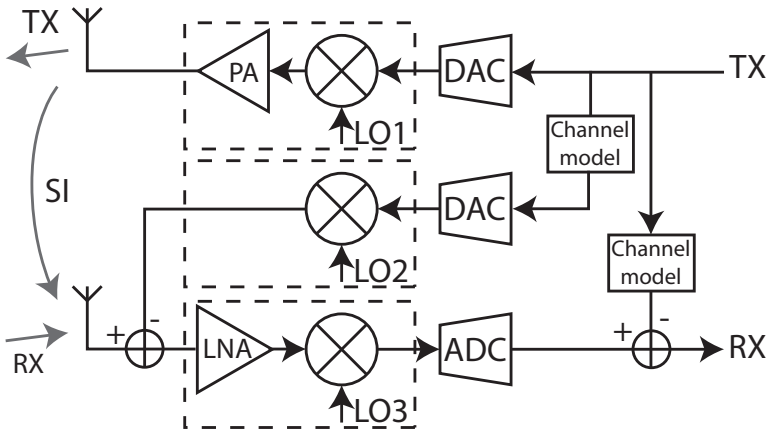

Figure 17: FD radio topology implemented in [6], using an additional upconverter for SI-cancellation at RF. All 3 up- and downconverters are separate transceivers with their own PLL.

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 [11]; we will evaluate this effect and examine its implications.

To exploit the maximum SINDR = 69.5dB of our RX and thus its maximum digital cancellation potential, the receiver should operate under its full 27 dB analog cancellation while receiving -18dBm SI. For higher SI powers, digital cancellation is impaired by SI-induced RX-distortion; for lower SI powers, dynamic range is limited by the RX noise floor. Hence, targeting the maximum link budget, the PN impact will be evaluated under 27dB analog cancellation of -18dBm SI at the receiver input. The remaining input-referred SI after analog cancellation will be -18dBm – 27dB = -45 dBm. The RX noise figure (NF_{DSB}) in its full bandwidth (50kHz - 12 MHz) is 11.7 dB (slightly more than reported in Table 3 since flicker noise is included down to 50kHz). Therefore, its input-referred DSB noise floor is -174dBm/Hz + $10*\log_{10}$(24MHz) + 11.7dB = -88.5dBm. This leaves room for -45 – (-88.5) = 43.5dB[9] 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.

The measurement setup is shown in figure 6. 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.

[9] Slightly more than the 42.5 dB mentioned earlier, since the influence of SI-induced distortion is omitted here for simplicity.

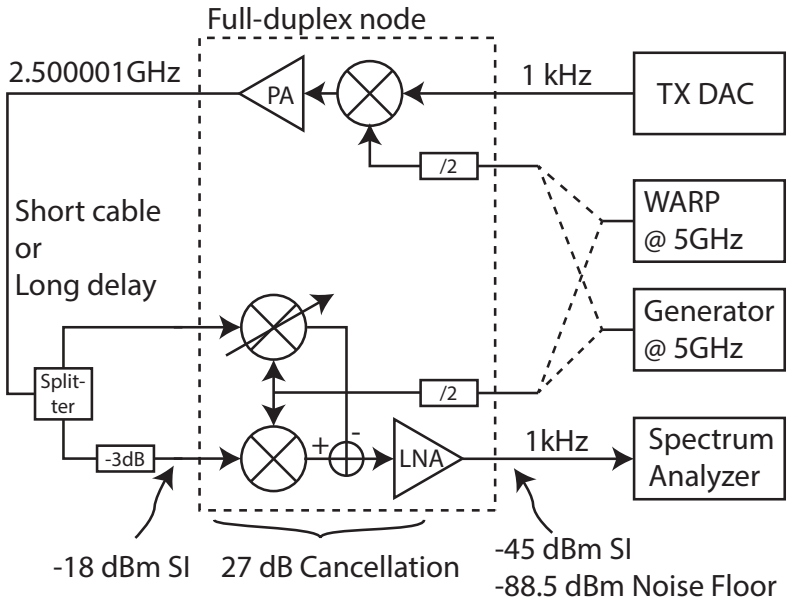

Figure 18: Simplified set-up used to measure the effect of correlated and non-correlated phase noise between TX and RX/VM, for short and long self-interference paths. The power levels are all referred to the RX input.

