

RF Spectrum Sensing in CMOS Exploiting Crosscorrelation



Mark Oude Alink

RF Spectrum Sensing in CMOS

Exploiting Crosscorrelation

Members of the dissertation committee:

prof. dr. ir.	A.J. Mouthaan	University of Twente (chairman and secretary)
prof. dr. ir.	B. Nauta	University of Twente (promotor)
dr. ir.	A.B.J. Kokkeler	University of Twente (assistant-promotor)
dr. ing.	E.A.M. Klumperink	University of Twente (assistant-promotor)
prof. dr. ir.	G.J.M. Smit	University of Twente
prof. dr. ir.	F.E. van Vliet	University of Twente / TNO
dr. ir.	M.J. Bentum	University of Twente / ASTRON
prof. dr. ir.	P.G.M. Baltus	Eindhoven University of Technology



This research is supported by the Dutch Technology Foundation STW, which is part of the Netherlands Organisation for Scientific Research (NWO) and partly funded by the Ministry of Economic Affairs (08081).

CTIT

CTIT Ph.D. Thesis Series No. 12-236
Centre for Telematics and Information Technology
P.O.Box 217, 7500 AE Enschede, the Netherlands

Copyright © 2013 by Mark S. Oude Alink, Enschede, The Netherlands.



This work is licensed under the Creative Commons Attribution-Non-Commercial 3.0 Netherlands License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/nl/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California 94105, USA.

Typeset with L^AT_EX. Cover photo by AGCuesta, YAY Images.

Printed by Gildeprint Drukkerijen, Enschede, The Netherlands.

ISBN 978-90-365-3497-0

ISSN 1381-3617 (CTIT Ph.D. Thesis Series No. 12-236)

DOI 10.3990/1.9789036534970

<http://dx.doi.org/10.3990/1.9789036534970>

RF SPECTRUM SENSING IN CMOS EXPLOITING CROSSCORRELATION

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
prof. dr. H. Brinksma,
volgens besluit van het College voor Promoties
in het openbaar te verdedigen
op vrijdag 24 mei 2013 om 16.45 uur

door

Mark Stefan Oude Alink

geboren op 20 juni 1984
te Hengelo (Ov.)

Dit proefschrift is goedgekeurd door:

prof. dr. ir.	B. Nauta	(promotor)
dr. ir.	A.B.J. Kokkeler	(assistent-promotor)
dr. ing.	E.A.M. Klumperink	(assistent-promotor)

Lente

Witruimtes, de blanco zijden
waarop onze breedband dromen
draaien, zonder heldere ontvanger.
Over een boodschap maar te zwijgen.

Luisterend tot ver onder het ruisniveau,
hoor ik op de volle maan
geen woorden maar signaaltonen,
die ik met zekerheid kan onderscheiden,
als de aarde spreekt.

Liefst breide ik
een groot rood hart voor jou,
of vrijde iets
ontegenzegbaars.

(Egbert van Hattem, UT-Nieuws, 24 maart 2011)

ABSTRACT

The introduction of new wireless services, the demand for higher data rates, and higher traffic volumes call for a more efficient use of the RF spectrum than what is currently possible with static frequency allocation. Dynamic spectrum access offers a more efficient use by allowing unlicensed users to opportunistically use locally and temporarily unoccupied licensed bands ('white space').

To prevent harmful interference to the licensed users, unlicensed users need to make sure the band is free before they are allowed to transmit. This means that, if resorting to databases is not possible or desired, the unlicensed users have to be able to detect very weak signals from the licensed users by means of spectrum sensing. Different types of spectrum sensing exist, but it is preferable to use one that does not require knowledge of the signals to be detected, as it can then be employed in arbitrary frequency bands. Such a solution is energy detection (ED). The first step of ED is similar to what a spectrum analyzer (SA) does: measure the power in a frequency band. The second step is to distinguish between measuring only noise, or noise plus a signal. Due to inaccuracies in the noise level estimation, there is a certain minimum signal-to-noise ratio (SNR), the SNR-wall, below which signals cannot be reliably detected. Several analog impairments, such as phase noise, nonlinearity, and limited harmonic rejection (HR), can also hamper the detection process by causing false alarms or missed detections.

To reduce the SNR-wall and the influence of analog impairments on sensing performance, crosscorrelation (XC) spectrum sensing, as a generalization of autocorrelation (AC) (the standard form of ED), is proposed. XC multiplies and integrates the outputs of two receivers, each processing the same signal, to obtain the signal power, while the noise (ideally) averages out. The noise uncertainty is removed at the cost of measurement time, and the SNR-wall reduces. A mathematical model is developed that predicts that (1) a lower noise correlation between the two receivers lowers the SNR-wall, and (2) resistive attenuation at the input of each receiver does not influence the sensitivity of XC. This allows a design to be optimized for high linearity without affecting the detection capabilities. By employing a separate oscillator for each receiver, XC can also reduce phase noise. A frequency offset between the two oscillators, in combination with some digital signal processing, also allows XC to improve HR.

A first mostly-discrete prototype is developed, employing a mixer-first architecture for high linearity. It demonstrates (phase) noise reduction and an attenuation-independent noise floor using XC, but suffers from external frequency-dependent coupling between the receivers, crosstalk between the mixers, and a poor HR. A second prototype tackles these disadvantages by integrating two RF-frontends into a single 1.2 V 65 nm CMOS-chip, with a novel distortion-cancellation technique in the attenuators for high linearity.

Measurements show that XC achieves 22 dB of phase noise reduction (limited by measurement time), and up to 25 dB of improvement in HR (limited by crosstalk). At 10 dB attenuation, the SNR-wall is found to be -184 dBm/Hz, which is 10 dB below the thermal noise floor, and even 12 dB below the measured SNR-wall of AC. XC achieves an attenuation-independent noise floor < -169 dBm/Hz from 0.3–1.0 GHz, with an IIP3 of +25 dBm at 10 dB attenuation, which makes the spurious-free dynamic range higher than that of high-end commercial SAs. Furthermore, it is experimentally shown that XC can be much faster and more energy-efficient than AC.

Overall, XC is shown to enable the integration of SAs with high sensitivity, good resilience to strong interferers, and with speed and, at low SNR, energy consumption benefits compared to AC. This not only makes sensitive spectrum sensing attainable in a hostile radio environment, but also paves the way for low-cost, low-power, and high-quality (mobile) measurement equipment. Furthermore, it may enable the integration of (many) small SAs inside other chips for built-in self-test (BIST), reducing on pin count and test time during manufacturing, as well as more reliable and stable performance during operation.

SAMENVATTING

De introductie van nieuwe draadloze diensten, de vraag naar hogere snelheden, en de continu toenemende datavolumes vereisen een efficiënter gebruik van het RF spectrum dan nu behaald wordt met de vaste toewijzing van frequentiebanden. Dynamisch spectrum gebruik biedt meer efficiëntie door het ongelicenseerde gebruikers toe te staan gelicenseerd spectrum te gebruiken, zolang deze maar niet in gebruik is door de gelicenseerde gebruiker.

Om verstoring van de gelicenseerde gebruikers te voorkomen, dienen ongelicenseerde gebruikers zich ervan te vergewissen dat de band vrij is alvorens ze mogen zenden. Dit betekent dat, als het gebruik maken van databases niet mogelijk of wenselijk is, de ongelicenseerde gebruikers in staat moeten zijn om zeer zwakke signalen van gelicenseerde gebruikers te detecteren. Er zijn verschillende manieren om de detectie te bewerkstelligen, maar het verdient de voorkeur om geen kennis nodig te hebben van het te detecteren signaal, zodat het in een willekeurige frequentieband te gebruiken is. Een techniek die hieraan voldoet is energie-detectie (ED). De eerste stap van ED is vergelijkbaar met wat een spectrum analyzer (SA) doet: het meten van het vermogen in een bepaalde frequentieband. De tweede stap in ED is om te bepalen of het gemeten vermogen alleen ruis, of ruis plus een signaal bevat. Door onnauwkeurigheden in de bepaling van het ruisvermogen is er een bepaalde minimale signaal-ruis verhouding (SNR) nodig, de SNR-muur, om nog met enige betrouwbaarheid signalen te kunnen detecteren. Verschillende niet-idealiteiten in het analoge deel van de implementatie, zoals faseruis, niet-lineariteit, en harmonische onderdrukking (HR), kunnen het detectie-proces ook nadelig beïnvloeden.

Om deze SNR-muur te slechten, en de invloed van analoge niet-idealiteiten op het beslissingsproces te verminderen, wordt kruiscorrelatie (XC) als generalizatie van autocorrelatie (AC) (de standaard methode van ED) voorgesteld. XC vermenigvuldigt en integreert de uitgangen van twee ontvangers, die elk hetzelfde ontvangen signaal bewerken, om tot het signaalvermogen te komen. De ruis middelt uit ten koste van meettijd: de ruisonzekerheid, en daarmee de SNR-muur, wordt lager. Een in dit proefschrift ontwikkeld wiskundig model voorspelt dat (1) een lagere ruis-correlatie een lagere SNR-muur geeft, en dat (2) resistieve verzwakking aan de ingang van elke ontvanger de gevoeligheid van XC niet beïnvloedt. Een systeem kan dan ontwikkeld worden voor hoge lineariteit zonder dat

dat ten koste hoeft te gaan van de detectie-mogelijkheden. Door een aparte oscillator in elke ontvanger te gebruiken, kan XC ook de faseruis verminderen. Verder kan een klein frequentie-verschil tussen de oscillatoren, in combinatie met enige digitale signaalbewerking, het voor XC ook mogelijk maken om de HR te verbeteren.

Een eerste, vooral discreet, prototype is ontwikkeld, met een mixer als eerste trap voor hoge lineariteit. XC-metingen tonen reductie van (fase)ruis en een verzwakkings-onafhankelijke ruisvloer, maar ook externe frequentie-afhankelijke koppeling tussen de ontvangers, overspraak tussen de mixers, en een slechte HR. Een tweede prototype verbetert dit door integratie van twee lineaire RF-ontvangers op één 1.2 V 65 nm CMOS-chip. Hierbij is een nieuwe techniek gebruikt om verzwakkers vervormingsarm te maken ten behoeve van een hoge lineariteit.

XC-metingen geven minimaal 22 dB faseruis-reductie (meettijd-gelimeerd), en tot 25 dB verbetering in HR (overspraak-gelimeerd). Met 10 dB verzwakking is een SNR-muur van -184 dBm/Hz gemeten, wat 10 dB onder de thermische ruisvloer ligt, en zelfs 12 dB onder de gemeten SNR-muur van AC. Met XC wordt een verzwakkingsonafhankelijke ruisvloer < -169 dBm/Hz behaald van 0.3–1.0 GHz, met $+25$ dBm IIP3 op 10 dB verzwakking, waardoor de spurious-free dynamic range hoger is dan die van moderne commerciële SAs. Bovendien wordt aangetoond dat XC veel sneller en energie-efficiënter kan zijn dan AC.

In dit proefschrift is aangetoond dat XC het mogelijk maakt om geïntegreerde SAs met hoge gevoeligheid te maken, terwijl ze tegelijkertijd goed overweg kunnen met sterke verstoorders, én voordeel bieden in snelheid en, voor lage SNR, energieverbruik ten opzichte van AC. Dit maakt het niet alleen mogelijk om gevoelig signalen te detecteren in drukbezette draadloze omgevingen, maar effent ook het pad naar goedkope, energiezuinige, en kwalitatief hoogstaande (mobiele) meetinstrumenten. Bovendien zouden (meerdere) kleine SAs in andere chips geïntegreerd kunnen worden voor zelftests, waarmee bespaard kan worden op meettijd en het aantal externe pinnen benodigd voor verificatie, maar waarmee het ook mogelijk wordt om meer betrouwbare en stabielere functionaliteit te bieden tijdens het gebruik.

DANKWOORD

Zoals het Twentse spreekwoord luidt: “*Alns hef ’n end, (mer de metworst hef der twee)*”, zit nu ook mijn AIO-tijd erop. Wat u hier voor u ziet is het werk van de afgelopen vier jaar. Dit boekje was er echter niet geweest zonder de hulp van velen.

Allereerst wil ik graag Bram Nauta en Gerard Smit bedanken om het mogelijk te maken mijn tanden in dit onderwerp te kunnen zetten bij hun vakgroepen. Mijn dagelijkse begeleiders Eric Klumperink en André Kokkeler hebben mij met veel raad en daad bijgestaan, en mijn manuscripten van vele nodige pennenstrepen of ballonnetjes voorzien. Ook wisten ze mij altijd scherp te houden met vragen of nieuwe ideeën. Ik denk dat iemand zich geen betere begeleiders kan wensen.

Binnen het ADREM project, gefinancierd door STW, heb ik vele zinnige discussies kunnen voeren met mijn collega’s Dlovan, Saqib en Amir. De leden van de gebruikerscommissie wil ik bedanken voor de nuttige opmerkingen en de positieve en kritische feedback. Ik wil NXP en ST bedanken voor het doneren van silicium, en CMP voor het faciliteren daarvan. Door deze donaties is het mogelijk geweest ideeën daadwerkelijk te kunnen realiseren. *Merci beaucoup!*

Henk de Vries, Gerard Wienk en de twee rechterhanden van Jasper Velner waren onmisbaar voor het ontwerpen van een werkende chip, PCB en meetopstelling. Zonder jullie waren vele metingen op niks uitgelopen. De goede en soms intensieve samenwerking op inhoudelijk vlak met Wei Cheng, Wim Korevaar, Nadia Haider, Zhiyu Ru en Michiel Soer wil ik hier ook zeker niet onvermeld laten. Mark Bentum en Frank van Vliet hebben af en toe als klankbord gediend voor wilde ideeën of fundamentele vragen; dit is van onschatbare waarde geweest voor mij in dit project. Ook de zinvolle discussies met Niels Moseley, Arjan Meijerink, Anne Hendrikse en Marnix Heskamp hebben inhoudelijk bijgedragen aan dit werk.

Met mijn kamergenoten Michiel, Jasper, Bram, David, Erik, Dirk-Jan, Harish, Johan en Kasra heb ik vele zinvolle en zinloze discussies gevoerd, grappen en grollen uitgehaald, goede lunch-barbecues gehad, en veel foute muziek geluisterd op vrijdagmiddag. ‘Intelligente’ gesprekken aan de koffietafel zijn bijna niet te voeren zonder Tom (+Tom), Ray en Rob; in de rappa doen Koen, Mark en Maurice er niet voor onder. De SC-dinners zorgden ook voor veel gezelligheid, wat zonder de organisatie van Giulia en Buket zeker niet zo goed gelukt zou zijn.

Ik heb tijdens deze vier jaar menig student begeleid, zowel voor de bachelor als de master. Dit was niet alleen erg leerzaam, maar ook erg leuk. Wim, Arnout, Arjan, Peter, Boris, Chiel, Matthijs en Marijn, ik hoop dat jullie dat ook vonden. Mijn onderwijstaken bleven niet beperkt tot eindopdrachten; ik wil hierbij Sabih Gerez, Jan Kuper, Ed van Tuijl, Mark Bentum, Anne-Johan Annema, André Kokkeler en Eric Klumperink bedanken voor de fijne samenwerking tijdens het verzorgen van vele onderwijsvakken. Van geheel andere categorie, maar niet minder belangrijk, waren de 6- tot 12-jarigen van het Leonardo-onderwijs in Hengelo die in het kader van hun project 'radio' de UT een dagje bezocht hebben. ASTRON heeft gezorgd voor mooie discolamp-setjes waar driftig aan gesoldeerd werd. Moge dat nog vele studenten opleveren.

Het is een bevoorrechte positie om in en (vlak)bij meerdere vakgroepen te zitten: meer taart, en meer secretaresses. Zonder hun hulp waren vele zaken een stuk lastiger, dan wel onmogelijk, geweest. Daarom: Gerdien, Annemiek, Satie, Marlous, Nicole, Thelma, Lilian en Remke, bedankt!

Dit proefschrift is in L^AT_EX geschreven. Zonder de hulp en tips van Philip en Pascal, of het template van Philip/Albert/Vincent/... dat als basis heeft gediend, zou het er allemaal een stuk minder uitzien. Het zorgvuldige doorlezen van Fleur en Wim hebben voor vele verbeteringen in de uiteindelijke tekst en plaatjes gezorgd. Alle fouten die er nog in staan, zijn volledig aan mijzelf te wijten.

Omdat dit boekje al dik genoeg is, en AIO's volgens Bram Nauta beter een beetje lui kunnen zijn, eindig ik hier met het bedanken van alle andere leden van ICD, CAES, SC en TE voor de gezellige tijd.

Naast werk is ontspanning ook nodig om het hoofd helder te krijgen. Ik wil daarom ook eenieder met wie ik heb gevolleybald, gebiljart en gesnookerd bedanken voor de gezellige tijd. Alfons, het was echt een gemis dat we het laatste jaar niet meer op het groene laken konden strijden.

Mijn ouders wil ik graag bedanken voor hun onvoorwaardelijke steun en toeverlaat. Hetzelfde geldt voor mijn opa en oma, die helaas allebei net niet meer hebben mogen meemaken dat hun kleinzoon eindelijk 'klaar is met school.'

Fleur, mijn grote liefde, zonder jou had ik dit project nooit tot een goed einde kunnen brengen. Als ik een deadline had klaagde je nooit als ik weer eens een avond of weekeinde doortrok, zelfs als je het zelf nog zo zwaar had. Jij hebt deze vier jaar echt de moeite waard gemaakt.

Bedankt allemaal, Mark

CONTENTS

	ABSTRACT	· vi
	SAMENVATTING	· viii
	DANKWOORD	· x
	1 INTRODUCTION	· 1
	1.1 <i>Problem Statement</i>	· 5
	1.2 <i>Thesis Outline</i>	· 6
2	DYNAMIC SPECTRUM ACCESS & SPECTRUM SENSING	· 7
	2.1 <i>Regulations & Standardization</i>	· 7
	2.2 <i>Spectrum Sensing</i>	· 11
	2.3 <i>Implementation Considerations</i>	· 15
	2.4 <i>Motivation for Energy Detection Spectrum Sensing</i>	· 30
	2.5 <i>Conclusions</i>	· 32
	3 CROSSCORRELATION SPECTRUM SENSING	· 33
	3.1 <i>Energy Detection</i>	· 33
	3.2 <i>Crosscorrelation</i>	· 39
	3.3 <i>Mitigating Analog Impairments</i>	· 50
	3.4 <i>Conclusions</i>	· 56
	4 A DISCRETE PROTOTYPE	· 57
	4.1 <i>Receiver Mismatch</i>	· 57
	4.2 <i>Prototype Design</i>	· 59
	4.3 <i>Expected Performance</i>	· 63
	4.4 <i>Measurements</i>	· 67
	4.5 <i>Conclusions</i>	· 72
	5 A HIGHLY LINEAR ATTENUATOR	· 75
	5.1 <i>Design Considerations of Attenuators</i>	· 75
	5.2 <i>Two-port Attenuators</i>	· 77
	5.3 <i>IM3-reduction in Discrete-step Resistive Attenuators</i>	· 78
	5.4 <i>Prototype Implementation & Measurements</i>	· 83
	5.5 <i>Conclusions</i>	· 89

6	AN INTEGRATED PROTOTYPE: UTSFINX	· 91
	6.1	<i>Design</i> · 91
	6.2	<i>Measurements</i> · 97
	6.3	<i>Benchmarking</i> ·113
	6.4	<i>Conclusions</i> ·113
7	LOWERING THE SNR-WALL WITH CROSSCORRELATION	·115
	7.1	<i>The SNR-wall of Crosscorrelation</i> ·115
	7.2	<i>Experimental Verification</i> ·122
	7.3	<i>Conclusions</i> ·128
8	CONCLUSIONS & RECOMMENDATIONS	·129
	8.1	<i>Summary & Conclusions</i> ·129
	8.2	<i>Original Contributions</i> ·131
	8.3	<i>Recommendations</i> ·132
	A	SFDR OF A UNIFORM QUANTIZER ·135
	A.1	<i>Quantization of a Sinusoid</i> ·136
	A.2	<i>Quantization of a Sinusoid with Noise</i> ·140
	A.3	<i>Multitone Quantization</i> ·141
	A.4	<i>Example</i> ·142
	A.5	<i>Conclusions</i> ·142
	B	CROSSCORRELATION WITH TWO ANTENNAS ·143
	B.1	<i>Noise in an Antenna</i> ·144
	B.2	<i>Noise Correlation with Two Antennas</i> ·145
	B.3	<i>Signal Correlation with Two Antennas</i> ·145
	B.4	<i>Conclusions</i> ·148
	C	RECEIVER TERMINOLOGY ·149
	D	LIST OF ACRONYMS ·151
	E	LIST OF SYMBOLS ·153
	E.1	<i>Mathematical Definitions</i> ·153
	E.2	<i>Variables</i> ·154
		BIBLIOGRAPHY ·157
		LIST OF PUBLICATIONS ·163
		<i>Peer-reviewed</i> ·163
		<i>Other</i> ·164

INTRODUCTION

When Heinrich Hertz experimentally verified the existence of electromagnetic (EM)-waves, as theoretically deduced by Maxwell a few years earlier, he supposedly said “*It’s of no use whatsoever*” [1]. He couldn’t have been more wrong: telegram services, radio broadcasting and ship-to-shore communications soon followed his discovery. Nowadays, TV broadcasting, radar, mobile telephones and Wireless Local Area Networks (WLANs) are just a few examples of the myriad of wireless services employed since then.

All of these services have been assigned specific frequency bands in the spectrum, usually by national authorities such as the FCC and NTIA (USA), Agentschap Telecom (Netherlands), MIC (Japan), ACMA (Australia), and Ofcom (UK). Often, harmonization with neighboring countries is required. This has led to complex and crowded charts. A simplified example is shown in fig. 1.1, where the 1.5–3 GHz band is divided in many bands with specified frequency ranges, all dedicated to specific services.

Lower frequencies propagate better and suffer from less attenuation due to walls and other obstacles, which is why FM-radio, TV-broadcasting and the first GSM-network all operate in the 0.1–1 GHz band. Several




Figure 1.1: Spectrum allocation between 1.5 GHz and 3 GHz in the USA. (Source: NTIA)




Figure 1.2: Wireless data demand in exabytes per month for recent history and near future (1 EB= 10^{18} bytes, or a billion Gbytes). (Source: [3])

ISM-bands are defined internationally, and are free to use as long as certain requirements are met. ISM-bands are often noisy and crowded environments; the 2.4 GHz ISM-band (indicated in fig. 1.1) is used by microwave ovens and standards such as Bluetooth and the 802.11-family (WLAN).

Virtually all of the usable spectrum has been assigned, which makes it very hard for new services to squeeze in. In 2008, operators paid 19.6 billion dollar to obtain their own 52 MHz of bandwidth in the 700 MHz range in the USA: that amounts to 400 million dollar per MHz! It seems to be a worthwhile investment: the annual worldwide revenue in mobile wireless is estimated at 3 trillion dollar [2].

Predictions are that the demand for wireless data communications will increase 18-fold from 2011 to 2016, see fig. 1.2. The FCC-chairman remarked in 2010 [4]: *“Our data shows there is a looming crisis. We may not run out of spectrum tomorrow or next month, but it is coming and we need to do something now.”* IEEE Spectrum even refers to it as *“the great spectrum famine”* [5]. The datarate can be increased in three ways: (1) more spatial reuse, (2) more bandwidth and (3) higher spectral efficiency. Spatial reuse is aggressively employed by operators using cell splitting [6]. Of the other two, LTE (also known as 4G) uses both approaches: from 5 MHz of spectrum per channel for UMTS to 20 MHz for LTE (and 40 MHz in the near future), and from QPSK for UMTS to 64-QAM for LTE.

Unfortunately, this is still not enough. Therefore, US President Obama ordered the FCC to free up 500 MHz of spectrum for mobile and fixed wireless broadband [7]. Part of this is being realized by the ‘digital dividend’, the part of the spectrum that has become available due to the switch from analog to digital TV (DTV) broadcasting. Similarly, in Europe, the Radio Spectrum Policy Programme, approved by European Parliament in 2012, aims at enhancing the efficiency and flexibility of spectrum use [8]. Concrete actions are to identify at least 1200 MHz of spectrum to address the demand for wireless data traffic, and to allow much more flexible use of spectrum via spectrum trading and spectrum sharing.

Measurements show that only 2% to 20% of the spectrum is actively being used at any given time and location, depending on the definition of ‘unused’ [9]. Thus, although spectrum is a scarce resource, it is used




Figure 1.3: DSA opportunistically uses unused spectrum.

in a very inefficient way. Dynamic Spectrum Access (DSA) promises a much more efficient use of the spectrum by letting unlicensed users, the secondary users (SUs), opportunistically use bands that are licensed to primary users (PUs), but are temporarily and/or locally unused, see fig. 1.3. Hence, the term ‘opportunistic spectrum access’ is sometimes also used instead of DSA.

The available spectrum bands depend on time and location, and are referred to as ‘spectrum hole’ [10] or ‘white space’ [11]. Whenever the PU starts communicating again, the SU has to back off. This is known as ‘vertical sharing:’ there is a hierarchy among the users. The PUs can come back at any time, so the SU has to frequently check whether the PU has returned. This is different from ‘horizontal sharing’ in the ISM-bands, where all users have equal rights. Both principles are illustrated in fig. 1.4.

Very often, DSA and CR are interchangeably used. The IEEE defines DSA as [13]: “*The real-time adjustment of spectrum utilization in response to changing circumstances and objectives,*” and CR as: “*A type of radio in which communication systems are aware of their environment and internal state and can make decisions about their radio operating behavior based on that information and predefined objectives,*” which is similar to the definitions given by the ITU [14] and FCC [15]. If one narrows ‘environment’ in the CR-definition




Figure 1.4: Vertical sharing in CR/DSA as opposed to horizontal sharing in the ISM-bands (Based upon: [12]).




Figure 1.5: WiFi coverage (simulations) in 1 km² in Notting Hill, London, UK at several frequencies. (Source: [18])

to ‘radio spectrum’, DSA can be considered the principle, and a CR the physical device implementing DSA, which is the point of view used in this thesis. Note that this is a much narrower interpretation of ‘environment’ than envisioned by Joseph Mitola III, the inventor of the CR-concept [16].

Radios which perform many functions in the digital domain, such that new standards and options can be accommodated by a change of software, are called software-defined radios. A software-defined radio can be seen as an all-in-one radio, for example to receive FM-radio, to communicate with a WLAN-router and to make a phone call over GSM. Mitola, as well as many others, consider a CR to be built on top of a software-defined radio. It would enable the CR to adapt its modulation type to the environment, which makes it more flexible and allows it to use the spectrum more efficiently. In addition, a CR needs to perform spectrum sensing, which also requires a flexible receiver and digital signal processing (DSP).

The basic idea of DSA has several significant advantages. It would allow any wireless connection to offload peak demands to other parts of the spectrum, thereby allowing it to obtain a higher data rate when it needs it (a useful feature for the GSM-network in the first hour of the new year). Moreover, with a CR built on top of a software-defined radio, the number of transceivers in a device, and thus the cost, may be considerably reduced.

Public safety services tend to have a very high peak-to-average utilization, which means that their dedicated spectrum remains largely unused for most of the time. When they do need spectrum, they actually need more than what is assigned to them, as multimedia transfer, such as photos and video, could be of great help in saving lives. A CR can provide this in a spectrum-efficient way, and also enable rescue workers inside buildings to switch to more favorable frequencies before losing connection [17, Ch. 16].

One of the main drivers of DSA is the possibility to get broadband wireless internet to areas that are too remote or geographically difficult to reach via a wired network. By using the sub-GHz frequencies, one can easily cover tens of kilometers at reasonable transmission powers.

British Telecom is thinking about providing WiFi coverage in London using TVWS-frequencies (in the UK, roughly 470 to 800 MHz). Fig. 1.5 shows the simulated coverage given the locations of the WiFi access points: it is much better in the TVWS-frequencies than in the 2.4 GHz or 5.8 GHz ISM-bands. Other applications of DSA include smart grids [19], machine-to-machine-communication [20], military communications, intelligent transport systems and wireless medical sensors.

The flexibility of CR does come at a cost: the analog hardware and the digital physical layer cannot be optimized for one band or modulation, thus power consumption increases. Another, perhaps less well known, issue is the loss of usable bandwidth for radio-astronomy that CR may cause; the coexistence of radio-astronomy and CR is discussed in [21].

1.1 Problem Statement

With the large amounts of DSP and software needed for all the different functionality of a CR, a power-efficient and fast technology for the digital circuitry is required. Since CR has a potential to be used in many different consumer devices, it should also be cheap to manufacture. These are precisely the strengths of CMOS technology. CMOS can also be used for analog circuits, such that the CR can potentially consist of one IC, which makes it very attractive for mobile applications due to the low cost, low power consumption, small form factor, and the simplified PCB-design.

Unfortunately, the analog performance does not scale nearly as well as the digital performance, partly due to the decreasing supply voltages. Moreover, receivers for CR face more stringent requirements than conventional receivers, since they have to be able to receive signals over a wide frequency band in a highly dynamic and a priori unknown environment.

The principle of DSA is applicable to any frequency band, where each band may contain many different types of modulation schemes. For ad-hoc, decentralized operation all over the world, it is therefore desirable to have some form of spectrum sensing that does not require any information of the signals to be detected. Yet, it must be sensitive enough to avoid interference to the PUs and at the same time not overlook too many opportunities. Furthermore, it should be fast enough to adhere to the different regulatory requirements and minimize the sensing overhead.

Fig. 1.6 illustrates the goal: a CR needs a transmitter and receiver for communication, and a 'scanner' or SA for spectrum sensing. All that functionality should be suitable for integration into a single CMOS-IC. In this work, the focus is on the spectrum sensing part.

In summary, the problem statement of this thesis is: *to implement some form of spectrum sensing that can, without prior knowledge, reliably and swiftly detect weak signals in a hostile radio environment, while being compatible with analog design in a low-voltage CMOS process.*




Figure 1.6: A SA in CMOS is desired for low-cost, low-power, fast and reliable spectrum sensing in a hostile radio environment.

1.2 Thesis Outline

DSA is discussed in more detail in chapter 2. An overview of the current regulations on DSA in the TV-bands will be given, followed by a brief discussion of different proposed spectrum sensing algorithms. Some major issues regarding the implementation of DSA in CMOS will be identified, and an overview of state-of-the-art in CMOS implementations targeting DSA will be given.

In chapter 3 the focus is on energy detection (ED) and its inability to detect signals below a certain signal-to-noise ratio (SNR), the SNR-wall. The concept of a crosscorrelation spectrum analyzer (XCSCA) is introduced, which is a SA that uses two receivers and performs crosscorrelation (XC) on their outputs to arrive at a spectrum estimate. Based on some assumptions, it is hypothesized that XC can reduce the SNR-wall, as well as mitigate several of the CMOS-integration issues. The theoretical sensing performance and the DSP-requirements of a XCSCA are analyzed and compared to ED.

Chapter 4 discusses the implementation and measurements of a mostly discrete (i.e., non-integrated) prototype to verify a few aspects of the proposed XCSCA-concept. A novel technique for very linear attenuators in bulk CMOS and a 0.16 μm CMOS-prototype are presented in chapter 5. This attenuator technique is used in a second prototype XCSCA, which contains two integrated highly-linear radio frequency (RF)-frontends in a single 65 nm CMOS-IC. The implementation is discussed in chapter 6, along with measurement results. Finally, the predicted sensitivity improvements using a XCSCA are verified experimentally in chapter 7.

A summary of this work with main conclusions is presented in chapter 8, along with a list of original contributions, and recommendations for future work. Appendices A and B contain some further relevant results with respect to a XCSCA.

DYNAMIC SPECTRUM ACCESS & SPECTRUM SENSING

In chapter 1, Dynamic Spectrum Access (DSA) was identified as a promising way for a more efficient use of the spectrum. In this chapter, the regulations regarding DSA will be discussed in more detail. Readers unfamiliar with receiver terminology are referred to appendix C for an overview.

The main approaches of spectrum sensing in order to find unoccupied spectrum are briefly discussed. It is argued that, from a flexibility perspective, it is most desirable to measure the power in a certain frequency band to determine whether it is occupied or not. However, it will be shown that the requirements imposed by current regulations are extremely tough for an integrated low-power device; in fact, some requirements are not even met by commercial state-of-the-art spectrum analyzers (SAs). Nevertheless, it will be argued that spectrum sensing functionality gives additional advantages to a cognitive radio (CR) and even other types of radio, which justifies the development of spectrum sensing functionality.

2.1 Regulations & Standardization

As a part of DSA, the secondary users (SUs) are supposed to be looking for free spectrum or *white space*, which can be defined as “a band of frequencies that are not being used by the primary user (PU) of that band at a particular time in a particular geographic area” [15]. A more accurate definition is a band of frequencies in which “a pair of SUs can communicate successfully without violating the interference constraint imposed by the primary network” [17]. In other words, the SU should not cause harmful interference to the PU: the SU should be (almost) orthogonal to the PU in its transmission parameters in “the theoretical hyperspace occupied by radio signals” [15]. This hyperspace

Parts of this chapter have appeared in [MOA:8]

consists of dimensions such as space, frequency, time, code, angle of arrival, and possibly others.

In all cases, however, there will be some residual interference. Even when the PU is very far away, it will still receive a tiny fraction of the transmitted power of the SU, which will look like noise to the PU, especially when many SUs are transmitting. Therefore, when multiple SUs are transmitting simultaneously, the PU may experience such degradation in its signal-to-noise ratio (SNR) that one can consider it to be harmfully interfered with. This phenomenon can be quantified as interference temperature, similar to the concept of noise temperature [22].

An SU operating at a different frequency can still degrade the performance of a PU-receiver because the PU cannot perfectly filter the received SU-power, resulting in increased noise or distortion. This is the reason the FCC asked the question “*rules are designed to control interference by regulating transmitter performance (e.g., power, emissions limits, and field strength) but not receiver performance. Should [we] adopt receiver standards?*” [23]. In fact, in [24] it is shown that PUs having some interference tolerance could greatly enhance the overall capacity of PUs and SUs combined. This interference tolerance discussion is still ongoing, and (to my best knowledge) is not yet incorporated in any regulation or standardization.

In any case, deployment of CR requires proper regulations by national and global authorities to protect PUs. Some authorities (such as the FCC and Ofcom) have already given the green light to deploy CR-networks in the TV-bands. Furthermore, many organizations are currently involved in defining worldwide and local standardization of DSA and CR-networking, such as the ITU, ETSI, COST, and IEEE, which is important for interoperability. Some standards in the draft phase feature DSA-capabilities: IEEE 802.22 for rural broadband in TV white space (TVWS) (from roughly 50 to 900 MHz) defining the physical layer and medium access control, IEEE 802.18 to define co-existence in TVWS, CogNeA/ECMA392 for home network applications such as HDTV-streaming in TVWS (physical layer- and medium access control-layer), and IEEE 802.11af for WiFi operation in TVWS (physical layer and medium access control). For more details on regulations and standardization, the interested reader is referred to [25–29].

In 2008, the FCC conducted field trials to assess attainable sensing performance. The DSA-prototypes that were used in these trials were provided by industrial players. They were able to reliably detect very weak signals in an otherwise clean spectrum, but failed in the presence of a strong interferer [30]. Based on these results, as well as many public and internal discussions, the FCC came with a new publication in 2010 [31], favoring a database approach with location and channel availability information over spectrum sensing. It states: “*Our actions here are expected to spur investment and innovation in applications and devices that will be used not only in the TV band but eventually in other frequency bands as well.*” The FCC decided that “*eliminating the requirement that TV bands devices that incorporate geo-location*

Table 2.1: CR requirements, assuming mobile devices that rely on spectrum sensing.

	FCC [31]	Ofcom [32]	802.22 [33]	Unit
Max. transmit power (TP)	50	50		mW
Max. adjacent-to-TV channel TP	40	2.5		mW
Adjacent-channel emission	-72.8 dBc ^a	-46 dBm	-49 dBm	
Modulation	free	free	OFDM	
DTV sensing limit	-114	-120	-116	dBm
DTV bandwidth	6	8	6,7,8	MHz
Microphone sensing limit	-107	-126	-107	dBm
Microphone bandwidth	0.2	0.2	0.2	MHz
Monitoring interval	≤ 60	1		s
Backoff time	2	< 1	2	s
P _{FA}			10	%
P _{MD}			10	%

^a Measured in 100 kHz out-of-band with reference total power in 6 MHz

and database access must also listen (sense) to detect the signals of TV stations and low power auxiliary service stations” was the best way to go forward.

Although spectrum sensing in TVWS has thus become unnecessary, regulations for spectrum-sensing-only devices are still included, because “we are encouraging continued development of [spectrum sensing] because we believe it holds promise to further improvements in spectrum efficiency in the TV spectrum in the future and will be a vital tool for providing opportunistic access to other spectrum bands” [31]. Sensing will be necessary for ad-hoc networks of SUs, where such geolocation and database infrastructure is not available, or where a central database is rendered useless due to the fast rate of change in the spectrum usage. Furthermore, even when databases are available, a spectrum sensing capability can be useful for other reasons as well, as will be discussed in section 2.4.

Some important requirements for mobile CRs in the TV-bands that incorporate spectrum sensing, as set by the FCC, Ofcom and the IEEE 802.22 standard, are listed in table 2.1. Note that some parameters are different for fixed devices, and that the numbers may change in the future. Although the concept of CR applies to other frequency bands as well, where e.g. emission and sensing parameters may be quite different, these numbers provide for a good starting point.

The maximum transmission power and maximum out-of-band emissions are a strict requirement, as well as the obligation to use transmission power control. The FCC specifies the adjacent-channel emission limit as the power in 100 kHz bandwidth, relative to the power transmitted in the 6 MHz channel. Thus, at full power (17 dBm), the limit is -55.8 dBm in 100 kHz, or -38 dBm in the total adjacent channel. Below full power, the absolute Ofcom and 802.22 regulations become easier to satisfy, while the FCC’s relative requirement remains equally strict.

The FCC and Ofcom mainly care about the protection of other users, and do not define modulation requirements. As a communication stan-

dard, 802.22 does define the type of modulation, in this case Orthogonal Frequency Division Multiplexing (OFDM).

The most important PU in the TV-bands is digital TV (DTV), which is implemented in different ways around the world. The USA has 6 MHz bandwidth with 8-VSB modulation, while Europe has 8 MHz bandwidth with OFDM. Another important PU is the wireless microphone. A wireless microphone occupies a bandwidth of roughly 200 kHz (not standardized), and there may be several of them in a single TV-channel. Many of them use analog FM, some AM, but several digital modulation schemes are also in use. Other PUs in the TV-bands are analog TV and some medical equipment, but for the sake of brevity they are not considered here.

The FCC, Ofcom, and 802.22 define the minimum power of each of these PU-signals that SUs should be able to detect. The spectrum sensing device should thus be able to cope with the different bandwidths of the different PUs. The detection threshold of -114 dBm for DTV as set by the FCC is based on the following reasoning. All receivers within the service contour (at -84 dBm) of the DTV broadcasting station should be able to properly demodulate it. The roof antennas often employed for TV-receivers have a directive gain of around 10 dB, plus a gain of 7 dB because of their elevation. A CR is assumed to have an isotropic antenna (0 dB gain). It may be inside a house, and thus experience 13 dB of attenuation due to the walls (which seems to be a rather arbitrary value). Adding all numbers results in -114 dBm [11]. The other limits are set by similar reasoning.

A typical integrated wideband CMOS-receiver noise figure (NF) is 5 dB (discussed in section 2.3.3). The noise power in 200 kHz will then be -116 dBm if the received noise power spectral density (PSD) is equal to -174 dBm/Hz. So, the SNR for detecting a wireless microphone in 200 kHz bandwidth is around 9 dB for FCC and 802.22, and around -10 dB for Ofcom. For DTV, the detection SNR ranges from -13 dB (FCC for 6 MHz bandwidth) to -20 dB (Ofcom for 8 MHz bandwidth). The spectrum sensing device should thus be able to detect signals in *highly negative SNR*.

The back off time is the time a CR may take to move to another band when a PU returns; Ofcom requires a CR to sense at a minimum interval of 1 s, hence it is inferred that the back off time should be shorter than that. With the Ofcom regulations being the most demanding for a sensing interval of 1 s, spectrum sensing should take less than 100 ms in order not to reduce the overall data rate by more than 10%.

A false alarm is the situation where a CR wrongfully concludes that a channel is occupied by a PU, whereas a missed detection occurs when a CR wrongfully concludes that a channel is free. A false alarm wastes an opportunity, while a missed detection causes harmful interference. P_{FA} and P_{MD} denote the respective probabilities that these events happen, and they are upper bounded in the 802.22 standard, which places further constraints on the sensing performance. Related to P_{MD} is the probability of detection P_D , with $P_D \triangleq 1 - P_{MD}$, which is sometimes more convenient to use.

2.2 Spectrum Sensing

One can distinguish between several main categories of spectrum sensing. There is no general consensus on the categorization; different categorizations can be found e.g. in [15, 17, 34]. In the following descriptions, the goal is to decide whether a channel is free or occupied, using the following notation: Y is the decision metric, λ the decision threshold, $r[k]$ the k -th sample at the output of the receiver, and K the total number of samples taken.

2.2.1 Energy Detection

energy detection (ED) is the most straightforward way of spectrum sensing. It measures the amount of energy (or power) in a frequency band, and compares the result to a threshold λ to determine whether a band is occupied or not:

$$Y = \frac{1}{K} \sum_{k=0}^{K-1} |r[k]|^2 \underset{\text{free}}{\overset{\text{occupied}}{\geq}} \lambda. \quad (2.1)$$

Due to variance in the measurement, λ should be somewhat higher than the noise power, but its exact value for given desired P_{FA} and P_{MD} depends on the distribution of Y in both cases. The required number of samples scales with $\mathcal{O}(1/\text{SNR}^2)$ [35], which means that 3 dB lower SNR requires 4 times more samples for the same detection performance.¹

2.2.2 Coherent Sensing

Many modulated signals contain certain characteristics (features) that can be used to an advantage by a spectrum sensing device. Examples are the strong carrier that is present in an ATSC-signal (analog TV in the US), the packet preamble of a IEEE 802.11b packet, or the spreading sequence of a waveform. One can correlate the received signal with the known waveform at the sensing device, which is similar to the operation of a matched filter. Basically, the following metric is calculated:

$$Y = \Re \left(\sum_{k=0}^{K-1} \overline{y_{\text{ref}}[k]} r[k] \right) \underset{\text{free}}{\overset{\text{occupied}}{\geq}} \lambda, \quad (2.2)$$

where $\Re(\cdot)$ takes the real part, y_{ref} is the reference signal, and $\overline{y_{\text{ref}}}$ its complex conjugate. If a signal is present and synchronization is somehow established, $r[k]$ is equal to $y_{\text{ref}}[k] + n[k]$, with $n[k]$ additive noise, and thus $\mathbb{E}[Y] = \mathbb{E}[|y_{\text{ref}}[k]|^2]$. Without signal, $\mathbb{E}[Y] = 0$.

Theoretically, the required number of samples scales with $\mathcal{O}(1/\text{SNR})$ [35], which is significantly better than the $\mathcal{O}(1/\text{SNR}^2)$ scaling for ED. However, there is usually an initial penalty in SNR as the signal characteristic

¹If $f(x) = \mathcal{O}(g(x))$, then $\limsup_{x \rightarrow \infty} |f(x)/g(x)| < \infty$ (Landau-notation)

contains only a fraction of the total signal power. For example, in ATSC, the pilot tone contains only about 10% of the total power. This means that for medium to high SNR, ED may actually be faster than coherent sensing.

Coherent sensing requires synchronization, and specific knowledge of the signals to be detected. Moreover, it requires a different matched filter for each type of signal, which can make it very complex and less suitable for use in arbitrary frequency bands.

2.2.3 Cyclostationary Sensing

Cyclostationary sensing exploits the cyclostationary features of modulated signals, and is sometimes also referred to as feature detection. These features are caused by periodicity in the signal or in its statistics, such as the symbol rate, the use of a cyclic prefix (CPr), or the use of pilot sequences. Cyclostationary sensing is quite involved mathematically. Here only an example will be shown; more information can be found in [34] and references therein.

An OFDM-signal $r(t)$ with CPr is cyclostationary with period T_0 equal to the OFDM-symbol time including CPr. Then

$$\gamma_r(t, \tau) = \gamma_r(t + T_0, \tau) \quad \gamma_r(t, \tau) \triangleq \mathbb{E} \left[r(t) \overline{r(t + \tau)} \right], \quad (2.3)$$

with $\gamma_r(t, \tau)$ the autocorrelation function of $r(t)$.

Assume the signal is being sampled during L symbols, with exactly N samples per symbol. Of these N samples, S samples cover the data symbol, and thus $N - S$ the CPr. As the CPr is a copy of the last samples of the OFDM-symbol, the discretized version of $\gamma_r(t, \tau)$, $\gamma_r[k, \tau_d]$, is definitely nonzero for $\tau_d = S$ for some k . This $\gamma_r[k, S]$ can be estimated as

$$\hat{\gamma}_r[k] = r[k] \overline{r[k + S]}, k \in [1, (L - 1)N]. \quad (2.4)$$

Due to the cyclostationary nature of the signal, $\mathbb{E}[\hat{\gamma}_r[k]] = \mathbb{E}[\hat{\gamma}_r[k + N]]$, thus a better estimator is

$$\tilde{\gamma}_r[k] = \frac{1}{L} \sum_{l=0}^{L-1} \hat{\gamma}_r[k + lN], k \in [1, N]. \quad (2.5)$$

The CPr usually takes between 1% and 20% of an OFDM-symbol. The absolute start time of the CPr is not known in advance, so it is not certain for which $k \in [1, N]$ the estimator $\tilde{\gamma}_r[k]$ will yield non-zero value(s). A decision metric could be (alternatives are given in [34])

$$Y = \max_{\substack{\text{occupied} \\ \text{free}}} |\tilde{\gamma}_r[k]| \stackrel{\text{occupied}}{\geq} \lambda. \quad (2.6)$$

If there is only white noise present, $\mathbb{E}[Y] = 0$, which makes this technique robust to uncertainty in the noise level. The periodicity itself may be unknown, which would require an additional search.

Table 2.2: Qualitative comparison of spectrum sensing categories.