A very low offset-frequency single tone of 1 kHz is used from the TX. This permits filtering out deterministic TX imperfections like LO-radiation, image and distortion by AC-coupling toward the spectrum analyzer, and this allows 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 [8]) in contrast to phase noise.

Table 5 shows for short SI channels how the PN deteriorates the RX noise floor under cancellation, and Figure 19 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.

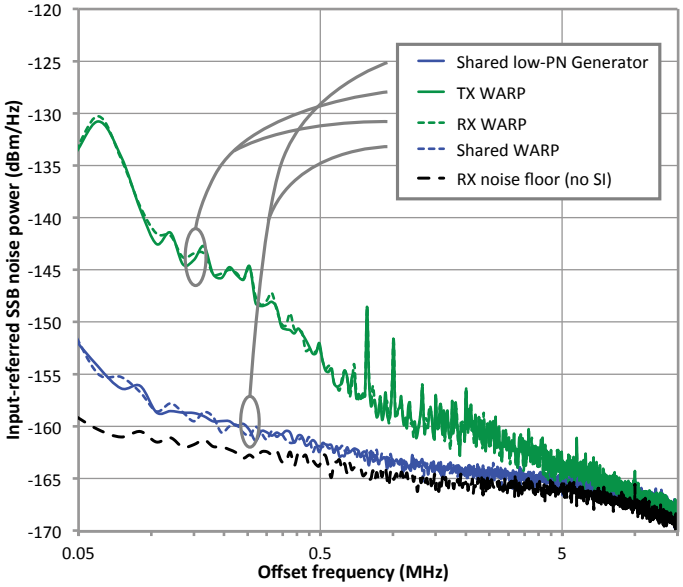

Figure 19: Measured input-referred noise spectra for the four clocking configurations with a short SI-path: A shared WARP clock shows comparable PN performance as a low-PN generator.

Table 5: Effect of phase noise on RX noise floor, short SI path

TX clock	RX / VM clock	DSB noise figure	Noise floor degradation
Disabled	WARP	11.7 dB	-
Shared low-PN generator		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

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 [11]. An experiment was set up with an available delay line, providing again -18 dBm of SI at the RX input and 27dB analog cancellation, but with 115ns of delay in the SI path. The results are shown in Figure 20 and Table 6. The effect of PN 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 6.2dB improvement compared to clocking one side with WARP, consistent with further analysis.

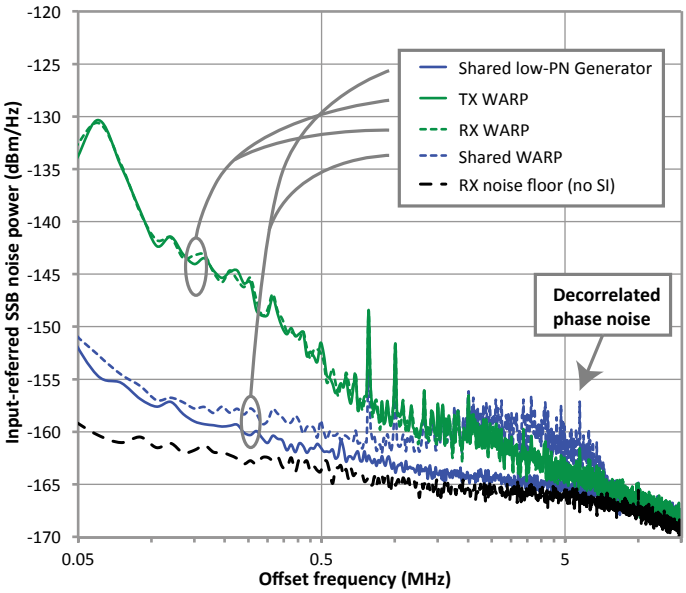

Figure 20: Measured input-referred noise spectra for the four clocking configurations with 115ns delay in the SI-path, resulting in phase noise decorrelation at high carrier offsets.

Table 6: Effect of phase noise on RX noise floor, long SI path

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

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 still be well below the noise floor. Although this is a discrete measurement point at 115ns delay, further analysis shows that for a path loss exponent of 2, no delay exists where decorrelated PN effectively raises the RX noise floor.

In conclusion, *by virtue of a common clock for all mixers in the system, the phase noise arising from a commercially available PLL with -38 dBc phase noise does not degrade the 89.5 dB link budget potential* calculated previously.

4.4. Antenna experiments

To verify the claims of 20dB as a representative worst-case antenna isolation and -40 to -50dB as the level where reflections start dominating the SI in 16.25 MHz BW, some real-world experiments were performed with the transceiver, using a crossed pair of commercial WLAN dipoles as a simple FD antenna [19]. The measurement setup is depicted in Figure 21a. A 20dB attenuator is used to set the maximum cancellation range of the VM. Since the VM can only cover a full phase circle over a reduced amplitude (as shown in Figure

6), the VM itself also attenuates another 3dB. Therefore, the 20dB attenuator configures the system for 23dB worst-case isolation from the antenna. Measurements were performed in a lab environment without special precautions.

Representative measured spectra are shown in Figure 21b. The TX power was close to 0dBm, and therefore the system is configured for 27dB cancellation of at most -23dBm SI at the RX input. In this particular situation, the SI entered the RX at about -29dBm, which is 6dB below the VM maximum range, set by the attenuator. This leaves room for at most 27 - 6 = 21dB analog cancellation. The actual achieved cancellation was roughly (-29 – (-46)) = 17dB. The remaining SI is clearly frequency-selective, indicating the presence of delayed SI-components that limit analog cancellation for this particular combination of environment, bandwidth and antenna solution. Note that the cancellation was limited because of the frequency selectivity, not because of imperfections in the design.

In conclusion, delayed SI, reflected by the environment comes in at 46dB below the transmitted power level, which is in the range of 40 to 50dB as expected. Such reflections can be subsequently handled in the digital domain. In comparable experiments, it was also confirmed that 20dB is a reasonable worst-case isolation value, e.g. when heavily influencing the near-field of this antenna with a hand. Further characterization of the transceiver in real-world scenarios and implementing digital cancellation is part of ongoing research.

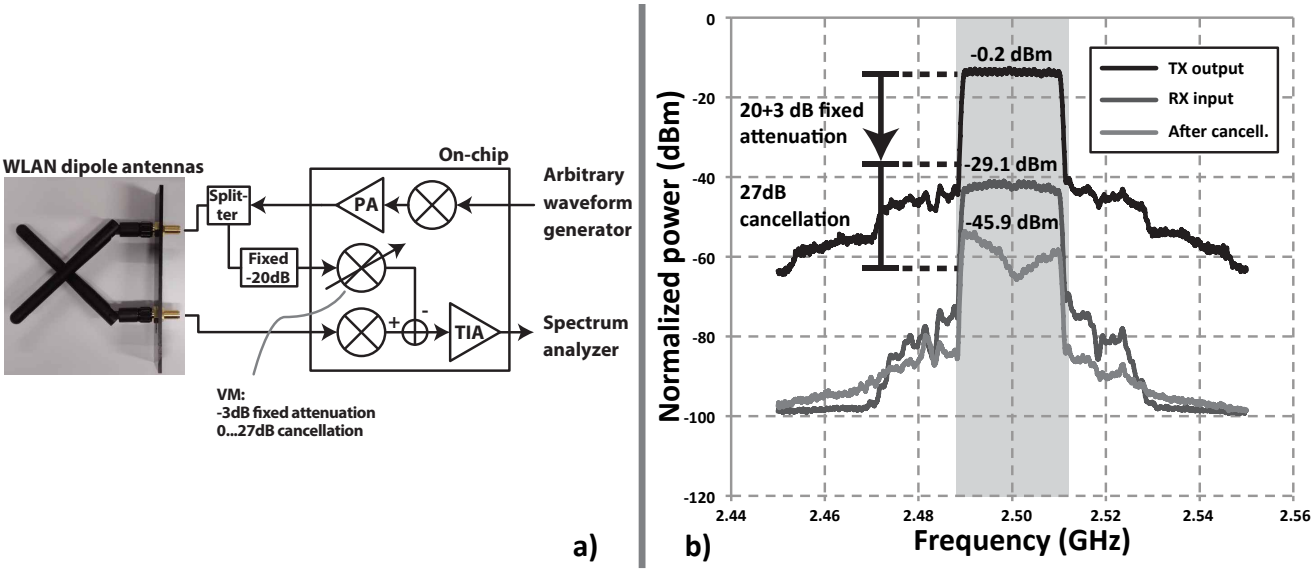

Figure 21: a) Simplified measurement setup using two crossed WLAN dipole antennas as a simple FD antenna, b) representative measured spectra at the output of the transmitter, input of the receiver and after analog cancellation (referred back to the RX input).