	ED	Coherent sensing	Cyclostationary sensing
Requires signal knowledge	no	yes	no ^a
Can classify signals	no	yes	yes
Works for any signal	yes	no	no
Distinguish man-made noise	no	yes	yes
Number of samples required	$\mathcal{O}(1/\text{SNR}^2)$	$\mathcal{O}(1/\text{SNR})$	$\mathcal{O}(1/\text{SNR}^2)$
Computational complexity	low	high	high
Sensitive to	noise uncertainty	frequency offset	sampling clock offset

^a Knowledge does significantly help in reducing complexity

The benefit of cyclostationary sensing is that the obtained periodicities can be used as a fingerprint for signal classification. In order to detect these cyclic frequencies, the signal often needs to be significantly oversampled. Nonlinearities in the frontend and cyclostationary noise may introduce spurious cyclic frequencies [17]. Moreover, cyclostationary sensing is sensitive to frequency offset and jitter, and the available averaging time depends on the coherence time of the channel. Similar to coherent sensing, it also requires a significant amount of processing.

2.2.4 Cooperative Sensing

In cooperative sensing, several SUs combine their findings to arrive at a more reliable decision. This can be essential in severe fading environments: if the SUs are sufficiently far apart, it is much less likely that they are all in a fading dip. Hence, P_{MD} (and/or P_{FA}) decreases significantly. The final decision can be based on hard decisions (e.g. a majority vote), or on soft decisions (including additional information). There will be some trade-off between the final decision quality, the required processing, and the required communication overhead.

Cooperative sensing can be added to further increase the detection performance of individual spectrum sensor nodes. Although its importance is noted, the development of cooperative sensing is more or less orthogonal to sensing at an individual node. This topic is considered outside the scope of this work; the interested reader is referred to [15, 17] and references therein.

2.2.5 Discussion

The different categories are qualitatively compared in table 2.2. The regulations only require a decision that a channel is either ‘free’ or ‘occupied’. However, according to [15], spectrum sensing “also involves determining what types of signals are occupying the spectrum including the modulation, waveform, bandwidth, [and] carrier frequency”. This additional information potentially allows a more efficient use of white space.




Figure 2.1: Mean man-made noise levels measured at different locations, with as reference the thermal noise power level $k_B T_K$ (Adapted from: [37]).

It can be advantageous to perform some preselection of which bands to scan, how frequently, and with what duration. These parameters could be optimized based on (inferred) knowledge or learning of PU-behavior, with regulations, such as set by the FCC, as constraints. This optimization is however considered outside the scope of this thesis.

Of the categories discussed, ED is the only option if a solution is sought that can be applied in arbitrary frequency bands without having any knowledge of the signals that could be present. However, in TVWS, man-made noise can be significant [36], and ED cannot distinguish between a PU and man-made noise. Besides black-body radiation (the source of the common reference $k_B T_K$ which is equal to -174 dBm/Hz at room temperature), there are many active radiators present. Some radiators are natural, such as lightning or stellar objects (an antenna aimed at the sun yields tremendous amounts of noise). Many man-made noise sources exist as well, either intentional (e.g. transmitters), or unintentional, such as lossy cables, spark plugs and electric fences. The sum of all these different man-made sources appears as additional thermal noise to the receiver. The power of this man-made noise is often described as a NF in dB, with as reference the thermal noise power level $k_B T_K$, using the symbol NF_α .

Man-made noise can dominate thermal noise: ITU measurements showed that around 5 MHz, NF_α can have a time-average value of 55 dB in a city environment, and in the TV-bands around 200 MHz and 400 MHz, it can average 30 dB [38]. Peak values can even be much higher than these average values [39, 40]. Due to a variety of reasons (propagation characteristics, transmitter power levels, spectral occupation, etc.), man-made noise levels decrease with increasing frequency. Judging from [40], the mean man-made noise level is below $k_B T_K$ for frequencies above 2 GHz. Some measurement results of another campaign are shown in fig. 2.1, which corroborates those observations. These results suggest that ED for a PU in




Figure 2.2: Power of the strongest TV-signal per 6 MHz TV-channel at locations in the USA. Channels 2–6 are in the VHF-low band (54–88 MHz), channels 7–13 in the VHF-high band (174–216 MHz) and channels 14–51 in the UHF-band (470–698 MHz). (Sources: Microsoft [41] and Shared Spectrum Company [42–44])

TVWS (or any frequency below a few GHz for that matter) at the sensing limits as required by the FCC, does not make any sense at all. Exploitation of known signal characteristics is thus required to meet FCC-regulations.

This does not mean that ED should therefore be discarded. In fact, the focus of this thesis is on ED: measuring the spectral content of the incoming signals, only considering time and frequency as degrees of freedom. In fact, the focus of this thesis is on ED: measuring the spectral content of the incoming signals, only considering time and frequency as degrees of freedom. The detection limit (as set by regulations) to decide between ‘occupied’ and ‘free’ implicitly takes care of the space dimension, as the received signal power decreases for higher distance. In order to better motivate the choice for ED in section 2.4, implementation considerations are discussed first.

2.3 Implementation Considerations

An implementation of ED for DSA should comply with regulations as discussed in section 2.1. Typically, with static frequency allocation, a standard also defines the type of blockers to expect at certain frequency offsets. In CR, the radio spectrum is unknown a priori. However, in order to derive specifications, it is important to get at least a rough idea of what to expect.

Fig. 2.2 shows the power of TV-signals at several locations in the USA. The data from Microsoft Research is obtained from their online White-FiFinder, which is a web-interface that calculates signal strength based on

transmission tower information (March 2010) in combination with Longley-Rice propagation models and terrain elevation data [41]. The data from Shared Spectrum Company (SSC) is based on measurements in the field, performed in the years 2004 and 2005. The measurements from SSC were performed with a standard SA, and the results in fig. 2.2 have been obtained by a rough estimation based on the graphs presented in [42–44]. The noise floor of the SA obscures signals below roughly -100 dBm in 6 MHz. The measurement and simulation results should therefore not be directly compared; what is of interest here is the general picture.

Major cities have many occupied TV-channels, including some with high power levels. In New York, simulated TV-signals reach levels around -15 dBm, while the strongest measured signal is at -20 dBm. In Chicago, measurements and simulations show several channels above -10 dBm, with the highest measured power at 0 dBm. At the other extreme, measurements in the West-Virginia National Radio Astronomy Observatory (NRAO) radio-quiet zone show a very low occupancy, and only a single channel with a power above -80 dBm. This -54 dBm signal may be an erroneous measurement, as the spectrum (see [44]) does not look like a TV-signal.

The good agreement in maximum power and overall occupancy suggests that the data obtained from the WhiteFiFinder provides a good picture of what to expect. For example, in a small city like Ithaca, NY, which has no major cities in the vicinity, signal powers between those of the major cities and NRAO are received. Simulation results from other towns (not shown here) confirm this trend.

2.3.1 Specifications

The sensing device is allowed to occasionally make mistakes (see table 2.1), so the dynamic range (DR) requirements can be reduced by ignoring situations that rarely occur, such as a 0 dBm input power. However, from fig. 2.2 one might wonder what could then be a reasonable upper limit. Since most people live in a city, it could be reasonable to take -15 dBm or -20 dBm as maximum channel power. To show the impact of this choice of maximum channel power, the performance requirements, based on the FCC-regulations from table 2.1, of an (integrated) SA using ED are derived for several maximum channel powers and shown in table 2.3. As a reference and to show that some of these requirements are really tough, the performance of some state-of-the-art commercial SAs are also listed.

The useful frequency range for operation would be roughly from 50 MHz to 6 GHz. Below 50 MHz, antennas become too large or inefficient, while above 6 GHz, line-of-sight-communication becomes a must, which is not desirable for mobile communications.

PU in TVWS occupy 200 kHz to 8 MHz in bandwidth, but PUs in other bands may use any amount of bandwidth. Because a lower bandwidth

Table 2.3: Desired performance of the integrated SA to comply with FCC-regulations, and comparison with state-of-the-art commercial SAs (typical values given for 0 dB input attenuation, and (optional) pre-amplifier off).

Spectrum Analyzer ^a	Model ^b	Type	Frequency			IIP3 [dBm]	DANL [dBm /Hz]	FoM SFDR [dB]	Phase Noise [dBc/Hz] ^d	HR [dB]	RBW		Price k€
			Low [kHz]	High [GHz]	Power [W]						Low [Hz]	High [MHz]	
D	0 dBm	IC				57	57	-169	114				
D	-10 dBm	IC				42	42	-159	104				
D	-15 dBm	IC	50000	6.0	0.05	35	-174	35	-154	99	10000	20	0.001
D	-20 dBm	IC				27	27	-149	94				
D	-30 dBm	IC				12	12	-139	84				
A	PXA-N9030A	B	0.003	3.6	450	22	-162 ^e	10	-146	?	1	8	42
R	FSU3	B	0.02	3.6	130	25	-157	8	-146	?	10	20	39
T	RSA6106B	B	9	6.2	450	20	-151	-3	-133	?	0.1	8	58
A	N9342C	H	100	7.0	15	10	-145	-19	-119	?	10	3	10
R	FSH3	H	100	3.0	7	14	-144	-16	-120	?	100	1	7
T	SA2600	H	10	6.2	?	7	?	?	-110	?	10	3	20
S	HF-60100 V4	H	1000	9.4	6	?	-155	?	?	?	1000	50	1.5
C	RFeye Node	M	10000	6.0	15	20	-164	10	-115	?	18	1.2	15
H	USB-SA44B	P	0.001	4.4	2.2	3	-144	-27	-135	?	0.1	5	0.9

^a Agilent, CRFS, Desired, Rohde & Schwarz, SPECTRAN, Signal Hound, Tektronix

^b For the desired implementation, the maximum assumed channel power is denoted here

^c Benchtop, Handheld, single IC, (outdoor) Monitoring, PC required for processing

^d At 1 MHz offset (from specification or calculated assuming $1/f^2$ dependency).

^e With NFE enabled (otherwise, DANL = -154 dBm/Hz).

means a lower datarate or more power per bit, and requires a longer sensing time, a practical lower limit on resolution bandwidth (RBW) could be 10 kHz. At the other extreme, 20 MHz covers most commonly used standards, with Ultra Wideband (UWB) and radar as notable exceptions.

The sensitivity of a SA is often thought to be equal to its displayed average noise level (DANL). This is not true: with averaging and a reasonable estimate of the DANL, it is possible to detect signals below the noise floor, as will be shown in chapters 3 and 7. If it is assumed that signals of 8 dB below the noise floor can be detected (such a power will raise the measured power by 0.6 dB), the required DANL will be -174 dBm/Hz. The distortion products in the SA should not be larger than -114 dBm to avoid obscuring the weak signals. If two signals of -20 dBm enter the re-



Figure 2.3: Examples of SAs (scale roughly 1:20): (a) Agilent PXA-N9030A, (b) Rohde & Schwarz FSH3, (c) Spectran HF-60100V4, and (d) Desired (scale 1:1)

ceiver, this means the input-referred third-order intermodulation intercept point (IIP3) must be at least $-20 + (-20 - (-114))/2 = +27$ dBm. Similar calculations for other input powers result in the values shown in table 2.3.

Phase noise manifests itself mostly around strong input signals [45]; the phase noise side skirts around a -20 dBm signal must not have more power than -114 dBm integrated over 6 MHz bandwidth to avoid obscuring the signal to be detected. If it is assumed the TV-signal can be modeled as a sine wave in the middle of the TV-channel, the phase noise has to be integrated from 3 MHz to 9 MHz offset from the carrier. At this frequency offset, it may be reasonable to assume the skirts fall off at 20 dB per decade frequency offset. Therefore, at 6 MHz offset, this requires a phase noise level of less than $-114 - (-20) - 10 \log_{10}(6 \cdot 10^6) - 2.5 = -164$ dBc/Hz. The 2.5 dB term originates from the $1/f^2$ behavior over the range of integration. Equivalently, it corresponds to -149 dBc/Hz at 1 MHz offset. Note that this calculation assumes that the strong signals at the input are not filtered out prior to downconversion.

It is instructive to compare this phase noise requirement with that of some commonly used communication standards. Referred back to 1 MHz offset, -100 dBc/Hz is required for DECT, -110 dBc/Hz for Wireless Local Area Network (WLAN) and Bluetooth, -122 dBc/Hz for UMTS, and -130 dBc/Hz for GSM [46, 47]. All of these requirements are significantly lower than stated in table 2.3.

The -20 dBm signal could also be present at any harmonic frequency of the local oscillator (LO). This means that the harmonic rejection (HR) of all harmonics should be better than $114 - 20 = 94$ dB.

The spurious-free dynamic range (SFDR), as will be explained in more detail in section 2.3.2, combines the DANL, which limits the minimum channel power due to noise contributions, and IIP3, which limits the maximum channel power due to distortion contributions, into a single number, and is a measure of the range of signal powers that can be detected at the same time. The DANL, and therefore the SFDR, depends on the RBW. Since different PUs occupy different bandwidths, it is perhaps easier to compare implementations based on SFDR by abstracting from the RBW. Therefore, $\text{FoM}_{\text{SFDR}} \triangleq \text{IIP3} [\text{dBm}] - \text{NF} [\text{dB}] = \text{IIP3} [\text{dBm}] - (\text{DANL} + 174) [\text{dBm/Hz}]$.

To get an idea of the size of modern SAs, fig. 2.3 shows a few examples. As a part meant to be included in many consumer devices, it is very desirable to have the SA integrated as part of a larger IC, and it should have a low cost, say below €1. It should be noted that for example the Agilent PXA contains roughly 20 ICs, of which the analog ones use GaAs-technology for high linearity (due to higher supply voltage) and low noise (intrinsic to the technology). Moreover, it uses mechanical attenuators to handle high power, to have low feedthrough and to have a high linearity.

The handheld SAs have a very limited thermal and power budget, resulting in lower performance. These budgets are even tighter for a single

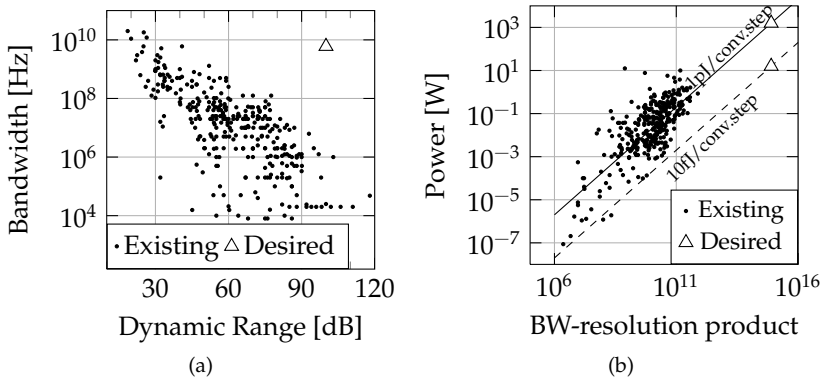


Figure 2.4: State-of-the-art ADC-performance and comparison with the desired DR = 100 dB and $W = 6$ GHz. (Source: [48])

integrated SA. To maximize battery operation, an integrated SA should not consume too much power. A regular receiver consumes roughly 50 mW and is therefore set here as the power consumption target of the SA.

2.3.2 Design Considerations

A SA processes the received signal to condition it for the power estimation process. In commercial SAs, this conditioning typically includes amplification, filtering, and frequency conversion. Apart from the final power estimate, these are the same functions that a regular radio receiver also has to perform. The power estimation can be done in the analog or the digital domain; both approaches are used in commercial SAs. In both cases, but especially when the power estimation is performed in the digital domain, the analog functions of a regular receiver and a SA are very similar, except perhaps for the specifications. For example, SAs typically have a higher linearity and NF than radio receivers.

This section provides a brief overview of receiver design to introduce some concepts related to receivers and SAs, and to show that a receiver is not simply an antenna with an analog-to-digital converter (ADC).

Direct Digitization

With the tremendous performance improvement of digital circuits over the years, one would like to do as much as possible in the digital domain. An ADC placed directly behind the antenna can convert the received analog signal to the digital domain, after which the demodulation or spectrum sensing algorithms can be executed. The ADC would then require 17 bits at a sample rate of 6 GS/s (assuming quadrature receivers), which is far from the specifications of ADCs that exist today, see fig. 2.4a. If historical trends are extrapolated, e.g. the observed doubling of bandwidth-resolution prod-




Figure 2.5: Block diagram of a direct-conversion receiver.

uct of ADCs every four years [49], such ADCs will become available only as early as the year 2055.

Even if such an ADC would exist, the power consumption would be somewhere between 16 W (assuming a very optimistic 10 fJ per conversion step) and 1.6 kW (extrapolation of the current 1 pJ per conversion step), see fig. 2.4b. This is far more power than desired (see table 2.3). Hence, filtering and frequency conversion are generally used to alleviate the requirements on the ADC.

Direct-conversion Receiver

A simplified diagram of a traditional receiver is depicted in fig. 2.5. After the antenna, an external filter (often a surface acoustic wave filter) blocks or significantly attenuates all signals outside the band of interest. The weak signal to be received is then first amplified by a low-noise amplifier (LNA) to make the noise contributions of subsequent stages insignificant; in other words, the LNA tries to keep receiver NF low. Many integrated solutions employ a direct-conversion architecture, so the LNA is followed by a quadrature mixer to mix the signal to baseband. A low-pass filter removes the high-frequency components, and the variable-gain amplifier amplifies the signal to use as much of the input range of the ADC as possible.

Unlike a traditional receiver, a CR-receiver (and the SA) needs a wide-band radio frequency (RF)-section and widely tunable internal frequency generation, and has to operate without an external RF filter bank for cost and size reasons. This makes the analog/RF part of a receiver significantly more challenging. The spectrum sensing implementation may use a separate receive chain and perhaps even a separate antenna (see e.g. [35]), but it may also simply reuse the existing receiver. In any case, its performance will be degraded by the same mechanisms that degrade the performance of a CR-receiver.

Spurious-Free Dynamic Range

For faster spectrum sensing, one would like to sense multiple channels at once. White space may be located just next to a very strong signal, so the SA should have both a high linearity and a low NF. An LNA can ensure low NF, but the amplified strong signal may cause the receiver to enter the nonlinear region, or even clip to the supply voltage.

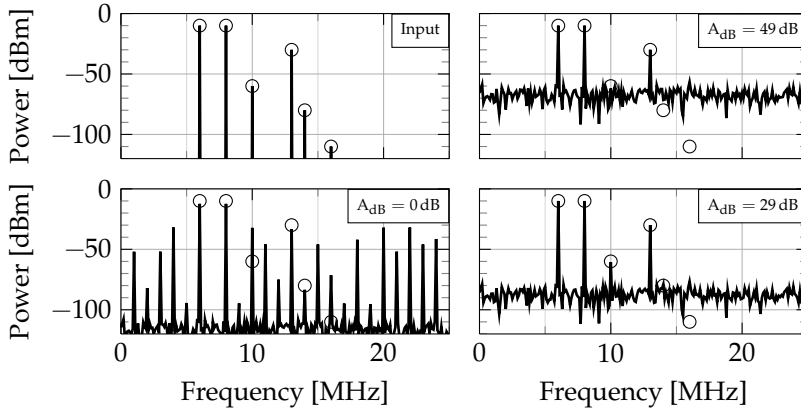


Figure 2.6: Simulation of a SA with $NF=5$ dB and $IIP3=+1$ dBm. The RBW is 100 kHz for a SFDR of 80 dB. Because of the limited SFDR it is not possible to detect all input signals at the same time.

The SFDR defines the difference in decibels between the strongest and weakest signal that can be detected *at the same time* [50].² Thus, the SA requires a high SFDR. The SFDR is limited by non-linearity and noise, but also by any spurious components (‘spurs’), e.g. from the phase-locked loop (PLL) or from harmonic downmixing. If the spurs can be reduced to negligible levels, and $IIP3$ is the limiting linearity factor, the following equation for SFDR can be derived [50, 51]:

$$SFDR = \frac{2}{3} (IIP3 - NF - 10 \log_{10} RBW + 174) \text{ [dB]}. \quad (2.7)$$

A lower RBW means that less noise power will be present in such a band, and hence the SFDR increases. The useful increase in SFDR by lowering RBW is limited by the bandwidth of the signals to be detected; at some point, the signal power will also drop, such that the SNR is not further increased when RBW is lowered. (2.7) assumes that the input signal can be attenuated to a level where the intermodulation components are at the same level as the noise floor.

The balancing of noise and intermodulation products is illustrated in fig. 2.6. Here it is assumed that the gain control is implemented by an attenuator at the input of the SA. When the linearity is limited by the LNA, it can only be improved by attenuating the signal in front of the LNA. Assuming a matched system and an ideal attenuator, x dB of attenuation raises both NF and $IIP3$ by x dB. At the input of the SA (top-left), six sine waves are present (circles indicate their power levels for easy reference). At the output, the spectrum looks quite different and depends on the attenuation (the power levels are referred to the antenna input). At low

²Definitions of SFDR differ between research fields and even between authors in the same field. Here the definition for the SFDR in SA-datasheets is used.

attenuation (bottom-left), the strongest signals generate many intermodulation products, which may lead to false alarms. At high attenuation (top-right), the increased noise obscures weak signals, which may generate missed detections. Even at the optimum attenuation (bottom-right; 29 dB here) where the noise and distortion products are at the same level and the SFDR of (2.7) is obtained, some signals cannot be detected.

The optimum attenuation $A_{\text{dB,opt}}$ can be calculated quickly and easily in the digital domain, as the SA roughly knows its NF and IIP3, and the strong input signals are readily detected. Assuming two input signals with power $P_{\text{in,max}}$ dBm, the strongest third-order intermodulation (IM3)-component is at $P_{\text{in,max}} - 2(\text{IIP3} + A_{\text{dB}} - P_{\text{in,max}} \text{ dBm})$, while the noise floor is at $-174 + \text{NF} + \text{RBW} + A_{\text{dB}}$ dBm. Equating these two results in:

$$A_{\text{dB,opt}} = P_{\text{in,max}} - \frac{2}{3}\text{IIP3} - \frac{1}{3}(-174 + \text{NF} + \text{RBW}) \text{ [dB]}. \quad (2.8)$$

If signals several dB below the noise floor can be detected, or intermodulation products can be predicted and (partially) corrected for, $A_{\text{dB,opt}}$ may be different. Nevertheless, the principle of balancing between the noise and distortion peaks remains the same.

Harmonic Rejection

Ideally, the desired signal is downconverted by multiplication with an exponential signal (the LO, with $\text{LO}_{\text{ideal}} = \exp(j2\pi f_{\text{LO}}t)$). In practice, the mixer output will also contain the signals that were present at higher harmonics of the LO, because the mixing is performed with non-linear components and/or time-variant systems.

In modern receivers, a square-wave-like digital LO is often used. One reason is that a digital square wave results from employing digital techniques to generate an LO. Another reason is the good performance of hard-switching passive mixers, which are simply switches in the signal path that are turned on and off by a digital LO. MOS-devices are very good at switching, and will only get better as technology improves. Moreover, passive mixers do not carry DC-current, and therefore suffer far less from flicker noise.

Mathematically, for a quadrature square wave with amplitude ± 1 :

$$\begin{aligned} \text{LO}_{\text{sq}}(t) &= \frac{4}{\pi} \sum_{h \in \mathbb{Z}} \frac{1}{4h+1} e^{j(4h+1)2\pi f_{\text{LO}}t} \\ &= \frac{4}{\pi} \text{LO}_{\text{ideal}}(t) + \frac{4}{\pi} \sum_{h \in \mathbb{Z} \setminus \{0\}} \frac{1}{4h+1} e^{j(4h+1)2\pi f_{\text{LO}}t}. \end{aligned} \quad (2.9)$$

The LO has essentially become a parallel combination of multiple LOs. The 3rd harmonic is only 10 dB weaker than the fundamental; in other words, the HR of the 3rd harmonic, HR_3 , is only 10 dB.

The lack of an external filterbank results in signals present at those harmonic frequencies to be downconverted as well. This is known as

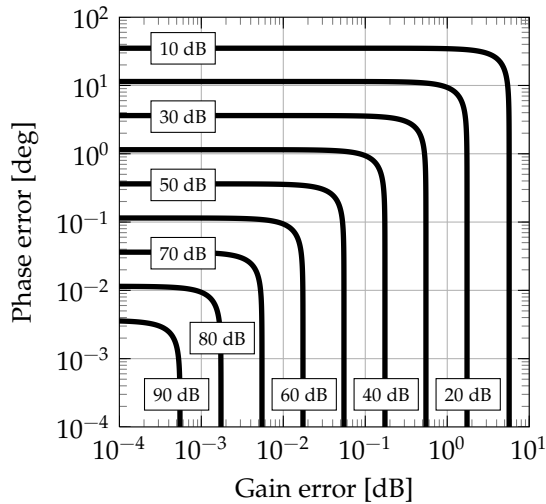


Figure 2.7: Required gain and phase accuracy for a certain amount of IR.

harmonic downmixing. A SA with poor HR trying to sense at 300 MHz will mistakenly identify it to be occupied when a strong signal is present at 900 MHz, leading to a false alarm.

The image (the -1^{st} harmonic) is rejected by properly combining the I- and Q-signals after downconversion. In practice, the I- and Q-signals are not perfectly 90° out of phase and have not experienced exactly the same gain. If this imbalance is referred back to the LO, it can be modeled as

$$\begin{aligned} \text{LO}_{\text{IQ-imbalance}}(t) &= \cos(2\pi f_{\text{LO}}t) + j(1 + \Delta g) \sin(2\pi f_{\text{LO}}t + \Delta\phi) \\ &= \text{LO}_{\text{ideal}}(t) + \frac{1 - (1 - \Delta g) e^{-j\Delta\phi}}{1 + (1 + \Delta g) e^{j\Delta\phi}} e^{-j2\pi f_{\text{LO}}t}. \end{aligned} \quad (2.10)$$

The image rejection (IR) can be calculated as a function of gain error ($20 \log_{10}(1 + |\Delta g|)$) and phase error ($|\Delta\phi| \cdot \pi/180$), as shown in fig. 2.7. For 90 dB IR, a maximum gain error of 0.0005 dB and phase error of 0.003° is required.

In direct-conversion receivers, the 30 dB to 40 dB of IR ($= \text{HR}_{-1}$) that is readily obtained is often sufficient, as the image has the same power as the signal. However, when a SA analyzes several channels at once to speed up the detection process, the image of a channel can be another channel, which could contain a much stronger signal. Therefore, the IR should be as good as the rejection of the other harmonics.

Frequency Generation

Frequency conversion requires an LO, which needs to be generated somehow. Most standards, such as GSM and WLAN (but with the notable exception of UWB), have an RF-bandwidth less than 10% of their center frequency. A low-power LC-oscillator can then be implemented with good

phase noise characteristics, and made tunable with some variable capacitors to cover the whole band. Integrated inductors are, however, relatively large; an inductor easily occupies $200\ \mu\text{m}$ by $200\ \mu\text{m}$, an area which can also house a (simple) microprocessor in modern CMOS-technology.

For these oscillators, a tuning range of more than 15% is possible, but there is a trade-off between tuning range and phase noise. A wideband receiver with many tunable LC-oscillators in parallel would become too bulky and expensive. Moreover, integrated inductors do not work well below 1 GHz and have issues related to magnetic coupling. A possibility is to make a very high frequency oscillator, and divide this frequency by 2, 3, 4, 5, etc. to cover all the desired frequency bands. With 10% tuning range, the oscillator must then be running in the order of 100 GHz to cover all frequencies below 10 GHz. This is challenging and tends to result in high power consumption. Ring-oscillators, on the other hand, are more easily tunable and can operate at low frequencies. However, they require an order of magnitude more power for the same noise performance.

Due to process spread, both types of oscillators require a reference frequency (usually an external crystal with a frequency below 100 MHz) for sufficient absolute frequency accuracy. One usually employs a PLL to connect the low-frequency crystal with the higher-frequency voltage-controlled oscillator (VCO). It has been shown that doubling the power in the PLL can reduce the PLL noise by 3 dB [52]. This suggests that with enough power available, the phase noise can be reduced to arbitrary low levels. However, at some point the noise of the reference becomes the limiting factor. The noise from the reference is reduced (filtered) outside of the bandwidth of the PLL, so there may be a possibility to reduce the PLL-bandwidth. This idea is, however, not further explored in this thesis.

At lower frequencies (up to several GHz), a direct digital synthesizer can be used, which consists of a frequency reference and digital logic. Followed by a digital-to-analog converter (DAC) and reconstruction filter, any frequency can be generated with a resolution set by the reference frequency. It has the additional advantage that it can change its frequency almost instantaneously. The downside is the relatively high power consumption and the presence of spurious frequency content due to the digital implementation.

2.3.3 State-of-the-art Integrated Solutions

The concept of DSA is relatively new, and only a few publications target the analog implementation of spectrum sensing functionality. These publications will be briefly discussed here.

Spurious-Free Dynamic Range

With respect to SFDR, reducing the gain of the receiver when the sensing performance is limited by distortion, and increasing it when it is limited

Table 2.4: Wideband CMOS-receiver performance overview

Reference	Node [nm]	RF Freq. [GHz]	Power [mW]	Gain [dB]	NF [dB]	IIP3 _{ib} [dBm]	IIP3 _{ob} [dBm]	OIP3 _{ib} [dBm] ^a	FoM _{SFDR} [dB]
Soer [53]	65	0.2–2.0	67	19	6	11		30	5
Ru [54]	65	0.4–0.9	60	34	4	4	16	38	0
Greenberg [55]	80	0.04–1.0	440		3	-15			-18
Murphy [56]	40	0.08–2.7	78	70	2	-22 ^b	14	48 ^b	-24
Youssef [57]	65	0.1–2.5	62	30	8	-20	20	10	-28
Andrews [58]	65	0.1–1.0	50	70	4	-67	25	3	-71

^a A dBm-scale may not be the most appropriate for out-of-band IIP3 as the impedance differs from 50 Ω . However, it can be defined as the power available from a source with 50 Ω output impedance: a receiver with high-ohmic input impedance then sees double the voltage.

^b Measured just outside IF-bandwidth; true IIP3_{ib} and OIP3_{ib} expected to be worse.

by noise, seems the obvious solution. In the extreme, the LNA is bypassed entirely, which allows very linear receiver implementations at the cost of NF, as e.g. in [53] (NF = 6.5 dB, IIP3_{ib} = +11 dB).³ An alternative is to stay in the current domain for as long as possible, as done in [54] (NF = 4 dB, IIP3_{ib} = +4 dBm, IIP3_{ob} = +16 dBm)³, such that the first significant voltage swing occurs at IF where linearity can be improved using feedback. In both cases, final linearity is limited by IF-circuitry.

Table 2.4 gives an overview of the performance of some wideband CMOS-receivers in terms of NF and IIP3, and the SFDR calculated from NF and IIP3_{ib}. Since a completely integrated solution requires some variable gain, it is hard to directly compare NF and IIP3: the change in NF and IIP3 depends on where the gain is changed. Therefore, it can be useful to also compare OIP3_{ib}. Furthermore, many solutions are targeted towards single-channel selection, where IIP3_{ib} may be less of a concern than IIP3_{ob}. Nevertheless, these numbers show that no implementation comes even close to the FoM_{SFDR} derived in table 2.3.

Harmonic Rejection

Weldon [59] tackles the problem of HR by combining several mixers in such a way that the total conversion wave looks more like a sine wave. The basic idea is depicted in fig. 2.8. It separately amplifies the (differential) input signal with three (differential) amplifiers, with a relative gain of 1, $\sqrt{2}$, and 1. It then mixes the amplifier outputs down with a multi-phase LO, such that the (even harmonics and the) 3rd and 5th harmonics are canceled, leaving the 7th harmonic as the first uncanceled one. Mismatch limits the rejection of the 3rd and 5th harmonics to typically 30–40 dB [54, 56, 58]. Nevertheless, the suppression strongly relaxes RF-filter requirements, allowing a few

³The in-band IIP3 is denoted with IIP3_{ib}, and out-of-band IIP3 with IIP3_{ob}. In-band means both signals and IM3 products fall within the intermediate frequency (IF) or baseband range, while out-of-band means that both signals are outside of the IF or baseband range, with the IM3-product falling in-band.




Figure 2.8: A HR-mixer removes several harmonics of the LO, including the 3rd and 5th, to mitigate harmonic downmixing.

external wideband filters, or lower-Q integrated filters, to further reduce the signals at these higher-order harmonic frequencies.

A two-stage technique based on the same principle is presented in [54], robustly obtaining 60 dB HR. To accurately define the different LO phases, the power consumption is high. A HR of 80 dB is obtained by adaptive interference cancellation, as also demonstrated in [54]. Alternatively, a tracking filter can be used in combination with a HR-mixer [55, 60]. Background calibration in combination with calibration signals can further improve overall HR and IR [55].

An overview of the HR-performance of several state-of-the-art receivers is shown in table 2.5. Virtually all of them are zero-IF receivers and do not care about IR, while only one of them explicitly cares about the 7th and 9th harmonic, all of which are important for a wideband CR. Clearly, all results are still far away from the HR-requirements of table 2.3.

Frequency Generation

As discussed, LC-oscillators exhibit significantly lower phase noise than ring-oscillators, but have only limited tuning capability. To save chip area,

Table 2.5: HR-performance of several state-of-the-art CMOS-receivers

Reference	Node [nm]	Cal ^a	Freq [GHz]	Power [mW]	IR [dB]	HR ₂ [dB]	HR ₃ [dB]	HR ₄ [dB]	HR ₅ [dB]	HR ₆ [dB]	HR ₇ [dB]	HR ₈ [dB]	HR ₉ [dB]
Greenberg [55]	80	Y	0.04–1.0	440	65	65	65	65	65	65			
Cha [60]	180	Y	0.05–0.9	140			72		70		60		60
Ru [54] ^b	65	N	0.4–0.9	60		62	60	62	63	62			
Ru [54] ^c	65	N	0.4–0.9	60		64	36/80 ^d	64	38/80 ^d	64			
Zhang [61]	65	N	0.05–0.2	8/7 ^e			56/76 ^e		56/75 ^e				
Andrews [58]	65	N	0.1–1.0	50			35		43				
Ryu [62]	130	N	0.04–0.9		60		69						
Murphy [56]	40	N	0.08–2.7	78			35		43				

^a Calibrated: yes (Y) or no (N)

^b In the two-stage HR-scheme

^c In the one-stage HR-scheme with adaptive interference cancellation

^d Either the 3rd or the 5th harmonic is suppressed by 80 dB

^e Without/with tracking filter enabled

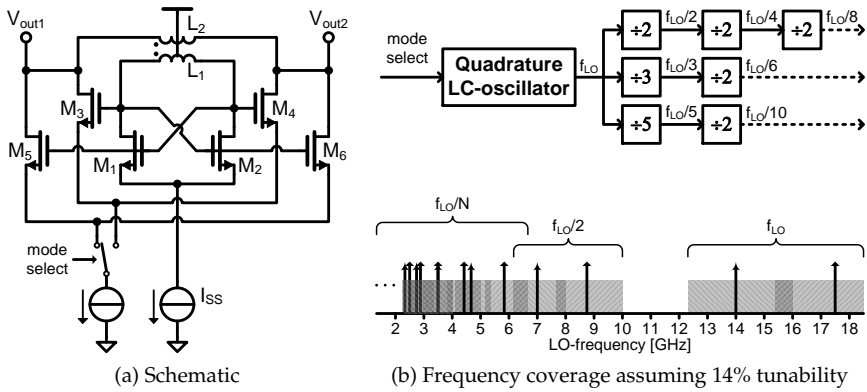


Figure 2.9: Nested inductors with sign-change option of the mutual coupling is area-efficient and allows dividers to run at a lower frequency [63].

without sacrificing frequency coverage, one can use two LC-oscillators, each tuned to such a frequency that any desired frequency can be obtained by selecting the appropriate oscillator in combination with an appropriate division ratio. In [63], this idea is taken one step further by using nested inductors, as shown in fig. 2.9a. By changing the sign of the mutual coupling between the inductors via a mode select bit, the LC-oscillator oscillates at either 14 GHz or 17.5 GHz. With 14% tuning range (not implemented in [63], but the principle is shown in fig. 2.9b) and integer frequency division, this bimodal oscillator covers all frequencies below 10 GHz. Even without this tuning, though, the phase noise is between -91 dBc/Hz to -120 dBc/Hz at 1 MHz offset, which is far from the specifications listed in table 2.3.

Phase noise and spurs impair the spectrum sensing process, and both should therefore be kept sufficiently low [64]. One of the best PLLs in terms of phase noise and spurious tone performance is presented in [65]. It uses an integer- N division for low spurs, adds spur-reduction techniques to suppress spurs below -76 dBc, and consumes 3.8 mW at $f = 2.21$ GHz for an integrated jitter of 0.3 ps [52]. It exhibits a constant phase noise of -120 dBc/Hz up to 6 MHz frequency offset.

Dividing the output frequency by N reduces the noise power by N^2 [52]. If a center frequency of 550 MHz is assumed (in the middle of TVWS), the phase noise of [65] reduces to -136 dBc/Hz. Integrating this over 6 MHz bandwidth gives -68 dBc, which for the -20 dBm input signal scenario of table 2.3 is still 26 dB short of the requirements. Assuming power scaling is possible (which is highly questionable for an integrated LC-oscillator, due to the parasitic resistance and the three-dimensional magnetic field of an inductor), the required power consumption becomes a prohibitive 1.5 W.

It must be noted, that even an expensive high-performance crystal oscillator (with buffer), e.g. the 55 MHz \$500 Wenzel Sprinter, exhibits

–170 dBc/Hz of phase noise in the MHz-range offset. Integrating this over 6 MHz at 550 MHz gives –82 dBc of phase noise power in the adjacent channel, which is still 12 dB short of the –20 dBm scenario. However, the noise may be caused by the buffer electronics, so there may be a way to reduce this noise further by spending more power. Clearly, it would be very desirable to have an alternative solution to reduce the phase noise requirements listed in table 2.3.

Analog Solutions for DSA

In this section, solutions proposed in literature, aiming to improve the analog part of ED, are briefly discussed.

RSSI-detectors have been around for quite some time, and are mainly used in a feedback loop for automatic gain control. Received Signal Strength Indication (RSSI)-detectors usually consist of a cascade of gain stages (limiting amplifiers), each followed by a rectifier. The outputs of the rectifiers are summed and low-pass filtered, resulting in a (near)-DC voltage indicating the input power. In principle, RSSI-detectors can be considered single-channel SAs, with an output equal to the total input power of the detector. RSSI-detectors often use logarithmic amplification for high DR. The RBW is usually determined by the receiver frontend. The typical DR is 60 dB to 80 dB with a relative error of ± 1 dB.

An interesting variant in the context of DSA is proposed in [66]. Here, the rectifiers are programmable to provide several gain settings. In this way, the DR and sensitivity of the detector can be traded off depending on input conditions. With all the different settings of the detector and the frontend of the receiver, signals up to +18 dBm can be detected. The RBW is determined by the frontend and is selectable between 0.2 MHz and 30 MHz. At minimum RBW, the maximum sensitivity is –83 dBm, limited by the maximum gain of the detector.⁴ The receiver NF of 5–8 dB suggests there is more sensitivity to be obtained by adding more amplifier stages in the RSSI-detector.

Analog windowing is proposed in [67] to reduce the requirements on analog filters. Ideally, an ADC captures a large chunk of spectrum, after which e.g. a fast Fourier transform (FFT) is used to process all the channels inside this chunk in parallel to minimize measurement time. As is well known, an FFT suffers from spectral leakage due to the truncation in the time domain. Therefore, the digital samples are often windowed to reduce this effect: windowing has a filtering effect. Windowing the input signal can also be done in the analog domain using an analog multiplier or correlator, which is proposed and implemented in [67].

⁴Personal communication with Masaki Kitsunezuka.

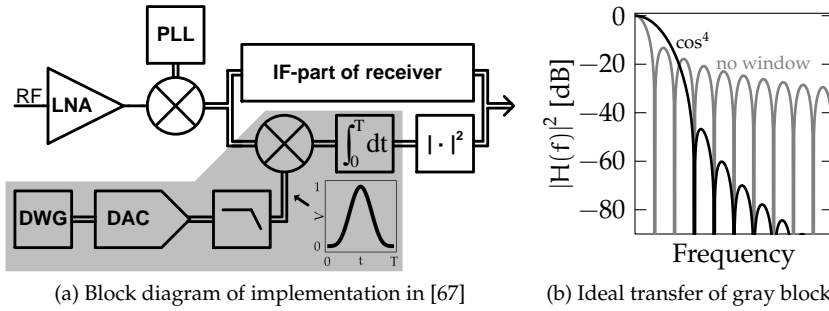


Figure 2.10: Analog windowing to reduce analog filtering requirements. The receiver and SA share the RF-part, but have a separate baseband implementation. (Adapted from [67])

A block diagram is shown in fig. 2.10a. Integrated together with a CMOS UHF-band receiver (also presented in [67]), the SA shares the RF-part with this UHF-receiver, but uses a separate baseband path including the analog windowing. The output of the digital window generator (DWG) (a digital memory) is converted to the analog domain by a DAC and low-pass filter, as shown in fig. 2.10a. This digital generation allows the characteristics, such as bandwidth, window length, and out-of-band suppression to be very flexible, reducing or eliminating the use of bulky analog filters. The measured suppression of adjacent channels is 35 dB (as opposed to the 13 dB that would be expected if no filtering was used, as shown in fig. 2.10b). The measured sensitivity is -74 dBm with a 100 kHz \cos^4 window. The sensitivity is likely limited by the noisy analog correlation process. The DWG consumes 29 mW and the correlator 14 mW.

Channelized receivers are discussed by [68] in the context of UWB. By handling several separate frequency channels in parallel in the analog domain, the requirements per path and per ADC are relaxed. When each path is used for single-channel sensing and uses a zero-IF architecture, each channel becomes its own image: IR problems are relaxed. This approach is taken in [69], where a 16-point analog FFT is implemented exploiting charge-sharing between capacitors. Mismatch limits the achieved signal-to-noise-and-distortion ratio to around 45 dB. The performance improvement will be limited due to the poor scalability of this approach: each path requires a different LO-frequency and additional chip area for filtering.

Detection by VCO-locking is proposed by [70]. The VCO is locked by the input signal, and subsequent circuitry allows the locking to be detected, effectively detecting the input signal. In this way, the prototype implemented in $0.18 \mu\text{m}$ CMOS is able to detect sine-like signals down to -75 dBm. However, it is unclear how well it can detect more noise-like sig-

nals such as OFDM. Moreover, weak signals will probably not be detected in the presence of a strong one, as the VCO will lock to the strongest signal.

Mimicking of the human cochlea is proposed by [71], inspired by the 120 dB of DR at only 14 μ W power consumption of the human cochlea. The prototype covers the spectrum from 600 MHz to 8 GHz. It has exponentially spaced output channels based on a tapered transmission line, and processes all channels in parallel. Where the power consumption and number of stages in the FFT scale with $N \log N$, N being the number of output channels, they scale with N for the cochlea. Implemented in 0.13 μ m CMOS, the tapered transmission line is approximated using lumped elements. The chip consumes 300 mW, has 50 output channels, and occupies 4.5 mm². Due to the lumped-element implementation of transmission lines, the implementation is very inflexible w.r.t. RBW, number of stages, etc. Also, the measured channel selectivity is very low: low frequencies give -20 dBc responses in the high-frequency stages.

Low-power low-resolution autocorrelation for built-in self-test (BIST) is patented by [72]. The goal is an integrated CMOS-SA for BIST to “1) minimize the amount of on chip analog signal processing and 2) minimize the complexity of a spectrum analyzer possibly at the expense of time required to measure the spectrum.” The way to achieve this is “based on that of digital autocorrelation of coarsely quantized signals.” The autocorrelation function thus obtained can be internally or externally Fourier-transformed to obtain a spectrum estimate. Coarse quantization tends to introduce spurious components (see appendix A), but the nonlinear quantization process can be inverted when all input signals are known to have a Gaussian distribution, see e.g. [73]. This latter requirement is something that seems to have been missed by the authors of [72].

Concluding Remarks

It is highly unlikely that a CMOS-implementation can achieve the desired performance of table 2.3. Even the benchtop SAs, which are optimized for performance, do not achieve any of the requirements besides RBW and RF-range. The most important identified shortcomings are the SFDR, the phase noise performance, and the HR.

2.4 Motivation for Energy Detection Spectrum Sensing

Current regulations in TVWS impose extremely tough requirements on the spectrum sensing receiver, which are unlikely to be met by an integrated CMOS-solution employing ED. Furthermore, it was argued that man-made noise in TVWS will make detection of very weak signals impossible using ED. Each argument by itself may be enough to refrain from using ED for

DSA, so why is this still pursued in this thesis? A counterargument is that some form of hierarchical sensing may be used: strong signals are easily detected with ED, and, assuming they exist, algorithms more resilient to the non-idealities of the spectrum sensing receiver may be used to zoom in on certain parts of the spectrum.

There are, however, much stronger arguments. It is important to stress that the regulations in TVWS are the only regulations currently available for DSA. Since DSA is a rather new concept, the PUs actively object against any threat of interference, which has resulted in these extreme requirements.

It is, however, not unlikely that DSA will be employed in many more bands in the future. To refrain from infrastructure demands, spectrum sensing then seems to be the only solution. ED is able to work in any band without prior knowledge of the signals in that band. With the ever-increasing demand for wireless capacity and the notion that interference tolerance of PUs greatly enhances overall capacity, the detection limits can be significantly relaxed. At multi-GHz frequencies, where most of the usable radio spectrum is, man-made noise is significantly less important, especially in the context of higher detection limits.

Moreover, modern communication signals are getting closer and closer to maximum entropy, and thus to a Gaussian probability density function (pdf) according to information theory, for which ED performs very close to the optimum detector [34]. Finally, the low computational complexity of ED is advantageous for the power consumption (although a comparison in energy consumption would be more fair, but this will highly depend on the spectrum occupancy and implementation details).

Besides DSA-applications, an integrated SA may have many more advantages. For example, in a transmitter, digital predistortion and spur reduction can be tuned on-the-fly by measuring its output spectrum. Another example is BIST: with higher integration levels and higher frequencies, it becomes more and more difficult to externally measure internal signals for test or calibration purposes, as each requires high-speed analog I/O connections. An integrated SA can enable BIST by connecting it to multiple internal nodes. The results can be internally used, or a digital output can be provided to the outside world for monitoring purposes.

Furthermore, even a SA that cannot detect extremely weak signals will have many advantages. The SA can alleviate the requirements of the CR-receiver, for example by finding strong interferers (blockers) in order to remove them with a tunable notch filter [74]. In [75], it is argued that filtering is essential to be able to use white space with reasonable receiver linearity requirements. Similarly, [76] shows that many white spaces pose such significant linearity requirements on the receiver, that the receiver would be better off not using them at all (for example when the IM3-product of two strong TV-signals would fall in this white space). This can only be determined if a true picture of the spectrum is obtained, without distortion components.