5. INTEGRATION AND VALIDATION POTENTIAL

This section briefly elaborates on further integration and validation of the *existing* prototype for research purposes. The next section describes potential improvements to be made in *future* designs.

- First of all, the presented results describe how the raw performance of the analog cancellation *theoretically* affects digital cancellation. Although many aspects have been studied (distortion, EVM, noise, phase noise) it remains to be confirmed if the design allows the full 42.5 dB digital cancellation potential that is expected. To this end, integration with a digital baseband section and digital cancellation algorithms is required. Implementing digital cancellation may reveal new, unexpected bottlenecks and opportunities in the link budget that give valuable insights for further research.
- When digital cancellation is implemented, further experiments can be performed with actual antennas in real-world environments. This may not only yield more information about the prototype, but also provide more insight in the statistics of the self-interference channel that an FD radio has to deal with.
- So far, all detailed experiments have been performed around 2.5 GHz carrier frequency to comply with the license-free ISM-band, but also the broad operation frequency range of the transceiver can be studied in detail.
- The prototype has a degree of freedom that has not yet been exploited: the virtual ground nodes of the baseband amplifiers are exposed as pins to the outside of the chip. These nodes could be useful to inject cancellation signals generated in the digital domain and converted to analog through DAC's. This could also cancel (part of) the delayed SI reflected by the environment already before the BB amplifiers. However, care should be taken since this node is very sensitive to noise.

6. FUTURE OPPORTUNITIES

This section lists some recommendations for follow-up designs using the proposed topology.

Some basic improvements to the design are:

- The VM currently complies with the 50-ohm measurement standard, but this makes the VM a large contributor of noise when the cancellation path is enabled. If the VM impedance is scaled up (which is feasible, since it has to provide attenuation of the TX signal), its noise contribution could be lowered significantly.
- The fixed attenuator was kept external to the prototype for versatility in working with different worst-case antenna isolation values. Integrating this attenuator on-chip saves an RF connection to the outside of the IC, making board lay-out much more convenient. Also, it becomes easier to move the TX tap and VM inputs away from the 50-ohm interfacing standard. Furthermore, making the on-chip attenuator variable allows configuring the system for various antenna solutions in the field.
- The baseband amplifiers could be redesigned, with special care for minimizing their area. The negative conductance linearizing the baseband amplifiers is also an area-consuming block, which can be removed since the baseband amplifiers are not the linearity bottleneck when dealing with strong self-interference.

If the topology is to be taken into a commercial transceiver, features such as automatic gain control (AGC) and harmonic rejection should be added to the design. It is also useful to re-think the transmitter topology, since the class-A power amplifier limits the efficiency of the TX, and its open-drain structure relies on external RF-chokes and is easily damaged. Furthermore, the current prototype uses a conservative, fully differential design. This is done in order not to run into unexpected second-order distortion bottlenecks while measuring. However, several components of both the transmitter and receiver are suitable for single-ended implementation, saving internal or external balun components. It is useful to investigate this.

7. SUMMARY AND CONCLUSIONS

This deliverable presented work on an integrated self-interference (SI) cancelling transceiver, aiming to bring in-band full-duplex wireless communication to compact low-power devices. The DUPLO project was among the first to present such an integrated design targeted specifically at in-band full-duplex. The work presented here was well received by the scientific community, in which FD is an unconventional scheme and meets quite some scepticism. The work contained in this deliverable has been presented at major events, such as the IEEE ISSCC [2], RFIC [3] and VTC [4] conferences and was invited to appear in IEEE Journal of Solid State Circuits.

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 transceiver takes an attenuated copy of the transmit signal, and provides simultaneous phase shift, amplitude scaling and downmixing using a vector modulator downmixer, for SI-cancellation in the RX analog baseband. The main receiver and vector modulator 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 potential and link budget.

For the sliced vector modulator, the cancellation performance was derived as a function of the number of slices. We also show 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 vector modulator 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.

With only 20dB isolation from the antenna, the receiver can achieve 27dB cancellation at 2dBm transmit power, without introducing distortion above the RX noise floor. Given its 12.3dB worst-case noise figure with cancellation enabled, this results in 89.5dB 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 entire system offers frequency-agile operation and cancellation from 0.15 to 3.5GHz LO frequency.