2.5 Conclusions

DSA is viewed by many, including regulatory authorities, academia and industry, as an important way towards more efficient use of the spectrum. Authorities are currently working on regulations in the TV-bands, and several communication standards incorporating DSA are in the draft phase. The regulations on spectrum sensing are very challenging, requiring weak signals of -114 dBm to be detected in negative SNR with strong signals (which can be as strong as -10 dBm) in adjacent channels. To benefit from the advantages of DSA in the near future, the most recent regulations favor a database approach over spectrum sensing. It is nevertheless still possible to use a sensing-only device (with or without using the database), as long as the regulations are complied with. More importantly, the authorities do recognize that sensing may be the only solution if, in the future, the concept of DSA is to be applied to other frequency bands as well.

Several categories of spectrum sensing have been discussed: ED, coherent sensing, cyclostationary sensing and collaborative sensing. The latter is a higher layer concept, relying on combining the results of individual nodes. Of the other categories, ED is the most general solution as it is the only one that does not require knowledge of the signals to be detected. This is a desirable feature in bands that may contain many different modulation schemes, such as in the ISM-bands. Moreover, coherent and cyclostationary sensing are more computationally complex.

Several analog impairments, such as the noise floor of the receiver, phase noise of the LO, limited linearity of the analog frontend, and limited HR, have been identified that hamper the detection process by causing false alarms and/or missed detections. These impairments are more severe for wideband receivers, such as a CR. It has been shown that the requirements in TVWS are so strict that it is highly unlikely that an integrated CMOS-solution employing ED for spectrum sensing is feasible with the specifications derived from the FCC-requirements.

Nevertheless, several other advantages have been identified when spectrum sensing can be performed, for CR-receivers as well as regular receivers. These advantages include the classification of white space in terms of receiver requirements, as well as the opportunity to tune receiver settings based on the radio environment. Therefore, even though the specifications in table 2.3 are unlikely to be achieved, it is certainly beneficial to improve the performance parameters as far as possible.

CROSSCORRELATION SPECTRUM SENSING

From chapter 2 it is clear that there are significant challenges involved in implementing a spectrum analyzer (SA) in CMOS. The main concerns are the noise floor, nonlinearity, phase noise, spurs, and limited harmonic rejection (HR). In this chapter, an energy detection (ED)-technique is proposed to mitigate many of these issues by employing two receivers in parallel and crosscorrelating their outputs. This technique will be referred to as crosscorrelation (XC) spectrum sensing, and the device implementing this as a crosscorrelation spectrum analyzer (XCSA).

In order to understand the way XC spectrum sensing works, ED is first discussed in greater detail. The principle of XC is elaborated next, and its detection performance is compared to standard ED. The way XC can alleviate some of the implementation challenges discussed in chapter 2 is described in section 3.3, while conclusions are drawn in section 3.4.

3.1 Energy Detection

A SA measures the total power in a certain bandwidth, which is also the first step in ED. For this discussion, it is assumed that the bandwidth (or a filtered version or subband thereof) is captured by the complex baseband receiver output $r[k]$. First, the energy detector, also referred to as a radiometer, performs the following operation:

$$Y = \frac{1}{K} \sum_{k=0}^{K-1} |r[k]|^2. \quad (3.1)$$

The $1/K$ normalization makes this an average power estimate rather than an energy estimate. Define $\hat{P}_{ac} \triangleq Y$ as the estimator of the received power.

Parts of this chapter have appeared in [MOA:1, 2, 6, 9, 12].

Where the SA in the lab displays the measured power, the detector also has to decide whether a signal $s[k]$ is present or not. In other words, it has to decide between two hypotheses:

$$\mathcal{H}_0 : r[k] = n[k] \quad (\text{only noise}) \quad (3.2)$$

$$\mathcal{H}_1 : r[k] = s[k] + n[k] \quad (\text{signal plus noise}). \quad (3.3)$$

In a general situation, $s[k]$ may be different from the transmitted signal due to the wireless channel: $s[k] \triangleq h[k]s_{\text{tx}}[k]$, with $s_{\text{tx}}[k]$ the actual transmitted signal, and $h[k]$ the complex channel gain to account for fading and power loss due to the distance. It is generally assumed that the signal and noise are independent, and that their statistics are at least (jointly) wide-sense stationary. Furthermore, it is often assumed that the noise samples are uncorrelated, i.e. $\mathbb{E}[n[k]\overline{n[l]}] = 0$ for $k \neq l$.

The problem now is to set a threshold λ_{ed} (from here on, a subscript is used to distinguish between different thresholds) such that

$$P_{\text{FA}} \triangleq \mathbb{P}(Y > \lambda_{\text{ed}} | \mathcal{H}_0) \leq P_{\text{FA,des}} \quad P_{\text{D}} \triangleq \mathbb{P}(Y > \lambda_{\text{ed}} | \mathcal{H}_1) \geq P_{\text{D,des}}, \quad (3.4)$$

where $\mathbb{P}(\cdot)$ is a conditional probability, and the subscript “des” indicates desired values. It can be seen that the probability of false alarm P_{FA} only depends on the noise statistics, while the probability of detection P_{D} also depends on the distribution (and thus the power) of the input signal. Since the power of the signal is not known a priori, one usually determines λ_{ed} by setting P_{FA} equal to $P_{\text{FA,des}}$. This implies that the obtained P_{D} is always at least equal to $P_{\text{D,des}}$, because whether a signal is present or not, the detector will occasionally decide a signal is present.

Many papers (e.g. [77–79]) assume that $s[k]$ is deterministic with power σ_s^2 , which is for example true for an additive white Gaussian noise-channel with M -ary phase shift keying-signals. If the noise is modeled as circularly symmetric zero-mean complex Gaussian noise, i.e. as $n[k] \sim \mathcal{CN}(0, \sigma_n^2)$, and the signal experiences block fading ($h[k] = h$), (3.1) will have a central χ^2 -distribution for \mathcal{H}_0 and a non-central χ^2 -distribution for \mathcal{H}_1 , both with $2K$ degrees of freedom [79]. One can then find (with $\text{SNR} \triangleq \sigma_s^2/\sigma_n^2$ as the instantaneous signal-to-noise ratio (SNR)) [79]:

$$P_{\text{FA}} = \frac{1}{\Gamma(K)} \Gamma\left(K, \frac{K\lambda_{\text{ed}}}{\sigma_n^2}\right) \quad P_{\text{D}} = Q_K\left(\sqrt{2K \cdot \text{SNR}}, \sqrt{2K\lambda_{\text{ed}}/\sigma_n^2}\right), \quad (3.5)$$

where $\Gamma(\alpha, \beta)$ is the upper incomplete gamma function defined as

$$\Gamma(\alpha, \beta) \triangleq \int_{\beta}^{\infty} t^{\alpha-1} e^{-t} dt, \quad \Gamma(\alpha) \triangleq \Gamma(\alpha, 0). \quad (3.6)$$

$Q_K(\alpha, \beta)$ is the generalized Marcum-Q function, defined as

$$Q_K(\alpha, \beta) \triangleq \frac{1}{\alpha^{K-1}} \int_{\beta}^{\infty} t^K e^{-\frac{1}{2}(t^2 + \alpha^2)} I_{K-1}(\alpha t) dt, \quad (3.7)$$

with $I_n(x)$ the n -th order modified Bessel function of the first kind. Under more general fading conditions, the SNR itself becomes a stochastic variable, and P_{FA} and P_D have to be weighted over all channel realizations. This results in complicated expressions, see e.g. [78, 79]. In this thesis, it is assumed that $h[k] = 1$.

Although exact expressions are very useful, approximations are often easier to work with. An approximation can be obtained by realizing that, according to the central limit theorem, the probability density function (pdf) of Y converges to a Gaussian distribution for large K . Knowing the first two moments then suffices for determining P_{FA} and P_D , and are found to be (subscript 0 for \mathcal{H}_0 , 1 for \mathcal{H}_1) [80]

$$\mu_0 \approx \sigma_n^2 \quad \sigma_0 \approx \frac{\sigma_n^2}{\sqrt{K}} \quad \mu_1 \approx \sigma_n^2 + \sigma_s^2 \quad \sigma_1 \approx \frac{\sqrt{\sigma_n^4 + 2\sigma_s^2\sigma_n^2}}{\sqrt{K}}. \quad (3.8)$$

If both the noise and the signal are modeled as circularly symmetric zero-mean complex Gaussian noise, which may be more accurate for e.g. Orthogonal Frequency Division Multiplexing (OFDM)-signals, simpler expressions can be obtained. For \mathcal{H}_0 , nothing changes, but for \mathcal{H}_1 , $r[k]$ will be zero-mean Gaussian. Y now has a central χ^2 -distribution with $2K$ degrees of freedom for both hypotheses. Thus,

$$P_{FA} = \frac{1}{\Gamma(K)} \Gamma\left(K, \frac{K\lambda_{ed}}{\sigma_n^2}\right) \quad P_D = \frac{1}{\Gamma(K)} \Gamma\left(K, \frac{K\lambda_{ed}}{\sigma_n^2 + \sigma_s^2}\right). \quad (3.9)$$

For the same P_{FA} , the probability P_D in (3.9) is somewhat lower than P_D in (3.5) (for $P_D > 0.5$), because now the additional variance caused by the randomness of the signal is taken into account. A Gaussian approximation for Y in this model yields again (3.8), but now with $\sigma_1 \approx (\sigma_n^2 + \sigma_s^2)/\sqrt{K}$. Using the Taylor series approximation $\sqrt{1+x} \approx 1 + \frac{1}{2}x$, these two models are seen to be almost equivalent for low SNR. This makes sense, because the signal-induced variance becomes negligible compared to the noise-induced variance.

Since the low-SNR regime is of most interest in this thesis, derivations will be made using (3.8). The resulting expressions for P_{FA} and P_D using the Gaussian approximations are

$$P_{FA} = Q\left(\frac{\lambda_{ed} - \mu_0}{\sigma_0}\right) \quad P_D = Q\left(\frac{\lambda_{ed} - \mu_1}{\sigma_1}\right), \quad (3.10)$$

where

$$Q(x) \triangleq \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-z^2/2} dz. \quad (3.11)$$

For constant P_{FA} , P_D will improve for larger K (and thus a longer measurement time). Increasing $P_{FA,des}$ by lowering the threshold leads to higher P_D , but general consensus is that simultaneously $P_{FA} < 0.5$ and $P_D > 0.5$ are required to speak of reliable spectrum sensing [81].

The threshold λ_{ed} can explicitly be derived from (3.10) as

$$\lambda_{ed} = \mu_0 + \sigma_0 Q_{FA,des}^{-1} \quad (3.12)$$

with $Q_{FA}^{-1} \triangleq Q^{-1}(P_{FA})$, $Q_{FA,des}^{-1} \triangleq Q^{-1}(P_{FA,des})$, $Q_D^{-1} \triangleq Q^{-1}(P_D)$, and $Q_{D,des}^{-1} \triangleq Q^{-1}(P_{D,des})$ to simplify notation. As a result (with $SNR \triangleq \sigma_s^2/\sigma_n^2$)

$$P_D = Q \left(\frac{Q_{FA,des}^{-1} - SNR\sqrt{K}}{\sqrt{1 + 2SNR}} \right). \quad (3.13)$$

3.1.1 Measurement Time

For given SNR and detection performance, the number of samples K can be calculated from (3.13):

$$K \approx \frac{\left(Q_{FA,des}^{-1} - Q_{D,des}^{-1} \right)^2 (1 + 2SNR)}{SNR^2}, \quad (3.14)$$

which scales with $\mathcal{O}(1/SNR^2)$ for small SNR. Assuming white noise, independent samples can be obtained at the Nyquist rate, i.e., $f_s = W$ for quadrature receivers, with W the (RF-)bandwidth of interest. Oversampling results in consecutive samples to be correlated in such a way that the sensing process is not improved [82]. The measurement time T is thus equal to K/W . To comply with IEEE 802.22 regulations ($SNR = -16$ dB for digital TV (DTV) with $W = 6$ MHz, $P_{FA} = 0.1$, $P_D = 0.9$, see section 2.1), $K \approx 1.1 \cdot 10^4$, resulting in $T \approx 1.8$ ms. This measurement time does not include any initial settling time of filters, phase-locked loops (PLLs), etc.

In the analog domain, the measurement time will be approximately the same [77]. However, there are certain advantages to do the processing in the digital domain. It is easier to design filters with desirable characteristics (such as bandwidth, out-of-band suppression, passband gain, and shape factor) in the digital domain, and it is much easier to change these characteristics. Furthermore, the dynamics of these filters are precisely known, and some imperfections (such as passband ripple) can be corrected for.

Handling N channels in parallel speeds up the process by a factor N . As CMOS scales better for digital than for analog implementations, it will become cheaper to handle several channels in parallel in the digital domain rather than in the analog domain, *provided* the analog frontend and the analog-to-digital converter (ADC) have sufficient bandwidth and dynamic range (DR). One must note, however, that the power consumption of digital signal processing (DSP) may not be negligible, and numerical errors may limit performance as well.

3.1.2 Comparing Detectors

The performance of a detector can be quantified by the receiver operating characteristic (ROC)-curve, which plots P_D as a function of P_{FA} for a given




Figure 3.1: ROC-curves for ED: (left) for several SNR with $K = 10^3$ and (right) for several K with $\text{SNR} = -20$ dB.

K and SNR. Algorithms may be compared based on a figure of merit (FoM) defined by the area under the ROC-curve (higher is better) [79]. Fig. 3.1 shows a few curves for ED based on (3.9). Clearly, more samples and higher SNR improve the performance.

The ROC-curves do not take into account the hardware power consumption. For spectrum sensing, it is probably useful to compare *energy consumption*. Given a certain measurement time, (3.13) can be rewritten to find the minimum SNR, SNR_{\min} , for the desired sensing performance:

$$\begin{aligned} \text{SNR}_{\min} &= \frac{Q_{\text{FA,des}}^{-1}}{\sqrt{K}} + \frac{Q_{\text{D,des}}^{-1}}{K} \times \left(Q_{\text{D,des}}^{-1} - \sqrt{\left(Q_{\text{D,des}}^{-1} \right)^2 + K + 2\sqrt{K}Q_{\text{FA,des}}^{-1}} \right) \\ &\approx \frac{Q_{\text{FA,des}}^{-1} - Q_{\text{D,des}}^{-1}}{\sqrt{K}}, \end{aligned} \quad (3.15)$$

where the last step follows for $Q_{\text{FA,des}}^{-1}, Q_{\text{D,des}}^{-1} \ll K$ and with K large. A receiver with 3 dB higher noise figure (NF) may consume half the power, but requires 4 times the number of samples for the same detection performance, draining twice the amount of energy from the battery.

3.1.3 SNR-wall

When K goes to infinity, (3.15) implies that any signal can be detected. The threshold is based on the noise power, which is composed of noise from the physical channel and noise from the receiving device. This noise level can be estimated, but with limited accuracy, e.g. due to the fact that the antenna noise varies as a function of weather, the noise of the receiver may vary over frequency and time (due to temperature and amplifier gain fluctuations), and the noise level estimation itself will always have some error [35, 80, 81]. The *noise uncertainty* in the receiver, i.e. the accuracy to which the noise power can be estimated, is estimated at 1 dB [36, 83].

Apart from this noise uncertainty, the nonidealities of the receiver, such as nonlinearity and harmonic downconversion, also contribute to measured power. These contribution are not known a priori, and may or may not be estimated (e.g. by an initial wide spectrum scan). In any case, they will introduce additional noise uncertainty.

What happens to the detection performance when there is noise uncertainty? Assume that the noise power estimate $\hat{\sigma}_n^2$ lies in an interval around the true noise power σ_n^2 , i.e. $(1 - \epsilon_1)\sigma_n^2 \leq \hat{\sigma}_n^2 \leq (1 + \epsilon_2)\sigma_n^2$, with $0 \leq \epsilon_1 < 1$ and $\epsilon_2 \geq 0$ [80]. As stated in (3.4), the obtained P_{FA} and P_D should not be worse than $P_{FA,des}$ and $P_{D,des}$. The worst-case situation for P_{FA} occurs when the noise level is underestimated, i.e. when $\hat{\sigma}_n^2 = \sigma_n^2(1 - \epsilon_1)$. To guarantee $P_{FA} \leq P_{FA,des}$, the threshold must then be purposely set higher by a factor $1/(1 - \epsilon_1)$ (compare with (3.12)):

$$\lambda_{ed,biasd} = \frac{1}{1 - \epsilon_1} \hat{\sigma}_n^2 + \frac{1}{1 - \epsilon_1} \frac{\hat{\sigma}_n^2}{\sqrt{K}} Q_{FA,des}^{-1}. \quad (3.16)$$

This biased threshold leads to a degraded P_D , but $P_D \geq P_{D,des}$ must also be guaranteed. The worst-case situation for P_D occurs when the noise level is actually overestimated, i.e. when $\hat{\sigma}_n^2 = (1 + \epsilon_2)\sigma_n^2$. So, for a guaranteed $P_{FA,des}$, $P_D \geq P_{D,des}$ can be guaranteed by using as threshold [80]:

$$\lambda_{ed,biasd} = \frac{1 + \epsilon_2}{1 - \epsilon_1} \left(\sigma_n^2 + \frac{\sigma_n^2}{\sqrt{K}} Q_{FA,des}^{-1} \right). \quad (3.17)$$

Overall, the threshold is multiplied by a factor equal to the peak-to-peak uncertainty U [80]

$$U \triangleq \frac{1 + \epsilon_2}{1 - \epsilon_1}. \quad (3.18)$$

Substituting $\lambda_{ed,biasd}$ for λ_{ed} in (3.10), and solving for SNR to find SNR_{min} , results in

$$\begin{aligned} SNR_{min} &= (U - 1) + \frac{U Q_{FA,des}^{-1}}{\sqrt{K}} + \frac{(Q_{D,des}^{-1})^2}{K} \\ &\quad - \frac{Q_{D,des}^{-1}}{K} \sqrt{(Q_{D,des}^{-1})^2 + (2U - 1)K + 2\sqrt{K} Q_{FA,des}^{-1}} \\ &= (U - 1) + \mathcal{O}\left(\frac{1}{\sqrt{K}}\right). \end{aligned} \quad (3.19)$$

So, regardless of $P_{FA,des}$ and $P_{D,des}$, SNR must be at least $(U - 1)$ to guarantee detection, even for an infinite number of samples.¹

¹If the support of $\hat{\sigma}_n^2$ is not bounded, e.g. when there are no bounds on the estimation error, $U \rightarrow \infty$. According to (3.19), this makes detection impossible for *any* SNR [80]. An example where this is the case is when the noise estimation error has a normal distribution. Intuitively, the obtained P_{FA} will not be significantly impacted compared to the bounded peak-to-peak uncertainty in (3.18) if there is negligible probability that the estimation error is very large ($\mathbb{P}(\hat{\sigma}_n^2/\sigma_n^2 \leq 1 - \epsilon_1) \ll P_{FA,des}$). The same holds for the obtained P_D when $\mathbb{P}(\hat{\sigma}_n^2/\sigma_n^2 \geq 1 + \epsilon_2) \ll P_{MD,des}$.

In the context of cognitive radio (CR), this minimum SNR is called the ‘SNR-wall’, a term coined by Tandra & Sahai [84]. The reason for this name is that if one calculates the required number of samples for desired detection performance as a function of SNR, the graph seems to hit a wall. They showed that an SNR-wall exists for cyclostationary sensing as well; in fact there is an SNR-wall for any type of sensing, although the SNR-wall may be caused by something else than noise uncertainty [81].

3.1.4 Other Energy Detection Schemes

Many recent publications have sought to overcome the SNR-wall for ED in one way or another. In [85], oversampling is proposed to distinguish between noise and a signal. The signal components of subsequent samples are highly correlated, while the noise samples (if white) are not. The same assumption is used in [86] to be able to distinguish between the autocorrelation functions of the signal and the noise. However, the noise itself may not be perfectly white, which results in correlation in subsequent samples, or, equivalently, an autocorrelation function that can be non-zero over large time-shifts, and hence obscure the signal to be detected in low SNR conditions. Moreover, both approaches require a more wideband frontend and faster ADCs to accommodate oversampling, resulting in a higher power consumption even when the system is not performing spectrum sensing.

An alternative technique is based on eigenvalues calculated from a covariance matrix [87]. The technique can handle multiple antennas and receivers, but it is always explicitly assumed that the noise samples of the different receivers are independent. However, [88] shows that this is (approximately) true only under certain conditions. Moreover, the technique of [87] is very computationally complex. An extension to the work of [87] is provided in [89], but again independent noise samples are assumed. The requirement on independence of the noise samples in the different receivers is undesirable, as these are difficult to meet in practice (see [88] and appendix B).

Cooperative sensing is often mentioned to improve performance by combining data from several nodes using some algorithm to obtain a decision whether the band is free or not. Fundamentally, the individual detectors limit the performance of cooperative sensing, so improving the detection capabilities of individual receivers is still desired.

3.2 Crosscorrelation

In order to mitigate the noise uncertainty of the receiver, XC spectrum sensing is proposed here. It will be shown that it can mitigate some of the analog impairments that were identified in chapter 2 as well. Where ED, as discussed in section 3.1, uses a single receiver, XC uses two receivers.

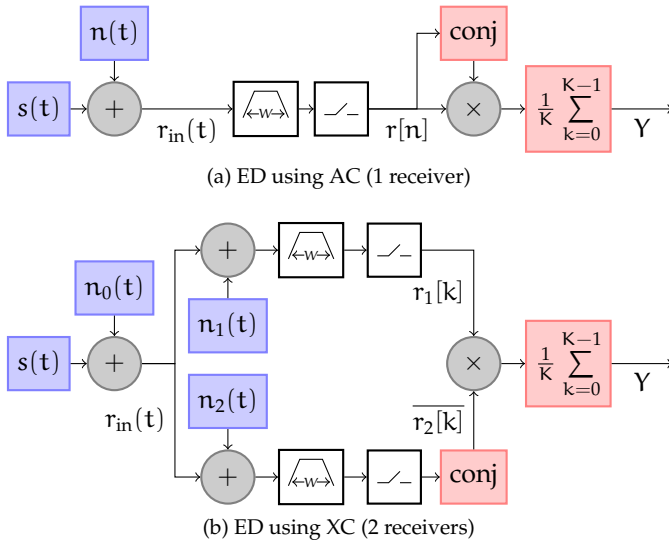


Figure 3.2: ED using XC is a generalization of ED using AC.

Instead of correlating the output of a receiver with itself to arrive at a spectrum estimate (as in (3.1)), the outputs of the two receivers are correlated with each other, which reduces the system noise.

3.2.1 Basic Principle

A block schematic of the XC-principle is shown in fig. 3.2b. The receiver chain is modeled here as a device that only adds some noise.

Each receiver contributes noise (n_1 and n_2 , respectively), with the contribution of each receiver independent of the other receiver (or at least uncorrelated). The expected value of the product of these noise contributions is then 0. The variance of this product can be reduced by repeating the process and averaging the results (similar to the averaging in (3.1)). After sufficient averaging, the role of n_1 and n_2 becomes insignificant in the detection process, even when the threshold is initially biased to account for uncertainty in the noise power.

This only leaves noise contribution n_0 , which represents noise that is correlated in both receivers. This includes the noise received by the antenna, all noise contributed by components before the signal is split, and (if present) part of the noise generated in the splitter. Furthermore, noise generated in one of the receivers may have some transfer unequal to 0 to the output of the other receiver, for example due to crosstalk between the receivers. Both principles introduce noise correlation. In fig. 3.2b this correlated noise is input-referred and simplified to additive noise as n_0 . Although this is a somewhat limited model (the noise correlation may be complex due to frequency-dependent transfers, etc.), it seems to be consistent with many models in literature [90–93], and it allows a direct

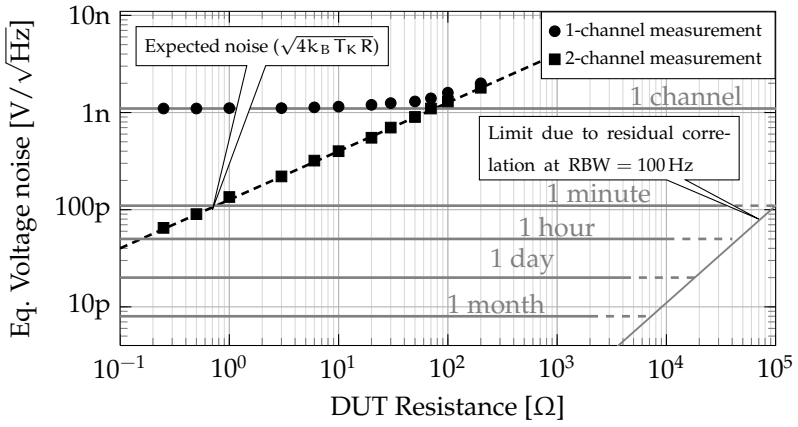


Figure 3.3: Resistor noise measurements. In the 1-channel measurement, the instrument noise obscures the resistor noise. (Adapted from: [94])

comparison with the traditional way of ED, which is shown in fig. 3.2a. If the contributions of the receivers to n_0 can be made zero, the output spectrum will eventually be exactly what arrives at the antenna. Intuitively, the noise uncertainty of the receiver then also disappears.

3.2.2 Related Work

The use of XC to lower system noise is not a new idea. For example, it is in widespread use in radio-astronomy to detect and measure weak sources. Modern phase noise measurement systems exploit XC for greater sensitivity (discussed in section 3.3.2), and rudimentary baseband SAs are proposed in [94]. The thermal noise of resistors can be measured to greater accuracy with XC [91, 92, 94], which can be useful to define the Boltzmann constant k_B [95]. An example measurement of resistor noise using XC is shown in fig. 3.3: the XC-technique allows much higher sensitivity because it removes virtually all of the system noise.

One of the advantages of XC is its higher robustness against gain variations and other analog imperfections compared to AC [96]. Sonnenschein [80] remarks, without any explanation, that XC is more robust than AC in the presence of noise uncertainty. Sousa [97] mathematically tackles the performance of XC under noise uncertainty, where he concludes that XC can indeed outperform AC under noise uncertainty. The results, however, are not easily interpreted quantitatively. The performance of XC under noise uncertainty will be dealt with in chapter 7.

In fact, XC has been proposed before as a possible means for Dynamic Spectrum Access (DSA) or spectrum analysis. Heskamp [93] discusses (without implementation) XC in terms of noise performance, where it is concluded that it “*may provide a viable way to relieve the requirement of the analog part of the spectrum sensing receiver*”. Moseley [98] proposes to

mitigate the effects of harmonic downmixing and finite image rejection (IR) on spectrum sensing by employing XC. This technique will be discussed in more detail in section 3.3.3.

Besides spectral analysis, XC (as a mathematical operation) is used in many digital receivers for synchronization, demodulation and error correction coding, and is an important tool in radar systems. All these different applications led to a lot of research on the statistics of XC-estimators, see e.g. [90, 91, 99–103]. Some will be discussed in section 3.2.4.

3.2.3 Mathematical Background

The crosscorrelation function of two jointly wide-sense stationary stochastic processes A and B is defined as

$$\gamma_{AB}(\tau) \triangleq \mathbb{E} \left[A(t) \overline{B(t+\tau)} \right]. \quad (3.20)$$

The XC theorem states that the cross power spectrum (or simply cross-spectrum) $\Gamma_{AB}(f)$ and the crosscorrelation function $\gamma_{AB}(\tau)$ form a Fourier transform pair [104]:

$$\begin{aligned} \Gamma_{AB}(f) &= \mathcal{F}[\gamma_{AB}(\tau)] = \int_{-\infty}^{\infty} \gamma_{AB}(\tau) e^{-j2\pi f\tau} d\tau \\ \gamma_{AB}(\tau) &= \mathcal{F}^{-1}[\Gamma_{AB}(f)] = \int_{-\infty}^{\infty} \Gamma_{AB}(f) e^{j2\pi f\tau} df. \end{aligned} \quad (3.21)$$

The better-known Wiener-Khinchin theorem, which states that the autocorrelation function and the spectrum of a wide-sense stationary stochastic process form a Fourier transform pair, is a special case of the XC theorem where $A = B$. Unlike the spectrum of a signal, which is by definition real, the cross-spectrum of two stochastic processes may be complex.

Assuming jointly ergodic processes, the time average of one realization equals the ensemble average²

$$\gamma_{AB}(\tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T a(t) \overline{b(t+\tau)} dt, \quad (3.22)$$

where T is the measurement time and a and b are realizations of the processes A and B. This means that an estimate of the crosscorrelation function (and thus the cross-spectrum) can be made by observing one realization for a certain amount of time. Similarly, by observing that the right-hand side of (3.22) resembles a convolution, one can equivalently formulate the cross-spectrum as the product of the Fourier-transforms of a and b :

$$\Gamma_{AB}(f) = \mathcal{F}[a(t)] \overline{\mathcal{F}[b(t)]}. \quad (3.23)$$

²Any individual realization may not converge to the ensemble average. Therefore, strictly mathematically speaking, one should take the expected value on the right-hand side of (3.22) [and (3.23)] [104]. Since this has no practical significance, the expectation is omitted.

Note that even though the Fourier transform of a stochastic process does not exist, the Fourier transform of any realization does exist. In practice the observation time is limited, so the Fourier transforms in both cases can only be approximated.

3.2.4 Statistics of Crosscorrelation

The two mathematically equivalent ways to arrive at the cross-spectrum, (3.21) and (3.23), suggest two different ways to arrive at a cross-spectrum estimate. In accordance with (3.21), the crosscorrelation function is first estimated for a number of discrete lags τ , and then the cross-spectrum is estimated by calculating the discrete Fourier transform (DFT) of the estimated crosscorrelation function. This is called an XF-correlator (XFC), because XC (the X in XF) is performed prior to taking the Fourier transform (the F in XF). Using (3.23), the frequency-domain representations of the signals from the individual receivers are first estimated via a DFT (F), after which the corresponding frequency bins are correlated (X). This is called an FX-correlator (FXC).

Although (3.21) and (3.23) are mathematically equivalent, it turns out that arriving at a cross-spectrum via either of these ways results in (slightly) different results due to the finite observation time in any practical measurement [102]. However, when only the total power in the band is required, using Parseval's theorem, only $\gamma_{AB}(0)$ is required, and the DFT-operator becomes an identity-operator, which makes the two estimation methods equivalent. $\gamma_{AB}(0)$ can be estimated as

$$\hat{\gamma}_{AB}(0) = \int_{(W)} \hat{\Gamma}_{AB}(f) df = \frac{1}{K} \sum_{k=0}^{K-1} r_1[k] \overline{r_2[k]}. \quad (3.24)$$

To a first-order approximation, the statistics to be discussed also apply when spectral estimates are made simultaneously at multiple frequencies within the received band, i.e. when the filtering is performed in the digital domain.³

The model of fig. 3.2b, which depicts the cross-spectrum estimation given in (3.24), is used for further derivations. It is assumed that the noise sources are independent complex zero-mean white Gaussian noise processes within the band of interest. The real and imaginary parts of each process have equal average power and are independent. The same assumptions are used for $s(t)$, as this is a good approximation of the signals transmitted by OFDM systems. OFDM is often proposed as the modulation

³The spectral estimates at different frequencies in the cross-spectrum are actually correlated (e.g. due to leakage in the fast Fourier transform (FFT)) [102]. In other words, estimates at frequencies close to a large step in power spectral density (PSD) may be significantly biased. One can use windowing or polyphase filterbanks to reduce this to negligible levels [105]. In fact, polyphase filterbanks have also been proposed in the context of XC for radio-astronomy [106]. The correlation between spectrum estimates at different frequencies is ignored in this thesis.

technique to use in CR, and is used by many wireless standards, including DVB, LTE, IEEE 802.11 (WiFi), IEEE 802.16 (WiMAX) and IEEE 802.22. For the output of each receiver, so for r_1 and r_2 , one can define $\text{SNR} = \sigma_s^2 / \sigma_n^2$. For simplicity, assume that the PSDs of n_1 and n_2 are equal, which would be the case for two equal receivers: $\sigma_n^2 = \sigma_{n_0}^2 + \sigma_{n_1}^2 = \sigma_{n_0}^2 + \sigma_{n_2}^2$. Define

$$\sigma_{n_0}^2 = \rho \sigma_n^2 \quad \sigma_{n_1}^2 = \sigma_{n_2}^2 = (1 - \rho) \sigma_n^2, \quad (3.25)$$

where ρ denotes the noise correlation between the two receivers.⁴ When $\rho = 1$, the noise is fully correlated, and fig. 3.2b simplifies to fig. 3.2a. When $\rho = 0$, the noise in one receiver is completely uncorrelated with the noise in the other receiver. Without loss of generality, define $\sigma_n^2 = 1$ to simplify notation.

For a total power estimate, the decision metric Y is equal to $\widehat{Y}_{AB}(0)$, see (3.24). Using the notation for n and r_{in} in fig. 3.2b, it follows that (for notational convenience, the subscripts “re” and “im” are used instead of the $\mathcal{R}(\cdot)$ and $\mathcal{J}(\cdot)$ operators, i.e. $Y = Y_{re} + jY_{im}$):

$$Y_{re} = \frac{1}{K} \sum_{k=0}^{K-1} \left(r_{in,re}(n_{1,re} + n_{2,re}) + r_{in,im}(n_{1,im} + n_{2,im}) \right. \\ \left. + n_{1,re}n_{2,re} + n_{1,im}n_{2,im} + \underbrace{r_{in,re}^2 + r_{in,im}^2}_{\text{contains signal power}} \right) \quad (3.26)$$

$$Y_{im} = \frac{1}{K} \sum_{k=0}^{K-1} (r_{in,re}(n_{1,im} - n_{2,im}) + r_{in,im}(n_{2,re} - n_{1,re}) \\ + n_{1,im}n_{2,re} - n_{1,re}n_{2,im}).$$

The distributions of Y_{re} and Y_{im} converge to a Gaussian distribution for a large number of samples, according to the central limit theorem. Assume there is no oversampling, such that the samples are independent identically distributed (the more general case is discussed in section 3.2.5).

Y_{im} can be discarded, as it does not contain any of the signal power that is supposed to be estimated. Thus, $\widehat{P}_{xc} \triangleq \mathcal{R}(Y) = Y_{re}$ can be used as an estimator for the signal power. The first two moments suffice to completely describe a Gaussian distribution. These are derived in [102, p378] based on the assumption that K is large enough. The results translated to the system model of fig. 3.2b are:

$$\mathbb{E}[\widehat{P}_{xc}] \approx \rho + \text{SNR} \quad (3.27)$$

$$\text{var}[\widehat{P}_{xc}] \approx \frac{1}{2K} (2\text{SNR}^2 + (2 + 2\rho)\text{SNR} + 1 + \rho^2).$$

⁴In radio-astronomy literature, ρ is usually defined as the correlation coefficient between r_1 and r_2 , which includes both the signal to be detected and correlated noise. In this work, it is also necessary to distinguish between the signal to be detected and correlated noise, hence a different definition of ρ is used.

Note that due to the variance, realizations of \widehat{P}_{xc} might actually be negative. This is similar to the concerns of [104], where certain FFT-windows may cause PSD estimates using AC to become negative. For screen displays, it may suffice to use $\max(\widehat{P}_{xc}, 0)$ to avoid these clearly erroneous results. This subtlety will not be further discussed here.

There can be a phase offset $\Delta\phi$ between the signal components at the output of the two receivers for several reasons, e.g. due to unequal phase behavior of parasitics or filters, or a timing difference or phase offset between the mixers. This phase offset can also model (as a narrowband approximation) a timing offset between the sampling clocks of the ADCs in the different receivers. When $\Delta\phi \neq 0$, part of the signal power is present in Y_{re} and part in Y_{im} .

When $\Delta\phi$ is negligibly small, or can be estimated and corrected for, (3.27) can be used. For arbitrary $\Delta\phi$, the estimator $\widetilde{P}_{xc} \triangleq |Y|$ can be used (instead of \widehat{P}_{xc}), as it makes the result independent of phase offset [90]. \widetilde{P}_{xc} has the additional practical advantage that the result will always be positive, so it can be displayed on a screen on a logarithmic scale, similar to regular SAs. The exact distribution for \widetilde{P}_{xc} has been calculated in [90], which translated to the model of fig. 3.2b results in:

$$\text{pdf}_{\widetilde{P}_{xc}}(z) = \frac{4K}{\eta\Gamma(K)} \left(\frac{zK}{1 + \text{SNR}} \right)^K \times \mathcal{K}_{K-1} \left(\frac{2zK}{\eta} \right) I_0 \left(\frac{2zK(\text{SNR} + \rho)}{\eta(1 + \text{SNR})} \right), \quad z > 0, \quad (3.28)$$

where $\mathcal{K}_n(x)$ is the n -th order modified Bessel function of the second kind, and η is defined as

$$\eta \triangleq (\text{SNR} + 1) - \frac{(\text{SNR} + \rho)^2}{\text{SNR} + 1}. \quad (3.29)$$

Since this estimator involves both Y_{re} and Y_{im} , it contains more noise than \widehat{P}_{xc} , which has a detrimental effect on the detection performance.

Unfortunately, (3.28) is not very insightful; [91] provides for approximations to the first two raw moments. These moments can be converted to central moments, which for the system model of fig. 3.2b results in:

$$\begin{aligned} \mathbb{E}[\widetilde{P}_{xc}] &\approx \sqrt{(\text{SNR} + \rho)^2 + \frac{\beta_K}{K}((1 - \rho^2) + 2(1 - \rho)\text{SNR})} \\ \text{var}[\widetilde{P}_{xc}] &\approx \frac{1}{K} (\text{SNR}^2 + (2 - 2\beta_K + 2\beta_K\rho)\text{SNR} + 1 - \beta_K + \beta_K\rho^2), \end{aligned} \quad (3.30)$$

where β_K is an interpolation function, defined as

$$\begin{aligned} \beta_K &= \frac{\pi}{4} \frac{1}{K} \left(\frac{\Gamma(K + \frac{1}{2})}{\Gamma(K)} \right)^2 \left(1 - \frac{(\rho + \text{SNR})^2}{\mathbb{E}[\widetilde{P}_{xc}^2]} \right) + \frac{1}{2} \frac{(\rho + \text{SNR})^2}{\mathbb{E}[\widetilde{P}_{xc}^2]} \\ \mathbb{E}[\widetilde{P}_{xc}^2] &= \frac{K+1}{K} (\text{SNR} + \rho)^2 + \frac{1}{K} (2(1 - \rho)(\text{SNR} + \rho) + (1 - \rho)^2). \end{aligned} \quad (3.31)$$




Figure 3.4: Comparison of the pdfs according to (3.28) (dashed gray) and the Gaussian approximation based on (3.30) (solid black) for several noise correlations ρ and K , for $\text{SNR} = -20$ dB.

β_K varies between $\frac{1}{2}$ (when the signal and/or correlated noise dominates) and $\frac{\pi}{4}$ (when the uncorrelated noise dominates) [91].

Fig. 3.4 shows a comparison of the pdfs of the exact solution in (3.28) and the approximation in (3.30) for several situations. Small ρ gives a larger error, as the true distribution does not include negative values; it is closer to a Rayleigh-distribution than a Gaussian distribution. The difference vanishes for larger K , which is required anyway for detecting small signals.

3.2.5 Performance of the Crosscorrelation Detector

Using the approximations of (3.30), the performance of the estimators \hat{P}_{xc} (using the real part) and \tilde{P}_{xc} (using the absolute value) can be compared to each other and to the performance of the standard energy detector \hat{P}_{ac} . This is illustrated in fig. 3.5, where the estimators are compared for several situations using ROC-curves. To make the plots more readable, the left plot compares \hat{P}_{ac} and \hat{P}_{xc} , while the right plot compares \hat{P}_{xc} and \tilde{P}_{xc} .

The left plot shows that \hat{P}_{xc} and \hat{P}_{ac} have equal performance for fully correlated noise ($\rho = 1$), which should be as the second receiver for XC yields exactly the same information as the first, and the estimators become identical. For $\rho = 0$ and $K = 10^3$, the performance of \hat{P}_{xc} is equal to the performance of \hat{P}_{ac} using $K = 2 \cdot 10^3$: with independent noise, the XC system obtains twice the amount of information per unit of time as compared to the AC system. The performance of \hat{P}_{xc} gets worse for higher ρ , although, as can be seen in the zoom-in, the difference between $\rho = 0$ and $\rho = 0.1$ is very small. This is an indication that a low ρ is desirable.

The right plot shows that \tilde{P}_{xc} always performs (slightly) worse than \hat{P}_{xc} , because it adds extra noise to the estimation (except for $\rho = 1$). It can be observed that, at least in the situations depicted, there is almost no difference in the ROC-curves. From numerical evaluation of the ROC-curves, a useful rule-of-thumb can be extracted that for $K\rho^2 > 1$, the performance of \hat{P}_{xc} and \tilde{P}_{xc} is nearly identical. The only situation for which this is not the case




Figure 3.5: ROC-curves for the discussed estimators of XC spectrum sensing for SNR = -15 dB. Left: \hat{P}_{xc} versus \hat{P}_{ac} . Right: \hat{P}_{xc} versus \tilde{P}_{xc} .

in fig. 3.5 is for $K = 3 \cdot 10^3$ and $\rho = 0.01$ ($K\rho^2 = 0.3$), and the performance of \hat{P}_{xc} is indeed visibly better than \tilde{P}_{xc} . To detect a signal at SNR = -15 dB, $K \approx 7 \cdot 10^3$ for $P_{FA} = P_{MD} = 0.1$, which means that \hat{P}_{xc} only outperforms \tilde{P}_{xc} significantly for $\rho < 0.012$.

Assuming the receivers have equal noise performance, the measurement time of a system exploiting XC is 50% to 100% that of a system using AC for the same detection performance. On the other hand, the XCSA will require roughly twice the power consumption compared to the AC-system. Overall, to reach a decision, the energy consumption of an XCSA will be 100% to 200% of a standard SA. Note that these results are based on the assumption of perfect knowledge of the noise level; as will be shown in chapter 7, in the presence of noise uncertainty, XC can outperform AC significantly in both time and energy consumption.

Other Types of Noise

All equations are based on wide-sense stationary white Gaussian noise assumptions. In practice, signals can deviate from this assumption, and it is important to know the performance of the XC-detector in these situations. For example, signals that do not occupy a certain band for 100% of the time, such as used in GSM and WiMAX, are not stationary. The nonstationary behavior changes the statistics, so a new derivation will be required to properly set the threshold for given $P_{FA,des}$ and $P_{D,des}$. Intuitively, as XC is a form of ED, it will detect the average energy of the signal. When the SNR is interpreted as the average SNR over time, simulations indicate that (3.27) and (3.30) give a good approximation for these 'duty-cycled' signals, see the duty-cycled white Gaussian (DCWG) line in fig. 3.6.

The statistics also change for input distributions other than Gaussian. Regardless of the input distributions, the central limit theorem states that




Figure 3.6: Simulated (10^5 independent simulations) mean μ_1 and variance σ_1^2 (signal present) for white Gaussian (WG), white uniform (WU), oversampled Gaussian (OG), 50% duty-cycled white Gaussian (DCWG) sources, and phase-noise-like (PN) sources, with $\rho = 0.01$, $\text{SNR} = -10$ dB, using \hat{P}_{xc} . The WG-line matches (3.27).

the distributions of \hat{P}_{xc} and \tilde{P}_{xc} still converge to a Gaussian distribution, thus the noise reduction is obtained at the same asymptotic rates for any kind of noise source. However, the initial values of the moments and the convergence to the asymptotic rates depend on the type of noise, which will have an impact on the detection process.

A relevant example is ADC quantization noise, which is often modeled as a white noise source with a uniform pdf between $-\text{LSB}/2$ and $\text{LSB}/2$. A simulation has been performed where the signal and all noise sources in fig. 3.2b have been replaced by white uniformly distributed sources, see the WU line in fig. 3.6. For the WU-simulation (using the same power of the noise sources), the simulated μ and σ^2 are very close to those of Gaussian sources, see fig. 3.6. By using the threshold calculated for white Gaussian sources, P_{FA} and P_{MD} for white uniform sources are found to be almost identical to those for white Gaussian sources. This indicates that the results obtained here are also relevant for uniformly distributed noise.

Colored noise, such as phase noise (partly white, partly $1/f^2$ dependency) and flicker noise ($1/f$ dependency), will have a higher σ^2 for the same K , because samples in each receiver are correlated in time, which limits the reduction in variance. The simulated variance for Gaussian noise that is partly white, partly $1/f^2$, is shown in fig. 3.6 as PN. It shows similar μ (as should be the case, because the power of all sources is kept constant in all simulations) but quite different σ as compared to the WG-simulation. One can define an effective sample size K_{eff} as [82]

$$\frac{K_{\text{eff}}}{K} \approx \frac{\left(\int_{\langle W \rangle} S(f) df \right)^2}{\int_{\langle W \rangle} S^2(f) df}, \quad (3.32)$$

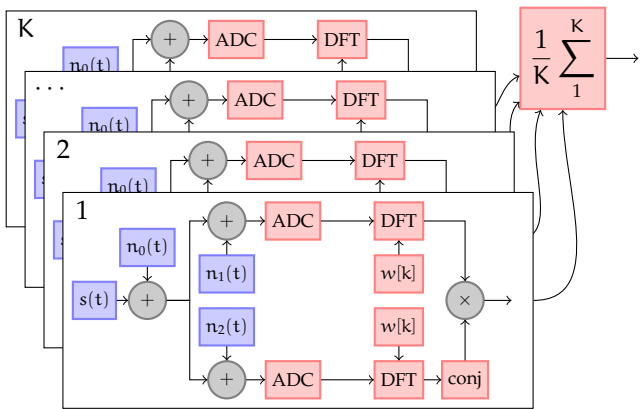


Figure 3.7: The FXC first estimates the (windowed) Fourier transforms of the incoming samples, and then multiplies them to arrive at a cross-spectrum estimate. Several estimates are averaged to lower the variance.

where $S(f)$ is the PSD of the total signal; for a uniform PSD, $K_{\text{eff}} = K$. This K_{eff} should then be used instead of K for setting the threshold and calculating the detection performance. As a verification for this claim, bandlimited Gaussian noise with 10 times oversampling (such that it covers the range from 0 to $f_s/10$) is simulated, see the line marked 'OG' in fig. 3.6. When the threshold is calculated using $K_{\text{eff}} = K/10$, the desired P_{FA} and P_{D} can be found.

3.2.6 Digital Signal Processing

Fig. 3.2b shows one band-pass filter per receiver, after which the outputs are multiplied with each other. As explained in section 3.1.1, it is advantageous for measurement time to process several bands in parallel.

An FFT is an efficient implementation of a DFT, provided samples are taken at a constant sample rate, and one can live with estimation of equally-spaced points on the frequency axis. From a receiver perspective, the FFT can be thought of as an efficient approximation of a bank of zero-IF downconverters, low-pass filters, and decimators. Each output can then be regarded in the same way as the receiver outputs in fig. 3.2b. Thus, for an M -point spectral estimate, M -point (possibly windowed) FFTs are needed. The computational complexity of the FFT is $\mathcal{O}(M \log_2 M)$, while that of the straightforward application of the DFT is $\mathcal{O}(M^2)$. Including averaging to reduce the variance in the measurement, the DSP for an FXC is shown in fig. 3.7. The number of samples K as used in e.g. (3.32) corresponds to the number of samples available *per bin* of the FFT, which is equal to the number of FFTs being averaged (which is indicated as K in fig. 3.7).

Even though the complexity largely favors the FXC over the XFC, there are some differences in performance with respect to word-length and a small loss in SNR, which has made the XFC the most popular in radio-

astronomy. Where many pairs of antennas have to be correlated, the FXC is nowadays favored for reasons of computational complexity [106]. In this thesis, the processing takes place on a general purpose processor, which makes the FXC by far the most attractive option due to its lower computational complexity. Furthermore, OFDM is often mentioned in the CR-context as the modulation to use, due to its flexibility in bandwidth. It extensively uses FFTs for modulation and demodulation, so it is likely that an efficient implementation of FFTs will be available in CRs.

Compared to a single-receiver SA, an FXC requires the same number of multiply-accumulates (MACs), but twice the number of FFTs (and windowing multiplications). Since the FFTs dominate in computational complexity for $M \geq 4$, one can say that the power consumption of the DSP for XC is doubled compared to the DSP of a single-receiver SA. To get an estimate of the power consumption of such DSP, the number of bits used in the calculations is important. A full-scale sine wave input, processed by 16-bit fixed-point 1024-point FFTs, gives a maximum spurious-free dynamic range (SFDR) of 86 dB [107]. In practice, full-scale inputs will rarely occur, so more bits are required to not limit the SFDR. Therefore, 24-bit FFTs are assumed as a worst-case estimate. It is assumed that enough additional bits at the accumulators are available to allow for integration without overflow.

The dynamic power consumption of multipliers scales quadratically with the number of bits and linear with the operating frequency [108]. CMOS-nodes below 100 nm all have very similar supply voltage, while simulations show that the input capacitance of the transistors scales roughly with the node size. Interconnect capacitance barely improves, but the interconnect length also scales roughly with the node size. The power consumption thus scales with CV_{DD}^2 , and therefore approximately linearly with the node size.

The figures of [109] can be used to arrive at an estimate for the DSP power consumption assuming a 65 nm implementation. Scaling from 90 nm to 65 nm CMOS ($\div 1.4$), 8-bit to 24-bit ($\times 9$), from 2.4 GS/s to 20 MS/s ($\div 120$), from 1 channel to 2 channels ($\times 2$), and allowing 10% more for the MACs, the DSP power consumption is estimated at 18 mW when it is integrated on chip. Using the 65 nm results of [110], and scaling from 16-bit to 24-bit ($\times 2\frac{1}{4}$), from 200 MS/s to 20 MS/s ($\div 10$), from 1 channel to 2 channels ($\times 2$), plus 10% for the MACs, only 4 mW is expected. Overall, it can be concluded that an optimized implementation in 65 nm CMOS does not have to consume more than 10 mW for 1024-point FFTs in 20 MHz bandwidth, and is likely to further scale down with newer CMOS processes.

3.3 Mitigating Analog Impairments

Fig. 3.2b is an idealized abstraction of a XCSA, where the only nonideality is the addition of noise. A practical implementation suffers from many

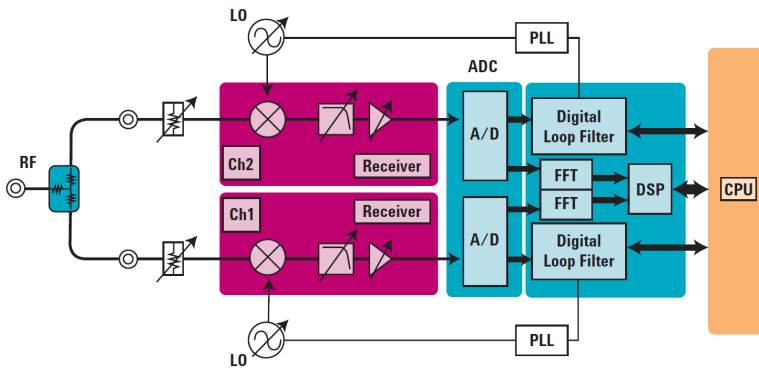


Figure 3.8: Some phase noise analyzers use XC (Source: Agilent).

more nonidealities, such as nonlinearity, limited HR, etc., as discussed in chapter 2. The increased sensitivity of XC [92, 111] is an advantage by itself, but XC can be exploited to mitigate other analog impairments as well.

3.3.1 Improving Linearity

Since the uncorrelated noise can be lowered through XC, the receiver can be designed for high linearity rather than low noise. However, the uncorrelated noise power is only reduced by 1.5 dB per doubling of measurement time. Due to this exponential increase in measurement time, the noise penalty to obtain a high linearity should not be too large.

As shown in table 2.4, if one compares the implementation of [53] to that of [54–56], a big improvement in input-referred third-order intermodulation intercept point (IIP3) can be made at the cost of only a few dB of NF. This indicates that spectrum sensing could significantly benefit from the XC-technique in terms of both linearity and noise performance. Furthermore, as will be shown in chapter 4, an attenuator can be used to improve IIP3 further at the cost of NF, while XC is then used to reduce the noise floor again.

3.3.2 Improving Phase Noise

As is well known from literature [112–115], XC offers a unique opportunity to reduce phase noise. It is actually available in some professional phase noise analyzers, such as the Rohde & Schwarz FSUP and Agilent E5052B. A basic block diagram of a phase noise analyzer based on XC is shown in fig. 3.8; note the similarity with fig. 3.2b.

It is possible to reduce phase noise in a system for spectrum sensing using XC. The local oscillators (LOs) in both receivers need to be frequency-locked, so that their phase noise is uncorrelated. The phase noise of a common external reference, such as a crystal, is usually much lower than internal voltage-controlled oscillators (VCOs). In this way, the

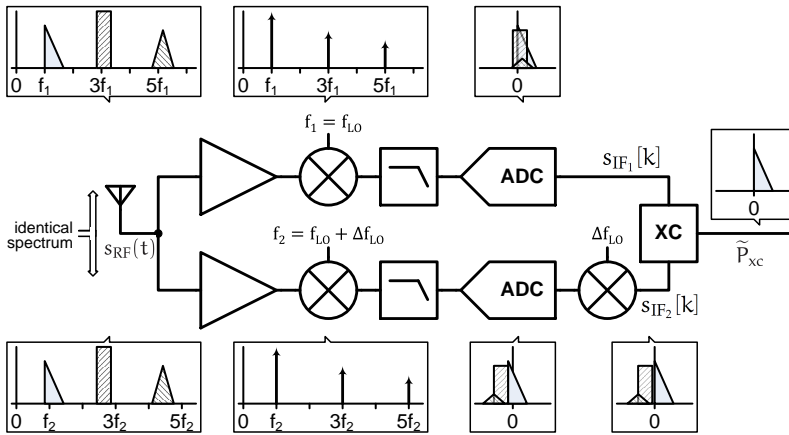


Figure 3.9: System architecture and basic principle to improve HR by XC.

phase noise may be reduced by 1.5 dB per doubling of measurement time. As LO phase noise is directly related to its power consumption, the phase noise reduction realized by XC can reduce LO power consumption. It could even save on total energy consumption by lowering the power (at the cost of an increase in phase noise) of the LOs. Alternatively, it allows for an LO with good phase noise performance to be improved even further.

3.3.3 Improving Harmonic Rejection

The problem with limited HR for spectrum sensing was discussed in section 2.3.2. In [98], a technique is proposed to improve HR using XC, which will be discussed here using fig. 3.9. In the upper receiver, the LO has a frequency f_L , while in the lower receiver, the LO has frequency $f_2 = f_1 + \Delta f_{LO}$. For now, the focus is on the first (upper) receiver.

As discussed in section 2.3.2, passive mixers are often used for their high linearity, but require a square-wave LO-signal $s_{LO_1}(t)$, which does not only contain the fundamental frequency f_1 , but also higher harmonics:

$$s_{LO_1}(t) = \sum_{h=-\infty}^{\infty} c_h e^{-j2\pi h f_1 t}. \quad (3.33)$$

Here c_1 is the desired coefficient, i.e. the fundamental frequency of the LO that is supposed to downconvert the desired signal. Without loss of generality, define $c_1 = 1$. The other c_h are the weight factors for the (undesired) harmonics and lead to harmonic downmixing, as shown at the output of the ADC in fig. 3.9. Particularly, $c_{-1} \neq 0$ leads to finite IR.

The radio frequency (RF) input signal $s_{RF}(t)$ may contain signals over a very wide range of frequencies. Define $s_{RF}(t)$ as a superposition of non-overlapping baseband-equivalent signals $z_h(t)$, centered around the

LO-harmonics (if $s_{\text{RF}}(t)$ is real, $z_{-h}(t) = \overline{z_h(t)}$) [98]:

$$s_{\text{RF}}(t) = \sum_{h=-\infty}^{\infty} z_h(t) e^{j2\pi h f_1 t}, \quad (3.34)$$

with $P_h = \mathbb{E} [z_h(t)^2]$ the power of $z_h(t)$, and \hat{P}_h the *estimated* power of $z_h(t)$. \hat{P}_h and P_h will differ due to the measurement variance, but will converge for infinite measurement time (neglecting noise contributions of the receiver). Note that P_1 , the power of $z_1(t)$, is the desired power to measure. Assuming direct-conversion, the mixer output is

$$s_{\text{IF}_1}(t) = \sum_{h=-\infty}^{\infty} c_h z_h(t) + \text{higher frequency components}. \quad (3.35)$$

The higher frequency components are filtered out by a low-pass filter; what remains is a superposition of the signals that were previously located around the LO-harmonics.

After the ADC, AC would estimate the power as in (3.1):

$$\hat{P}_{\text{ac}} = \frac{1}{K} \sum_{k=0}^{K-1} |s_{\text{IF}_1}[k]|^2 = \frac{1}{K} \sum_{k=0}^{K-1} \sum_{h=-\infty}^{\infty} |c_h z_h[k]|^2 = \sum_{h=-\infty}^{\infty} |c_h|^2 \hat{P}_h. \quad (3.36)$$

Here it is assumed that all z_h are uncorrelated and thus add in power. The estimated power highly depends on signals present at harmonics of the LO. The same result would be obtained by a XCSA if $f_2 = f_1$.

Filtering to reduce z_h before downconversion helps, but is difficult to implement for wideband systems. Therefore, many receivers also reduce c_h by differential implementations ($c_h = 0$ for all even h) and HR-mixers ($c_h = 0$ for some odd h). Due to mismatches, the odd harmonics are typically suppressed by 40 dB ($|c_h/c_1| \approx 0.01$), denoted as $\text{HR}_h = 40$ dB.

A frequency offset between the two receivers in combination with XC can improve the HR. The second mixer output is $\sum_h c_h z_h(t) e^{-j2\pi h \Delta f_{\text{LO}} t}$: the h -th harmonic is shifted by $h \Delta f_{\text{LO}}$ at baseband. To align the desired signals ($h = 1$), the resulting signal can be multiplied by $e^{j2\pi \Delta f_{\text{LO}} t}$. This results in the output of the second receiver as:

$$s_{\text{IF}_2}(t) = \sum_{h=-\infty}^{\infty} c_h z_h(t) e^{-j2\pi(h-1)\Delta f_{\text{LO}} t}. \quad (3.37)$$

The frequency difference between the same signals at the outputs of the two different receivers, that were present around the h -th harmonic of the LO at the antenna, is now $(h-1)\Delta f_{\text{LO}}$. The downconverted signals are no longer aligned with themselves, except for $h = 1$. Under the assumption that Δf_{LO} is larger than the bandwidth of each of these signals, or that there is no correlation between different frequency components of these signals, they are decorrelated. They therefore behave like uncorrelated noise contributions in each receiver, and can be reduced at the cost of measurement

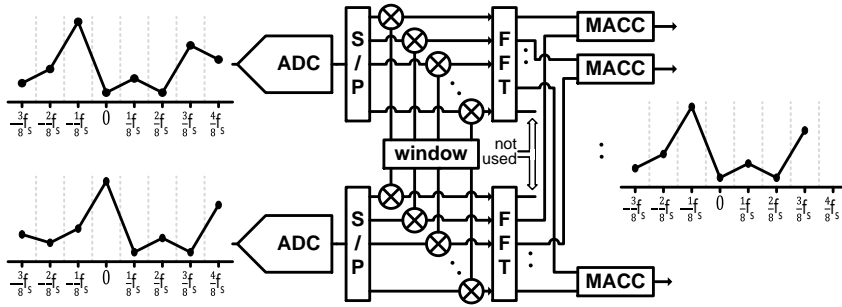


Figure 3.10: Efficient DSP of XC to improve HR using a frequency offset. Here, 8-point FFTs are assumed with $\Delta f_{LO} = W/M$ (1 FFT-bin).

time. The $e^{j2\pi\Delta f_{LO}t}$ -term should be implemented in the digital domain, as it should not suffer from additional harmonic downconversion. The power in the received band is then estimated by XC as (with $k = tf_s$):

$$\begin{aligned} \tilde{P}_{xc} &= \left| \frac{1}{K} \sum_{k=0}^{K-1} s_{IF_1}[k] \overline{s_{IF_2}[k]} \right| = \left| \frac{1}{K} \sum_{k=0}^{K-1} \sum_{h=-\infty}^{\infty} |c_h z_h[k]|^2 e^{j2\pi(h-1)\Delta f_{LO} \frac{k}{f_s}} \right| \\ &= \hat{P}_1 + \frac{1}{K} \sum_{h=-\infty, h \neq 1}^{\infty} |c_h|^2 \hat{P}_h \left| \sum_{k=0}^{K-1} e^{j2\pi(h-1)\Delta f_{LO} \frac{k}{f_s}} \right| \triangleq \hat{P}_1 + \epsilon, \end{aligned} \quad (3.38)$$

where ϵ is a bias (error made) in the estimation, with

$$|\epsilon| = \sum_{h=-\infty, h \neq 1}^{\infty} |c_h|^2 \hat{P}_h \left| \sum_{k=0}^{K-1} \frac{e^{j2\pi(h-1)\Delta f_{LO} \frac{k}{f_s}}}{K} \right|. \quad (3.39)$$

For $\Delta f_{LO}/f_s \notin \mathbb{Z}$ (always true when Δf_{LO} is a fraction of f_s), $\lim_{K \rightarrow \infty} \epsilon = 0$. Thus, for large K , $\tilde{P}_{xc} \rightarrow \hat{P}_1$: $z_1(t)$ can eventually be found without contamination from harmonic downmixing.

The above derivation does not include the variance of \tilde{P}_{xc} , which will depend on the specific spectrum conditions. In any case, a higher HR of the individual receivers will be beneficial for the output variance, so it is best to use the technique of [98] in combination with a HR-receiver.

In the derivation, it is assumed that the bandwidth of $z_1(t)$ is less than $W - |\Delta f_{LO}|$, with W the bandwidth of the filter + ADCs, such that after the frequency shift in the digital domain, z_1 from the two receiver paths can still fully overlap. If Δf_{LO} is chosen as kW/M , with $k \in \mathbb{Z}, |k| < M$, the multiplication with $e^{j2\pi\Delta f_{LO}t}$ can be implemented as a shift in FFT-bins with negligible processing overhead [98], see fig. 3.10.

3.3.4 Reducing ADC Requirements

ADCs perform two operations: sampling and quantization. Like mixing, sampling also suffers from harmonic downconversion. For sampling, the

LO is referred to as the sampling clock, and the harmonic downconversion is usually referred to as ‘aliasing’. An important difference between sampling and mixing is that for an ideal sampler, higher harmonics are just as strong as the fundamental, which is unlike a mixer where higher harmonics reduce in power. To reduce aliasing to satisfactory levels, anti-alias filters are generally employed in combination with an oversampling ratio.

Quantization, however, is an inherently nonlinear process. It limits a signal to a finite set of discrete values, which are usually uniformly spaced at a distance known as the least significant bit (LSB). The quantizer output generally does not exactly represent the input, and the difference is defined as the quantization error. It is usually assumed that the quantization error is uniformly distributed between $-\frac{1}{2}\text{LSB}$ and $\frac{1}{2}\text{LSB}$, with an equal amount of power for each frequency, i.e. it is modeled as white noise and is independent of (and hence uncorrelated with) the input signal. With a full-scale sinusoid as the input signal, it can be shown that the SNR at the output is equal to $\text{SNR} = 6.02b + 1.76 \text{ dB}$, with b the number of bits [116].

For a sine wave input without noise, uniform quantization results in pure harmonic distortion rather than noise. With multiple noiseless sine waves at the input, intermodulation distortion is also generated. Define SFDR_{ADC} as the difference in power between the desired signal and the most powerful distortion component, which is the usual definition of the SFDR of an ADC. The noise power at the output is not taken into account because the noise level depends on the resolution bandwidth (RBW). For spectrum sensing, SFDR_{ADC} should be better than the SFDR of the analog frontend, such that the ADC is not the limiting factor for distortion. Increasing the resolution of the quantizer increases SFDR_{ADC} , but the resolution of ADCs is limited by the required sampling rate and the maximum allowable power consumption [117]. It can be shown (see appendix A) that SFDR_{ADC} for an ideal quantizer with a single noisy sine wave input can be accurately approximated by

$$\text{SFDR}_{\text{ADC}} \approx 8.07b + 3.29 + 171.5\sigma_n^2 \text{ [dB]}, \quad (3.40)$$

where σ_n is the standard deviation of the noise in LSBs at the input of the (otherwise noiseless) quantizer. Dithering, i.e. intentionally adding noise, increases SFDR_{ADC} , but also raises the total noise power.

When the receivers are noiseless, their ADCs see exactly the same input signal, and thus the quantization error in both receivers is fully correlated. As a result, the quantization error is not removed through XC. The power of the correlated components are upperbounded by (3.40), thus with a sufficient amount of uncorrelated noise per receiver, the ADCs in both receivers see a different input signal. This decorrelates the quantization error, which allows XC to remove the quantization error as well.

This can be exploited in a number of ways in a system design. For thermal noise performance of the ADCs, the input range is usually relatively large; the typical LSB-size for a medium-resolution ADC in a modern

CMOS-process is in the order of 1 mV. Using (3.40), the ADC of a XCSA is allowed to be more noisy, and thus more power-efficient, to increase linearity. Secondly, the input range of the ADC (and thus each LSB) can be made smaller: a lower gain typically makes receivers more linear and power-efficient. Thirdly, the smaller swings have much less effect on the gate capacitance of MOS-transistors, which reduces the dynamic nonlinearities of the ADC. Fourthly, the distortion caused by integral nonlinearity of the ADC, which often limits the linearity performance in higher-resolution ADCs, rather than the quantization process, can be reduced as well.

3.4 Conclusions

The standard energy detector (radiometer) squares and sums the input samples. It then compares the result to a threshold, which is based on the receiver noise power, and decides whether the band is occupied or that it is free. A larger sample size K allows the detector to obtain a smaller P_{FA} and P_{MD} for the same SNR. Alternatively, increasing K by a factor 4 allows signals at 3 dB smaller SNR to be detected for the same P_{FA} and P_{MD} . Unfortunately, this simple view breaks down if the noise power is not exactly known. It turns out that with uncertain noise power, there is a certain SNR, the SNR-wall, below which signals cannot be reliably detected, regardless of K and desired P_{FA} and P_{MD} .

ED also suffers from several analog impairments, such as thermal noise, phase noise, nonlinearity, and harmonic downmixing. Linearity can be improved at the cost of noise (by design and/or by employing an attenuator), but that decreases the SNR. To mitigate the above mentioned issues, XC is proposed as a solution. It requires two receivers, and about twice the amount of DSP as compared to the standard energy detector. Analysis of the XC-process shows that, without noise uncertainty and for the same detection performance, it is at least as fast, and up to two times faster than ED, at roughly twice the power consumption. From an energy consumption point of view, XC thus performs somewhat worse.

The benefit of XC lies in the fact that it reduces the noise level, and thereby also the noise uncertainty. As a result, the SNR-wall of XC is lower than that of ED. This also alleviates the tradeoffs between noise performance and other system parameters. Very different tradeoffs for system performance are found for XC compared to traditional designs. For example, a system may be designed for linearity rather than low noise, as XC reduces the uncorrelated noise contributions. One can even think of applying attenuation for higher linearity, and reduce the increased noise by XC. Furthermore, XC offers opportunities to improve phase noise, HR, and ADC-performance. However, these advantages require additional measurement time: a 3 dB higher NF (3 dB lower SNR) requires the measurement time to be increased by a factor 4.

A DISCRETE PROTOTYPE

A prototype is developed to show the principle of crosscorrelation (XC) spectrum sensing in practice. XC should allow a spectrum analyzer (SA) with higher linearity and lower noise floor, and thus increased spurious-free dynamic range (SFDR).

As two practical receivers are generally not perfectly identical, the effect of receiver mismatch is discussed first in section 4.1. The design, expected performance and measurement results of the prototype are discussed in section 4.2, section 4.3 and section 4.4, respectively. Finally, conclusions are drawn in section 4.5.

4.1 Receiver Mismatch

The proposed XC-system requires two receivers. Even when identically designed, the receivers will show at least some difference in actual behavior, e.g. due to mismatch. For the signal component, this can be captured as a difference in transfer function. Denote the transfer functions of the individual receivers by H_1 and H_2 , with $\mathbb{E}[H_1(f)] = \mathbb{E}[H_2(f)] \triangleq H_{\text{rcv}}(f)$, with $H_{\text{rcv}}(f)$ the nominal transfer function. This mismatch in transfer function affects the estimated signal power. When it is assumed that the signal has a power of 1 and a constant power spectral density (PSD) in the bandwidth of interest, the power estimation results in (with f_c the center frequency):

$$\tilde{P}_{\text{xc}} = \frac{1}{W} \left| \int_{f_c - W/2}^{f_c + W/2} \frac{H_1(f)\overline{H_2(f)}}{|H_{\text{rcv}}(f)|^2} df \right|. \quad (4.1)$$

As discussed in section 3.2.4, \tilde{P}_{xc} remains unaffected if there is only a phase offset between the two receivers. To calculate the effect of gain mismatch, assume $H_1(f) = H_{\text{rcv}}(f)$ and $H_2(f) = 10^{\times/20} H_1(f)$ (i.e., receiver

Parts of this chapter have appeared in [MOA:2, 6].




Figure 4.1: Simulation results for the power estimation error with $f_c = f_{3\text{dB}} = 20$ MHz, $\text{RBW} = 1$ MHz, $\sigma_\delta = 5\%$ for an AC and XC system.

2 has x dB more gain than receiver 1). From (4.1), $\mathbb{E}[\tilde{P}_{\text{XC}}] = 10^{x/20}$: a gain error of x dB in one receiver results in $\frac{x}{2}$ dB error in the power estimate.

When $H_{\text{rcv}}(f)$ is a low-pass RC-filter, both the gain and the phase mismatch vary across the bandwidth of interest. Assume that the RC-product has a Gaussian distribution around its mean, i.e.

$$H_1 = \frac{1}{1 + j2\pi f(1 + \delta_1)RC} \quad H_2 = \frac{1}{1 + j2\pi f(1 + \delta_2)RC'} \quad (4.2)$$

with $\mu_{\delta_i} = 0$ and $\sigma_{\delta_i} = 0.05$. One can imagine that the mismatch in the transfer will be worst around the edge of the passband, where gain and phase vary the most. Assume $f_c = f_{3\text{dB}} = 1/2\pi RC$ Hz = 20 MHz and $W = 1$ MHz. Numerically evaluating (4.1) with the transfer of (4.2) for random independent δ_i yields the probability density function (pdf) of the power estimation shown in fig. 4.1. The pdfs for both AC and XC are close to Gaussian.

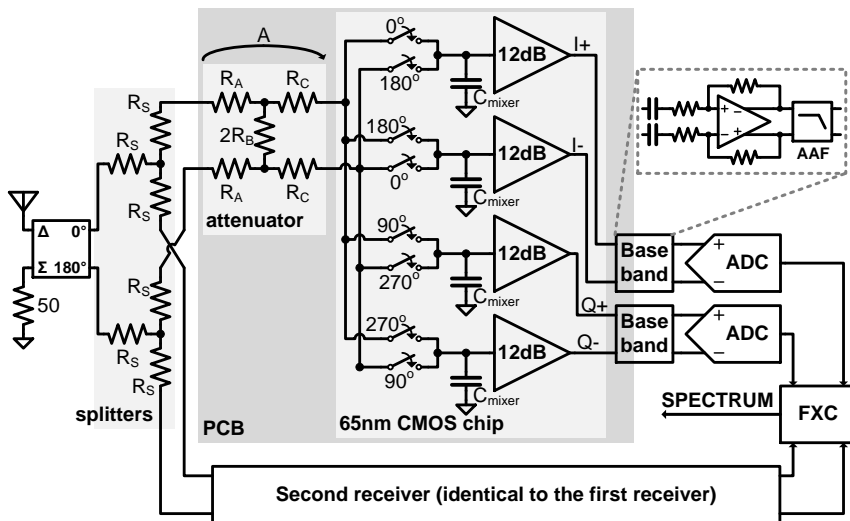


Figure 4.2: Implementation overview of the prototype XC system.




Figure 4.3: Schematic of a passive quadrature mixer [53].

Clearly, the XC-system *outperforms* the AC-system in estimation accuracy. This can be intuitively understood by realizing that for AC (one receiver), $\delta_1 = \delta_2$, while for XC, δ_1 and δ_2 tend to ‘cancel out’. Indeed, simulations indicate that the standard deviation of the power estimation error for AC is about $\sqrt{2}$ larger than for XC.

4.2 Prototype Design

The basic philosophy for the design of the prototype is (1) to have a high linearity by design, (2) to improve linearity by attenuation, (3) to reduce the increased noise by XC, and (4) to make it suitable for integration on one single CMOS-chip. The block diagram of the system is shown in fig. 4.2.

4.2.1 Mixer-first Receiver

The mixer-first architecture of [53] provides a good starting point for high-linearity spectrum analysis, with $\text{IIP3}_{\text{ib}} = +11$ dBm. The 65 nm CMOS chip consists of a mixer and IF-amplifiers. It avoids amplification at RF, where linearization techniques such as feedback are inadequate to achieve very high linearity. Instead, the first stage is a passive mixer with only switches and capacitors, which are both very linear in CMOS.

The passive double-balanced quadrature sampling mixer is shown in its single-balanced form in fig. 4.3. It achieves $\text{IIP3} = +26$ dBm [53, 118]. The switches are controlled by a four-phase non-overlapping square-wave LO with 25% duty cycle. The conversion loss of the mixer is 0.9 dB. Due to noise folding and 1 dB loss due to the switch-resistance, the double sideband noise figure (NF) is 1.9 dB [53, 119, 120].

At IF, the signal is amplified, where a high loop-gain can be achieved to attain good linearity using resistive feedback ($\text{IIP3} = +11$ dBm). The IF-amplifiers provide 13.5 dB voltage gain at a simulated 50 Ω -based NF of 3 dB, which brings the total NF of the chip to a measured 5.5 dB (in [53], the NF fluctuates somewhat). In total, the receiver obtains $\text{SFDR} = 79$ dB in 1 MHz resolution bandwidth (RBW) [53].

4.2.2 Impedance Matching & Attenuation

The input impedance of the passive quadrature mixer in fig. 4.3 at the local oscillator (LO)-frequency for non-overlapping LO-phases has been

analyzed in [119], resulting in the following equation (more elaborate analyses exist, see e.g. [121], but are not required here)

$$Z_{\text{mixer}} \approx (X - 1)R_{\text{left}} + XR_{\text{switch}}, \quad (4.3)$$

where R_{left} is the resistance seen from the mixer looking towards the source, R_{switch} is the on-resistance of the switches in the mixer, and

$$X = \frac{4}{D} \left(1 - \left(\frac{\sin \pi D}{\pi D} \right)^2 \right)^{-1}, \quad (4.4)$$

where D denotes the duty cycle ($0 < D < 1$). For a 25% duty cycle, this simplifies to $X = \pi^2/(\pi^2 - 8) \approx 5.28$. One could state that this mixer is ‘fundamentally’ unmatched; for $R_{\text{left}} = 50 \Omega$ (a typical RF-source) and $R_{\text{switch}} = 5 \Omega$, $Z_{\text{mixer}} \approx 240 \Omega$.

It is desirable to match to 50Ω for at least two reasons. First of all, each receiver will have its own PCB, and requires cables to connect them. These cables, as well as the on-PCB transmission lines, have a characteristic impedance of 50Ω . If the system has a significantly different input impedance, non-negligible reflections will occur, with a resulting input signal that heavily depends on the actual frequency and the physical dimensions of transmission lines and cables. Furthermore, this system will require external RF-filters to mitigate harmonic downmixing, and these filters usually require a 50Ω load to operate properly.

Besides impedance matching, R_{left} and R_{switch} also influence the bandwidth of the mixer at the IF-frequency: together with C_{mixer} it forms a lowpass RC-filter. With a duty cycle of 25%, the resistance seen from the capacitor is effectively quadrupled [53], making the cutoff-frequency

$$f_{3\text{dB,IF}} = \frac{1}{8\pi(R_{\text{left}} + R_{\text{switch}})C_{\text{mixer}}}. \quad (4.5)$$

It is possible to provide matching by using baseband components, which, translated by the mixer switches, form an RF-impedance [118]. However, in the XC-case, impedance matching and attenuation can be easily combined using resistors. It is desirable to keep $f_{3\text{dB,IF}}$ equal at different settings of the attenuator, so that the specific setting does not influence the circuitry behind it. The interface between the attenuator and the mixer does not need to be matched, provided they are located closely enough together.

Overall, there are three requirements for the attenuation/matching network: provide 50Ω -matching, keep $f_{3\text{dB,IF}}$ constant, and provide a certain attenuation A . Since there are three independent requirements, at least three components are needed to fulfill these requirements, which in fig. 4.2 are represented by R_A , R_B , and R_C . Implementing the attenuator with resistors guarantees wideband operation, high linearity and easy CMOS-integration. It will be shown in chapter 5 that such an attenuator can be implemented in CMOS with $\text{IIP3} > +30 \text{ dBm}$.

Table 4.1: Some attenuator component values for $R_{src} = Z_{in} = 50 \Omega$, $f_{3dB,IF} = 18 \text{ MHz}$, $C_{mixer} = 64 \text{ pF}$, and $R_{switch} = 5 \Omega$.

A	A_{dB} [dB]	R_A [Ω]	$2R_B$ [Ω]	R_C [Ω]	used
1.00	0	0	147	0	\checkmark^a
0.89	1	5	126	1	
0.79	2	10	108	2	
0.71	3	14	94	3	
0.63	4	17	82	5	
0.50	6	24	63	8	\checkmark^b
0.35	9	31	44	13	
0.25	12	36	30	17	
0.20	14	39	24	20	

^a Referred to as network 1 or N1 ^b Referred to as network 2 or N2

Due to the required 50Ω impedance per PCB and the low operating frequencies, a resistive power splitter is used to split the source signal and route it to the two receivers. The splitter consists of three resistors with $R_{split} = R_{src}/3 \approx 17 \Omega$, and provides a fixed voltage attenuation of 6 dB from the voltage source to each output of the splitter.

Define the attenuation A *without* this 6 dB (see fig. 4.2), so

$$A = \frac{Z_{mixer} R_B}{R_B (Z_{mixer} + R_C) + R_A (Z_{mixer} + R_B + R_C)}. \quad (4.6)$$

The input impedance of the system is

$$Z_{in} = R_{split} + \frac{1}{2} (R_{split} + R_A + R_B \parallel (R_C + Z_{mixer})), \quad (4.7)$$

where $x \parallel y$ means that x and y are in parallel. R_{left} generally yields a long expression, but when input matching is achieved, it simplifies to

$$R_{left} = R_C + R_B \parallel (R_A + R_{src}). \quad (4.8)$$

Given Z_{in} , $f_{3dB,IF}$, and A , the values of R_A , R_B and R_C can be found.

Table 4.1 shows some component values for $Z_{in} = R_{src} = 50 \Omega$, $f_{3dB,IF} = 18 \text{ MHz}$, $C_{mixer} = 64 \text{ pF}$, and $R_{switch} = 5 \Omega$, which correspond to the prototype design. These values imply $R_{left} = 30 \Omega$. For simple reference, define $A_{dB} \triangleq -20 \log_{10} A$. The final column indicates which attenuators are used in the measurements. The two attenuator networks will be referred to as network 1 (for $A = 1.00$) and network 2 (for $A = 0.50$) in the remainder of this chapter. The choice for these attenuations is based on obtaining a considerable improvement in linearity, while not increasing integration time too much for displayed average noise level (DANL) measurements or practical applicability.

4.2.3 Baseband

A PCI-card with 4 14-bit differential analog-to-digital converters (ADCs) (PMC66-14HSA14) is used to convert the differential IQ-pairs of the receiver




Figure 4.4: Noise floor of, and correlated noise between, the ADC-pairs in the PMC66-14HSA14 (sampling at 10 MS/s).

outputs to the digital domain, see fig. 4.2. The maximum sample rate that can be handled by the operating system is slightly above 11 MS/s per channel, but it is experimentally found that 10 MS/s offers the cleanest spectrum (harmonics of an internal 10 MHz reference clock then all alias to DC). To reduce the noise folding due to sampling, a discrete fifth-order passive Butterworth-filter is implemented that in combination with the ADCs has a cut-off frequency of 12 MHz. This cut-off frequency is chosen to allow this fixed filter to be used at higher sample rates as well, while the additional noise due to noise folding in this prototype is less than 1 dB. The ADCs suffer from a large amount of correlated noise from the many nearby electronics on the board itself and within the PC, as shown in fig. 4.4.

This high correlated noise floor requires a sufficiently large voltage gain to be able to measure the correlated noise from the frontend itself. This voltage gain is provided by a cascade of low-noise TI-THS4130 opamps, in total providing 52 dB gain (11 dB, 21 dB and 20 dB, respectively). The high gain causes the DC-offset of the zero-IF receiver (≈ 7 mV) to exceed the input range of the ADCs. Therefore, the final two stages of the opamp cascade are AC-coupled with a cutoff frequency of 1 kHz.¹ The whole cascade of opamps is used for noise measurements (to overcome the correlated noise of the ADCs), while only the first stage is used for linearity measurements (to prevent the ADCs from saturating). This is also the reason why the SFDR of the frontend cannot be directly shown in a single measurement, and has to be derived from the separate NF and input-referred third-order intermodulation intercept point (IIP3)-measurements.

As discussed in section 3.2.6, an FX-correlator (FXC) is used. To filter out spurious tones present in the lab environment, and to get a good visualization of the spectrum, 1024-point fast Fourier transforms (FFTs) are used. To eliminate finite wordlength effects as much as possible, XC is performed on a regular PC in MATLAB using double floating-point precision. To be more robust for receiver mismatch and differences in LO phase, \tilde{P}_{XC} is used as the spectral estimator.

¹More sophisticated techniques exist to mitigate DC-offset, which may be applied here as well, but they are not important for the experiments to be performed.

4.3 Expected Performance

The available power at the input of a SA is -174 dBm/Hz (room temperature), and the SA adds noise to it. The power displayed on the screen is equal to -174 dBm + NF. Due to the stochastic nature of noise, actual values will fluctuate around these mean values with a certain standard deviation. The DANL is the value that will be displayed *on average*, and therefore the DANL of the SA is equal to -174 dBm + NF.

The true PSD can be more closely approximated by averaging, such that the mean value does not change, but the variance reduces. Many publications (e.g. [92, 94, 111, 122]) claim that a lower noise floor allows higher sensitivity, even when averaging is allowed. When the spectra are displayed on a logarithmic scale, as is generally done on SAs, this makes sense: a signal far below the (smooth) noise floor will not raise the displayed line even a single pixel, so it will not be recognized. When the SA has a much lower NF, and thus a lower DANL, the signal will visibly raise the displayed trace. In fact, after averaging long enough to reduce the variance to negligible levels, there is no difference between the two situations at all when viewed on a linear scale, except for a positive offset in the situation with a higher DANL.

In chapter 7, it will be shown that in the presence of noise uncertainty, a lower noise correlation ρ leads to detection of smaller signals. A lower ρ directly translates to a reduction in DANL for \tilde{P}_{xc} . This may explain the claims of [92, 94, 111, 122] that a lower DANL allows better sensitivity. In any case, the connection between DANL and ρ is the reason why the expected performance of this prototype will be expressed using the DANL in the calculations that follow.

With XC, the DANL depends not only on the NF, but also on the averaging time, which is denoted by DANL_{eff} (DANL_{eff} is always defined in dBm/Hz in this thesis). Equivalently, from DANL_{eff} one can define an effective NF, NF_{eff} , as $\text{NF}_{\text{eff}} \triangleq \text{DANL}_{\text{eff}} + 174$. After infinitely long averaging, the DANL is set by the correlated noise. This DANL is defined as $\text{DANL}_{\text{corr}}$ (always in dBm/Hz), with corresponding $\text{NF}_{\text{corr}} \triangleq \text{DANL}_{\text{corr}} + 174$.

The system performance parameters of interest for the prototype are IIP3 and NF, where the NF (or DANL) is interesting before, during and after the XC process. IIP3 and NF without XC are derived in section 4.3.1. $\text{DANL}_{\text{corr}}$ will be derived in section 4.3.2, and DANL_{eff} as a function of measurement time in section 4.3.3.

4.3.1 Performance without Crosscorrelation

Resistive impedance matching adds about 3 dB NF (and IIP3) compared to the unmatched case (NF = 5.5 dB, IIP3 = +11 dBm) in [53], because in this chapter a resistive network is used to implement the matching. Including the 6 dB loss of the resistive splitter, the NF and IIP3 will be increased by 9 dB using network 1, and 15 dB using network 2 (as defined

in table 4.1). The opamps add 2.1 dB to the NF (calculated using datasheet and feedback resistor values), and the ADCs another 0.7 dB due to noise folding (calculated using opamp datasheets and measured transfer of anti-alias filter and ADCs). Thus, the total for network 1 (network 2) is $NF = 5.5 + 9 + 2.1 + 0.7 = 17.3$ dB (23.3 dB) and $IIP3 = +20$ dBm (+26 dBm).

4.3.2 Noise floor due to correlated noise

In this prototype, the correlated noise power originates from the antenna (or source), from the resistors in the splitter and attenuator, and from noise folding (of the noise from the previous two effects) in the mixer and the ADCs. The noise from the antenna limits $DANL_{\text{corr}}$ to -174 dBm/Hz.

As the splitter is composed of resistors, and there is no buffer between the attenuator networks in both paths (see fig. 4.2), noise generated by these passive components induces a voltage *in both paths*. Since XC multiplies the voltage outputs of the two receivers (with the result in the power domain), this results in *correlated* noise power. The noise correlation due to the splitter and attenuators is analyzed using fig. 4.5 for the single-ended case. The differential implementation with a balun at the input yields identical results [123]. In the analysis, it is assumed that there is no parasitic coupling between the two receivers, and that there is enough gain to make the correlated noise of the ADCs negligible.

The input impedance Z_{mixer} of the mixer is assumed noiseless. The antenna noise is modeled by R_1 , the splitter by R_2 , R_3 and R_7 , and the attenuation networks by the other resistors. The on-resistance of the switches is not explicitly modeled, but the generated noise can be treated in the same way as the noise of R_6 and R_{10} ; it turns out to have a negligible impact on the noise correlation.

The transfer functions α_i to node x and β_i to node y of the noise voltage v_{n_i} of each resistor R_i can be calculated in a straightforward manner, resulting in $v_x = \sum_{i=1}^{10} \alpha_i v_{n_i}$ and $v_y = \sum_{i=1}^{10} \beta_i v_{n_i}$. Note that except for $i \in \{1, 2, 5, 9\}$, α_i and β_i have opposite sign. From symmetry of the circuit it immediately follows for $i \in [3, 6]$ that $\alpha_i = \beta_{i+4}$, $\beta_i = \alpha_{i+4}$ and $R_i = R_{i+4}$.

The crosscorrelation function of noise voltages v_x and v_y is

$$\begin{aligned} \gamma_{XY}(\tau) &= \mathbb{E} \left[v_x(t) \overline{v_y(t + \tau)} \right] = \mathbb{E} \left[\left(\sum_{i=1}^{10} \alpha_i v_{n_i}(t) \right) \left(\sum_{i=1}^{10} \beta_i \overline{v_{n_i}(t + \tau)} \right) \right] \\ &= \mathbb{E} \left[\sum_{i=1}^{10} \alpha_i \beta_i \overline{v_{n_i}(t) v_{n_i}(t + \tau)} \right] = \sum_{i=1}^{10} \alpha_i \beta_i \mathbb{E} \left[\overline{v_{n_i}(t) v_{n_i}(t + \tau)} \right] \\ &= \sum_{i=1}^{10} \alpha_i \beta_i \overline{\gamma_{n_i}(\tau)}. \end{aligned}$$

The second line uses the fact that the noise of a resistor is independent of the noise of the other resistors. Using the linearity of the Fourier transform,




Figure 4.5: Schematic for calculating the effect on NF by the splitter plus attenuation network, before and after correlation.

the cross-spectrum is

$$\Gamma_{XY}(f) = \sum_{i=1}^{10} \alpha_i \bar{\beta}_i \Gamma_{n_i n_i}(f). \quad (4.9)$$

In the prototype, $R_1 = 50 \Omega$ and $R_2 = R_3 = R_7 = 17 \Omega$. Using table 4.1, $R_{\text{left}} = 30 \Omega$, and thus $Z_{\text{mixer}} \approx 150 \Omega$. The correlated noise power from the resistor network, normalized to the power delivered by the source, is equal to F_{corr} , with $\text{NF}_{\text{corr}} \triangleq 10 \log_{10} F_{\text{corr}}$, with

$$F_{\text{corr}} = \frac{\sum_{i=1}^{10} \alpha_i \bar{\beta}_i \langle v_{n,i}^2 \rangle}{|\alpha_1|^2 \langle v_{n,1}^2 \rangle} = 1 + \frac{|\alpha_2|^2 \langle v_{n,2}^2 \rangle + 2 \sum_{i=3}^6 \alpha_i \bar{\alpha}_{i+4} \langle v_{n,i}^2 \rangle}{|\alpha_1|^2 \langle v_{n,1}^2 \rangle}, \quad (4.10)$$

where $\langle v_{n,i}^2 \rangle = 4k_B T_K R_i V^2 / \text{Hz}$ is the one-sided spectral density of the thermal noise of R_i . In total, this yields $\text{DANL}_{\text{corr}} = -174 + \text{NF}_{\text{corr}} = -169.7 \text{ dBm/Hz}$ for all networks listed in table 4.1. Simulations using a combination of transient noise in SPECTRERF and processing in MATLAB verify these results. *For arbitrary attenuation, $\text{DANL}_{\text{corr}}$ is limited by the splitter and attenuator to at least -169.7 dBm/Hz .*

Driven by a square-wave LO, the mixer suffers from harmonic down-conversion. Since the paths to both mixer inputs are equal and, to a first-order approximation, frequency-independent, the correlation between the noise at both mixer inputs is equal for all frequencies. Hence, noise folding in the mixers increases NF and $\text{DANL}_{\text{corr}}$ by the same amount, which is 0.9 dB in the prototype. In total, the derived theoretical minimum $\text{DANL}_{\text{corr}}$ is -168.8 dBm/Hz . Simulations (SPECTRERF + MATLAB) confirm these results within $\pm 0.7 \text{ dB}$; the frequency translation of the mixer also influences its input in a complicated way, which is not taken into account in (4.10).

Noise folding also occurs in the ADCs (which adds another estimated 0.7 dB to the NF), but due to the frequency-selective operation of the mixer and the limited bandwidth of the IF-circuitry, the noise folding of correlated noise is expected to be negligible. In other words, noise folding at the ADCs is expected to contribute only to the NF, but not to NF_{corr} .




Figure 4.6: The measurement setup with the two PCBs identically oriented and about 1 cm apart. To reduce coupling between them, another setup uses 50 cm lateral separation between the PCBs.

4.3.3 Noise Floor and Measurement Time

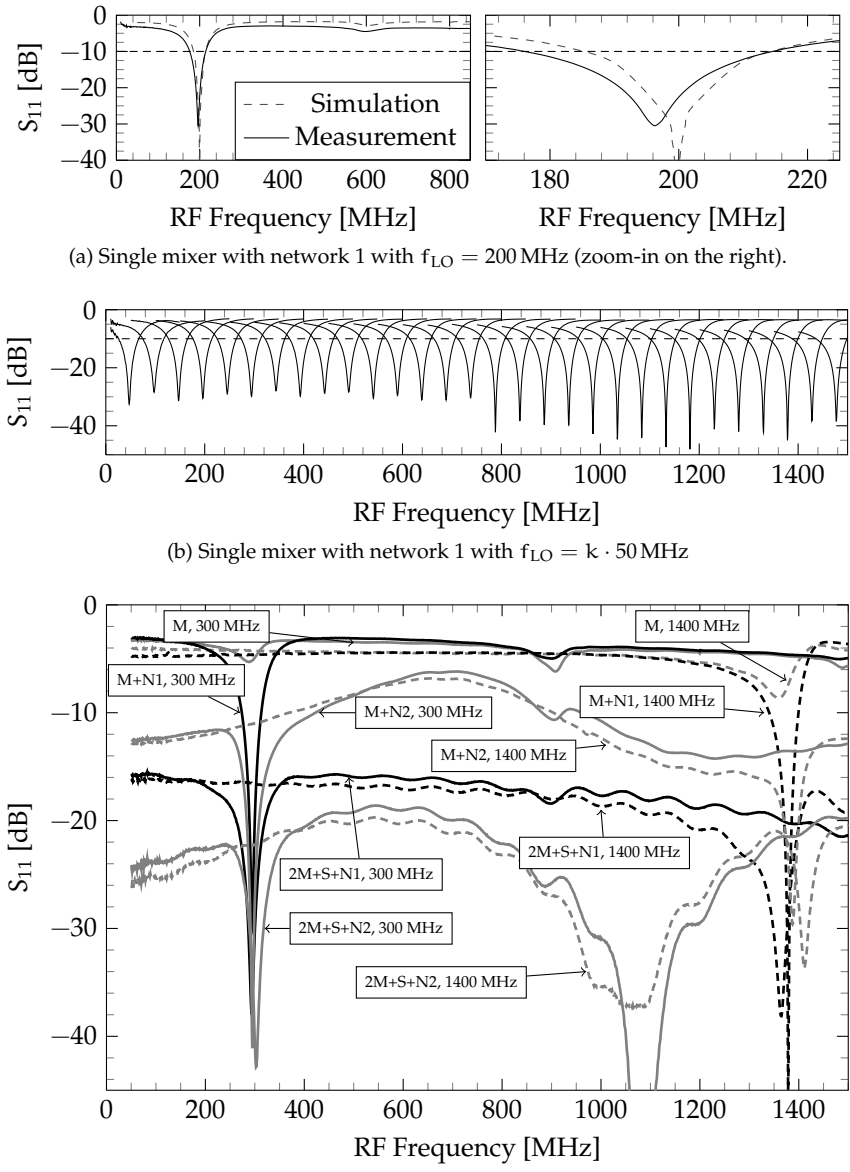
With \tilde{P}_{xc} , the DANL_{eff} is lowered as a function of measurement time. The uncorrelated noise will be reduced with longer measurement time, but the correlated noise will remain, see section 3.2.4.