The integrated TX has -40dB EVM performance at 0dBm TX power, requiring only minor improvement by e.g. pre-distortion to allow the full 89.5dB 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 WARP transceivers, allowing operation from the commercial WARP 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.

Initial antenna experiments confirm the claims of 20dB as a worst-case antenna isolation for a simple FD-antenna, and -40 to -50dB as the level where reflections limit the efficacy of a phase / amplitude based canceller in 16.25MHz bandwidth. This validates the initial system-level aspects of the design, although further research is required.

Several hints were included to proceed research based on the current design, as well as how to take the proposed transceiver topology towards a product.

This deliverable 2.3.2 concludes the work performed within DUPLO by task 2.3.

8. REFERENCES

1. van den Broek, D.-J., *Transceiver Circuits Simulation, Implementation and Measurement Report.* DUPLO - Report D2.3.1, 2014.
2. van den Broek, D.-J., E.A.M. Klumperink, and B. Nauta, *19.2 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. p. 1-3.
3. van den Broek, D.-J., E.A.M. 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, 2015 IEEE*. 2015 (Accepted for publication).
4. van Liempd, B., et al., *RF self-interference cancellation for full-duplex*, in *Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM), 2014 9th International Conference on*. 2014. p. 526-531.
5. Tapio, V., *System Scenarios and Technical Requirements for Full-Duplex Concept, INFSO-ICT 316369.* DUPLO - Report D1.1, 2013.
6. Debaillie, B., *Design and Measurement Report for RF and Antenna Solutions for Self-Interference Cancellation, INFSO-ICT- 316369 - Report D2.1.* DUPLO - Report D2.1, 2014.
7. Debaillie, B., et al., *Analog/RF Solutions Enabling Compact Full-Duplex Radios.* IEEE J. Sel. Areas Commun., 2014. **32**(9): p. 1662-1673.
8. Bharadia, D., E. McMilin, and S. Katti, *Full Duplex Radios.* SIGCOMM Comput. Commun. Rev., 2013. **43**(4): p. 375-386.
9. Zhou, J., et al., *19.1 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. p. 1-3.
10. Duarte, M., C. Dick, and A. Sabharwal, *Experiment-Driven Characterization of Full-Duplex Wireless Systems.* IEEE Trans. Wireless Commun., 2012. **11**(12): p. 4296-4307.
11. Sahai, A., et al., *On the Impact of Phase Noise on Active Cancelation in Wireless Full-Duplex.* IEEE Trans. Veh. Technol., 2013. **62**(9): p. 4494-4510.
12. Yang, D., H. Yuksel, and A. Molnar, *A Wideband Highly Integrated and Widely Tunable Transceiver for In-Band Full-Duplex Communication.* IEEE J. Solid-State Circuits, to be published.
13. Everett, E., A. Sahai, and A. Sabharwal, *Passive Self-Interference Suppression for Full-Duplex Infrastructure Nodes.* IEEE Trans. Wireless Commun., 2014. **13**(2): p. 680-694.
14. Mahrof, D.H., et al., *Cancellation of OpAmp Virtual Ground Imperfections by a Negative Conductance Applied to Improve RF Receiver Linearity.* IEEE J. Solid-State Circuits, 2014. **49**(5): p. 1112-1124.
15. Soer, M.C.M., et al., *3.5 A 1.0-to-2.5GHz 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. p. 66-67.
16. Andrews, C. and A.C. Molnar, *Implications of Passive Mixer Transparency for Impedance Matching and Noise Figure in Passive Mixer-First Receivers.* EEE Trans. Circuits Syst. I, 2010. **57**(12): p. 3092-3103.
17. Xin, H. and J. van Sinderen. *A 45nm low-power SAW-less WCDMA transmit modulator using direct quadrature voltage modulation*. in *Solid-State Circuits Conference - Digest of Technical Papers, 2009. ISSCC 2009. IEEE International*. 2009.
18. Korpi, D., et al., *Widely Linear Digital Self-Interference Cancellation in Direct-Conversion Full-Duplex Transceiver.* IEEE J. Sel. Areas Commun., 2014. **32**(9): p. 1674-1687.
19. Debaillie, B., et al., *RF Self-Interference Reduction Techniques for Compact Full-Duplex Radios*, in *Vehicular Technology Conference (VTC), 2015 IEEE*. 2015 (Accepted for publication).