A higher frequency resolution (higher M) requires a longer measurement time to obtain the same noise floor (in dBm/Hz). To obtain expressions independent of different M and sample rates, the term normalized measurement time (NMT) is introduced. One NMT is the time required to acquire enough samples for one FFT. The number of independent samples in a single bin of the FFT then is equal to NMT.

The NF of each receiver is $\text{NF} = 17.3 + A_{\text{dB}}$. As was calculated in section 4.3.2, $\text{DANL}_{\text{corr}}$ is independent of the attenuation, which implies that ρ is lower for higher attenuation. By noting that ρ denotes the fraction of the noise power that is correlated, it is equal to the ratio of the noise power after performing XC infinitely long to the noise power without XC. In other words, ρ can be rewritten as $\rho = F_{\text{corr}}/F$. $\mathbb{E}[\tilde{P}_{xc}]$ can then be found by substituting $K = \text{NMT}$, $\text{SNR} = 0$, and $\rho = F_{\text{corr}}/F$ in (3.30), which in its turn can be used to calculate DANL_{eff} :

$$\text{DANL}_{\text{eff}} = -174 + \text{NF} + 10 \log_{10} \mathbb{E}[\tilde{P}_{xc}] \quad [\text{dBm/Hz}]. \quad (4.11)$$

As $0 \leq \mathbb{E}[\tilde{P}_{xc}] \leq 1$, and $\mathbb{E}[\tilde{P}_{xc}]$ decreases for larger K , the effective noise floor, DANL_{eff} , decreases with increasing measurement time.



(c) Measurement results for different situations with $f_{LO} = 300$ MHz and $f_{LO} = 1400$ MHz. M denotes mixer, N1 denotes network 1, N2 denotes network 2, and S denotes the splitter.

Figure 4.7: Measurement results for S_{11} .

4.4 Measurements

Several measurements are performed with the measurement setup in fig. 4.6: impedance matching (section 4.4.1), noise reduction (section 4.4.2), and reduction of oscillator phase noise (section 4.4.3). The overall system performance is summarized in section 4.4.4.




Figure 4.8: Measured spectrum as function of correlation time. The test tone of -100 dBm was inserted to verify correct operation. The RBW in this measurement is 10 kHz.

4.4.1 Impedance Matching

To check whether the calculations on resistor values for the unusual input impedance of the mixer (see (4.3)) are correct, S_{11} is measured for several situations. The measured frequency range is from 50 MHz (limited by the mixer) to 1.5 GHz (limited by the measurement setup). Fig. 4.7a shows the measured and simulated S_{11} of a single mixer plus network 1 (N1) for one specific f_{LO} (200 MHz). Fig. 4.7b shows the measured S_{11} of the same configuration for many different f_{LO} , where, for clarity, only the measurement results around f_{LO} are shown. There is clearly good matching around f_{LO} .

Fig. 4.7c shows the measured results for two different f_{LO} , for several different situations (only mixer (M), mixer and attenuation network (M+N1, M+N2), and splitter with two parallel mixers and attenuation networks (2M+S+N1, 2M+S+N2)). The mixer itself is unmatched, but the attenuator networks provide matching. It is interesting to note that the best matching is obtained at somewhat lower frequencies than at the LO-frequency. This is caused by parasitic capacitances in front of the mixer (bondpad, PCB) [121], which also explains why it is more pronounced at higher frequencies.

4.4.2 Noise Figure

An output spectrum for different measurement times is shown in fig. 4.8, where it can be seen that the noise floor is lowered for longer measurement time. Some peaks are visible in the spectrum, which may originate from




Figure 4.9: Measured NF_{eff} (markers) and theoretical fits (black) using (3.30) as a function of measurement time for attenuation network 1. The dashed line shows the expected results using section 4.3.3 with $NF = 17.3$ dB and $NF_{corr} = 5.2$ dB.

PCs, measurement equipment, supplies or other nearby labs. They are discarded in the determination of the DANL. Eventually, the noise floor remains at a certain level, $DANL_{corr}$. The spectra with low NMT look smoother because they have been averaged for better visibility of the trend in noise reduction. The higher noise around DC is DC-offset and flicker noise of the final opamp stage and ADCs (the AC-coupling removes the DC-offset and flicker noise of the mixer-first CMOS chips).

The theoretical $DANL_{eff}$ from (4.11) for network 1 ($NF = 17.3$ dB, $DANL_{corr} = -168.8$ dBm/Hz) is plotted as a function of measurement time in fig. 4.9, together with some representative measurement results. The measurement results are based on the noise floor such as shown in fig. 4.8, with a test tone to determine the total gain of the system. Clearly, the measured $DANL_{corr}$ is sometimes significantly lower and sometimes significantly higher than predicted. For network 2, similar results are obtained (not shown), with roughly the same $DANL_{corr}$ at the same RF-frequency.

The deviation between the predicted and measured $DANL_{corr}$ may be explained by several factors:

- The transfers calculated using fig. 4.5 assume no time delay between the noise from a resistor arriving at v_x and v_y , while in the measurement setup the time difference for some resistors is 0.50 ns due to the PCB transmission lines, which changes the values of α_i in (4.10);
- The mixers perform frequency translation, which affects their inputs, resulting in complicated relations between v_x and v_y due to the different phase-shifts of each frequency component;

- The two receivers are stacked on top of each other, such that the transmission lines on the PCB run in parallel, isolated only by a centimeter of air (see fig. 4.6), which introduces additional frequency-dependent coupling;
- External interference introduces positive or negative correlation when it couples to the two receivers, depending on the phase difference.

The two parameters NF and $DANL_{\text{corr}}$ in (4.11) can be fitted to match the *measured* $DANL_{\text{eff}}$ for different RF-frequencies using network 1, see fig. 4.9. Using only these two parameters, the match with the measurements is excellent for any NMT (for 400 MHz: $NF = 16.8$ dB, $DANL_{\text{corr}} = -172.2$ dB; for 600 MHz: $NF = 17.2$ dB, $DANL_{\text{corr}} = -164.7$ dBm/Hz; for 860 MHz: $NF = 15.3$ dB, $DANL_{\text{corr}} = -168.3$ dB). Similar agreement is obtained with network 2 (not shown). Hence, it is believed that the measurements are performed correctly, and that the aforementioned list explains the differences between calculated and measured $DANL_{\text{corr}}$.

For further verification, the measurements are repeated with the PCBs physically separated. In this situation, $DANL_{\text{corr}}$ is almost flat over the entire RF-bandwidth of the system and averages -170.5 dBm/Hz (see fig. 4.11). This is almost 2 dB lower than expected, but may be explained by the decorrelation of the noise from the resistors due to the long cables.

As this prototype is likely to have different crosstalk mechanisms from an integrated SA, which is the eventual goal, no further effort is put into finding the true cause of the discrepancies. Both measurement results are reported (boards in close proximity and separated); further work on an integrated version has to show what can practically be achieved.

4.4.3 Phase Noise Reduction

In this prototype, the external oscillators have good phase noise performance. Therefore, phase noise is emulated by applying phase modulation. The modulation is done by independent white noise sources, such that the phase noise of both oscillators is independent. The oscillators are frequency-locked using one of the two as reference for the other; each mixer is then driven by one of the oscillators. The phases are synchronized manually before the start of each experiment, which is necessary, because otherwise the impedance seen by one mixer at its input is affected by the other mixer.

The obtained spectra are shown in fig. 4.10. The AC spectra are the spectra as observed at the output of each single receiver. Because the LO-signal generators are of different brands, it is likely that the phase modulation input is handled in a different way, which explains why the spectra of the two receivers differ. This is, however, not a key concern for demonstrating the principle. For XC with $NMT = 1$, the observed spectrum is roughly halfway (in dBs) between the AC spectra. Clearly, a longer measurement reduces the phase noise.




Figure 4.10: Measurement results for phase noise reduction using two phase-modulated frequency-locked oscillators for (left) network 1 and (right) network 2. For network 1, the XC-line has not been averaged.

Some crosstalk can be observed between the two receivers, as the phase noise is only reduced by 11 dB (network 1) and 16 dB (network 2). In another experiment where only one of the two oscillators is modulated, phase noise shows up in the other path as well. This coupling is expected to scale with the attenuation: 1 dB additional attenuation provides an additional 2 dB isolation between the mixers, and thus 1 dB reduction in the correlated phase noise. Hence, 6 dB difference in phase noise reduction between network 1 and network 2 is expected, which is close to the measured 5 dB.

4.4.4 Overall Performance

The NF and IIP3 are measured over the entire bandwidth of operation (with a common low phase-noise oscillator that, in contrast with the previous subsection, is not modulated), of which the results are shown in fig. 4.11. The measured NF corresponds with the predicted NF (17.3 dB and 23.3 dB, see section 4.3.1) ± 1.5 dB, while the measured IIP3 is about 1.5 dB lower than calculated in section 4.3.1 (20 dBm and 26 dBm).

The additional 6 dB of attenuation increases both NF and IIP3 by 6 dB; the additional noise caused by the attenuation circuitry is completely reduced by XC. This agrees with the calculations of section 4.3.2, showing that the lower bound on NF after XC is independent of the attenuation. As a result, using (2.7), the 6 dB of attenuation increases the SFDR by 4 dB. For both attenuation networks, $DANL_{corr}$ is around -170 dBm/Hz, which is somewhat lower (better) than the DANL of the original frontend without attenuator. For network 1, the SFDR is increased by 5 dB as compared to the single frontend, while for network 2 the improvement is 9 dB.




Figure 4.11: NF and IIP3 measurement results. Here, ‘separated’ denotes the situation where the two receivers are physically separated by about 50 cm, while ‘close’ denotes the situation where the two receivers are physically very close, as shown in fig. 4.6.

Each frontend (mixer + IF-amplifier) consumes 61 mW ($f_{LO} = 50$ MHz) to 83 mW ($f_{LO} = 1.5$ GHz), while the digital signal processing (DSP) power consumption is estimated at 10 mW (see section 3.2.6). Table 4.2 compares the results to several other architectures, showing the high linearity and low noise of the prototype. For the SAs, the values given at 0 dB attenuation are used, without any (optional) preamplifiers turned on. For the integrated solutions, the reported NF or sensitivity values are converted to $DANL_{Hz}$. For NF, $DANL_{Hz} = NF - 174$ (dBm/Hz), and for sensitivity, $DANL_{Hz}$ is set equal to the sensitivity in dBm/Hz. Even though the Noise Floor Extension (NFE) in the PXA seems to reduce DANL by 8 dB immediately (noise level calibration is done at start-up), the variance is still the same as at the level of -154 dBm. It takes an estimated 40 averages to get a variance that belongs to a -162 dBm/Hz noise level.

4.5 Conclusions

A prototype is developed to test some of the predictions made in chapter 3. The core of the prototype is a high-linearity mixer-first receiver with moderate noise performance, presented in [53]. To provide both matching and attenuation, a resistor network at the input is used. Attenuation results in more linearity at the cost of extra (uncorrelated) noise.

The prototype confirms that duplication of the frontend and XC of the outputs of the two frontends significantly reduces the noise. It shows

Table 4.2: Comparison with other designs

Architecture	CMOS-node [nm]	RF-range [GHz]	Power [mW]	NF [dB]	DANL _{corr} [dBm/Hz]	Time penalty [NMT] ^f	IIP3 [dBm]	FoM _{SFDR} [dB]
This work (network 1)	65 ^b	0.05–1.5	176 ^c	17	-169	2.7·10 ²	17	11 ^d
This work (network 2)	65 ^b	0.05–1.5	176 ^c	23	-169	4.3·10 ³	24	18 ^d
Soer et al. [53]	65	0.2–2.0	67	5.5	-168		11	5
Ru et al. [54]	65	0.4–0.9	60	4	-170		4	0
Park et al. [67]	180	0.4–0.9	180	50 ^e	-124		-17 ^f	-67
Kitsunezuka et al. [124]	90	0.03–2.4	37	39 ^e	-135		-11	-50
Pollin/Ingels et al. [125, 126]	40	0.1–6.0	100 ^g	3	-171		-12	-15
Murphy et al. [56]	40	0.08–2.7	35–78	2	-172		-22 ^h	-14
Greenberg et al. [55] ⁱ	80	0.04–1.0	440	3	-171		-15	-18
Agilent PXA N9030A-503		0–3.6	450 · 10 ³	18	-166 ^j	40 ^k	22	14
Rohde & Schwarz FSH4		0.01–3.6	12 · 10 ³	29	-145		15	-14
CRFS RFeye Node		0.01–6.0	15 · 10 ³	10	-164		20	10
Signal Hound USB-SA44B		0.00–4.4	2.2 · 10 ³	30	-144		3	-27

^a Calculated NMT to have an effective NF within 1 dB of NF_{corr}

^b The RF-core is 65 nm, but many other components are discrete

^c Only integrated part and estimated DSP power consumption

^d Using the value 174 + DANL_{corr} + 1 as NF

^e Calculated based on sensitivity given in dBm in a certain bandwidth

^f P_{1dB} given in paper as -27 dBm

^g Analog frontend + 10-bit ADCs

^h IIP3_{ib} close to, but outside IF-bandwidth; true IIP3_{ib} expected to be worse

ⁱ Receiver in maximum gain setting

^j Using noise figure extension

^k Calculated based on 8 dB lower expected value without lower variance

that most of the noise introduced by each frontend is uncorrelated. The theoretical expectation of the noise floor reduction (section 3.2.4) fits well to the measured results. Moreover, the theoretical prediction that the final noise floor after XC is independent of the attenuation is experimentally verified. Therefore, passive attenuation can be used to improve linearity, while the loss in noise performance is compensated by XC and a longer measurement time.

Experiments with the two receivers spaced closely together and far apart show that large fluctuations in the measured DANL_{corr}, the DANL after XC, are (at least to a large part) caused by crosstalk between the two receivers, which is most likely the result of coupling between the transmission lines on the PCB. Since this effect will not be present in a single-chip implementation, no further efforts have been put into investigating this effect in more detail.

Overall, a significant reduction in noise level is obtained within acceptable measurement time. The achieved IIP3 of +24 dBm is 13 dB better than what is achieved using a single frontend, while the final noise floor is similar. In total, the SFDR increases from 79 dB to 88 dB in a 1 MHz RBW, which compares favorably to other designs.

Phase noise reduction by XC is also experimentally verified. The obtained reduction is limited due to the lack of isolation (by design) between the two mixers, which also prevents the use of a frequency offset between the two receivers. Therefore, improvement of harmonic rejection (HR) with XC (see section 3.3.3) cannot be tested with this prototype. An isolating component between the two mixers, such as an amplifier, could solve these issues, as will be further discussed in chapter 6.

A HIGHLY LINEAR ATTENUATOR

As discussed in chapter 4, an attenuator can be used to improve the linearity of the receiver, while crosscorrelation (XC) can be used to reduce the increased noise floor. However, the attenuator itself should not be the linearity bottleneck.

In this chapter, some design considerations and topologies are discussed, after which it is concluded that a two-port attenuator is desirable for power and area reasons. Unfortunately, current implementations of two-port attenuators do not simultaneously satisfy the linearity, bandwidth, and insertion loss (IL)-requirements. A technique to improve the input-referred third-order intermodulation intercept point (IIP3) of discrete-step resistive attenuators is introduced, which mainly alleviates the tradeoff between IIP3 and bandwidth. A prototype of this topology is designed in 0.16 μm CMOS and measurement results are discussed.

5.1 Design Considerations of Attenuators

An attenuator provides a scaled-down version of the input signal at its output. The ideal attenuator does this without distortion, at any input power and at any input frequency. Furthermore, the scaling factor is ideally tunable from 0 dB (no attenuation, no IL) to a very high attenuation. It should have no feedthrough (via parasitic paths) so that high attenuation at high frequencies can be obtained. Finally, it should occupy very little chip area and consume no DC power.

Assume that the attenuator and the receiver circuitry following it are in the weakly nonlinear region with only third-order distortion:

$$y_{\text{att}}(x) \approx a_1x + a_3x^3 \text{ [V]} \quad y_{\text{rcv}}(x) \approx b_1x + b_3x^3 \text{ [V]}, \quad (5.1)$$

Parts of this chapter have appeared in [MOA:5, 11].

with y_{att} and y_{rcv} the attenuator and receiver output, and x in V, as MOS-transistors are voltage-controlled devices. $\text{IIP}_{3,\text{att}}$ (in V²) is $\frac{4}{3} |a_1/a_3|$, and $\text{IIP}_{3,\text{rcv}}$ is $\frac{4}{3} |b_1/b_3|$ [127]. Then

$$y_{\text{rcv}}(y_{\text{att}}(x)) \approx a_1 b_1 x + (b_1 a_3 + b_3 a_1^3) x^3. \quad (5.2)$$

For the cascade, $\text{IIP}_{3,\text{casc}} = \frac{4}{3} |a_1 b_1 / (b_1 a_3 + b_3 a_1^3)|$. An ideal attenuator has $a_3 = 0$, such that $\text{IIP}_{3,\text{casc}} = \frac{4}{3} |(b_1/b_3)/a_1^2|$. Thus, $\text{IIP}_{3,\text{casc}} = \text{IIP}_{3,\text{rcv}} + A_{\text{dB}}$ dB, with $A_{\text{dB}} \triangleq -20 \log_{10} A$ the attenuation in dB and $A = a_1$. As nonlinearity contributions add up, $\text{IIP}_{3,\text{casc}}$ can be lower for a non-ideal attenuator. To calculate the allowable a_3 , and thus the required $\text{IIP}_{3,\text{att}}$, a compressive nonlinearity is assumed, where a_3 and b_3 have the opposite sign from a_1 and b_1 , respectively. With α the allowable deterioration of $\text{IIP}_{3,\text{casc}}$ as compared to the cascade with an ideal attenuator, one can find

$$\frac{a_1 b_1}{b_1 a_3 + b_3 a_1^3} > \alpha \frac{b_1}{a_1^2 b_3} \Rightarrow \frac{a_1}{a_3} > \frac{1}{a_1^2} \cdot \frac{b_1}{b_3} \cdot \frac{\alpha}{1 - \alpha}. \quad (5.3)$$

For 1 dB of allowed deterioration ($\alpha \approx 0.8$), $\text{IIP}_{3,\text{att}} > \text{IIP}_{3,\text{rcv}} + A_{\text{dB}} + 6$ [dBm]. Thus, if $\text{IIP}_{3,\text{rcv}} = 15$ dBm and $A_{\text{dB}} = 10$ dB, $\text{IIP}_{3,\text{att}} > 31$ dBm is required.

Integrated resistors can be extremely linear, provide input matching, handle large input powers, occupy a small area, have low feedthrough and low parasitics, and can be cascaded for high attenuation. A straightforward implementation would therefore be a resistor ladder with an amplifier or buffer attached to each node, as in [128], but this requires several amplifiers, which may occupy a large area and complicates the overall design.

The attenuator could also be a two-port device with tunable attenuation, thus requiring only one amplifier at the output. This type of attenuator will be discussed in more detail in section 5.2. One important trade-off is that a low IL at 0 dB attenuation requires large transistors, which increases area, parasitic capacitance, and feedthrough, and lowers achievable bandwidth.

A compromise for the number of amplifiers can be made when two amplifiers are used: one is connected to the input without a switch, so there is no IL, and another is connected via a switchable resistor ladder. An example is shown in fig. 5.1. Exactly one of the two amplifiers is on at any time. Matching can be obtained by feedback [129, 130] or by the ladder / amplifier [131, 132].

As resistors add noise, many implementations use capacitive ladders rather than resistive ones to lower the noise penalty of the attenuator. Integrated capacitors, however, are relatively large in area when they have to AC-couple sub-GHz frequencies, and often have a relatively large parasitic capacitance. Moreover, the improvement in noise figure (NF) over resistive attenuators is not as much as one might expect. Assume a resistor is used for matching, followed by a capacitive divider. Due to the loss of the attenuator, the noise of subsequent circuitry will play a more significant role. Using Friis' equations (assuming matching), it can be derived that for a receiver with $\text{NF} = 4$ dB and an attenuation of 6 dB, there is only 1 dB




Figure 5.1: Example of a structure using an attenuator with two amplifiers in parallel, of which only one is on at any time. (Adapted from: [131])

difference in cascade NF. For 20 dB attenuation, the difference in cascade NF is still only 2 dB.

A two-port attenuator is preferred, as the amplifier may consume a significant amount of power or chip area, and may require external components. Ideally, the attenuator itself is small and low-power, with low IL at minimum attenuation.

5.2 Two-port Attenuators

The Π -attenuator and T-attenuator shown in fig. 5.2 are widely-used two-port topologies [133–135]. For linear-in-dB controllability, the transistors are used as voltage-controlled resistors. By properly changing the gate voltages between V_{SS} and V_{DD} , a continuously-controlled signal attenuation level can be achieved while maintaining input/output matching [133].

For the Π -attenuator, higher attenuation is achieved by increasing the resistance of the series device M_1 . Simultaneously, the control voltage of the shunt devices M_2 and M_3 adjusts their channel resistance for input/output matching. At high attenuation settings, the channel resistance of M_1 is large, so that a large part of the input signal drops across M_1 . As a result, the nonlinear channel resistance of M_1 generates relatively high levels of distortion [135].

Similarly, for the T-attenuator, higher attenuation is achieved mainly by decreasing the resistance of the shunt device M_3 , while series devices M_1 and M_2 together with M_3 provide input/output matching. The channel resistance of M_3 is small to short the signal to ground, resulting in less




Figure 5.2: Basic schematics of continuously-tunable Π - and T-attenuators.




Figure 5.3: A single branch of a discrete-step attenuator is turned on (off) by connecting the gates to V_{DD} (V_{SS}).

distortion by M_3 . Consequently, T-attenuators are claimed to have a higher IIP3 than the Π -attenuators at higher attenuation settings [135].

Alternatively, for better linearity, the transistors can be replaced by passive resistors (which are usually much more linear than transistors) together with transistor switches as shown in fig. 5.3. Several of these branches in parallel can then provide discrete-step attenuation [134].

A lot of effort has been devoted to improving the linearity and power handling capability of continuously-tunable attenuators [133, 136, 137] and discrete-step attenuators [134]. Adaptive bootstrapped body biasing [136] is used in a cascaded Π -attenuator to suppress the body-related parasitic effects and to improve the 1-dB compression point (CP). The stacked-FET technique used in [137] reduces the third-order intermodulation (IM3) distortion by distributing the voltage swing over many FETs in series to reduce the drain-source voltage swing per FET. However, the large parasitic capacitances of the large transistors required by this technique lower the bandwidth and increase the minimum IL at high frequencies. Moreover, the capacitive nonlinearities will limit the highest achievable IIP3. Therefore, this technique is mainly effective in silicon-on-insulator (SOI) CMOS [137].

The continuously-tunable T-attenuator in [133] uses two stages to reduce the swing per stage, and obtains a worst-case IIP3 of +20 dBm. For a discrete-step implementation with low switch-on resistance to minimize distortion, the switches have to be large, resulting in less bandwidth and larger active area. In [134], a Π -attenuator with parallel branches is designed achieving a worst-case IIP3 of +23 dBm in the TV band. Overall, these IIP3 numbers are not good enough according to table 2.3, so a more linear implementation is desirable.

5.3 IM3-reduction in Discrete-step Resistive Attenuators

As a more linear implementation of a two-port attenuator is desirable, the cause of the limited IIP3 in current implementations needs to be identified. Fig. 5.4 shows a signal source, a Π -attenuator, and its load, where the input power source is modeled as a voltage source $v_s = 2v_{in}$ with source




Figure 5.4: Illustration of IM3 cancellation principle in a Π -attenuator.

impedance R_s .¹ Assuming perfect matching, the input voltage for the attenuator is v_{in} and the gain is defined by $A = v_{out}/v_{in}$, with $A_{dB} \triangleq -20 \log_{10} A$. Large resistors in series with gate and bulk of M_1 cause the gate and bulk voltages to follow the average of source and drain voltage (the transistors are in deep triode): with sufficiently large resistors these voltages are purely AC-coupled via the parasitic capacitors of the transistors [134], resulting in $v_{gs,M_1} \approx v_{bs,M_1} \approx v_{ds,M_1}/2$. These resistors extend the bandwidth of the attenuator, and minimize the distortion caused by all nonlinearities related to v_{gs,M_1} and v_{bs,M_1} [134]. For M_2 and M_3 , they are not always necessary as M_2 and M_3 have less impact on linearity and bandwidth than M_1 .

All transistors are assumed to have the minimum length for maximum bandwidth. As the switched-on transistors stay in the deep triode region, the third-order output admittance nonlinearity is dominant. As an approximation, assume that parasitic capacitances do not play a role. Thus, the transistor distortion current (defined from drain to source) is dominated by third-order output conductance nonlinearity. Define²

$$G_{nml} \triangleq \frac{1}{n!m!l!} \left. \frac{\partial^{n+m+l} I_{ds}}{\partial V_{gs}^n \partial V_{ds}^m \partial V_{bs}^l} \right|_{\substack{V_{gs}=V_{GS}, \\ V_{ds}=V_{DS}, \\ V_{bs}=V_{BS}}} [A/V^{n+m+l}]. \quad (5.4)$$

Then $r_{on} \triangleq 1/G_{010}$ is the (linear) small-signal on-resistance of the transistor, and G_{030} is the third-order output conductance nonlinearity. With fixed channel length, G_{010} and G_{030} are, to a good approximation, proportional to transistor width W ; define $r_{on} \triangleq K_r/W$ and $G_{030} \triangleq K_G W$. For the $0.16 \mu\text{m}$ CMOS-process of NXP, as used in this chapter, $K_r \approx 640 \Omega \cdot \mu\text{m}$ and $K_G \approx -0.001 \Omega^{-1} \mu\text{m}^{-1} \text{V}^{-2}$.

Assume the source outputs two tones with amplitude $2v_{in}$ (thus the power delivered to the attenuator is v_{in}^2/R_s). At frequency f_{IM3} , the distortion voltage at the output $v_{out}^{f_{IM3}}$ is the sum of the distortion voltages contributed by the individual transistors, which are created by the distortion currents i_{d,M_k} of each transistor via the transimpedance transfer

¹The T-attenuator is not discussed here, but all results naturally apply to T-attenuators as well. They are explicitly described in [MOA:11].

²The following notation is used: $V_{gs} = V_{GS} + v_{gs}$, with V_{GS} the DC bias voltage, v_{gs} the (small-signal) AC voltage, and V_{gs} the total voltage between gate and source terminals.

function G_{M_k} . In its turn, i_{d,M_k} is induced by the output conductance nonlinearity G_{030,M_k} and the ‘linear’ voltage v_{ds,M_k} across the drain-source terminal via the transfer H_{ds,M_k} from v_{in} to v_{ds,M_k} . Thus

$$v_{out}^{f_{IM3}} \approx \frac{3}{4} \sum_{k=1}^3 G_{M_k} \underbrace{G_{030,M_k} v_{ds,M_k}^3}_{i_{d,M_k}} = \frac{3}{4} K_G v_{in}^3 \sum_{k=1}^3 G_{M_k} W_{M_i} H_{ds,M_k}^3 [V]. \quad (5.5)$$

The origin of the factor $\frac{3}{4}$ is the scaling factor of the component at f_{IM3} of a two-tone input with equal amplitudes [127]. G_{M_k} and H_{ds,M_k} are linear transfer functions, and can thus be easily calculated using network theory (with $R_x \triangleq R_1 + r_{on,M_1}$, $R_y \triangleq R_2 + r_{on,M_2} = R_3 + r_{on,M_3}$, and $R_{load} = R_s$):

$$\begin{aligned} G_{M_1} &= \frac{A r_{on,M_1}}{1+A} [V/A] & H_{ds,M_1} &= \frac{r_{on,M_1}}{R_x + R_y \parallel R_s} [V/V] \\ G_{M_2} &= \frac{(A-1)r_{on,M_2}}{2(1+A)} [V/A] & H_{ds,M_2} &= \frac{r_{on,M_2}}{R_y} \frac{R_y \parallel R_s}{R_x + R_y \parallel R_s} [V/V] \\ G_{M_3} &= \frac{A(A-1)r_{on,M_3}}{2(1+A)} [V/A] & H_{ds,M_3} &= \frac{r_{on,M_3}}{R_y} [V/V]. \end{aligned} \quad (5.6)$$

With input/output matching, $R_x = R_s(1-A)(1+A)/2A$ and $R_y = R_s(1+A)(1-A)$, resulting in (after straightforward, but tedious algebra)

$$v_{out}^{f_{IM3}} \approx \frac{3v_{in}^3 K_r^4 K_G}{8R_s^3(1+A)^4} \left(\frac{16A^4}{W_{M_1}^3} - \frac{A^3(1-A)^4}{W_{M_2}^3} - \frac{A(1-A)^4}{W_{M_3}^3} \right) [V]. \quad (5.7)$$

As expected, (5.7) indicates that IM3-distortion decreases with larger transistors (the voltage swings across the transistors become smaller). Interestingly, (5.7) also shows that low IM3-distortion can be achieved without requiring wide transistors, as long as the right widths are chosen to minimize $|v_{out}^{f_{IM3}}|$. This latter point breaks the tradeoff between linearity and bandwidth. $IIP_{3,att}$ can be found as

$$IIP_{3,att} = \frac{1}{2} \left(20 \log_{10} \frac{v_{in}}{v_{out}^{f_{IM3}}} - A_{dB} \right) + 10 \log_{10} \frac{v_{in}^2}{2R_s} + 30 \text{ [dBm]}. \quad (5.8)$$

Solving (5.7) for W_{M_1} to minimize $|v_{out}^{f_{IM3}}|$, the switch width for maximum IIP3 is independent of the process constants K_r and K_G :

$$W_{M_1,opt} \approx \frac{2^{\frac{4}{3}} A}{1-A} \left(\frac{(1-A)A^2}{W_{M_2}^3} + \frac{1-A}{W_{M_3}^3} \right)^{-\frac{1}{3}} [\text{m}]. \quad (5.9)$$

The optimum combination of W_{M_1} , W_{M_2} , and W_{M_3} depends on A , so different transistors are needed at different A . Hence, exploiting this IM3-cancellation mandates the use of a discrete-step attenuator.




Figure 5.5: Simulated (marker) and calculated (line) IIP3 versus W_{M_1} for two attenuators ($W_{M_2} = 20 \mu\text{m}$ and $W_{M_3} = 40 \mu\text{m}$) at different frequencies.

Fig. 5.5 shows simulation results for two Π -attenuators for signal frequencies up to 50 GHz by sweeping W_{M_1} for fixed $W_{M_2} = 20 \mu\text{m}$ and $W_{M_3} = 40 \mu\text{m}$. W_{M_2} is chosen smaller than W_{M_3} because M_2 contributes less distortion than M_3 (see (5.7)) at the output of the attenuator. To keep 50Ω impedance matching and the desired attenuation, R_1 , R_2 and R_3 are set accordingly (thus R_1 is swept along with W_{M_1}). IIP3 is extrapolated for an input power of -10 dBm with tones at $f_{RF} \pm 1.6$ MHz (thus 3.2 MHz spacing). For center frequencies up to 50 GHz, the simulation results agree very well with the simple model of (5.7), which validates the assumption for deriving (5.7) that the drain-source nonlinearity is dominant. For small W_{M_1} , M_1 is dominant for the IM3 output. As W_{M_1} increases, its distortion decreases and hence IIP3 increases until it is dominated by M_2 and M_3 .³

The optimum width is well predicted, but the calculated IIP3 near the optimum is too optimistic as it neglects other nonlinearities. For $A_{dB} = 6$ dB, the optimum W_{M_1} is 88 μm according to (5.9) while the optimum W_{M_1} according to simulations is between 84 μm and 90 μm , depending on frequency. For $A_{dB} = 18$ dB, the optimum W_{M_1} according to (5.9) is 15 μm while simulations show an optimum W_{M_1} between 14 μm and 15 μm .

Similar analyses may be performed for second-order intermodulation (IM2)-distortion, which are likely to yield similar equations, but with a different optimum W_{M_1} . In this thesis, the focus is on improving IIP3, as input-referred second-order intermodulation intercept point (IIP2) is usually high enough (using differential implementations) to make IIP3 the distortion bottleneck. Therefore, it will not be further discussed here.

³All simulations are performed in SPECTRERF, using the PSP compact model of the 0.16 μm CMOS process of NXP that is used for the implemented prototype. The PSP model correctly fits derivatives up to the third order [138, 139] and satisfies the so-called Gummel symmetry test (see [140, 141]), which is essential for accurate simulation of distortion.




Figure 5.6: The Π -attenuator ($A_{dB} = 12$ dB) used for PVT-simulations.

5.3.1 Limitations in IIP3-improvement

At least three factors limit the improvement that can be obtained using this IM3-cancellation technique. Process/Voltage/Temperature (PVT)-variation is the most obvious one. Using a realistic production variation model, 200 Monte Carlo simulations are performed on the Π -attenuator in fig. 5.6 at several different temperatures and two different V_{DD} to check the effect of PVT-variation. Fig. 5.7 shows the simulated range of IIP3 at 1 GHz as a function of temperature for $V_{DD} = 1.8$ V (nominal supply) and $V_{DD} = 1.5$ V. The IIP3 remains above 30 dBm for a wide temperature range, even when the supply voltage drops to 1.5 V. The latter can be explained by the fact that the switches operate in very deep triode, so the threshold voltage mismatch hardly plays a role. The effect of device mismatch around the IIP3 peaking region can be reduced by increasing the width of all the switches with the same factor, at the cost of reduced bandwidth.

In fig. 5.4 it is assumed that the distortion currents of the transistors have either 0° or 180° phase shift with respect to each other. This is a valid assumption at low frequencies, but at higher frequencies parasitic capacitances introduce a different phase shift for each distortion current, which leads to degraded distortion cancellation.

Finally, each attenuation setting is optimized for IIP3. During operation only one branch is enabled. The nonlinear parasitic capacitances of off-state transistor switches now set an upper bound on the maximum IIP3 that can be achieved by the enabled attenuator branch.




Figure 5.7: Simulated IIP3 range of 200 Monte Carlo simulations for mismatch and process spread at 1 GHz for the circuit in fig. 5.6.

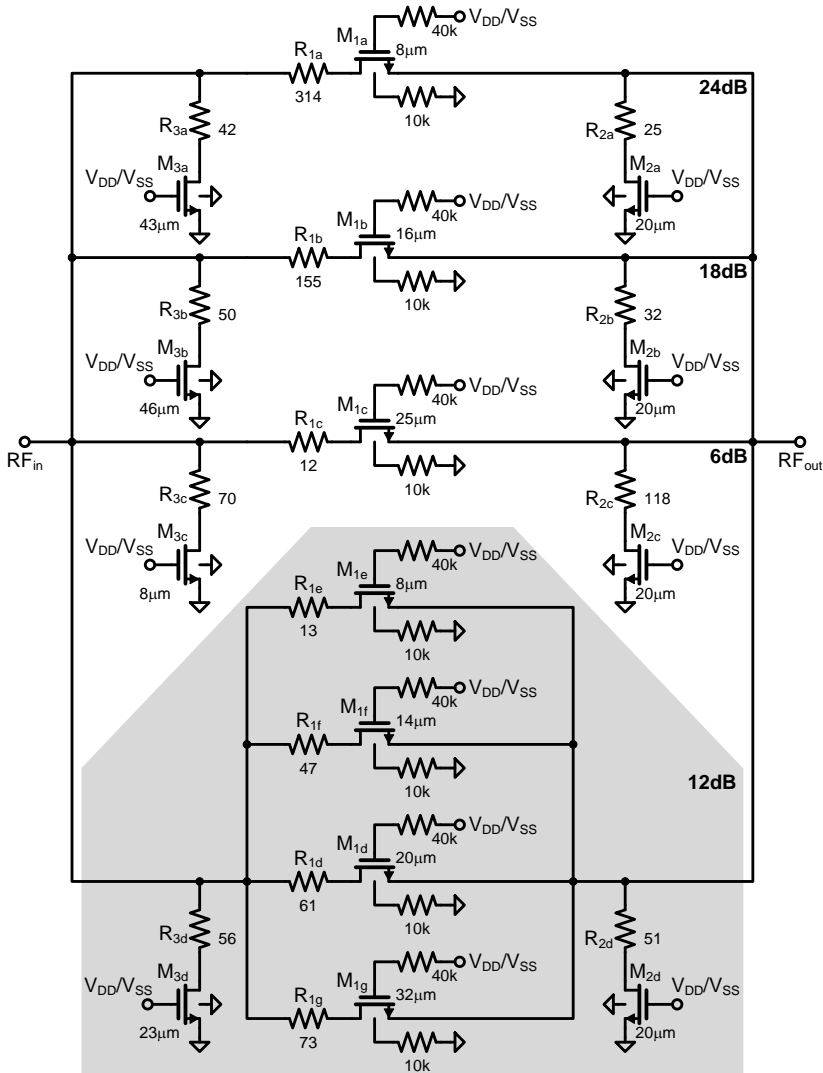


Figure 5.8: Schematic of the designed Π -attenuator network.

Overall, wide transistor switches reduce the sensitivity of IM3 cancellation over PVT-variations, but also limit the bandwidth and introduce two other factors that limit maximum achievable IIP3: capacitance nonlinearity and phase shift due to the parasitic capacitances. As a result, careful optimization is necessary.

5.4 Prototype Implementation & Measurements

To verify the proposed concept, the Π -attenuator of fig. 5.8 is implemented in a $0.16 \mu\text{m}$ bulk CMOS process. It contains two blocks for different measurement purposes: (1) a block ($A_{\text{dB}} = 12 \text{ dB}$, shown inside the gray area in




Figure 5.9: The chip micrograph for both attenuators fabricated in the $0.16\ \mu\text{m}$ bulk CMOS process of NXP. The IC is shared with another design, hence the additional lines and bondpads.

fig. 5.8) for demonstrating the validity of (5.9) and (2) a four-step attenuator with 6 dB, 12 dB, 18 dB and 24 dB attenuation. In the 12 dB attenuation block, each of the four branches is designed for 12 dB attenuation, but has different W_{M_1} to mimic a Π -attenuator with selectable W_{M_1} for fixed W_{M_2} ($20\ \mu\text{m}$) and W_{M_3} ($23\ \mu\text{m}$).

The system contains the upper three attenuator branches in fig. 5.8 (all optimized for IIP3) and the 12 dB setting optimized for IIP3 (M_{1d} , M_{2d} , M_{3d}). During operation, only one branch is enabled. For isolation and bootstrapping purposes, the gate and bulk of M_1 are connected to the controlling voltage via $40\ \text{k}\Omega$ resistors; the gates and bulks of the shunt devices are connected directly to the controlling voltages to save area. For minimum signal attenuation, the transistors M_{1a} , M_{1b} , M_{1c} , M_{1d} are enabled, and the shunt transistors are disabled, yielding an additional 1.8 dB setting that sets the minimum IL of this system. Poly resistors are used for the series and shunt resistance in the attenuator because of their high linearity (IIP3 around +50 dBm according to simulations). A T-attenuator is designed as well, with a minimum attenuation equal to 1.2 dB. A digital decoder provides the controlling voltages ($V_{DD} = 1.8\ \text{V}$ for enabling and $V_{SS} = 0\ \text{V}$ for disabling), and is shared by the attenuators.

The chip micrograph is shown in fig. 5.9. The active areas of the digital decoder (not optimized), the Π -attenuator system and the T-attenuator system are $60 \times 65\ \mu\text{m}^2$, $50 \times 30\ \mu\text{m}^2$ and $54 \times 53\ \mu\text{m}^2$ respectively.




Figure 5.10: Measured and simulated IIP3 for $f_{RF} = 1$ GHz as a function of W_{M_1} . Sim-A refers to simulation with the nonlinear capacitance of the off-state switches taken into account, Sim-B to simulation where it is not taken into account.

The measurements are performed by on-wafer probing. The CP is extrapolated from an input power of -20 dBm, and IIP3 from an input power of -15 dBm with 3.2 MHz two-tone spacing. All simulations include the estimated bondpad capacitances (100 fF) at the input/output of the attenuator system in order to get a proper comparison with measurements. IIP2 is not measured; the simulated nominal IIP2 for the Π -attenuator (T-attenuator) is $+55$ dBm ($+45$ dBm) for all settings at $f_{RF} = 2.5$ GHz with a two-tone spacing up to 1 GHz.

5.4.1 Verification of IM3-Cancellation

The 12 dB branch of the Π -attenuator is used to demonstrate the IM3 cancellation theory. The measured, simulated and calculated (using (5.7)) IIP3 as a function of W_{M_1} at 1 GHz show reasonable agreement, see fig. 5.10. The




Figure 5.11: Measured (markers) and simulated (lines) IIP3 vs f_{RF} for the Π -attenuator for different W_{M_1} . Note the different y-axes for the measurement and simulation results.




Figure 5.12: Measured IIP3 curves vs input power ($f_{\text{RF}} = 1 \text{ GHz}$) for the Π -attenuator for the different W_{M_1} at 12 dB attenuation (gray lines: first/third-order extrapolations).

off-state switches have a significant influence on the linearity performance: While simulations without taking these switches into account (Sim-B) agree reasonably well with the calculations, measurements clearly agree more with the simulations that do take these switches into account (Sim-A). The difference between measured and simulated (Sim-A) IIP3 may be due to limited accuracy of transistor modeling and unaccounted parasitics.

IIP3 as a function of f_{RF} (up to 3 GHz) is shown in fig. 5.11. The optimum W_{M_1} of $20 \mu\text{m}$ provides higher IIP3 than the other W_{M_1} .

Fig. 5.12 shows the IIP3 curves for the different W_{M_1} at 1 GHz. As the attenuation is the same for all W_{M_1} , the power of the fundamental output powers is the same and the markers are thus on top of each other. However, the IM3-products are significantly different: the gray lines show a third-order extrapolation of these IM3-products. For the particular device




Figure 5.13: Measured S_{11} and S_{21} of the four-step Π -attenuator system. For the minimum-attenuation setting, simulation results are shown (markers).




Figure 5.14: Measured IIP3 curves versus input power at 1 GHz for different settings of the Π -attenuator.

sizing in the implementation, it shows that the IM3 improvement becomes less effective for input powers above approximately -8 dBm due to higher-order nonlinearities. As discussed in section 2.3, this is good enough for TV white space (TVWS). Nevertheless, the IM3-products at the optimum width of $20 \mu\text{m}$ remain (very close to) the lowest IM3-products of all four branches, even up to higher input powers. By scaling up the transistors, the voltage swing across them will be less, and thus the IM3-curve will follow the small-signal third-order behavior up to higher input powers.

The optimum sizing has very little effect on CP compared to the oversized switch: the measured CP is 5.6, 12.0, 15.7, and 16.0 dBm for W_{M_1} is 8, 14, 20, and $32 \mu\text{m}$, respectively. Note that IIP3 increases together with CP, except at the optimum size for IIP3.

5.4.2 Discrete-step Attenuator Measurements

For the 6, 12, 18 and 24 dB settings of the Π -attenuator, the measured and simulated S_{11} and S_{21} (50Ω reference) are shown in fig. 5.13. Due to a mistake in the decoder design, the 1.8 dB setting cannot be enabled, and therefore only simulated values are shown. S_{21} for the 24 dB setting increases with frequency due to feedthrough via parasitic capacitances.

The measured IIP3 curves as a function of input power at $f_{\text{RF}} = 1$ GHz are shown in fig. 5.14. The IIP3 for the Π -attenuator are 31 dBm, 33 dBm, 38 dBm, and 36 dBm for 6 dB, 12 dB, 18 dB and 24 dB attenuation. Again, for high input powers (> -8 dBm) higher-order nonlinearities kick in.

The measured IIP3 of ten dies in one wafer for $f_{\text{RF}} = 1$ GHz shows $< \pm 1.5$ dB IIP3 variation, as shown in fig. 5.15. This confirms the robustness of this IM3 cancellation technique.

Fig. 5.16 summarizes the measured IIP3 for various f_{RF} . Due to bandwidth limitations of the measurement setup, the IIP3 cannot be measured below $f_{\text{RF}} = 50$ MHz. The measured IIP3 is above +30 dBm in the TV bands (0.05–1 GHz), above +26 dBm for 0.05–5 GHz and above +24 dBm for 0.05–10 GHz. At higher f_{RF} , extra phase shifts, caused by the parasitic




Figure 5.15: Measured IIP3 for ten samples at 1 GHz.

capacitances, degrade the IM3 cancellation. The measurement results agree rather poorly with the simulation results. One explanation could be that due to the IIP3-cancellation in the design, other effects start to dominate, which may not be modeled well enough as they are significantly smaller than the term that is usually dominant. More research is required to find the cause of the discrepancy.

The CP is above +3 dBm from 0.05–10 GHz, see fig. 5.17. For the 18 dB and 24 dB attenuation settings, M_1 is quite small and hence experiences a relatively large voltage swing, causing $CP < 10$ dBm for $f_{RF} < 1$ GHz. Using wider M_1 in these settings can increase CP. At lower frequencies, the AC-coupling between source/drain and gate becomes less effective. This increases v_{gs} of the series transistors, and thus generates more distortion and decreases CP. Increasing the gate series resistor can alleviate this problem, but making it too large may increase the noise and/or slow down the transition when changing the attenuation setting.




Figure 5.16: Measured and simulated IIP3 vs f_{RF} for different settings of the Π -attenuator.




Figure 5.17: Measured (solid) and simulated (dashed) CP vs f_{RF} for different settings of the Π -attenuator.

5.4.3 Benchmarking

In table 5.1 the Π - and T-attenuator designs are compared with state-of-the-art two-port attenuators. Both designs using the proposed IM3 cancellation technique achieve very high linearity, and simultaneously, high bandwidth for a very low active area in standard bulk CMOS.

5.5 Conclusions

For maximum linearity, the input attenuator, which is followed by a receiver / spectrum analyzer (SA), will ultimately limit the linearity performance. Therefore, a highly linear attenuator implementation is desired. Attenuators based on simple resistive or capacitive ladders may offer the highest linearity, but require an amplifier or buffer attached to each node. A two-port attenuator is desirable as it does not require more amplifiers or buffers, but it suffers from a tradeoff between linearity and bandwidth.

In this chapter, a wideband IM3 cancellation technique has been presented for two-port attenuators that allows individual optimization of branches in a resistive discrete-step attenuator. Simple, yet accurate equa-

Table 5.1: Comparison with state-of-the-art attenuators.

	[136]	[134]	[133]	[137]	[142]	This (Π)	This (T)
CMOS	0.18 μm	65 nm	0.13 μm	SOI	0.18 μm	0.16 μm	0.16 μm
V_{DD} [V]	1.8	1.2	1.2	5	N/A	1.8	1.8
Area [mm^2]	0.28	0.05	0.7	N/A	0.5	0.0054	0.0067
RF-band [GHz]	0.4–3.7	0.4–0.8	0.0–2.5	0.05–4.0	0.0–14	0.05–5.0	0.05–5.6
IIP3 [dBm]	+15	+23	+10	+47	+29	+30 / +27	+30 / +27
(at f_{RF} [GHz])	0.7	0.4–0.8	10	1	10	0.05–1 / 5	0.05–1 / 5.6
CP [dBm]	+6	N/A	+2.5	+30	+15	+3 / +10	+11
(at f_{RF} [GHz])	0.7	N/A	N/A	1	10	0.05–1 / 1–10	0.05–10
Flatness [dB]	2.6	N/A	2.6	3.0	0.7	1.6	1.6
Max. att. [dB]	33	48	42	40	31.5	24	24
Min. att. [dB]	0.96–2.9	5.8	0.9–3.5	2.4–4.0	3.7–10	1.8–2.4 (sim)	1.3–2.2 (sim)
Return loss [dB]	> 9	> 12	> 8.2	> 14	> 9	> 14	> 10
Control mode ^a	L	D	L	L	D	D	D

^a L: linear-in-dB, D: discrete-step

tions for transistor width dimensioning have been introduced. For specific transistor widths, the dominant distortion currents of their nonlinear on-resistance cancel at the load, which results in a high IIP3, even for relatively small transistors. This ensures a high linearity, while the small transistors result in a lower chip area and higher bandwidth. It thus enables highly linear wideband CMOS attenuators with small active area.

A four-step Π -attenuator system designed in 0.16 μm CMOS using this IM3 cancellation technique achieves > 30 dBm IIP3 for the TV bands (0.05–1 GHz), > 26 dBm IIP3 for 0.05–5 GHz and a CP higher than 3 dBm for 0.05–10 GHz, with only 0.0054 mm^2 of active area. Both measurement and simulation results show good robustness of this IM3 cancellation technique.

The cancellation works up to input powers of -8 dBm, where higher-order nonlinearities become dominant. Note that almost all received signals in TVWS have lower power, as discussed in section 2.3. Therefore, this technique is very suitable for an integrated SA.

AN INTEGRATED PROTOTYPE: UTSFINX

In this chapter, an integrated prototype of a crosscorrelation spectrum analyzer (XCSA) is discussed. This prototype is dubbed **UTSFINX**: *University of Twente: Spurious-Free dynamic range Improvement by removing Noise through Xcorrelation*. Section 6.1 discusses the design of **UTSFINX**, while section 6.2 presents all the measurement results. The linearity and noise performance are compared to other implementation in section 6.3, and conclusions are drawn in section 6.4.

6.1 Design

UTSFINX is designed with several considerations in mind. First and foremost, it should be designed for high linearity, as crosscorrelation (XC) reduces the noise floor. Secondly, it should address the problems that were identified in chapter 4 with the discrete prototype: (1) limited phase noise reduction and the inability to use two different local oscillator (LO)-frequencies in the receivers due to crosstalk between the mixers, (2) crosstalk between the two PCBs, (3) a high noise figure (NF) due to the resistive splitter at the input, and (4) a discrete implementation of the attenuator. With an integrated prototype, there will be only one PCB, which solves problem (2).

The passive mixer-first architecture used in chapter 4 lacks isolation between the frontends, which limits the obtainable phase noise reduction. It also prohibits the use of a frequency offset between the two receivers, which is useful for the harmonic rejection (HR)-techniques employing XC of section 3.3.3, and when two bands should be received simultaneously.

Parts of this chapter have appeared in [MOA:7, 9, 12].




Figure 6.1: System diagram of the proposed system, with the dashed box indicating the parts integrated on-chip.

A linear receiver with reverse isolation between mixer and receiver input is used in [54]. High linearity is obtained by keeping voltage swings low at radio frequency (RF) and using feedback techniques at baseband. By using two identical receivers, each with $100\ \Omega$ input impedance, input matching can be obtained without using a resistive splitter. Fig. 6.1 shows the system diagram employing this topology in a XCSA, with the parts that are implemented on-chip in the dashed box. In this implementation, the attenuator is followed by a low-noise transconductance amplifier (LNTA) for input power to current conversion. The output current of the LNTA is then mixed down by a passive mixer employing a HR-architecture. Following circuitry is left off-chip to allow for more experimental freedom and to more easily measure the achieved RF linearity.

A transimpedance amplifier (TIA) converts the current to a voltage at baseband, while RC-feedback provides a first-order low-pass filter to limit the intermediate frequency (IF)-bandwidth. This low-pass filtering will attenuate blockers before they generate a large swing, thus improving overall linearity. The inputs of the TIA (implemented with a TI-THS4130 opamp) act as a virtual ground, reducing the swing at the output of the LNTA, improving its linearity. The TIAs are followed by more amplification to properly interface with the analog-to-digital converters (ADCs), which are preceded by anti-alias filters.

When not used for XC spectrum sensing, both receivers may be operating stand-alone. An (on-chip) switch is then required to (dis)connect them. This gives the configurations in fig. 6.2a and 6.2b. To provide input matching in both configurations, the receivers need to change their input impedance, the implementation of which will be discussed later. Furthermore, a mode is implemented where both receivers have a high-ohmic input impedance, referred to as highZ, as shown in fig. 6.2c. The reason for this mode will be explained in more detail in section 6.1.2.




Figure 6.2: The chip can be reconfigured to several modes.

6.1.1 Attenuator

The attenuator is used to improve the linearity of the receiver, while the additional noise is reduced through XC. Thus, the attenuator should not limit linearity itself. Therefore, a discrete-step attenuator employing the technique of chapter 5 is used. Several settings are implemented (single-ended) with 2 dB, 6 dB, and 10 dB attenuation. An additional large bypass switch is added as a 0 dB setting; it gives only 0.1 dB insertion loss (IL).

The input impedance needs to be 50 Ω when a single receiver is used, and 100 Ω when two receivers are put in parallel for XC spectrum sensing. Therefore, each setting is designed for 100 Ω input and output matching, and two of them are put in parallel to enable 50 Ω input matching.

6.1.2 Low-noise Transconductance Amplifier

The LNTA is based on [54], and consists of a differential input and seven identical differential slices, see fig. 6.3. The seven slices of the LNTA are combined in a 2:3:2 ratio to approximate the ideal 1:√2:1 ratio of a HR-mixer to suppress the 3rd and 5th harmonic of the square-wave LOs.

A CG-stage (M_1 and M_2) provides matching: M_1 and M_2 each have $g_m \approx 1.4$ mS, which amounts to $R_{in} \approx 1/(7 \cdot 2 \cdot 1.4 \cdot 10^{-3}) \approx 50 \Omega$ for seven




Figure 6.3: Circuit-level implementation of one slice of the LNTA.




Figure 6.4: General small-signal equivalent circuit for LNTA noise analysis in all operating modes.

slices in parallel. By selectively activating zero, one or two transistors per slice, a highZ, 100 Ω or 50 Ω input impedance can be obtained, respectively. The common mode feedback (CMFB)-circuit makes sure that the output remains biased at half the supply voltage to enable maximum output swing. A CS-stage [M_3 ($g_m \approx 6.8$ mS) and M_4 ($g_m \approx 4.8$ mS)] provides additional gain. M_3 provides the bias current for M_1 and M_2 , which is shunted by the external inductor of 100 nH that is shared by all slices.

The CG-CS-configuration has the ability to cancel the noise of the CG-stage if their gain has the same magnitude but opposite phase [143]. In this implementation $g_{m,CS} \gg g_{m,CG}$, which increases the overall transconductance, but makes the cancellation only partial. For the noise analysis, consider the transistors as ideal voltage-controlled current sources with noise power spectral density (PSD) equal to $4k_B T_K \gamma_{NEF} g_m A^2/\text{Hz}$, with γ_{NEF} the noise excess factor of the transistor (usually somewhere between 2/3 and 3/2). A small-signal equivalent circuit is shown in fig. 6.4.

Define $g_{m,CG} \triangleq 7(g_{m,M_1} + g_{m,M_2})$, $g_{m,CS} \triangleq 7(g_{m,M_3} + g_{m,M_4}) \approx 80$ mS, $g_{m,tot} \triangleq g_{m,CG} + g_{m,CS}$. In the XC-mode, two LNTAs will be in parallel: define $\alpha = 0$ when the two receivers are disconnected, and $\alpha = 1$ when they are in parallel. The output current of the first LNTA $i_{out}^{(1)}$ is equal to

$$\begin{aligned} i_{out}^{(1)} &= i_{out,+}^{(1)} - i_{out,-}^{(1)} \\ &= H_{n,src}^{(1)} i_{n,src} + H_{n,CG,1}^{(1)} i_{n,CG,1} + H_{n,CG,2}^{(1)} i_{n,CG,2} \\ &\quad + H_{n,CS,1}^{(1)} i_{n,CS,1} + H_{n,CS,2}^{(1)} i_{n,CS,2}, \end{aligned} \quad (6.1)$$

with $H_{xx}^{(1)}$ the transfer from noise current i_{xx} to $i_{out}^{(1)}$, which can be derived from the small-signal equivalent circuits. Note that with these approximations, $H_{n,CS,2} = 0$, and when $\alpha = 0$, $H_{n,CG,2} i_{n,CG,2} = 0$. The following expression for the noise factor can then be obtained:

$$\begin{aligned} F &= 1 + \gamma_{NEF} g_{m,CG} \frac{(1 + (\alpha g_{m,CG} - g_{m,CS}) R_s)^2}{g_{m,tot}^2 R_s} \\ &\quad + \gamma_{NEF} g_{m,CS} \frac{(1 + (1 + \alpha) g_{m,CG} R_s)^2}{g_{m,tot}^2 R_s} + \alpha \gamma_{NEF} g_{m,CG} R_s. \end{aligned} \quad (6.2)$$

For $\gamma_{\text{NEF}} \approx 1$, the NF in 50Ω mode ($g_{\text{m,CG}} \approx 20 \text{ mS}$) is 3.0 dB. In 100Ω mode ($g_{\text{m,CG}} \approx 10 \text{ mS}$), when the other receiver is connected and in 100Ω mode as well, the NF is 3.8 dB. In the highZ mode ($g_{\text{m,CG}} = 0$), it results in a NF of only 1.0 dB. This is all within 0.1 dB from simulation results.

Part of the LNTA and the following stages contribute independent noise in the individual receivers, so at this point the displayed average noise level (DANL) expected after XC can be calculated. The product of the receiver outputs is of interest here, as this is the result used for power estimation. By noting that cross-terms involving noise components from different devices have an expected value of 0, one can find

$$\mathbb{E} \left[\overline{i_{\text{out}}^{(1)} i_{\text{out}}^{(2)}} \right] = |H_{\text{n,src}}|^2 \mathbb{E} [i_{\text{n,src}}^2] + H_{\text{n,CG,1}}^{(1)} \overline{H_{\text{n,CG,1}}^{(2)}} \mathbb{E} [i_{\text{n,CG,1}}^2] + H_{\text{n,CG,2}}^{(1)} \overline{H_{\text{n,CG,2}}^{(2)}} \mathbb{E} [i_{\text{n,CG,2}}^2]. \quad (6.3)$$

Due to symmetry, $H_{\text{n,CG,1}}^{(1)} = H_{\text{n,CG,2}}^{(2)}$, $H_{\text{n,CG,2}}^{(1)} = H_{\text{n,CG,1}}^{(2)}$, etc. Using the fact that the transfers are in the idealized situation not frequency-dependent, and $\mathbb{E} [i_{\text{n,src}}^2] = 4k_{\text{B}} T_{\text{K}} / R_{\text{s}}$, $\mathbb{E} [i_{\text{n,CG,1}}^2] = \mathbb{E} [i_{\text{n,CG,2}}^2] = 4k_{\text{B}} T_{\text{K}} \gamma_{\text{NEF}} g_{\text{m,CG}}$, it is possible to find a ‘correlated noise factor’ similar to (6.2):

$$F_{\text{corr}} = 1 + 2 \frac{H_{\text{n,CG,1}}^{(1)} H_{\text{n,CG,1}}^{(2)}}{|H_{\text{n,src}}|^2} \gamma_{\text{NEF}} g_{\text{m,CG}} R_{\text{s}}. \quad (6.4)$$

The result is that for $\gamma_{\text{NEF}} \approx 1$, $\text{NF}_{\text{corr}} = 2.0 \text{ dB}$ in the 100Ω mode. In other words, after crosscorrelating many samples, the DANL_{Hz} should converge to -172 dBm/Hz . The residual noise correlation is caused by the CG-noise currents flowing between the two receivers, as is illustrated in fig. 6.2b.

From (6.4) it is clear that if $g_{\text{m,CG}} = 0$ (as is the case for the highZ-mode), $F_{\text{corr}} = 1$, and thus $\text{NF}_{\text{corr}} = 0$. Fig. 6.2c illustrates this effect by showing zero current flow between the two receivers, compared to non-zero current flow in the 100Ω -mode. XC removes *all* receiver noise, and only the noise from the source remains: DANL_{Hz} will be equal to -174 dBm/Hz .

6.1.3 Mixer & LO

The HR-mixer steers the LNTA output currents to a differential I/Q output. It is driven by an 8-phase LO with 1/8 duty cycle, which is generated from an externally applied clock at an eight times higher frequency.

The LO-generation circuitry is shown in fig. 6.5. It first converts a differential sine-wave input to a square wave, which is then used to drive a circular shift register. One flipflop is preloaded with a ‘1’, and all others with ‘0’. In this way, the eight outputs have a duty cycle of 1/8. The maximum input frequency is 8 GHz (limited by the reset-circuitry), so the maximum LO-frequency is 1 GHz.

The shift-register outputs are buffered by clock drivers to drive the mixer switches. A small overlap of the clock phases, which can occur due

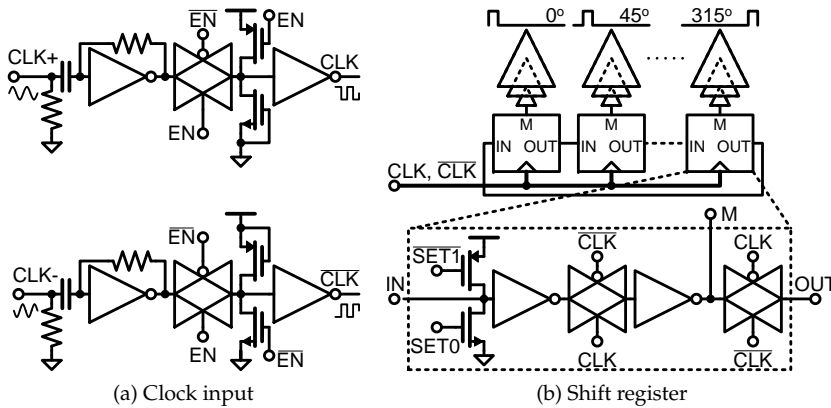


Figure 6.5: LO-generation circuitry to generate an 8-phase LO.

to mismatch, can result in significant noise degradation. Therefore, the clock drivers internally employ asymmetric rise and fall times to slightly reduce the duty cycle to nominally $1/9$.¹

Fig. 6.6 shows a plot of input-referred third-order intermodulation intercept point (IIP3) for the cascade of LNTA and passive mixer (driven by an ideal LO), where the width of the mixer switches is swept. The impedance from mixer output to TIA-input and the non-ideality of the TIA are simulated as an ideal $10\ \Omega$ resistance (IIP3 is determined at these $10\ \Omega$ resistors). Larger mixer switches require a more power-hungry clock driver, so a width of $25\ \mu\text{m}$ is used to obtain a simulated IIP3 of around $+16\ \text{dBm}$ for the cascade, in both $50\ \Omega$ and $100\ \Omega$ mode. In both cases, the choice is less than 1 dB from the simulated optimum at a width of $40\ \mu\text{m}$.

6.1.4 On-chip Receiver Connection

The two receivers on-chip are identical. For easier routing on the PCB (one differential transmission line per side), the second receiver is rotated 180° with respect to the other (no strict receiver matching is required, see section 4.1). This requires a long on-chip wire ($> 1\ \text{mm}$) to connect them.

To connect or disconnect the two receivers, a series-shunt configuration is used, as shown in fig. 6.7a. This configuration improves the isolation between the two receivers in stand-alone mode, and isolates the input from the parasitic capacitance of this long line. With the switches sized to have $2\ \Omega$ on-resistance each, and an ohmic loss of $2\ \Omega$ for the wire, $0.5\ \text{dB}$ IL is added for the second receiver.

¹It is worth mentioning that this implementation has some problems. At high frequencies, the reset-circuitry is not fast enough, so a reset has to take place at a lower frequency and the frequency should be slowly increased. At certain frequencies, however, the available signal generators seem to switch internally, which creates a transient, causing the ring of flipflops to lose state. Therefore, it may be better to use stateless dividers, or add some logic such that the ring automatically converges to a desired state, regardless of the initial state.




Figure 6.6: Simulated IIP3 LNTA and mixer cascade versus W of the mixer switches at $f_{LO} = 500$ MHz.

6.2 Measurements

The chip is implemented in the 65 nm low-power bulk CMOS technology of ST. It operates at 1.2 V, with an area (including bondpads) of 1×1 mm². An annotated photograph is shown in fig. 6.8a. The active area excluding decap is 0.15 mm². The chip is packaged in a QFN44-package, as shown in fig. 6.8c, where the top has been removed to show the IC inside it.

The measurement setup is shown in fig. 6.9. It is located inside a Faraday cage to minimize external interference. The IL of cables and hybrid, but not the PCB, are corrected for (de-embedded) from the measurement results. Off-chip baseband circuitry is not de-embedded. The feedback resistance of the TIA is 1 k Ω , and the feedback capacitance is 8 pF for 20 MHz bandwidth. The anti-alias filters, different from the ones used in chapter 4, have a cut-off frequency of 8 MHz, and the ADCs sample at 10 MS/s, resulting in a noise floor that is slightly higher near the Nyquist frequency in digital baseband due to aliasing.

The isolation in stand-alone mode is measured as follows. With both receivers in 100 Ω mode and connected on-chip, the output power of the second receiver is measured for a certain input power applied to the RF-input of the first receiver. Then the receivers are put in 50 Ω mode and discon-






Figure 6.7: The two receivers can be (dis)connected.



(a) Chip micrograph of UTFSINX,




(b) Zoom-in on logo

(c) UTFSINX inside a QFN44-package


Figure 6.8: UTFSINX is fabricated in the 65 nm bulk CMOS process of ST.

nected without changing the applied input power. The power difference between these two measurements is the isolation between the receiver inputs when they are disconnected, and the result is shown in fig. 6.7b for 0 dB attenuation. More than 80 dB of isolation is obtained up to 0.8 GHz, after which it drops to just below 60 dB at 1 GHz.

Fig. 6.10a shows the LO power consumption (per receiver): it scales almost linearly from 7.5 mW at $f_{LO} = 0.3$ GHz to 20.4 mW at $f_{LO} = 1.0$ GHz.



(a) PCB with IC and opamps



(b) Measurement setup in Faraday cage


Figure 6.9: Setup for UTSFINX-measurements. The PC with anti-alias filters and ADCs is located outside the Faraday cage to minimize interference.

The differential implementation ensures low LO-leakage from the mixer to the input, while the LNTA (and attenuator) reduce the leakage further. The measured LO-radiation (0 dB attenuation) is well below -70 dBm (see fig. 6.10b), and scales with the attenuation (not shown). The RF circuitry consumes 15.3/12.8/10.3 mW per receiver in the 50 Ω /100 Ω /highZ-mode.

6.2.1 Two Receivers in Parallel

Measurement results of two receivers in parallel, each with 100 Ω input impedance, are shown in fig. 6.11. Matching ($S_{11} < -9.6$ dB) at 0 dB attenuation is obtained from 150 MHz (limited by the external inductor) to 650 MHz (limited by capacitance of the attenuators, LNTAs, four bondpads, and the long interconnect). At higher attenuation, the inductor and part of the capacitance are shielded by the resistive attenuator, achieving matching from below 100 MHz to above 1 GHz.

At 0 dB attenuation, the gain is around 30 dB, and varies 2.2 dB over the whole band, which closely matches simulation. The NF is around 10 dB, which is almost 3 dB more than simulated. After extensive searching



(a) Power consumption (incl. drivers)

(b) Radiation at the antenna

Figure 6.10: Measurements of LO power consumption and radiation.




Figure 6.11: Measurement results of two receivers in parallel for all attenuation settings (black lines and circles for receiver 1; triangles for receiver 2). Some simulation results for 0 dB attenuation are shown as gray lines. In the bottom graph, the HR of the first receiver is shown for several harmonics.

and debugging, it is concluded that the cause of the discrepancy must be originating from inside the chip, but the origin has not been found. The gain of the other receiver follows that of the first, but is 0.2–0.6 dB lower, as expected from the IL. The attenuation lowers the gain and increases the NF, as expected.

For 0 dB attenuation, the 1-dB compression point (CP) increases from 0 dBm at 300 MHz to around +3 dBm at 700 MHz and above. The IIP3 is very close to the simulation results at around +15 dBm over the whole band, except below 400 MHz where it is slightly higher. Both CP and IIP3 increase 1 dB per dB with the attenuation (CP at 10 dB attenuation is above +12 dBm; the input power is not increased further due to risk of oxide/junction breakdown).

HR-measurements are performed with input frequencies up to 3.5 GHz; e.g. for measuring HR₇, the LO-frequency is increased only up to 0.5 GHz. The even-order harmonics are suppressed by 57 dB or more, indicating good matching and LO-balance. image rejection (IR) is relatively poor at 30 dB, which may be caused by the baseband components (3° difference in phase or 0.6 dB difference in gain already limits IR to 30 dB, see fig. 2.7). Theoretically, the 7th harmonic is 17 dB down for an ideal square wave and with infinite bandwidth; the measurements show HR₇ around 26 dB, with the difference most likely caused by limited matching bandwidth.

6.2.2 Single Receiver

With the receiver used for regular reception, the two receivers are disconnected on-chip. Each receiver turns to 50 Ω input impedance by reconfiguring the attenuator and LNTA as described in section 6.1.

Disconnecting the receivers largely reduces the parasitic capacitance at the input, which results in a much wider matching bandwidth, see fig. 6.12. With two CG-transistors activated for 50 Ω matching, the transconductance of the LNTA, and therefore overall gain, is expected to increase by about 1 dB, which is verified in the measurements. Due to the lower capacitance, the gain curve is more constant than for two receivers. NF at 0 dB attenuation is measured at around 8 dB, which is again almost 3 dB more than simulated. Compared to the two-receiver case, CP is lower by about the same amount as the gain is higher. IIP3, on the other hand, is very comparable, which is in agreement with fig. 6.6.

The HR-results (here measured up to 3 GHz input frequency) show again that the even harmonics are well suppressed, although HR₂ < 60 dB at higher frequencies. HR₅ is comparable to the two-receiver case, while HR₃ is somewhat worse. Based on the 2 : 3 : 2 ratio of the LNTA, HR₃ ≈ HR₅ > 30 dB is expected at all frequencies. It is not yet clear why HR₃ is worse, and why there is so much difference between HR₃ and HR₅.

6.2.3 High-impedance Mode

The noise correlation with two receivers in parallel limits the measured DANL_{Hz} to about -171 dBm/Hz due to the 100 Ω matching at the input. The inputs of the receivers can be made high-ohmic, so that in principle no noise currents can flow from one receiver to the other. This should not only




Figure 6.12: Measurement results of a single receiver for all attenuation settings. Some simulation results for 0 dB attenuation are shown as gray lines. In the bottom graph, HR is shown for several harmonics.

achieve a higher sensitivity [92], but is also very interesting for built-in self-test (BIST), where internal nodes are preferably not loaded.

The setup is exactly the same as in the matched case. The IL due to wiring and the external hybrid is assumed to be identical to the matched case. The source (signal generator) is modeled as a voltage source with $50\ \Omega$ output resistance. The signal generator assumes $50\ \Omega$ matching, and




Figure 6.13: Measurement results of two receivers in parallel with high-ohmic inputs (circles for receiver 1, triangles for receiver 2). Some simulation results are shown as gray lines.

internally generates twice the voltage required for the desired power at the load. Since the MOS-devices are voltage-controlled, this $2\times$ passive voltage gain improves gain and NF, and degrades IIP3 and CP.

The transfer of the used hybrid (Tyco H-183-4) is not specified for terminations other than 50Ω , so its effect is hard to de-embed. Here, the IL of the hybrid is de-embedded, measured with proper 50Ω terminations. The measurement results, shown in fig. 6.13 for two high-ohmic receivers in parallel, agree reasonably well with simulations with an ideal hybrid (50Ω input impedance on all ports, regardless of their loads). A more or less periodic ripple in measured gain is observed (also visible in NF, CP, and IIP3), which may indicate impedance mismatch effects. Note that resistive attenuation does not make sense in this case, and thus only the 0 dB setting is measured.

The measured return loss is less than 3 dB over the whole band; most of the power is reflected, as desired. At higher frequencies, parasitic input capacitance lowers the gain and increases the NF (which is again almost 3 dB higher than simulated). As one would expect, a higher gain gives a lower compression and linearity.




Figure 6.14: Crosscorrelation noise measurements at $f_{LO}=400$ MHz. Left: noise floor as function of NMT for 10 dB attenuation. Right: $DANL_{Hz}$ as function of NMT for all attenuation settings.

6.2.4 Crosscorrelation

Fig. 6.14 shows the measured DANL as a function of normalized measurement time (NMT) for 10 dB attenuation at $f_{LO} = 400$ MHz, where $NMT = 1$ equals the time required to obtain enough samples for one fast Fourier transform (FFT) per receiver (100 μ s for the 10 kHz resolution bandwidth (RBW) in fig. 6.14). It is worth explaining why the DANL of -155 dBm/Hz at $NMT = 1$ is about 1 dB less than the DANL of -154 dBm/Hz from autocorrelation (AC) (a NF of 20 dB in fig. 6.11). For $K = 1$, AC has a non-central χ^2 -distribution with 2 degrees of freedom (with mean 2), while XC has a Rayleigh-distribution (with mean $\sqrt{\pi/2}$); thus the difference is a factor $\sqrt{2/\pi} \approx 0.8$. This difference also follows from the approximation using (3.30): for AC ($\rho = 1$), $\mathbb{E}[\tilde{P}_{xc}] = 1$, while for XC ($\rho \approx 0$), $\mathbb{E}[\tilde{P}_{xc}] \approx \sqrt{\pi/4} \approx 0.8$. Both results correspond to -1 dB.

As an example of the improvement obtained with XC, consider the 10 dB attenuation setting with a RBW of 1 MHz (which is somewhere between the bandwidth of a wireless microphone and a digital TV (DTV)-signal). Obtaining enough samples for each FFT (independent of the actual ADC sample rate) then takes 1 μ s, so the use of 600 FFTs takes 0.6 ms, an acceptable time for many purposes. From fig. 6.14, $DANL_{Hz}$ then decreases from -155 dBm/Hz to -167 dBm/Hz, improving spurious-free dynamic range (SFDR) by $\frac{2}{3} \cdot 12 = 8$ dB.





Figure 6.15: Final $DANL_{Hz}$ obtained after $NMT = 2^{16}$ as a function of RF-frequency and attenuation setting.

Fig. 6.15 shows $DANL_{corr}$ over the RF-bandwidth of 0.3–1.0 GHz, measured by inserting a known tone of low power at 1 MHz IF (the image power is well below the noise) and determining the noise floor around -1 MHz IF. The IF-gain at 1 MHz is the same as at -1 MHz, so in this way an accurate gain estimate is obtained, and, therefore, an accurate input-referred noise level. Even though the receiver NF is almost 3 dB more than expected, $DANL_{corr}$ is close to the predicted -172 dBm/Hz for all attenuation settings. The DANL at 10 dB attenuation deviates somewhat from the other attenuation settings at higher frequencies; an explanation for this has not yet been found.

In the highZ-mode, there is ideally *no* correlated noise in the two receivers: $DANL_{corr}$ is expected to be close to -174 dBm/Hz. The measurement results indicate that this is almost achieved: a $DANL_{corr}$ around -173 dBm/Hz is found, again about 1 dB higher than predicted in section 6.1.2. Apart from small calibration errors, this difference may be



(a) XCSA-output ($A_{dB} = 10$ dB) of a modulated 1 MHz-wide 100-tone signal. The individual receiver spectra are also shown. (b) Measurement of phase noise reduction using two PM-modulated frequency-locked oscillators (0 dB attenuation).

Figure 6.16: Several other measurements performed with the XCSA.

explained by the shared low-ohmic ground that may introduce some noise correlation, and the IL of the PCB. This IL is shared by both receivers, and directly translates to a higher DANL.

To demonstrate the benefit of a high SFDR, fig. 6.16a shows the measured spectrum of a modulated multitone signal that was applied to the input of the system. As a reference, the spectrum as detected by the individual receivers is also shown. The signal is created by setting the generator to a high output power and using a passive attenuator at its output to lower the signal to around -50 dBm. The spectral widening due to non-linearity in the signal generator can be observed much more clearly and with greater accuracy in the XC-output.

In the presented system, the oscillators are external with good phase noise performance, which prohibits a visible improvement in phase noise performance using our ADC-board. Therefore, wideband PM-modulation is applied to two frequency-locked synthesizers from Agilent. In this way, the phase noise of both LOs is independent and detectable. The obtained spectra at 0 dB attenuation are shown in fig. 6.16b. As a reference, the spectra of the individual receivers are also shown. The phase noise is seen to be reduced by more than 20 dB. This improvement is independent of the actual phase noise performance of the LOs, so XC can also be used to make a good LO look even better.

6.2.5 Crosscorrelation and Harmonic Rejection

In section 3.3.3, the technique of [98] was discussed to improve HR by employing XC with a frequency offset between the two receivers. This frequency offset could not be applied to the prototype of chapter 4 because of the lack of isolation between the two mixer-first receivers. In UTSEINX, a frequency offset *can* be applied, and measurement results thereof are discussed next.

The measurements of HR_k ($k \neq 0, 1$) are carried out with the following settings: $f_2 = f_1 - 0.2$ MHz (thus $\Delta f_{LO} = -0.2$ MHz), $f_{in,1} = f_{LO} + 1.2$ MHz ($P_1 = -90$ dBm), $f_{in,k} = kf_{LO} + 1.9$ MHz ($P_k = -20$ dBm). The measured power at 1.2 MHz IF is set to equal the input power of -90 dBm, and the HR_k can then be found by evaluating the power of the downconverted harmonic compared to its original input power. For IR, only one input tone is applied at $f_{in,1} = f_{LO} + 1.2$ MHz at -50 dBm, with the image power measured at -1.2 MHz IF. Flat-top windows are used to accurately estimate the power of the sinusoids and to have good sidelobe suppression to maximize the maximum improvement in HR. This will be explained in more detail later in this section.

An example measurement for IR is shown in fig. 6.17 (only part of the usable baseband spectrum shown). In the spectrum of receiver 1 (receiver 2), the desired signal at 1.2 MHz (1.4 MHz) and the image at -1.2 MHz (-1.4 MHz) are clearly visible. The IR is around 29 dB for both




Figure 6.17: Example measurement for IR at $f_{LO} = 400$ MHz.

receivers. After correcting for Δf_{LO} in receiver 2, the desired signal lines up, while the images do not. After performing XC only once, the strongest spur present is at -48 dBc.² Unfortunately, this spur does not go down by correlating longer, because it is present in both receiver outputs.

The spurs that can be found in the spectrum of receiver 1 are located at -1.4 , -0.2 , 1.4 and 1.6 MHz. The spur at -1.4 MHz may be caused by crosstalk at IF: it is 50 dB below the desired signal in receiver 2 (due to the IQ-nature of the IF-signals, the signal at 1.4 MHz can be visible at -1.4 MHz in the other receiver). The other spurs may be explained by a mistake in the layout: although the analog and digital supplies of the two receivers are all separate (there are 4 supply domains, all at 1.2 V), their grounds were accidentally low-ohmically connected via the padding. Measurements of the on-chip ground (indirectly accessible via the RF-ground of the second RF-input, which was not used in the measurements) clearly show both LO-frequencies (and both clock frequencies at eight times the LO-frequencies), as well as many of their harmonics.

The LO-buffers are inverters, which respond to their input, but also to the supply and ground. As a result, the mixers are driven by an LO that is modulated by the other clock and LO. The inverter-like CS-stages of the LNTAs are connected to the same ground, and are also modulated by the LOs and their harmonics. Overall, all kinds of intermodulation products

²As they look similar, spurs and image components are collectively referred to as spurs.




Figure 6.18: Example measurement for IR at $f_{LO} = 1000$ MHz.

between the input signal and the LO-frequencies and their harmonics may arise in the signal and the LO-path, which after downconversion can result in even more terms. Further research should be carried out to determine the cause of these spurs, e.g. by modulating the input signal such that it can be clearly distinguished from the tones of the LOs.

Another measurement example (with higher RBW) of IR is shown in fig. 6.18, where the LO-frequency is now 1 GHz instead of 400 MHz. Again, the IR of the individual receivers is around 29 dB, and the spur after XC at -1.4 MHz is at -48 dBc. However, in this case the strongest spur is located at -1.6 MHz, likely caused by more crosstalk at the higher operating frequency. One can argue now about which frequencies constitute the image: is it just at -1.2 MHz and -1.6 MHz, or should any spurs that are (in)directly caused by finite IR be counted as well? In order to be on the safe side, the strongest spur in the whole frequency range around (more precisely: plus or minus Δf_{LO}) the two image frequencies is used to determine IR. Similarly, for HR_k , the strongest spur around the two downconverted components $\pm k\Delta f_{LO}$ is taken to determine HR_k . Thus, in this example, IR is found to be 35 dB, only about 5 dB better than the individual receivers.

The results are also shown when, instead of correlating only once, XC is also performed with $NMT = 8$. The noise floor goes down (as expected), but so are some of the spurs, such as the ones at -1.2 and 0.4 MHz. For the spur at -1.2 MHz, this is precisely what is predicted by (3.38): the power of the decorrelated image is reduced by XC. The initial power in the XC-spectrum is the geometrical mean of the power of receiver 1 (caused by the image) and that of receiver 2 (caused by noise). HR_k is thus not




Figure 6.19: Example measurement for HR_3 at $f_{LO} = 300$ MHz.

only a function of the HR_k of the individual receivers, *but also of their noise performance and the RBW*.

A measurement of HR_3 is shown in fig. 6.19, with $f_{LO} = 300$ MHz. The signal $f_{in,3}$ of -20 dBm at 901.9 MHz mixes down to -1.9 MHz for receiver 1 and to -2.7 MHz for receiver 2 (after frequency correction), both at -48 dBm. Thus, with an input power of -20 dBm, this results in $HR_3 = 28$ dB. For XC, the strongest remaining peak is -70 dBm at -2.7 MHz, resulting in $HR_3 = 50$ dB, an improvement of 22 dB. As was the case for IR in the previous examples, correlating longer does not lower this highest peak as it is correlated between the receivers.

Without crosstalk, the immediate HR-improvement would be higher, and further improved by increasing the measurement time (see also (3.38)). Nevertheless, even with crosstalk, a significant gain in HR is obtained.

HR_k is measured for different f_{LO} and k , and the results are shown in fig. 6.20. The results without XC are the geometrical mean of HR_k of the two receivers (the HR of the two receivers matches quite well). HR_k with XC is consistently determined using the strongest spur present (as explained before), which is always better than HR_k for the individual receivers, except in the single instance of HR_2 measured at 300 MHz, where it is a fraction of a dB worse.

The 4th and 6th harmonics are already 70 dB rejected in the individual receivers by the differential implementation, but XC still improves HR_4 and HR_6 to 75 dB, *without* calibration. With XC, HR_7 (the 7th harmonic is not suppressed by the HR-mixer) improves from 25 dB for the individual receivers to 45 dB, an improvement of 20 dB. For other harmonics, such as the image, the 3rd and the 5th, the improvement obtained using XC ranges from 5 dB to 25 dB. The IR decreases with frequency, because the spur at




Figure 6.20: Measured HR without (black markers) and with XC (white markers) for several relevant harmonics, measured up to 3.5 GHz input frequency (thus the 7th harmonic is only measured up to 500 MHz).

−1.6 MHz (also visible in fig. 6.17) increases with frequency (likely due to the crosstalk mechanism).

Discussion

The definition of HR becomes somewhat vague when a frequency offset in combination with XC is used, as there are now two fundamental components, each of which can generate responses in the XC-spectrum due to finite HR of the individual receivers. Crosstalk mechanisms, such as experienced in UTSFINX, may generate additional components.

Without averaging, where receiver 1 would estimate $|R_1[k]|^2$, and receiver 2 would estimate $|R_2[k]|^2$, XC using \tilde{P}_{xc} would estimate $|R_1[k]| \cdot |R_2[k]|$. In other words, the power in the XC-spectrum is the geometrical mean, or average power *in dB*, of the AC power estimates of the individual receivers. The power estimates of the individual receivers include the noise contributions, which are both functions of their receiver NF and the RBW. So, even when there is no crosstalk, the finite HR in a receiver will introduce a spur in the XC-spectrum. If a signal in a single receiver is 10 dB stronger, the spur in the XC-spectrum will be 5 dB higher. Similarly, if there happens to be another signal downconverted to the same baseband frequency, the spur in the XC-spectrum will also be higher. Thus, HR actually is a function of the input spectrum. If the noise is uncorrelated from the signals (which

is generally assumed), the spur will reduce for longer measurement time. This means that HR also is a function of measurement time.

Even if one assumes that there is no crosstalk, no correlated noise, infinite measurement time, and infinite precision in the digital domain, the achievable HR using the basic FX-correlator (FXC) is still limited due to limitations in the FFT-window.³ Consider an arbitrary window $w[m]$ of length M , normalized such that $\frac{1}{M} \sum_{m=0}^{M-1} w[m] = 1$ to correctly estimate the power of a sinewave. Its frequency response $W(f)$ is defined as

$$W(f) \triangleq \sum_{m=0}^{M-1} w[m] e^{-j2\pi m f}. \quad (6.5)$$

Thus, $W(f) = W(f + 1)$ and $W(0) = 1$. A simple case to evaluate is the situation where a sinewave at a certain frequency f_0 is present in the output of receiver 1, and a sinewave at a frequency $f_0 + \Delta f$ in receiver 2:

$$r_1(t) = e^{j2\pi f_0 t} \quad r_2(t) = e^{j2\pi(f_0 + \Delta f)t}, \quad (6.6)$$

with $t = m/f_s$ and $f_0 = \frac{f_s}{M}(k_0 + \Delta k)$. Here, Δk , with $-\frac{1}{2} \leq \Delta k < \frac{1}{2}$, represents the offset of f_0 with respect to the center of a bin. Note that when $\Delta f \neq 0$, there is no mathematical correlation at *any* frequency between $r_1(t)$ and $r_2(t)$. Writing the receiver outputs in the digital domain as

$$r_1[m] = e^{j2\pi m \frac{k_0 + \Delta k}{M}} \quad r_2[m] = r_1[m] e^{j2\pi \Delta f m} = e^{j2\pi m \frac{k_0 + \Delta k + \Delta f}{M}}, \quad (6.7)$$

the response of the FFT of r_1 and r_2 in bin k are

$$\begin{aligned} R_1[k] &= \sum_{m=0}^{M-1} r_1[m] w[m] e^{-j2\pi \frac{km}{M}} = \sum_{m=0}^{M-1} w[m] e^{-j2\pi(k - k_0 - \Delta k) \frac{m}{M}} \\ &= W\left(\frac{k - k_0 - \Delta k}{M}\right) \\ R_2[k] &= W\left(\frac{k - k_0 - \Delta k - \Delta f}{M}\right). \end{aligned} \quad (6.8)$$

Thus, using an FXC and \tilde{P}_{xc} , the result of XC is

$$\tilde{P}_{xc}[k] = \left| W\left(\frac{k - k_0 - \Delta k}{M}\right) \right| \left| W\left(\frac{k - k_0 - \Delta k - \Delta f}{M}\right) \right|. \quad (6.9)$$

With $\Delta f = 0$, $\tilde{P}_{xc}[k_0] = |W(-\Delta k/M)|^2$, which is only the true power 1 if $\Delta k = 0$. In practice, input signals can occupy any frequency, so the window response may be evaluated at arbitrary Δk . The deviation $|W(-\Delta k/M)|/W(0)$ is referred to as scalloping loss. It is thus desirable to have $|W(-\Delta k/M)| \approx W(0)$ for all Δk . Similarly, for $\Delta k = 0$ and $\Delta f \neq 0$, \tilde{P}_{xc} should converge to 0 as there is no correlation, which means

³Window-tradeoffs important for XC spectrum sensing, such as equivalent noise bandwidth and sidelobe fall-off rate, are left outside of the discussion.

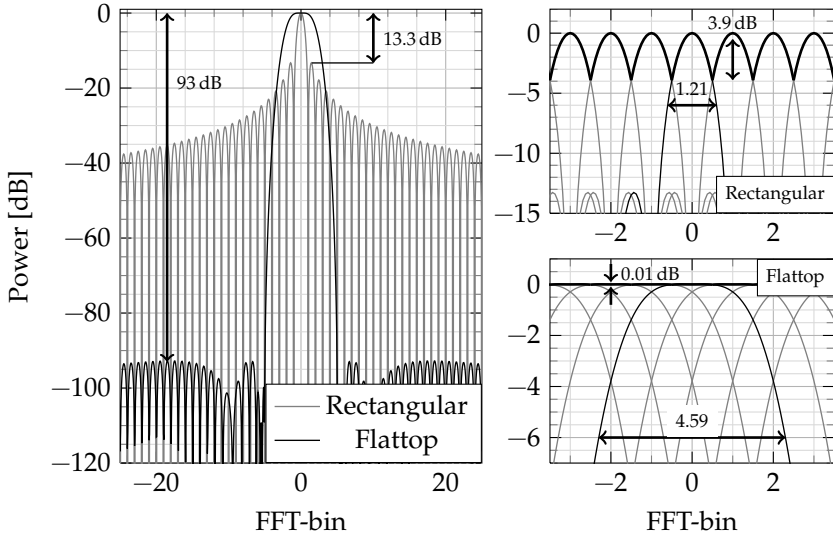


Figure 6.21: Illustration of sidelobe suppression (left), scalloping loss, and 6 dB-BW (right) of 1024-point rectangular and flattop windows.

$W(-\Delta f/M) = 0$ is desired $\forall |\Delta f| > \frac{1}{2}$: there should be no significant leakage. For finite M , there must be a transition region between $W(f) = 1$ and $W(f) = 0$. Usually a 6 dB-BW is defined where $|W(f)| \geq 0.5$.

Moreover, $|W(f)|$ may increase again due to sidelobes. The sidelobe suppression is the relative power of the largest sidelobe: it is equal to the inverse of the maximum of $|W(f)|^2$ after the first zero-crossing. The sidelobe suppression of a window is generally less for smaller M .

To properly decorrelate harmonic images, the sidelobe suppression should be high enough and Δf should be significantly larger than the BW of the window. This is graphically illustrated for the rectangular window and the flattop window in fig. 6.21. The rectangular window w_R is equal to applying no window at all:

$$w_R[m] = 1 \implies W_R(f) = e^{-j\pi f(M-1)} \frac{\sin \pi f M}{M \sin \pi f}, \quad (6.10)$$

while the flattop window w_F is defined as (rounded coefficients)

$$w_F[m] = 1 - 1.93 \cos\left(2\pi \frac{m}{M-1}\right) + 1.29 \cos\left(4\pi \frac{m}{M-1}\right) - 0.388 \cos\left(6\pi \frac{m}{M-1}\right) + 0.032 \cos\left(8\pi \frac{m}{M-1}\right). \quad (6.11)$$

For the rectangular window, $W_R(k/M) = 0$ for $k \in [1, 2, \dots, M-1]$. If $\Delta k = 0$, the response is 1 for $k = k_0$, and 0 otherwise (i.e. no leakage). The 6 dB-BW is 1.21 bins, the scalloping loss 3.9 dB, and the sidelobe suppression for a 1024-point FFT 13.3 dB. In the scenario with a single sinewave in each receiver output, this means that, if Δf happens to be the frequency

distance between the main and the first sidelobe, $\tilde{P}_{xc}[k_0]$ is only 6.6 dB lower than the power of the sinewave that is present in only one of the two receivers.

The flattop window has a scalloping loss below 0.01 dB, and a sidelobe suppression for a 1024-point FFT of 93 dB. These are the reasons to choose the flattop window in the measurements presented here. However, the 6 dB-BW is 4.59 bins, which means that the RBW is significantly reduced. Moreover, the flattop window gives a loss of 5.76 dB in signal-to-noise ratio (SNR). Both properties would significantly increase the required measurement time for a given $P_{FA,des}$ and $P_{MD,des}$, and desired RBW.

All in all, there is a lot more to consider with respect to improving HR with XC than what has been done in the measurements presented in this chapter. The measurement results should therefore be interpreted with care. What can be concluded with certainty based on these measurements is that (1) HR can be improved with this technique, (2) it works for all harmonics, (3) it can give significant improvements without additional measurement time, and (4) crosstalk can be a major limiting factor.

6.3 Benchmarking

Table 6.1 compares a number of spectrum analyzers (SAs), spectrum sensing solutions and wideband receivers from literature (identical to table 4.2) with UTSFINX. Without XC, the linearity and DANL are comparable to commercial SAs. With XC, a significantly lower noise floor is obtained, even lower than the PXA with Noise Floor Extension (NFE) enabled. Lowering the noise floor comes at the cost of measurement time. The 10 dB attenuation setting requires 1700 averages to get a $DANL_{Hz}$ of -168 dBm/Hz. The lower NF for the same IIP3 makes it 2.5 times faster than the prototype in chapter 4 with network 2.

Compared to spectrum sensing solutions in literature, a much higher linearity and better DANL is obtained, even without XC. As the external baseband components used in this experiment operate at 15 V, the comparison may not be entirely fair compared to [55, 56, 66, 67], which have integrated baseband components. Nevertheless, it still gives an indication that the performance achieved by this technique is very promising.

6.4 Conclusions

An integrated SA is desirable for software-defined radio, Dynamic Spectrum Access (DSA), and BIST, but the design faces many challenges. Spectrum analysis using XC can reduce thermal noise and phase noise, and improve linearity, at the cost of measurement time.

In this chapter, a high-linearity design in 65 nm CMOS operating at 1.2 V from 300 MHz to 1.0 GHz is discussed. It consists of two RF-frontends with attenuators, and achieves +25 dBm IIP3 and a DANL of better than

Table 6.1: Comparison with other designs

Architecture	CMOS-node [nm]	RF-range [GHz]	Power [mW]	NF [dB]	DANL _{corr} [dBm/Hz]	Time penalty [NMT] ^a	IIP ₃ [dBm]	FoM _{SFDR} [dB]
UTSFINX (0 dB attenuation)	65 ^b	0.3–0.7	41–54	11	-169	16	15	9 ^c
UTSFINX (2 dB attenuation)	65 ^b	0.3–1.0	41–66	13	-169	42	17	11 ^c
UTSFINX (6 dB attenuation)	65 ^b	0.3–1.0	41–66	17	-169	2.7·10 ²	21	15 ^c
UTSFINX (10 dB attenuation)	65 ^b	0.3–1.0	41–66	21	-169	1.7·10 ³	25	19 ^c
UTSFINX (highZ-mode)	65 ^b	0.3–1.0	36–61	6–10	-172	6–42	4–10	1–7 ^c
Chapter 4 (network 1)	65 ^d	0.05–1.5	176 ^e	17	-169	2.7·10 ²	17	11 ^c
Chapter 4 (network 2)	65 ^d	0.05–1.5	176 ^e	23	-169	4.3·10 ³	24	18 ^c
Soer et al. [53]	65	0.2–2.0	67	5.5	-168		11	5
Ru et al. [54]	65	0.4–0.9	60	4	-170		4	0
Park et al. [67]	180	0.4–0.9	180	50 ^f	-124		-17 ^g	-67
Kitsunezuka et al. [124]	90	0.03–2.4	37	39 ^f	-135		-11	-50
Pollin/Ingels et al. [125, 126]	40	0.1–6.0	100 ^h	3	-171		-12	-15
Murphy et al. [56]	40	0.08–2.7	35–78	2	-172		-22 ⁱ	-24
Greenberg et al. [55] ^j	80	0.04–1.0	440	3	-171		-15	-18
Agilent PXA N9030A-503		0–3.6	450 · 10 ³	18	-166 ^k	40 ^l	22	14
Rohde & Schwarz FSH4		0.01–3.6	12 · 10 ³	29	-145		15	-14
CRFS RFeye Node		0.01–6.0	15 · 10 ³	10	-164		20	10
Signal Hound USB-SA44B		0.00–4.4	2.2 · 10 ³	30	-144		3	-27

^a Calculated to have an effective NF within 1 dB of NF_{corr}

^b Baseband components are external

^c Using the value 174 + DANL_{corr} + 1 as NF

^d The RF-core is 65 nm, but many other components are discrete

^e Only integrated part and estimated DSP power consumption

^f Calculated based on sensitivity given in dBm in a certain bandwidth

^g P_{1dB} given in paper as -27 dBm

^h Analog frontend + 10-bit ADCs

ⁱ IIP_{3,ib} close to, but outside IF-bandwidth; true IIP_{3,ib} expected to be worse

^j Receiver in maximum gain setting

^k Using noise figure extension

^l Calculated NMT based on 8 dB lower expected value without lower variance

-169 dBm/Hz at a power consumption of around 50 mW. When matching is not required, the DANL can be reduced further to below -172 dBm/Hz.

The presented prototype also allows verification of the concept of combining XC with an LO frequency offset for one of the receivers to decorrelate the harmonic images, and thus improve HR for all harmonics. Improvements of up to 25 dB in HR have been observed without requiring additional measurement time, and up to 80 dB HR is obtained for some harmonics *without* calibration. The improvement in HR is limited by crosstalk due to sub-optimal layout; without the crosstalk, the improvement is expected to be much more. In spite of these imperfections, the technique gives significant improvements.

Overall, it can be concluded that XC with two linear frontends is promising to realize integrated SAs in CMOS with high linearity and sensitivity.

LOWERING THE SNR-WALL WITH CROSSCORRELATION

The existence of a lower bound on the signal power that can be detected, the signal-to-noise ratio (SNR)-wall, was explained for energy detection (ED) in section 3.1.3. In section 3.2.1, it was intuitively argued that crosscorrelation (XC) can lower this SNR-wall. In this chapter, this explanation is quantified using a theoretical model, and some of its predictions are compared to SNR-wall measurements performed with `UTSFINX`.

7.1 The SNR-wall of Crosscorrelation

The model for XC as discussed in chapter 3 is repeated here for convenience in fig. 7.1. It consists of two receivers, which are ideally identical, but which may have differences in terms of noise performance and phase, frequency and amplitude response.

To find the SNR-wall for this model, the derivations follow the procedure for ED in section 3.1.3. First, the SNR-wall is derived for the ideal situation with two identical receivers. This allows the use of the estimator \hat{P}_{xc} , which uses only the real part of the output decision metric Y , as discussed in section 3.2.4. The important situation of phase offset of the input signal at the output of the two receivers is considered next. In that case, \tilde{P}_{xc} , which is based on the absolute value of the XC-output rather than only the real part, can be used to be independent of phase offset.

When both receivers are in phase, i.e., $\Delta\phi = 0$, the first two central moments of \hat{P}_{xc} are given in (3.27). For \mathcal{H}_0 (SNR = 0) and for \mathcal{H}_1 (SNR > 0),

Parts of this chapter have appeared in [MOA:3] and may appear in [MOA:14].




Figure 7.1: The XC system model of fig. 3.2b, repeated for convenience.

these can be written as

$$\mu_0 = \rho \quad \sigma_0^2 = \frac{1 + \rho^2}{2K} \quad (7.1)$$

$$\mu_1 = \rho + \text{SNR} \quad \sigma_1^2 = \frac{1}{2K} (2\text{SNR}^2 + (2 + 2\rho)\text{SNR} + 1 + \rho^2). \quad (7.2)$$

If the noise power is not exactly known, the threshold must be purposely biased by scaling with the peak-to-peak uncertainty U , which is assumed to equally affect the correlated and uncorrelated noise parts:

$$\lambda_{xc} = U(\mu_0 + \sigma_0 Q_{\text{FA,des}}^{-1}). \quad (7.3)$$

The theoretical and simulated P_{MD} and P_{FA} without noise uncertainty (i.e., $U = 1$, or, equivalently, $U_{\text{dB}} = 0$ dB) are in good agreement in fig. 7.2.

Filling in (7.3) into (3.10), and approximating for small SNR ($\text{SNR} \ll 1$) and long averaging ($K \gg 1$) results in

$$P_{\text{D}} \approx Q \left(\frac{(\rho(U - 1) - \text{SNR})\sqrt{K} + U\sqrt{\xi}Q_{\text{FA,des}}^{-1}}{\sqrt{\xi + (\rho + 1)\text{SNR}}} \right), \quad (7.4)$$

where $\xi \triangleq \frac{1}{2} + \frac{1}{2}\rho^2$. For a certain desired P_{D} , $P_{\text{D,des}}$, (7.4) can be solved to find the minimum SNR, SNR_{min} , to obtain both $P_{\text{FA}} \leq P_{\text{FA,des}}$ and $P_{\text{D}} \geq P_{\text{D,des}}$:

$$\text{SNR}_{\text{min}} \approx \rho(U - 1) + \frac{\xi U Q_{\text{FA,des}}^{-1}}{\sqrt{K}} + \frac{(Q_{\text{D,des}}^{-1})^2}{K} \frac{\rho + 1}{2} - \frac{Q_{\text{D,des}}^{-1}}{K} \sqrt{K(\xi + (U - 1)(\rho^2 + \rho)) + \theta}, \quad (7.5)$$

where $\theta \triangleq (\rho + 1)^2 (Q_{\text{D,des}}^{-1})^2 + \xi U Q_{\text{FA,des}}^{-1} (\rho + 1) \sqrt{K}$. Simplifying this using $K \gg 1$ and $|Q_{\text{FA,des}}^{-1}|, |Q_{\text{D,des}}^{-1}| \ll \sqrt{K}$:

$$\text{SNR}_{\text{min}} \approx \rho(U - 1) + \frac{\xi U Q_{\text{FA,des}}^{-1} - Q_{\text{D,des}}^{-1} \sqrt{(\rho^2 + \rho)(U - 1) + \xi}}{\sqrt{K}}, \quad (7.6)$$

which for $\rho = 1$ (autocorrelation (AC)) simplifies to (3.15).




Figure 7.2: Comparison between theory (lines) and simulation (markers) for several noise correlations ρ , with $U_{dB} = 0$ dB and $SNR = -10$ dB, using \hat{P}_{xc} . The marker values are obtained using 10^5 independent simulations.

For infinitely long averaging, the minimum SNR below which a signal cannot be reliably detected, the SNR-wall, is found as:

$$SNR_{wall} = \lim_{K \rightarrow \infty} SNR_{min} = \rho(U - 1). \quad (7.7)$$

This result is independent of $P_{FA,des}$ and $P_{D,des}$. Clearly, when $\rho = 0$, there is *no SNR-wall, independent of noise uncertainty* U .

SNR_{min} according to (7.6) is plotted in fig. 7.3 for several ρ and U with $P_{FA} = P_{MD} = 10^{-5}$ (similar results are obtained for other P_{FA} and/or P_{MD} as the SNR-wall does not depend on $P_{FA,des}$ and $P_{D,des}$). Without noise uncertainty ($U_{dB} = 0$ dB), there is no SNR-wall, although it requires less samples for $\rho = 0$ than for $\rho = 1$ to detect signals at a certain SNR, as was




Figure 7.3: Relation between SNR_{min} and sample complexity for several noise correlations ρ and noise uncertainties U according to (7.6).




Figure 7.4: Theoretical (lines) and simulated (markers) P_{MD} for several U and ρ with $\text{SNR} = -10$ dB using \tilde{P}_{xc} . The simulated P_{MD} is obtained using 10^5 independent simulations. Note: for $U_{dB} = 2$ dB, $\rho = (0.5, 1)$ are on top of each other, and for $U_{dB} = 5$ dB, $\rho = (0.1, 0.5, 1)$ are on top of each other.

also established in section 3.2.5. With noise uncertainty, one can clearly observe that both a lower uncertainty U and a lower correlation ρ lower the SNR-wall, which allows smaller signals to be detected.

The theoretical and simulated (assuming worst-case noise estimates) P_{MD} are plotted in fig. 7.4 for different U and ρ , with $P_{FA,des} = 0.1$. For higher U , a lower ρ is required to keep the detector robust.

When a phase offset $\Delta\phi$ between the two receivers may exist, \tilde{P}_{xc} can be used. K is assumed large enough to justify Gaussian approximations for Y_{re} and Y_{im} . When $\rho = 0$ and $\text{SNR} = 0$, Y_{re} and Y_{im} have equal variance. The covariance between Y_{re} and Y_{im} can be found to be 0, based on mutually independent real and imaginary noise components. Hence, Y_{re} and Y_{im} are uncorrelated, and, as they are normally distributed, independent. Thus, \tilde{P}_{xc} is Rayleigh-distributed. This gives $P_{FA} = \exp(-\lambda_{xc}^2/2\sigma_0^2)$ with explicit threshold

$$\lambda_{xc} = U\sigma_0 \sqrt{2 \ln \frac{1}{P_{FA,des}}}. \quad (7.8)$$

For other situations, i.e. $\rho > 0$ and/or $\text{SNR} > 0$, the distribution of \tilde{P}_{xc} tends to a Gaussian distribution, and is hence adequately described by the first two moments [90], given in (3.30).

To find the SNR-wall, the decision threshold can again be set as in (7.3) (with the proper μ_0 and σ_0 from (3.30)):

$$\lambda_{xc} = U(\mu_0 + \sigma_0 Q^{-1}(P_{FA,des})). \quad (7.9)$$




Figure 7.5: Comparison between theory (lines) and simulation (markers) for several ρ , with $U_{\text{dB}} = 0$ dB and $\text{SNR} = -10$ dB, using \tilde{P}_{x_c} . The marker values are obtained using 10^5 independent simulations.

The theoretical and simulated P_{MD} and P_{FA} are plotted in fig. 7.5 for $U_{\text{dB}} = 0$ dB. Theory matches simulations very well, except for very small ρ , where the simulated P_{FA} and P_{MD} are somewhat higher than desired. This is due to the Gaussian approximation, while the actual probability density function (pdf) is close to a Rayleigh-distribution, which has a heavier tail.

In both cases [(7.8) and (7.9)], $\lambda_{x_c} = \mu_0 + \sigma_0 g(P_{\text{FA,des}})$, with $g(x)$ some monotonically decreasing function of x . This means the derivation as shown in the case without phase offset can be repeated from (7.3) to (7.6). With the central moments given in (3.30), SNR_{min} can then be found using straightforward, but lengthy calculations. The result is identical to (7.6), except that ξ should be replaced by $1 - \beta_K + \beta_K \rho$. Hence, the same SNR-wall as in (7.7) is found:

$$\text{SNR}_{\text{wall}} = \rho(U - 1). \quad (7.10)$$

For completeness, the theoretical and simulated P_{MD} are plotted using \tilde{P}_{x_c} , see fig. 7.6. The central moments of (3.30) and worst-case noise estimates for P_{MD} are used, and, as can be observed, theory matches simulations very well. Although the SNR-wall is the same, taking the absolute value of the output increases the variance, and therefore the measurement time needs to be longer for a certain $P_{\text{FA,des}}$ and $P_{\text{D,des}}$. This is clearly visible for $\rho = 0$ and $U_{\text{dB}} = 5$ dB in fig. 7.4 as compared with fig. 7.6: for $P_{\text{MD}} = 10^{-3}$ the required measurement time is roughly doubled.

7.1.1 Example

As an example to show the improved performance of the XC process, consider the IEEE 802.22 standard (detection of -116 dBm in 6 MHz bandwidth, with $P_{\text{FA}} = 0.1$, $P_{\text{D}} = 0.9$, and the sensing time limited to 2 s, see




Figure 7.6: Theoretical (lines) and simulated (markers) P_{MD} for several U and ρ with $SNR = -10$ dB using \tilde{P}_{xc} . The simulated P_{MD} is obtained using 10^5 independent simulations. Note: for $U_{dB} = 2$ dB, $\rho = (0.5, 1)$ are on top of each other, and for $U_{dB} = 5$ dB, $\rho = (0.1, 0.5, 1)$ are on top of each other.

also table 2.1). Assume a noise figure (NF) of 5 dB ($SNR = -14.8$ dB) and $U_{dB} = 1$ dB.

With AC ($\rho = 1$), the SNR-wall is -6 dB, so the system would not be able to detect the signal. With XC, however, the signal can be detected if $\rho \leq 0.125$. One can numerically solve for T using (3.10) and (7.3), where μ_0, μ_1, σ_0 and σ_1 are obtained from the statistics of \hat{P}_{xc} (as in this example) or \tilde{P}_{xc} . However, for engineering purposes, a first-order estimate of the minimum required measurement time T_{min} may be very useful. It can be found by substituting $K = TW$ in (7.6) and solving for T :

$$T_{min} \approx \frac{1}{W} \left(\frac{\xi U B - A \sqrt{(\rho^2 + \rho)(U - 1)} + \xi}{SNR_{min} - \rho(U - 1)} \right)^2, \quad (7.11)$$

provided $SNR_{min} > \rho(U - 1)$. Using this equation, $T_{min} \approx 9.8$ ms for $\rho = 0.1$, $T_{min} \approx 0.53$ ms for $\rho = 0.01$ and $T_{min} \approx 0.46$ ms for $\rho = 0$. The numerically calculated T_{min} is 14 ms, 0.77 ms, and 0.66 ms, respectively. The very simple estimate given by (7.11) is off by less than a factor of 2.

Fig. 7.7 shows the resulting spectrum of a simulation where the NF of each receiver is 5 dB, $\rho = 0.1$, $U_{dB} = 1$ dB, $T = 140$ ms, $f_s = 30$ MHz and one digital TV (DTV)-signal is present with total power -116 dBm, occupying the band from -2 MHz to 4 MHz. The noise uncertainty is modeled here as a non-white noise floor to get an intuitive feeling for the detection problem. To better visualize the spectrum, 1024-point fast Fourier




Figure 7.7: Simulation showing the clear benefit of XC as compared with AC in the presence of noise uncertainty.

transforms (FFTs) are used. The TV-signal cannot be observed in the AC spectra due to the noise uncertainty, but it is clearly present in the XC spectrum.

7.1.2 Discussion

Tandra and Sahai [81] suggested three options to get around or lower the SNR-wall for ED: (1) impose some structure on the primary signal (i.e. use knowledge of the signal to be detected), (2) force some diversity in fading to improve SNR (e.g. by moving around while sensing), and/or (3) “*somehow reduce the noise uncertainty*”. (7.7) shows that XC achieves the latter option: the SNR-wall can be reduced by minimizing ρ .

As XC removes uncorrelated noise, it can be considered to improve the SNR. So, one could either argue that the SNR-wall is lowered, or that the SNR-wall remains the same and the SNR is improved. Both lead to the same conclusion: XC allows smaller signals to be detected.

Fig. 7.8 shows the required number of samples (per receiver) for detection of a signal for different SNR as a function of ρ for $U_{dB} = 1$ dB. For SNR = -5 dB, XC with $\rho = 0.1$ reduces measurement time by a factor of 38 as compared to $\rho = 1$ (AC of a single receiver output). This allows more time for actual data transmission, as well as a reduction of the energy consumption for spectrum sensing by almost a factor 20. This more than makes up for the additional power temporarily required by turning on the second receiver, and may be enough incentive to spend the additional chip area and cost to include the second receiver. For single-receiver systems, the faster sensing and the ability to detect smaller signals may be enough to warrant the use of an additional receiver, solely used for spectrum sensing. When the higher sensitivity of XC is not required, it can be turned off or used to lower the variance in a more traditional ED process by acting as a second energy detector.

Just like the noise level estimation, the correlation factor has to be estimated as well [97]. Even with extreme care in the design of the two receivers to minimize coupling, there will be some correlation of the noise, such as black-body radiation from surrounding objects or man-made noise.




Figure 7.8: Required number of samples for signal detection as a function of ρ for several different SNR (annotated in the figure) with $P_{FA} = 0.1$, $P_D = 0.9$ and $U_{dB} = 1$ dB, using \hat{P}_{xc} .

Hence, it seems impossible to obtain $\rho = 0$. Moreover, the assumptions of Gaussianity and stationarity of the noise may fail for long measurement times, so a larger part of the uncertainty may be residing in the correlated noise, reducing the improvement provided by the XC technique.

In conclusion, the model predicts that XC has a lower SNR-wall than AC, where a lower ρ implies a lower SNR-wall. Close to the SNR-wall of AC, XC can be orders of magnitude faster in signal detection.

7.2 Experimental Verification

In this section, the theoretical results that lower signals can be detected using XC are experimentally verified. The prototype UTFSINX implementation with the external analog-to-digital converters (ADCs) (see chapter 6) provides a good starting point. The performance of XC can be compared directly with AC (the performance of the individual receivers), as the used samples, and thus any temperature and gain variations, are identical for both ED schemes. \tilde{P}_{xc} is used for the XC process, as it is robust against the local oscillator (LO) phase difference between the two receivers without requiring any calibration.

7.2.1 Experimental Setup

Based on the measurement results presented in chapter 6, the LO-frequency is chosen in the band at 400 MHz, where there is good input matching. At 10 dB attenuation, the noise correlation between the two receivers is rather low, which, according to section 7.1, should significantly lower the SNR-wall, and is therefore used in all measurements. The displayed average noise level (DANL) (see fig. 6.14) is reduced from -154 dBm/Hz to -169.5 dBm/Hz, which, using fig. 7.1, corresponds to $\rho \approx 10^{-15.5/10} \approx 0.028$.




Figure 7.9: The measured noise floor between 2 MHz and 3 MHz at different points in time during the measurements (top: the individual receivers, bottom: after XC).

The measurement process involves a number of steps:

1. Measure the noise floor and peak-to-peak uncertainty;
2. Find λ_{ed} and λ_{xc} ($P_{FA,des} = 0.1$) for various K and SNR;
3. Determine the output SNR for a given input power;
4. Find P_D for various K and SNR using λ from step 2.

These steps will now be explained in more detail.

The noise floor is determined by measuring the averaged output spectrum of the individual receivers and that obtained with XC without applying an input signal. The ADCs sample at 10 MS/s: the baseband frequency ranges from -5 MHz to 5 MHz. In every measurement, 1024-pt FFTs with rectangular windows are used (verifications with other FFT-sizes give similar results). The noise power is determined by adding the power in the 103 bins that have a center frequency between 2 MHz and 3 MHz.

Fig. 7.9 shows the output noise floors of 13 measurements, performed at the start, at the end, and in between measurements with a signal present. The first measurement was started about an hour after turning on all equipment to allow everything to reach thermal equilibrium. The last measurement was done almost two hours after the first measurement. The receiver noise floors are higher near 3 MHz due to noise aliasing. The difference in the noise floor in the two receivers is caused by the loss of the on-chip switches to connect the receivers.




Figure 7.10: Measured distribution of the noise power (bins), and Gaussian fits (lines), for several K for XC (left) and the individual receivers (right).

Clearly, the noise level fluctuates over time, which was also observed in [35]. The difference between the highest and lowest total power in these bands is an indication of the peak-to-peak uncertainty U . For receiver 1, $U_{\text{dB}} \approx 0.08$ dB (with $U_{\text{dB}} \triangleq 10 \log_{10} U$), and for receiver 2, $U_{\text{dB}} \approx 0.07$ dB. The noise uncertainty of the two receivers is approximately equal, as expected. For XC with just one average (not shown), the noise floor is the geometric mean of the spectra of the two individual receivers. Indeed, the measurements show $U_{\text{dB}} \approx 0.08$ dB. Note that U_{dB} is much lower than the 1 dB mentioned in section 3.1.3, due to the controlled environment and the use of a signal generator to generate the input (for comparison, [35] finds $U_{\text{dB}} \approx 0.03$ dB in a similar configuration).

The final noise floor obtained using XC is about 15.5 dB lower than that of the individual receivers. For the final noise floor of XC ($K = 1.35 \cdot 10^6$), $U_{\text{dB}} \approx 0.19$ dB. Interestingly, U for high K is higher, which probably means that some of the fluctuations are from an external source (the signal generator or interferers) and are not removed through XC. These results in combination with (7.6) suggest that the SNR-wall for the individual receivers ($\rho = 1$, $U_{\text{dB}} \approx 0.07$ dB) is around -18 dB, and for XC ($\rho \approx 0.028$, $U_{\text{dB}} \approx 0.19$ dB) the SNR-wall is expected to be -29 dB, an improvement of more than 10 dB. This will be experimentally verified later.

Some examples of the measured noise power distribution are given in fig. 7.10. Higher K reduces the variance, and for XC also the mean. As expected, (1) the initial distribution for XC is close to a Rayleigh-distribution, and (2) all distributions converge to a Gaussian for high K .

The threshold λ for given $P_{\text{FA,des}}$ and K is set such that in a $P_{\text{FA,des}}$ fraction of the cases the measured values exceed λ . It could be calculated by combining the NF of the receiver and the statistics given in chapter 3, but using the measured results avoids errors introduced by limitations of the model. The thresholds λ_{ed} and λ_{xc} are determined for $P_{\text{FA}} = 0.1$, see




Figure 7.11: Obtained λ for $P_{FA} = 0.1$. Left: AC, right: XC.

fig. 7.11. The noise fluctuations cause λ to be slightly different for each of the measurements; the maximum λ is used to guarantee $P_{FA} \leq P_{FA,des}$.

To mimic a noise-like input signal between 2 MHz and 3 MHz at IF, a signal generator outputs a 64-tone signal (spaced at 15.6 kHz, each with random initial phase) with center frequency 402.5 MHz. The SNR is determined at a relatively high input power, see fig. 7.12. The DC-offset at 0 MHz and the spurs visible at ± 4 MHz in receiver 2 are of no concern as they are outside the band of interest.

With 103 FFT-bins between 2 MHz and 3 MHz, the ADCs capture approximately 103 independent samples in the band of interest per 1024 samples. The total power of the signal is determined by adding the power of the bins between 2 MHz and 3 MHz, and subtracting the average measured power during the noise calibration. Since 103 bins are used to estimate the power, K increases in steps of 103 rather than 1.

With these numbers the SNR turns out to be 10.67 dB for receiver 1, and 9.77 dB for receiver 2. Regardless of this 0.9 dB difference, it will be referred to as the measurement with $SNR = 10$ dB. Reducing the output power of the signal generator by 1 dB reduces the SNR of each receiver by 1 dB. Smaller FFTs would give more information at higher SNR, but that is of no interest for SNR-wall measurements.

Finally, the input signal is applied, and the threshold values found earlier are used to determine P_D for several SNRs. In each measurement, ²²⁷




Figure 7.12: Output spectrum with a signal present at $SNR \approx 10$ dB.




Figure 7.13: Measured P_D for $P_{FA} = 0.1$ as a function of K in the bandwidth of interest for several SNRs above the SNR-wall. For $\text{SNR} = -17$ dB, the graphs for receivers 1 and 2 are on top of each other.

complex samples per receiver are captured (limited by computer memory), which results in 13.5 million independent samples available to detect the signal. To determine P_D , 10 independent realizations are used to compare to $P_{MD,des} = 0.1$ (although more would be better). This leaves only 1.35 million samples available per power estimate.

7.2.2 Measurement Results

Fig. 7.13 and 7.14 show the measured P_D as a function of K in the bandwidth of interest for several SNRs. As expected, P_D starts at $P_{FA,des}$ and increases with more samples. At 0.9 dB lower SNR, the second receiver theoretically requires $(10^{0.9/10})^2 \approx 1.5$ times more samples, which agrees with the measurements. XC is clearly faster than the individual receivers.

The SNR-wall for AC shows up at $\text{SNR} \approx -18$ dB. P_D of receiver 2 goes down rather than up, which is a clear sign of an overestimation (biased threshold) of the noise level. With a 0.9 dB higher SNR, P_D of receiver 1 still goes to 1. XC is now about two orders of magnitude faster than receiver 1, as predicted in fig. 7.8. At $\text{SNR} = -25$ dB, P_D of both receivers goes to 0, while XC still goes to 1. However, at $\text{SNR} = -32$ dB, the P_D for XC also stays at or below 0.1. More than the available 1.35 million are required to see what really happens, which is a subject for further research.

Fig. 7.15 summarizes the results by showing the required number of samples versus SNR for $P_{FA} = 0.1$ and $P_D = 0.9$. Here, the actual SNR for each receiver is used (e.g., -12.33 dB for receiver 1 and -13.23 dB for receiver 2, whereas both are indicated as $\text{SNR} = -13$ dB in fig. 7.13). To remove false positives, the required number of samples is taken to be the minimum K for which $P_D \geq 0.9$ for all K equal to or larger than this minimum. This means that at $\text{SNR} = -25$ dB, XC needs $6 \cdot 10^5$ samples for $P_D \geq 0.9$ (even though P_D was occasionally measured to be above 0.9 for lower K , see fig. 7.14). The theoretical curves for AC and XC (obtained by numerical evaluation) are also included, with $\rho = 0.028$ for XC. $U_{dB} =$




Figure 7.14: Measured P_D for $P_{FA} = 0.1$ as a function of K in the bandwidth of interest for several SNRs close to or below the SNR-wall.

0.13 dB seems to fit better than $U_{dB} = 0.08$ dB or $U_{dB} = 0.19$ dB, although the differences are small. The results clearly show that XC can detect significantly lower signal powers than AC, as predicted in section 7.1. Based on this $U_{dB} = 0.13$ dB, an SNR-wall of -30.7 dB is found, about 12.8 dB lower than the SNR-wall of AC.

The FCC-requirement of -114 dBm in 6 MHz bandwidth corresponds to -182 dBm/Hz. With the 20 dB receiver NF off UTSFINX in the 10 dB attenuation setting, the SNR-wall of -30.7 dB for XC results in a maximum sensitivity of -184 dBm/Hz. For -182 dBm/Hz, $SNR = -28$ dB, which requires approximately 10^7 samples (see fig. 7.15), equivalent to a measurement time of 1.7 s. In less adverse spectrum conditions, the 10 dB input attenuation may be removed, which, according to section 7.1, reduces the sensing time by a factor 100 to around 17 ms.

Fig. 7.16 shows the measurement time and energy consumption of XC ($U_{dB} = 0.13$ dB, $\rho = 0.028$) and AC ($U_{dB} = 0.07$ dB). For SNRs above -7 dB, AC is more energy-efficient than XC as \tilde{P}_{xc} has a higher variance for low




Figure 7.15: Sample complexity versus SNR for $P_{FA} = 0.1$ and $P_D = 0.9$.




Figure 7.16: Measured performance of XC ($U_{\text{dB}} = 0.13$ dB, $\rho = 0.028$) as compared to AC ($U_{\text{dB}} = 0.07$ dB) in terms of energy consumption and measurement time as a function of SNR.

K (see chapter 3). However, XC is significantly faster, which allows a cognitive radio (CR) to spend more time on actual data communications. Furthermore, for SNRs below -7 dB, XC is not only more than twice as fast as AC, but also more energy-efficient, as shown in fig. 7.16. So XC spectrum sensing can save both on battery life and spectral efficiency.

7.3 Conclusions

Based on a simple model for noise uncertainty, it is theoretically predicted that XC can detect weaker signals than AC. A lower noise correlation is expected to detect signals at smaller SNR. This theoretical result generalizes the result for the detection performance of AC in the presence of noise uncertainty. Furthermore, it predicts that the input attenuation in each receiver has no influence on the detection capabilities of a crosscorrelation spectrum analyzer (XCSA), except that the measurement time will go up.

Experiments performed with UTSFINX conclusively show that XC can detect smaller signals than AC, at least in a system where a lot of uncorrelated noise is present due to 10 dB input attenuation in each receiver. The measured noise uncertainty for AC is 0.08 dB, which theoretically gives an SNR-wall of -17.5 dB, close to the measured -17.9 dB. The noise uncertainty for XC turns out to be somewhat ill-defined in the theoretical model, as the measured peak-to-peak uncertainty in noise level is 0.08 dB for $\text{NMT} = 1$, and 0.19 dB for $\text{NMT} = 10^6$. The experimental SNR-wall for XC is found to be around -30 dB, which corresponds to a peak-to-peak uncertainty of 0.13 dB in the theoretical model. The model therefore needs some further refinement. In any case, *the SNR-wall is reduced by 12 dB using XC*. Furthermore, the measurement results confirm the theoretical prediction that XC is about twice as fast as AC for negative SNR, and even orders of magnitude faster close to the SNR-wall of AC.

The results indicate that (neglecting man-made noise), XC can comply with the sensitivity requirements of the FCC-regulations for spectrum sensing in the TV-bands, except that the measurement time is too long in the 10 dB attenuation setting of UTSFINX. Therefore, further work is required to perform SNR-wall measurements at different attenuation settings.

CONCLUSIONS & RECOMMENDATIONS

In this chapter, the most important conclusions presented in this thesis are summarized, followed by a presentation of the original contributions. Finally, recommendations for future work are discussed.

8.1 Summary & Conclusions

Dynamic Spectrum Access (DSA) is viewed as an important methodology towards more efficient use of the radio spectrum (chapter 1). The underlying principle is that (cognitive) radios search for locally unused spectrum and exploit this for their own communications. To prevent interference with primary users (PUs), signals below the noise floor (i.e., in negative signal-to-noise ratio (SNR) conditions) need to be detected, even in the presence of strong signals in adjacent channels (section 2.1). The main objective of this thesis is *“to implement some form of spectrum sensing that can, without prior knowledge, reliably and swiftly detect weak signals in a hostile radio environment, while being compatible with analog design in a low-voltage CMOS process”*.

Different types of spectrum sensing can be distinguished, but of the ones discussed, energy detection (ED) is the most general solution as it does not require prior knowledge of the signals to be detected (section 2.2). The first step of ED is similar to what a spectrum analyzer (SA) does: measure the power in a frequency band. The second step is to distinguish between measuring only noise, or noise plus a signal. Due to inaccuracies in the noise level estimation, there is a certain minimum SNR, the SNR-wall, below which signals cannot be reliably detected (section 3.1.3). Several other analog impairments, such as phase noise, nonlinearity, and limited

harmonic rejection (HR), can also hamper the detection process by causing false alarms or missed detections.

To lower the SNR-wall, as well as mitigate other analog impairments, crosscorrelation (XC) is proposed (chapter 3). A SA employing XC, referred to as a crosscorrelation spectrum analyzer (XCSA), has two receiver chains, which both process the same input signal. The outputs of the two receivers are multiplied to obtain a power estimate of the input signal. The uncorrelated noise in each receiver averages out, effectively reducing the noise floor. With a decreasing noise floor, small signals that are hidden under the noise will eventually appear. As correlated noise does not average out, it is essential that the correlation between the noise components resulting from the two receiver chains is minimal.

Calculations based on the noise and signal correlation show that, if each receiver is preceded by a matched attenuator, the sensitivity (i.e., the smallest signal power at the input that can be detected given enough measurement time) remains unaffected (section 4.3.2). This allows for trading off the effects of nonlinearity against measurement time: XC allows improvement of linearity without degrading sensitivity (section 3.3). When using separate local oscillators (LOs), their uncorrelated phase noise is also reduced by XC. Thus, XC reduces the effect of phase noise in the output spectrum. XC can also improve HR by decorrelating the harmonic images, while not affecting the signal itself. This is achieved by using a frequency offset between the two receivers, and correcting the offset in the digital domain.

A first prototype is developed with a resistive splitter and two PCBs, each containing a mixer-first frontend, followed by external baseband circuitry and analog-to-digital converters (ADCs) (chapter 4). Measurements confirm the theory set out in earlier chapters: a significant reduction of the noise floor and phase noise by XC is achieved. In addition, the measurements confirm that the input attenuation does not influence the final noise floor after XC. However, the noise floor fluctuates significantly due to coupling between the two PCBs if they are located close to each other. Moreover, the architecture has little isolation between the mixers of each receiver, which introduces phase noise correlation that limits the achievable reduction. The lack of isolation causes the mixers to fail when different LO-frequencies are used.

In order to mitigate these issues and evaluate isolation in an integrated design, a new prototype, UTSFINX, is developed in 65 nm CMOS (chapter 6). It consists of two RF-frontends in a single IC. Rather than a mixer-first architecture, it employs a low-noise transconductance amplifier (LNTA) as a first stage, followed by HR-mixers. Again, baseband components and ADCs are external. Just as in the first prototype, attenuators precede each receiver. They are now integrated on-chip as discrete-step attenuators. A third-order intermodulation (IM3)-cancellation technique is used that allows these attenuators to be wideband and very linear. This technique

was demonstrated separately in a prototype in $0.16\ \mu\text{m}$ CMOS, showing that $+30\ \text{dBm}$ input-referred third-order intermodulation intercept point (IIP3) can be robustly obtained (chapter 5).

The noise level after XC in UTSFINX is found to be almost constant over the whole band of operation (0.3–1.0 GHz). At least 22 dB of phase noise reduction is obtained. At 10 dB input attenuation, UTSFINX achieves an IIP3 of $+25\ \text{dBm}$ and a displayed average noise level (DANL) of better than $-169\ \text{dBm/Hz}$ at a power consumption of around 50 mW. This results in a spurious-free dynamic range (SFDR) based on DANL and IIP3 that outperforms commercial state-of-the-art SAs (section 6.3). When input-matching is not required, the DANL can be reduced even further to below $-172\ \text{dBm/Hz}$ by employing a high input-impedance setting.

Measurements with a frequency offset between the LOs show an improvement up to 25 dB in HR without requiring calibration or additional measurement time (section 6.2.5). The improvement is limited by crosstalk due to sub-optimal layout. Without crosstalk, it is expected that the improvement could be significantly more.

A mathematical model is developed to quantitatively assess the influence of noise correlation in a XCSA on the SNR-wall (chapter 7). It predicts that the SNR-wall scales with the noise correlation. Experimental results show good agreement with the theoretical model: XC is significantly faster than autocorrelation (AC) for the same detection performance, and XC has a lower SNR-wall than AC. In the 10 dB attenuation setting of UTSFINX, the receivers have 20 dB noise figure (NF). The measured SNR-wall in this setting is $-17.9\ \text{dB}$ for AC and $-30.8\ \text{dB}$ for XC. XC gives an improvement of almost 13 dB. Moreover, it shows that XC can actually detect signals 10 dB below the thermal noise floor in this setting.

Overall, XC is shown to enable the integration of SAs with high sensitivity, good resilience to strong interferers, and with both speed and (at low SNR) energy consumption benefits compared to AC. This may enable the integration of (many) small SAs inside other chips for built-in self-test (BIST), reducing on pin count and test time during manufacturing, as well as more reliable and stable performance during operation. For DSA, XC not only makes sensitive spectrum sensing attainable in a hostile radio environment, but also paves the way for low-cost, low-power, and high-quality (mobile) measurement equipment.

8.2 Original Contributions

The work presented in this thesis contains several original contributions:

- The introduction of XC as a means to lower the SNR-wall for ED, including a mathematical foundation and experimental proof; (*chapters 3 and 7 and [MOA:4]*)

- The introduction of attenuation in combination with XC as a means to improve linearity without affecting the detection sensitivity, provided enough measurement time is available;
(*chapters 4 and 6 and [MOA:6, 12]*)
- The experimental proof that XC can improve (phase) noise performance, HR, image rejection (IR), and linearity;
(*chapters 4 and 6 and [MOA:9, 12]*)
- The introduction of a high-impedance mode to further reduce correlated noise;
(*chapter 6 and [MOA:12]*)
- The introduction of an IM3-cancelling mechanism in CMOS discrete-step attenuators to allow more linear and wideband operation;
(*chapter 5 and [MOA:5, 11]*)
- The introduction of simple closed-form expressions for the SFDR of an ideal uniform quantizer with and without Gaussian noise.
(*appendix A and [MOA:1]*)

8.3 Recommendations

This thesis has only just begun to unravel the potential of XC for spectrum sensing or integrated SAs in general. The ultimate goal is a completely integrated XCSA, rather than a partially integrated version with many external components. Most of the work described in this thesis focuses on the implementation and the effects of the radio frequency (RF)-frontend, but the other components need to be investigated as well. For example, one should find a good architecture for a high-SFDR baseband section, and determine how gain could be optimally divided over RF and baseband for optimum SFDR, taking into account XC. The efficient implementation of XC in the digital domain also deserves further attention, including taking into account the effects of quantization on SFDR.

XC may prove useful in other areas as well, where XC can be used to mitigate analog impairments. One example is communication at low SNR, which was recently proposed in [144]. The system architecture proposed in [144] is very similar to the architecture proposed in this thesis for XC spectrum sensing.

A few words are said about the potential of XC on reducing the requirements and improving the performance of ADCs in section 3.3.4: they can be made less power-hungry and more linear. Some preliminary work has been done [145, 146], but the ideas still remain to be verified.

The wideband generation of LO-signals is a big challenge in itself, and perhaps even more when one considers two of the ideas used in this thesis: to have two LOs with uncorrelated phase noise, and, for improved HR, to have two LOs with a small frequency offset. It may be possible to use two

phase-locked loops (PLLs) referred to the same external reference (crystal) to get uncorrelated phase noise. A frequency difference between the two receivers may be obtained by mixing the PLL-output of one of the two receivers with a low-frequency signal using an image-rejection mixer. This low-frequency signal may be obtained by division of the external reference. In any case, spurs from the frequency generation, whether it originates from the PLL, from the mixers, or from crosstalk between the mixers of the two receivers (as in UTSFINX), need to be addressed as well.

The IM3-cancellation technique in discrete-step attenuators is shown to work, but the obtainable linearity is still limited by several factors, including non-linear capacitance of off-state switches. The use of silicon-on-insulator (SOI) CMOS may already reduce the capacitance and its nonlinearity, but it can be worthwhile to explore new circuit techniques to further reduce these effects. Furthermore, in wideband systems, second-order intermodulation (IM2)-components may also significantly impact the performance, but this issue has not yet been addressed. It seems likely that IM2-cancellation can be achieved with a technique similar to the proposed IM3-cancellation technique. It might even be possible to find an optimum where both IM2- and IM3-cancellation is obtained.

The prototype UTSFINX can operate with a single antenna, as used in this thesis, but can also use two separate antennas. This may be exploited in a XCSA, as it can further reduce the noise correlation between the two receivers, and thus further lower the SNR-wall. However, mutual coupling, spatial correlation of black-body radiation, and man-made noise can still cause noise correlation between the two receivers. Moreover, the two antennas will not see exactly the same input signal, due to time delay (although this may be electronically corrected by an inverse delay, and could even provide directivity in the sensing process), differential Doppler shift, and fading, all of which may reduce the (desired) signal correlation. Some preliminary work is given in appendix B, and some further analyses can be found in [147, 148]. However, the combination of using multiple antenna systems and XC spectrum sensing requires further research.

The mathematical model presented in chapter 7 predicts that the sensitivity of a XCSA is independent of its attenuation. Furthermore, it predicts that the high-impedant mode in UTSFINX, which has a significantly lower noise correlation than the matched mode, should have a higher sensitivity. It would be interesting to validate these conjectures, and thereby the theoretical model itself, by actual measurements.

This thesis is limited to the usage of only two receivers in a XCSA. An obvious question to ask is whether more than two receivers can be used. Multiplying more than two receiver outputs together provides a signal measure different from power, which may be useful for higher-order statistics. Restricting to power measurements, with N receivers (and maybe N antennas), there are $\frac{1}{2}N(N - 1)$ pairs that can be crosscorrelated and averaged, which can lower the measurement time for a given SNR, at the

cost of an increase in hardware and computational requirements [149]. Having more than two receivers may also allow additional techniques, such as the three-receiver technique discussed in [111] to estimate and subtract residual noise correlation. This may further lower the SNR-wall by reducing the overall uncertainty.


SPURIOUS-FREE DYNAMIC RANGE OF A UNIFORM QUANTIZER

Quantization is an important step in analog-to-digital converters (ADCs) and digital-to-analog converters (DACs). It limits a signal to a finite set of values, and thus introduces a difference with the original signal, the *quantization error*. Although the quantization error is often modeled as white noise, quantization actually leads to distortion of the signal. For an ADC, the difference in power between the desired signal and the most powerful distortion component (or spur, but that is not considered here) is termed the spurious-free dynamic range (SFDR). To distinguish it from the SFDR of spectrum analyzers (SAs), where the noise floor also plays a role, here it is referred to as $SFDR_{ADC}$. $SFDR_{ADC}$ is relevant, because in general one cannot distinguish between distortion components and true input signals.

Increasing the resolution of the quantizer increases $SFDR_{ADC}$, but at the cost of power consumption and maximum sample rate [117]. Analytical formulas for distortion components are available, but they do not easily translate to $SFDR_{ADC}$. Simple but accurate equations for $SFDR_{ADC}$ would be useful for the design of a crosscorrelation spectrum analyzer (XCSA), but also for the design of systems in general. These equations will be derived in this appendix. The results apply equally well to zero-order hold DACs (which retain their sample value until the following sample).

Sampling is not considered here as it is orthogonal to quantization. It is, however, important to realize that the spectrum of a quantized signal may contain components at any frequency. The components above the Nyquist frequency will alias when the signal is sampled.

Parts of this appendix have appeared in [MOA:1].



(a) Any quantization staircase (left) can be decomposed into a straight line and a repetitive quantization error (right). The black line ($\Delta = 0$) corresponds to a midriser quantizer. (b) Spectrum of a full-scale sinusoid after 8-bit quantization using (A.1) and simulation. Differences are due to numerical inaccuracies.

Figure A.1: Quantization is inherently nonlinear and distorts the spectrum.

A.1 Quantization of a Sinusoid

Multilevel quantization of a sinusoid was investigated by Blachman [150]. In a midriser quantizer (a threshold exactly at 0), only odd-order harmonics are produced due to the odd-symmetric nature of the quantization staircase. To simplify calculations, amplitudes are expressed in least significant bits (LSBs), with $\text{LSB} \triangleq 1$. The resulting output signal then is [150]:

$$A_p = \delta_{p,1}A + \sum_{m=1}^{\infty} \frac{2}{m\pi} J_p(2m\pi A), \quad (\text{A.1})$$

with A_p is the output amplitude of the p -th harmonic, $\delta_{i,j}$ the Kronecker delta function, and A the input amplitude.

Using the quantization staircase $q(x)$ as shown in fig. A.1a, (A.1) can be generalized to uniform quantizers by expressing it as a linear transfer plus the quantization error. This quantization error is periodic with a period of 1 LSB, with Δ representing the offset of the threshold in LSB (for a midriser quantizer, $\Delta = 0$). As a result, $q(x)$ can be written as the sum of x and the Fourier series of the quantization error:

$$q(x) = x + \sum_{m=1}^{\infty} a_m \sin(2\pi m x) - \sum_{m=1}^{\infty} b_m \cos(2\pi m x). \quad (\text{A.2})$$

By straightforward calculation, the coefficients are found to be

$$a_m = \frac{2 \cos^2(\Delta\pi m) - 1}{\pi m} \quad b_m = \frac{\sin(2\pi\Delta m)}{\pi m}. \quad (\text{A.3})$$

Using the same method as Blachman [150] the following result is found (which reduces to (A.1) for $\Delta = 0$):

$$A_p = \begin{cases} (2 - \delta_{p,0}) \sum_{m=1}^{\infty} b_m J_p(2\pi m A), & \text{for } p \text{ even} \\ 2 \sum_{m=1}^{\infty} a_m J_p(2\pi m A) + \delta_{p,1} A, & \text{for } p \text{ odd.} \end{cases} \quad (\text{A.4})$$




Figure A.2: Left: SFDR_{ADC} for a full-scale sinusoid as a function of the number of quantization levels for a midriser quantizer (points) and linear fit (line). Right: Error of linear fit.

With A expressed in LSB, the number of quantization levels n_q directly depends on the amplitude A of the sinusoid. Hence, n_q does not need to be a power of two, which is important because in practical situations sinusoids are not always full-scale.

Fig. A.1b shows the lower part of the spectrum of a full-scale sinusoid quantized with 8 bits ($A = 128$), obtained by numerical evaluation of (A.1) and by simulation (both in MATLAB). As sampling is not considered, a high sampling rate is used to make aliasing effects negligible. For the harmonic p with the highest power, \hat{p} , \hat{p} equals 795, which is close to $2\pi A \approx 804$ as derived in [150]. The approximation of \hat{p} being located roughly at $2\pi A$ is only valid for $A \geq 10$ ($n_q \geq 20$). In other cases, the third harmonic is the strongest, as obtained by numerical evaluation.

Pan & Abidi [116] simulated the effect of midriser-quantization of a sinusoid. They constructed two linear fits for the power of \hat{p} as a function of b , both with a slope of 9 dB/bit, but with different offsets. Although these fits were intuitively explained, it can be seen from their simulation results [116, fig. 3] that the true slope is somewhat less than 9 dB/bit.

SFDR_{ADC} is numerically evaluated (and verified by simulation) for a full-scale sinusoid with $\Delta = 0$ for all even values of n_q (to keep symmetry around zero) up to 13 bits, as shown in fig. A.2. A linear fit (also shown in fig. A.2) of these points results in:

$$\text{SFDR}_{\text{ADC}} \approx 8.07b + 3.29 \text{ [dB]}, \quad (\text{A.5})$$

where $b = \log_2 n_q$. Because n_q is not necessarily a power of 2, b is not necessarily an integer. This approximation has a mean absolute error of 0.25 dB and a standard deviation of 0.31 dB, with a maximum error of 1.56 dB occurring for $n_q = 58$. Numerical evaluation shows that this approximation holds up to at least 25 bits.

There is one ambiguity in the definition of SFDR_{ADC} that needs to be clarified. In the case of only a few quantization levels, the first harmonic (or fundamental component) is lower in amplitude at the output than at the input. For example, for 1-bit quantization of a full-scale sinusoid, the power of the output fundamental is 3.9 dB lower than the input power. One could define SFDR_{ADC} with respect to the input amplitude (which is the




Figure A.3: Difference in dB between SFDR_{ADC} related to the input power and related to the fundamental output power.

usual definition) or with respect to the output amplitude. The difference becomes negligible for 4 bits or more, as shown in fig. A.3. The linear fit described in (A.5) uses SFDR_{ADC} with respect to the output fundamental, which seems to be in accordance with the choice made in [116].

For an arbitrary value of Δ , both even and odd-order harmonics are present, which means the distortion power is distributed over more distortion components. Indeed, numerical evaluation shows that when a_m in (A.4) is set equal to b_m for $m = 1$ (for reasons to be discussed in the derivation), which is the case for $\Delta = \frac{1}{8}$ and $\Delta = \frac{5}{8}$, the SFDR_{ADC} increases by roughly 3 dB as compared to (A.5).

In practice, the amplitude will never be exactly full-scale. Numerical evaluation shows that SFDR_{ADC} changes randomly with a magnitude similar to the error shown in fig. A.2 if the amplitude is somewhere in the range between full-scale and full-scale minus 1 LSB. An example of this change is shown in fig. A.4. This suggests that approximations to the theoretical value may deviate by 1 or 2 dB without compromising practical relevance.

A.1.1 Derivation of 8 dB/bit

The harmonics for $p \ll 2\pi A$ decrease by 3.01 dB/bit, when the LSB remains 1 [150]. If the resolution is increased, but the amplitude stays the same, an extra bit means that the LSB is divided by two, which corresponds to




Figure A.4: Zoom-in on spectra for 8-bit quantization of (close to) full-scale sinusoids. SFDR_{ADC} is 67.51 dB, 68.49 dB, and 66.63 dB respectively. The ‘holes’ in the spectra contain values below the range shown.




Figure A.5: $-\hat{p}^{\frac{2}{3}}\zeta(\pi n_q/\hat{p})$ (black) and $\hat{p}/\pi n_q$ (gray).

another 6.02 dB/bit decrease. The net result is a decrease of 9.03 dB/bit, which is the conclusion drawn in [116].

However, the harmonics around $p \approx 2\pi A$ actually decrease only by about 2 dB/bit (with $\text{LSB}=1$). The exact analytic formula (A.4) can be used to give an indication of this trend (unfortunately, a thorough proof is still lacking). The Bessel-functions can be approximated by Airy-functions in the region where $p \approx 2\pi A$ [151, p366, eq. (9.3.6)]:

$$J_p(pz) \approx \left(\frac{4\zeta(z)}{1-z^2} \right)^{\frac{1}{4}} \frac{\text{Ai}\left(\hat{p}^{\frac{2}{3}}\zeta(z)\right)}{\hat{p}^{\frac{1}{3}}}, \quad (\text{A.6})$$

where

$$\zeta(z) \triangleq -\left(\frac{3}{2}\sqrt{z^2-1} - \frac{3}{2}\arccos\frac{1}{z} \right)^{\frac{2}{3}}. \quad (\text{A.7})$$

With $\hat{p} \approx 2\pi A$, (A.1) in combination with (A.6) yields:

$$A_{\hat{p}} \approx 2 \sum_{m=1}^{\infty} c_m \left(\frac{4\zeta(z)}{1-z^2} \right)^{\frac{1}{4}} \frac{\text{Ai}\left(\hat{p}^{\frac{2}{3}}\zeta(z)\right)}{\hat{p}^{\frac{1}{3}}}, \quad (\text{A.8})$$

where $c_m = a_m$ if \hat{p} is odd, $c_m = b_m$ if \hat{p} is even, and $z = 2\pi m A/\hat{p}$. The cases $m = 1$ and $m > 1$ of (A.8) are considered separately. It will turn out that only the $m = 1$ term in the summation contributes to the trend, while the other terms merely act as 'random' deviations from this trend.

Numerical analysis shows that \hat{p} is always slightly smaller than $2\pi A$, but tends to approach $2\pi A$ for larger A , see fig. A.5. For $m = 1$, numerical evaluation shows $z \in (1; 1.1)$. In this region, the factor $4\zeta(z)/(1-z^2)$ remains virtually constant, as does the parameter $\hat{p}^{\frac{2}{3}}\zeta(2\pi A/\hat{p})$ of the Airy-function (see fig. A.5).

For $m = 1$ and removing (approximate) constants, (A.8) reduces to

$$A_{\hat{p}}(m=1) \propto \frac{1}{\hat{p}^{\frac{1}{3}}}, \quad (\text{A.9})$$

which corresponds to a $\hat{p}^{-\frac{2}{3}}$ dependency in the power spectrum, equivalent to a decrease of 2.01 dB/bit. For $m \geq 2$, the Airy-function $\text{Ai}(\cdot)$ is

approximated using [151, p449, eq. (10.4.83)] (for $x \gg 1$)

$$\text{Ai}(-x) \approx \frac{\sin\left(\frac{2}{3}x^{\frac{3}{2}} + \frac{\pi}{4}\right)}{\sqrt{\pi x^{\frac{1}{4}}}}. \quad (\text{A.10})$$

Substituting (A.10) and $z = 2\pi m A / \hat{p} \approx m$ into (A.6) results in

$$J_p(pz) \approx J_p(pm) \approx \left(\frac{4}{m^2 - 1}\right)^{\frac{1}{4}} \frac{\sin\left(-\frac{2}{3}p\zeta(m)^{\frac{3}{2}} + \frac{\pi}{4}\right)}{\sqrt{\pi p^{\frac{1}{2}}}}. \quad (\text{A.11})$$

There seems to be no relation between n_q and the phase of the sinusoid in (A.11). Hence, the sine-term can be considered as a random variable that takes values between -1 and 1 with an expectation of 0.

In conclusion, only the first term in the summation of (A.8) is important for the overall trend, while the other terms provide more or less random deviations. This randomness explains the erratic behaviour around the trend of SFDR_{ADC} shown in Figure A.2. Combining the 2.01 dB/bit for the $m = 1$ term and the 6.02 dB/bit from halving the amplitude of the quantization error, a trend of 8.03 dB/bit increase in SFDR_{ADC} is expected, which is very close to the 8.07 dB/bit obtained from numerical evaluation.

A.2 Quantization of a Sinusoid with Noise

Adding noise to the input signal decorrelates it with the quantization error [152], and can be regarded as a form of dithering [153]. The effect depends only on the univariate probability density function (pdf) of the noise, and not on its spectrum [150]. The resulting signal frequency response is the product of the quantization error, as given in (A.1), and the Fourier transform of the pdf of the noise, e.g. as in [150, eq. (35)].

As thermal noise has a Gaussian pdf and is often the most important noise contribution, the effect of this noise on SFDR_{ADC} is studied in more detail. Flicker noise, assuming it has a Gaussian pdf, can be incorporated by realizing that frequency components below $1/T$ will constitute an apparent constant offset [154], which can be represented by Δ in (A.4). The exact resulting spectrum is found to be [150, 152, 155]:

$$A_p = \delta_{p,1} A + \sum_{m=1}^{\infty} c_m J_p(2m\pi A) e^{-2\pi^2 \sigma_n^2 m^2}, \quad (\text{A.12})$$

with σ_n the standard deviation of the noise in LSB, and p the p -th harmonic of the input frequency (which is unrelated to any noise frequency).

Using the previously derived result that only $m = 1$ contributes to the trend of SFDR_{ADC} as a function of n_q , the increase in SFDR_{ADC} in dB depends quadratically on the noise level in LSB:

$$\begin{aligned} \text{SFDR}_{\text{ADC}} &\approx 8.07\text{b} + 3.29 + 20 \log_{10} e^{-2\pi^2 \sigma_n^2} \\ &\approx 8.07\text{b} + 3.29 + 171.5 \sigma_n^2 [\text{dB}]. \end{aligned} \quad (\text{A.13})$$




Figure A.6: Difference between approximation (A.13) and theory (A.12).

Fig. A.6 shows the error in dB between the approximated (using (A.13)) and the calculated (using (A.12)) SFDR_{ADC} as a function of the two important parameters n_q (or b) and σ_n . The plot is generated by evaluating (A.12) and (A.13) for all combinations of σ_n (in steps of 0.01 LSB) and even n_q . A negative error means the approximation underestimates SFDR_{ADC} given by theory. Clearly, the approximation is quite close over the whole range of values shown. The difference of less than 2 dB in virtually all cases bears no practical relevance, as was discussed in appendix A.1.

Because the instantaneous amplitude of the noise can assume any value, practical quantizers will sometimes clip, resulting in (additional) distortion of the spectrum. Gaussian noise has an amplitude of less than 3σ for 99.87% of the time. Simulations show that if the amplitude is full-scale minus 3σ , clipping effects have no discernable influence on the spectrum.

A.3 Multitone Quantization

Using the derivation of the 8 dB/bit trend for quantization of a single sinusoid, it turns out to be quite straightforward to generalize it for an arbitrary number of sinusoids. Suppose N sinusoids of frequency f_i and amplitude A_i are present at the input, where i ranges from 1 to N . The output will contain peaks at frequencies $\sum_i p_i f_i$, with $\sum_i p_i > 0$ (note that any individual p_i can be negative), the amplitude of which is denoted by A_{p_1, \dots, p_N} . This amplitude can be derived using the identities [150]:

$$e^{jz \sin(\theta)} = \sum_{p=-\infty}^{\infty} J_p(z) e^{jp\theta} \quad J_{-p}(z) = (-1)^p J_p(z), \quad (\text{A.14})$$

resulting in

$$A_{p_1, \dots, p_N} = \sum_{i=1}^N \left(\delta_{p_i, 1} A_i \prod_{j \neq i} \delta_{p_j, 0} \right) + \sum_{m=1}^{\infty} \left(c_m \prod_{i=1}^N J_{p_i}(2\pi m A_i) \right). \quad (\text{A.15})$$

If some of the input frequencies are commensurate, i.e. $\exists k, l_j \in \mathbb{Z}$ such that $k f_i = \sum_{j \neq i} l_j f_j$, multiple components will occupy the same

frequency. Their respective phase relations are then required to obtain the total amplitude at that specific frequency. This is a much more complicated situation, and the solution is not pursued here.

In the derivation of the trend in $SFDR_{ADC}$ for a single sinusoid, a $p^{-\frac{1}{3}}$ amplitude dependence is found if the LSB is kept equal to 1. Since the summation now contains the product of N of these Bessel-functions, it results in a $p^{-\frac{N}{3}}$ amplitude dependency, or a $6.02+2.01N$ dB/bit increase in $SFDR_{ADC}$. Simulations using two non-commensurate tones confirm a trend of roughly 10 dB/bit. For more than two tones, simulations suffer from numerical problems, which may require further attention.

A.4 Example

Consider a flash ADC with differential-pair-based amplifiers at the input, amplifying the difference between the input signal and different reference voltages generated through a resistor ladder. Assume there are no mismatches and offsets, and that the input capacitance of each differential pair is 30 fF. Assume the thermal noise of the ladder, equal to $\langle v_n^2 \rangle = k_B T_K / C \approx 1.4 \cdot 10^{-7} \text{ V}^2$, dominates. Then $\sigma_n = \sqrt{1.4 \cdot 10^{-7}} \text{ V} \approx 0.38 \text{ mV}$.

The noise becomes important for $SFDR_{ADC}$ when $\sigma_n > 0.1$ LSB, because then (using (A.13)) it contributes 1.7 dB. Considering a realistic input voltage range of 500 mV, this will be the case for 7 bits, giving an $SFDR_{ADC}$ of $8.07 \cdot 7 + 3.29 + (0.38/3.9)^2 \cdot 171.5 = 61.4$ dB. For 5 bits, the noise contribution is negligible, and the $SFDR_{ADC}$ is found to be 43.7 dB, while the signal-to-noise ratio (SNR) will be 31.9 dB. For 9 bits, the noise increases the $SFDR_{ADC}$ by 28.0 dB to a total of 101.9 dB.

A.5 Conclusions

The SFDR of a uniform quantizer, $SFDR_{ADC}$, when quantizing a sinusoid without noise, increases by roughly 8 dB/bit, which is found by numerical evaluation of exact formulas and verified by simulations. This trend and the seemingly random deviations around it are mathematically explained. Adding noise decorrelates the quantization error from the input signal and therefore increases $SFDR_{ADC}$. For Gaussian noise, $SFDR_{ADC}$ (in dB) increases quadratically with the standard deviation of the noise in LSB. A numerical comparison between a simple approximation formula and the exact value obtained from existing but complicated analytical formulas shows an error of less than a few dB, which is often acceptable for practical purposes. For N tones, $SFDR_{ADC}$ is expected to increase by $6 + 2N$ dB/bit, which was derived in the same way as for the single tone situation.

These results can be applied to the design of systems without having to use the exact formulas. This can save a lot of time and effort. The results relate the SNR, the SFDR and the number of quantization levels; knowing two of them allows easy calculation of the third.

CROSSCORRELATION SPECTRUM SENSING WITH TWO ANTENNAS

The noise reduction using crosscorrelation (XC) is limited by the correlated noise between the two receivers. One immediate thought to reduce the noise correlation is to use a separate antenna for each receiver, so that there are no shared components to introduce noise correlation. The use of multiple antennas is getting more and more accepted, as it is used in diversity receivers, for beamforming and MIMO-systems, and has been proposed in many other works on spectrum sensing. Although this seems a straightforward next step, the consequences for the noise and signal powers in the XC spectrum sensing process are nontrivial, and a first exploration of these effects is discussed here.

Fig. B.1a shows an abstraction of a XC-system using one antenna and two receivers, as e.g. discussed in chapters 4 and 6. The noise voltages generated by resistors in the splitter and attenuators will be present at the input of both receivers. The measurement results in chapters 4 and 6 show residual correlated noise of approximately -170 dBm/Hz, which leaves room for improvement. The use of a separate antenna for each receiver, as depicted in fig. B.1b, has the direct advantage of reducing the noise figure (NF) of the system (before XC) because no form of splitter is needed. In the case of the resistive splitter in chapter 4, this reduces NF by 6 dB, in principle reducing measurement time by a factor 16.

It is, however, not trivial what the noise correlation and the signal correlation on the two antennas will be. Therefore, first the origin of noise in a single antenna is considered in appendix B.1, after which the noise correlation between antennas will be discussed in appendix B.2. The signal correlation is discussed in appendix B.3.

Parts of this appendix have appeared in [MOA:3].




Figure B.1: An XCSA may use a single antenna or two antennas.

B.1 Noise in an Antenna

Black-body radiation is the electromagnetic (EM)-radiation emitted by a black body at a given temperature. Real-life objects are not perfect black bodies (as can be observed in the visible range of the spectrum), but the approximation is very good in the radio frequency range [156]. From the Rayleigh-Jeans approximation (accurate for frequencies up to hundreds of GHz at room temperature), the radiation of a black body is (also referred to as brightness) is

$$B_b \approx \frac{2k_B T_K}{\lambda^2} [\text{W} \cdot \text{Hz}^{-1} \cdot \text{m}^{-2} \cdot \text{sr}^{-1}], \quad (\text{B.1})$$

with $\lambda = c/f$ the wavelength, and c the speed of light ($3.00 \cdot 10^8$ m/s).

The purpose of an antenna is to guide the energy of the incident EM field to the load. The maximum power that can be delivered by the antenna to the load is called the available power P_{av} , with S_{av} the available power per Hertz. This maximum power transfer is obtained under conjugate matching, which is assumed from here on. The effective area A_e of an antenna is defined as “the ratio of the available power at the terminals of a receiving antenna to the power flux density of a plane wave incident on the antenna from that direction, the wave being polarization-matched to the antenna” [157], and in general depends on direction, with $A_e(\theta, \phi) = \frac{\lambda^2}{4\pi} G(\phi, \theta)$ [m²], and $G(\phi, \theta)$ the directional antenna gain. Therefore, assuming the antenna is completely surrounded by black bodies in thermal equilibrium

$$S_{av} = \int_0^{2\pi} \int_0^{\pi} B_b A_e(\theta, \phi) \sin \theta d\theta d\phi = \frac{1}{2} \frac{2k_B T_K}{\lambda^2} \frac{\lambda^2}{4\pi} 4\pi = k_B T_K [\text{W}/\text{Hz}], \quad (\text{B.2})$$

where the factor $\frac{1}{2}$ in the second step is due to the unpolarized nature of the black-body radiation and the polarized reception of the antenna. The noise floor, taken as $k_B T_K$, with $T_K = 290^\circ\text{C}$, is then -174 dBm/Hz. So, even when the antenna itself is noiseless, which is assumed from here on, it will deliver noise power to the attached receiver.

B.2 Noise Correlation with Two Antennas

Many papers propose multiple antennas for spectrum sensing, and assume that independent noise samples at each antenna are available. However, several factors introduce correlated noise between two antennas.

The EM-field impinging on an antenna induces a time-varying current in it, which by itself generates an EM-field that may be received by the other antenna. This is known as mutual coupling, which introduces correlation of the external noise at the different antennas. It also introduces correlation of the noise generated by the receiver, as some noise generated in one receiver will be transmitted by its antenna, and then received by the other [88]. Hence, mutual coupling introduces correlated noise.

Assuming the antenna is surrounded by black bodies, it can be shown that, even when there is no antenna coupling, the thermal noise of two antennas can be correlated due to spatial correlation. The crosscorrelation function between the voltages induced on the antenna terminals by these noise sources depends on the distance between the antennas, the antenna patterns, and the polarization the antennas are sensitive to [158].

A receiver tries to receive information from one ‘wanted’ radiator, but there are many other radiators in that band, collectively combined into ‘man-made noise’ (see section 2.2.5). Man-made noise, just like the black-body radiation, originates from many sources and arrives from many directions, so it will also have spatial correlation.

B.3 Signal Correlation with Two Antennas

When both the noise and the signal correlation are reduced by the same factor, the signal-to-noise ratio (SNR) is not changed, so the sensitivity of a XCSA will not change. Therefore, it is equally important to consider the signal correlation in the two-antenna XCSA. To achieve insight into this effect, a few different scenarios are considered in an idealized situation where the antennas are identically polarized and have no mutual coupling.

B.3.1 Line-of-sight Connection

In line-of-sight-connections, there will be a time delay between the signal arriving on the first and the second antenna, depending on the orientation of the antennas with respect to the signal source. This time delay is unknown, because the angle of arrival of the signal is unknown. Mathematically, the output of the crosscorrelator is (neglecting noise and thereby removing expectation operators) $Y(f) = P_s(f)e^{-j2\pi f\tau_0}$, where τ_0 is the time delay. Hence, using $\tilde{P}_{xc} = |Y|$ removes the phase shift and the true signal power is detected.

However, the minimum bandwidth that can be estimated is inversely proportional to the measurement time, so one has to consider detection in a finite bandwidth f_b , as done in section 4.1. Each frequency within f_b




Figure B.2: Simple model of a moving transmitter and stationary two-antenna XCSA.

experiences a different phase shift between the two antennas. For simplicity, assume the signal is a white noise source in the band of interest; then

$$\tilde{P}_{xc} = \left| \int_{f_0 - \frac{f_b}{2}}^{f_0 + \frac{f_b}{2}} Y(f) df \right| = P_s \left| \left(\frac{\sin \pi f_b \tau_0}{\pi f_b \tau_0} \right) \right|. \quad (\text{B.3})$$

Clearly, for $f_b = 0$ and/or $\tau_0 = 0$, $\tilde{P}_{xc} = P_s$. For larger τ_0 and/or larger f_b , \tilde{P}_{xc} becomes lower. If maximally 1 dB loss of signal power is allowed, $f_b \cdot \tau_0 \lesssim 0.36$.

For TV white space (TVWS), with channels of 6 MHz, $\tau_0 \leq 60$ ns. In a handheld device, the maximum distance between the two antennas may be about 15 cm, so the maximum time delay is $\frac{d}{c} \approx 0.5$ ns, which gives negligible loss in signal power. Even when the antennas would be 10 m apart, the detected signal power will only be lowered by 0.3 dB. So, in static line-of-sight-connections, using two antennas for XC spectrum sensing makes sense.

B.3.2 Movement in LOS-connections

In a mobile context, the transmitter and/or the spectrum sensing device can move during the sensing process, which makes some parameters dynamic. Consider the situation depicted in fig. B.2, where the transmitter moves during the sensing process, such that at time instant t , it is located at p_t , while the two receive antennas, located at p_x and p_y , are stationary. The two receive antennas are separated by a distance d_a , and the transmitter moves at a uniform velocity v . Assume the signal power received by each antenna does not change, but the relative phase between the two antennas does (which can be considered a manifestation of the Doppler-effect). The largest change in phase is obtained if the transmitter moves as indicated: parallel to the baseline of the antennas, passing point p_0 (at distance $d_{tx,rx}$ from the baseline) halfway during the sensing process, which takes T seconds. The distance d_{p_1,p_2} between p_1 and p_2 covered by the transmitter during the sensing process is thus $d_{p_1,p_2} = vT$.

With $d_{p_t,p_0}(t) \triangleq \left(\frac{T}{2} - t\right) \cdot v$ and $\Delta d(t) \triangleq d_{p_t,p_x}(t) - d_{p_t,p_y}(t)$, where $d_{p_t,p_x}(t) \geq 0$ and $d_{p_t,p_y}(t) \geq 0$, the phase difference between the two

receivers becomes $\Delta\phi(t) = 2\pi f \frac{\Delta d(t)}{c}$. Considering a single frequency, Y after measuring for T seconds becomes

$$Y(f, T) = \frac{1}{T} \int_0^T P_s(f) e^{j\Delta\phi(t)} dt = P_s(f) \cdot \frac{1}{T} \int_0^T e^{j\Delta\phi(t)} dt. \quad (\text{B.4})$$

This integral is difficult to evaluate analytically for given $\Delta\phi(t)$, but it can be simplified for certain situations.

For small vT a Taylor-series approximation can be used for d_{p_t, p_x} and d_{p_t, p_y} : $\Delta d(t) \approx d_a d_{p_t, p_0}(t)/r$, with $r = \sqrt{(d_a/2)^2 + d_{t_x, r_x}^2}$. Then

$$Y(f) \approx \frac{1}{T} P_s(f) \int_{-T/2}^{T/2} e^{-j2\pi f \frac{d_a v t}{c r}} dt = P_s(f) \frac{\sin\left(\frac{\pi v T f d_a}{rc}\right)}{\left(\frac{\pi v T f d_a}{rc}\right)}. \quad (\text{B.5})$$

To maximally lose 1 dB of estimated signal power, $vT f d_a \lesssim 0.36rc$.

For very large vT , the integral is mainly determined by the transmitter being at the far left or the far right of the receiver, so $\Delta d(t) \approx \pm d_a$. Then

$$Y(f) \approx P_s(f) \cos\left(2\pi f \frac{d_a}{c}\right). \quad (\text{B.6})$$

To maximally lose 1 dB of estimated signal power, $d_a \lesssim 0.1 \frac{c}{f} = 0.1\lambda$.

The integral of (B.4) can be numerically evaluated for some practical cases. Consider the detection of a wireless microphone, operating at a center frequency of 500 MHz, at a distance $d_{t_x, r_x} = 20$ m. The two antennas of the XCSA are 10 cm apart ($d_a \approx 0.17\lambda$). With a velocity of 1.6 m/s (a person carrying the microphone), the maximum measurement time T_{\max} to lose not more than 1 dB of signal power is 43 s, which is much longer than required by IEEE 802.22. Next, consider the detection of a signal in the ISM-band at 2.4 GHz with a laptop that has two antennas at a distance of 30 cm ($d_a = 2.4\lambda$). The signal is transmitted from a car moving at 30 m/s with $d_{t_x, r_x} = 50$ m. Now $T_{\max} \approx 190$ ms (note that using the approximate (B.5), T_{\max} is found to be 250 ms). These practical cases show that only for large receive antenna distances and fast-moving transmitters and/or excessive measurement time, the signal can be decorrelated significantly. Under these line-of-sight-assumptions, it is clear that a two-antenna XCSA is applicable in many situations.

B.3.3 Fading

Not all connections are line-of-sight (or have a large specular component), but must rely on a multitude of multi-path reflections. Diversity is often used to improve the performance of a communications system in these fading channels, and is based on the idea that with multiple antennas, the probability that at least one of them is not in a deep fade is increased.

The simplest fading channel is a Rayleigh fading channel, with a Rayleigh distribution for the instantaneous amplitude $A = \sqrt{r_{\text{re}}^2 + r_{\text{im}}^2}$ and a uniform distribution for the phase. Both r_{re} and r_{im} are normally distributed with zero mean and variance $\sigma_s^2/2$. The output power for autocorrelation (AC) (or an XCSA with one antenna) will then be $\sigma_s^2 = P_s$.

The XC system is sensitive to phase variations during the measurement, as was established in appendix B.3.2. Therefore, a good channel model is required that describes the phase correlation over time and location, which may require further research. Here, for simplicity, a channel is assumed in which the amplitudes of the signal undergo (correlated) fading, but the phase difference of the signal between the two antennas remains constant during the measurement. Using \tilde{P}_{xc} removes any constant phase difference over frequency between the two antennas of the receiver, so a phase difference of 0 can be assumed.

With received signal amplitudes A_1 and A_2 at antenna 1 and 2, respectively, the result is $\tilde{P}_{\text{xc}} = \frac{1}{T} \left| \int_{\langle T \rangle} A_1(t) A_2(t) dt \right|$. In a fading environment, $A_1(t)$ and $A_2(t)$ are random processes, with $A_i \geq 0$. The average signal power received by both antennas is assumed equal ($\mathbb{E}[A_1^2] = \mathbb{E}[A_2^2]$, a well-accepted model [159]). With fully correlated fading ($A_1 = A_2$), $\mathbb{E}[Y] = \sigma_n^2 = P_s$, which is equal to the single-antenna case. For independent fading, $\mathbb{E}[A_1 A_2] = \mathbb{E}[A_1] \mathbb{E}[A_2] = \sqrt{\frac{\pi}{4}} \sigma_s \sqrt{\frac{\pi}{4}} \sigma_s = \frac{\pi}{4} \sigma_s^2$. These are two extreme cases, so it is expected that for arbitrary correlated fading the result will be somewhere in between: $\frac{\pi}{4} P_s \leq \mathbb{E}[Y] \leq P_s$. Maximally $10 \log_{10} \frac{4}{\pi} \approx 1$ dB of signal power is lost, under the optimistic assumption that the relative phase does not change during the measurements.

B.4 Conclusions

A single-antenna XCSA can significantly lower the noise level for energy detection (ED), but is limited by correlated noise. Using a separate antenna for each receiver may lower NF of each receiver, and reduce the noise correlation, thereby reducing measurement time and increasing sensitivity. Different mechanisms are identified that still cause residual noise correlation, which will ultimately limit the reduction of the noise level: mutual coupling, spatial correlation, and man-made noise.

Unfortunately, the use of two antennas may also significantly degrade the signal power that is detected. This is caused by decorrelation of the signal, due to the physical distance between the antennas, the presence of relative signal phase shift between the antennas due to Doppler-shift, and the presence of fading. The calculations indicate that in static and line-of-sight-connections the loss of signal power may be acceptable, making two-antenna XC spectrum sensing an interesting solution to more reliable ED that deserves further exploration.

RECEIVER TERMINOLOGY

This appendix is a brief introduction to receiver design and terminology.

The digital data transmitted over the air is modulated, i.e., (groups of) bits are transformed into an analog waveform, with its properties (amplitude, phase, frequency, etc.) denoting the transmitted bits. This waveform is transmitted at a certain power, often denoted in dBm, which is a dB-scale with 1 mW as a reference. Typical transmitted power levels of mobile devices are in the range of 0–20 dBm, or 1 mW to 100 mW. As the transmitter generally transmits in many or all directions, only a small fraction of the transmitted power is received by the receiver. The transmitted waveform is susceptible to noise, which is added in the wireless link (the channel) as well as in the receiver. This noise causes errors in the detection process, because it changes the waveform. The channel may also introduce frequency- and time-dependent changes to the waveform, which is known as fading, but that is not considered in this thesis.

The noise figure (NF) quantifies the noise performance of a receiver as a deterioration of signal-to-noise ratio (SNR) from input to output. The SNR is the ratio of the useful signal power divided by the noise power in the signal band. The higher the SNR, the less bit errors are introduced by noise. In case of wireless communications, depending on the type of modulation, the required SNR ranges from 8 dB to 25 dB (note that this is different for spectrum sensing, as discussed in chapter 2). This means that the NF of the receiver should be low enough. Receivers typically have a NF of 2 dB to 10 dB. The received signal is often so weak that it would be completely obscured by the receiver noise if it is not first amplified. To obtain a low NF, almost all receivers start with a low-noise amplifier (LNA), which amplifies the input signal without adding much noise, such that the noise added by successive stages becomes relatively less important.

Ideally, such an amplifier can be described mathematically as $y(t) = a \cdot x(t) + n(t)$, where $x(t)$ is the input signal (which can be a voltage,

current, or charge), a is the amplification, $n(t)$ is the noise added, and $y(t)$ is the output signal. Unfortunately (neglecting $n(t)$), amplifiers are better described by $y(t) \approx a_1 \cdot x(t) + a_2 \cdot x^2(t) + a_3 \cdot x^3(t)$, i.e., as a nonlinear device approximated using a Taylor-series. The values $|a_1/a_2|$ and $|a_1/a_3|$ are a measure for the linearity of the system. The explanation here will be brief; a more elaborate discussion can be found in [127].

Suppose two sine waves of equal power P at frequencies f_1 and f_2 are applied to the input of such a system. Using Fourier theory, it can be shown that a_3 will introduce frequency components at $2f_1 - f_2$ and $2f_2 - f_1$. These components increase in amplitude with the third power of the input amplitude. The input-referred third-order intermodulation intercept point (IIP3) is the (extrapolated) input power for which these undesired intermodulation components have the same magnitude as the desired components, and is usually expressed in dBm. Extrapolation is often necessary, because for such high input powers the Taylor-series approximation with only a few terms is usually not accurate anymore. Typical receiver IIP3 is in the range of -30 dBm to $+5$ dBm. The input-referred second-order intermodulation intercept point (IIP2) is defined in a similar way, where the intermodulation components at $f_2 - f_1$ and $f_1 - f_2$ (caused by a_2) are considered. Values for IIP2 typically range from 20 dBm to 70 dBm. The dynamic range (DR) is defined as the ratio, usually expressed in dB, of the maximum to the minimum signal input power levels over which a device can operate; the minimum level is usually determined by the noise, and the maximum level by nonlinearity.

The range of frequencies over which a receiver is able to receive information is its bandwidth, sometimes called radio frequency (RF)-bandwidth. It is usually defined as the band over which the receiver is matched to the antenna (e.g. the power received by the antenna is (almost) completely transferred to the receiver) and performance of the receiver is not degraded too much (e.g. NF increased by 2 dB or gain decreased by 3 dB). Matching is generally considered good if the voltage standing wave ratio is smaller than 2 , which means constructive interference due to reflections will never increase the voltage amplitude at any point by more than a factor 2 . An equivalent representation can be obtained using S-parameters: less than -9.6 dB of the incoming power is reflected, or $S_{11} < -9.6$ dB. The bandwidth that the analog-to-digital converter (ADC) eventually has to convert to the digital domain is often referred to as intermediate frequency (IF)-bandwidth, and in general is much smaller than the RF-bandwidth, because the signal to be demodulated only occupies a fraction of the RF-bandwidth.

LIST OF ACRONYMS

AC	autocorrelation	DWG	digital window generator
ADC	analog-to-digital converter	ED	energy detection
AM	amplitude modulation	EM	electromagnetic
AWGN	additive white Gaussian noise	ETSI	European Telecommunications Standards Institute
BIST	built-in self-test	FCC	Federal Communications Commission
CG	common gate	FFT	fast Fourier transform
CMFB	common mode feedback	FM	frequency modulation
CMOS	Complementary Metal-Oxide-Semiconductor	FoM	figure of merit
COST	European Cooperation in Science and Technology	FXC	FX-correlator
CP	1-dB compression point	GSM	Global System for Mobile Communications
CP_r	cyclic prefix	HDTV	High-Definition TV
CR	cognitive radio	HR	harmonic rejection
CS	common source	IC	integrated circuit
DAC	digital-to-analog converter	IEEE	Institute for Electrical and Electronics Engineers
DANL	displayed average noise level	IF	intermediate frequency
DECT	Digital Enhanced Cordless Telecommunications	IIP_x	input-referred x-th order intermodulation intercept point
DFT	discrete Fourier transform	IL	insertion loss
DR	dynamic range	IM_x	x-th order intermodulation
DSA	Dynamic Spectrum Access	IR	image rejection
DSP	digital signal processing	ISM	Industrial Scientific Medical
DTV	digital TV	ITU	International Telecommunication Union
		LNA	low-noise amplifier
		LNTA	low-noise transconductance amplifier

LO	local oscillator	ROC	receiver operating characteristic
LSB	least significant bit	RSSI	Received Signal Strength Indication
LTE	Long Term Evolution	SA	spectrum analyzer
MAC	multiply-accumulate	SFDR	spurious-free dynamic range
NF	noise figure	SNR	signal-to-noise ratio
NFE	Noise Floor Extension	SOI	silicon-on-insulator
NMT	normalized measurement time	SSC	Shared Spectrum Company
NRAO	National Radio Astronomy Observatory	SU	secondary user
NTIA	National Telecommunications and Information Administration	TIA	transimpedance amplifier
Ofcom	Office of Communications	TVWS	TV white space
OFDM	Orthogonal Frequency Division Multiplexing	UHF	ultra high frequency
pdf	probability density function	UMTS	Universal Mobile Telecommunications System
PCB	printed circuit board	UWB	Ultra Wideband
PLL	phase-locked loop	VCO	voltage-controlled oscillator
PM	phase modulation	VHF	very high frequency
PSD	power spectral density	VSB	vestigial sideband
PU	primary user	WiMAX	Worldwide Interoperability for Microwave Access
PVT	Process/Voltage/Temperature	WLAN	Wireless Local Area Network
QAM	quadrature amplitude modulation	XC	crosscorrelation
QPSK	quadrature phase shift keying	XCSA	crosscorrelation spectrum analyzer
RBW	resolution bandwidth	XFC	XF-correlator
RF	radio frequency		

LIST OF SYMBOLS

All symbols are explained at their first occurrence in the main text. As a reference, this table contains an explanation and/or definition of the symbols used at multiple locations throughout this thesis.

E.1 Mathematical Definitions

\triangleq	Defined as
\sim	Distributed as
γ_A	Autocorrelation function of stochastic process A; $\gamma_A \triangleq \gamma_{AA}$
γ_{AB}	Crosscorrelation function of A and B, with $\gamma_{AB}(t, \tau) \triangleq \mathbb{E} [A(t)\overline{B(t + \tau)}]$
$\Gamma(\alpha, \beta)$	Upper incomplete gamma function $\Gamma(\alpha, \beta) \triangleq \int_{\beta}^{\infty} t^{\alpha-1} e^{-t} dt$
$\Gamma(\alpha)$	Gamma function, with $\Gamma(\alpha) \triangleq \Gamma(\alpha, 0)$. For positive integers $\Gamma(\alpha) = (\alpha - 1)!$
$\Gamma_{AB}(f)$	Cross-spectrum of stochastic processes A and B; $\Gamma_{AB}(f) = \mathcal{F}[\gamma_{AB}(\tau)]$
μ	The mean of a stochastic variable (usually denoted by the subscript)
σ	The standard deviation of a stochastic variable (usually denoted by the subscript)
$\text{Ai}(\cdot)$	Airy-function
$\mathcal{CN}(0, \sigma^2)$	Circularly symmetric complex normal distribution, with the variance in the real and imaginary parts equal to $\sigma^2/2$
$\Re(\cdot)$	Operator to take the real part
$\Im(\cdot)$	Operator to take the imaginary part
$\mathbb{E}[\cdot]$	Expected value
$\mathcal{F}[\cdot]$	Fourier-transform, with $\mathcal{F}^{-1}[\cdot]$ its inverse; $\mathcal{F}[x(t)] \triangleq \int_{-\infty}^{\infty} x(t) \exp(j2\pi ft) dt$
$\mathbb{P}(x y)$	Probability of x given y
$\mathcal{O}(\cdot)$	If $f(x) = \mathcal{O}(g(x))$, then $\limsup_{x \rightarrow \infty} f(x)/g(x) < \infty$ (Landau-notation)
j	Imaginary unit; $j \triangleq \sqrt{-1}$
$I_n(\cdot)$	The n-th order modified Bessel function of the first kind
$J_n(\cdot)$	The n-th order Bessel function of the first kind
$\mathcal{K}_n(\cdot)$	The n-th order modified Bessel function of the second kind
$\text{pdf}_X(z)$	Probability density function of stochastic variable X
$Q(x)$	Probability of a standard-normal variable being larger than x, with $Q^{-1}(x)$ its inverse: $Q(x) \triangleq \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-z^2/2} dz$
$\text{var}[\cdot]$	Variance
\bar{X}	Complex conjugate of X
$ X $	Absolute value of X
$X Y$	Parallel application of impedances X and Y: $X Y = XY/(X + Y)$
X_{re}	Real part of X
X_{im}	Imaginary part of X
\hat{X}	Estimator of stochastic variable X
\mathbb{Z}	The set of all integers ($\dots, -2, -1, 0, 1, 2, \dots$)

E.2 Variables

β_K	Interpolation function defined in (3.31) on page 45
γ_{NEF}	Noise excess factor
Δf_{LO}	Frequency difference between LO-frequencies of the receivers: $\Delta f_{LO} \triangleq f_1 - f_2$
$\Delta\phi$	Phase difference between the signal components at the two receiver outputs
ϵ_1	Noise power underestimation error ($0 \leq \epsilon_1 < 1$)
ϵ_2	Noise power overestimation error ($\epsilon_2 \geq 0$)
λ	Threshold (general)
λ_{ed}	Threshold for energy detection
λ_{xc}	Threshold for crosscorrelation (XC)
$\lambda_{ed,biased}$	Biased threshold for energy detection
μ_0	Shorthand for μ of an estimator given that \mathcal{H}_0 applies
μ_1	Shorthand for μ of an estimator given that \mathcal{H}_1 applies
ρ	The noise correlation between two receivers
σ_0	Shorthand for σ of an estimator given that \mathcal{H}_0 applies
σ_1	Shorthand for σ of an estimator given that \mathcal{H}_1 applies
A	Attenuation (transfer) on linear scale ($0 \leq A \leq 1$)
A_{dB}	Attenuation in dB ($A_{dB} \triangleq -20 \log_{10} A$)
b	Number of bits in an ADC
c_h	Coefficient of the h -th harmonic of the LO; $c_1 \triangleq 1$
C_{mixer}	Capacitance at each output-phase of the passive mixer
D	LO duty cycle ($0 < D < 1$, or $0\% < D < 100\%$)
$DANL_{corr}$	DANL after XC-process (all uncorrelated noise removed): $DANL_{corr} \triangleq -174 + NF_{corr}$ dBm/Hz
$DANL_{eff}$	DANL with XC during the XC-process: $DANL_{eff} \triangleq -174 + NF_{eff}$ dBm/Hz
f	Frequency (Hz)
F	Noise factor: $NF \triangleq 10 \log_{10} F$
f_{3dB}	Cut-off frequency (amplitude transfer $1/\sqrt{2}$ compared to maximum transfer)
f_{LO}	Frequency of the LO
f_1	Frequency of the LO in the first receiver
f_2	Frequency of the LO in the second receiver
f_s	Sample rate
FoM_{SFDR}	Performance indicator for SFDR, with $FoM_{SFDR} \triangleq IIP3$ [dBm] – NF [dB]
g_m	Transconductance
h	General channel realization
\mathcal{H}_0	The null hypothesis (signal to be detected is not present)
\mathcal{H}_1	The alternative hypothesis (signal to be detected is present)
HR_h	HR of LO-harmonic h ; $HR_1 \triangleq 0$ dB
$IIP3_{ib}$	In-band IIP3 (both large tones and intermodulation products inside IF-bandwidth)
$IIP3_{ob}$	Out-of-band IIP3 (large tones outside, IM-product(s) inside IF-bandwidth)
K	Number of samples (per receiver) in a measurement
k_B	Boltzmann constant ($1.3806503 \cdot 10^{-23}$ J/K)
M	Number of points per FFT
n	General noise realization
n_0	That part of the noise that is fully correlated between the two receivers
n_1	That part of the noise that present in the first receiver, but not in the second receiver
n_2	That part of the noise that present in the second receiver, but not in the first receiver
NF_α	Man-made noise level in dB as compared to $k_B T_K$, with $T_K = 290$ °C
NF_{corr}	Effective NF of XC during the XC-process: $NF_{eff} \triangleq DANL_{eff} + 174$ dB
NF_{eff}	Effective NF after XC (all uncorrelated noise removed): $NF_{corr} \triangleq DANL_{corr} + 174$ dB
NMT	Normalized measurement time: time required to acquire samples for one FFT
P_h	The power of the signal present at the h -th harmonic of the LO
\hat{P}_{ac}	Spectrum estimator which uses AC (single receiver)
\hat{P}_{xc}	Spectrum estimator which uses the real part of the measured cross-spectrum
\tilde{P}_{xc}	Spectrum estimator which uses the absolute value of the measured cross-spectrum
P_D	Probability of detection
$P_{D,des}$	Desired probability of detection, i.e., $P_D \geq P_{D,des}$ is desired
P_{FA}	Probability of false alarm
$P_{FA,des}$	Desired probability of false alarm, i.e., $P_{FA} \leq P_{FA,des}$ is desired
P_{MD}	Probability of missed detection ($P_{MD} \triangleq 1 - P_D$)
$P_{MD,des}$	Desired probability of missed detection, i.e., $P_{MD} \leq P_{MD,des}$ is desired
Q_D^{-1}	Shorthand notation for $Q^{-1}(P_D)$ (similarly, $Q_{D,des}^{-1} \triangleq Q^{-1}(P_{D,des})$)
Q_{FA}^{-1}	Shorthand notation for $Q^{-1}(P_{FA})$ (similarly, $Q_{FA,des}^{-1} \triangleq Q^{-1}(P_{FA,des})$)
$r[k]$	The k -th sample of the receiver output
$r_1[k]$	The k -th sample of the first receiver output

$r_2[k]$	The k -th sample of the second receiver output
r_{in}	The total of signal and noise entering a receiver
R_{load}	Load resistance
R_{left}	Resistance seen from the passive mixer towards the source
r_{on}	Attenuator-switch on-resistance
R_s	Source impedance (always assumed 50Ω)
R_{switch}	Mixer-switch on-resistance
s	General signal realization
S_{11}	Input port voltage reflection coefficient
$SFDR_{ADC}$	Difference between fundamental and strongest distortion component in ADC
SNR_{min}	Minimum SNR (required or detectable)
t	Time (seconds)
T	Measurement time (seconds)
T_K	Absolute temperature in Kelvin
U	Peak-to-peak noise power estimation uncertainty; $U \triangleq (1 + \epsilon_2)/(1 - \epsilon_1)$
U_{dB}	Peak-to-peak noise power estimation uncertainty in dB: $U_{dB} \triangleq 10 \log_{10} U$
V_{DD}	Supply voltage
V_{gs}	Gate-source voltage (subscripts 'b' and 'd' for bulk and drain terminals): $V_{gs} = V_{GS} + v_{gs}$
V_{GS}	DC-bias gate-source voltage
v_{gs}	Small-signal gate-source voltage
V_{SS}	Ground voltage
W	Bandwidth (Hz)
W	Transistor width
Y	Decision metric
Z_{in}	Input impedance
Z_{mixer}	Input impedance of passive mixer

BIBLIOGRAPHY

- [1] A. Capri, *Quips, Quotes and Quanta: An Anecdotal History of Physics*. World Scientific Publishing Company, Inc., 2011.
- [2] E. Noam, "The economists' contribution to radio spectrum access: The past, the present, and the future," *Proc. IEEE*, vol. 100, no. Special Centennial Issue, pp. 1692–1697, May 2012.
- [3] Cisco, "Visual networking index: Global mobile data traffic forecast update, 2011–2016," Cisco, Tech. Rep., Feb. 2012.
- [4] M. Reardon. (2010, Jan.) FCC chairman pushes policy agenda. [Online]. Available: http://www.cnet.com/8301-31045_1-10430991-269.html
- [5] M. Lazarus, "The great spectrum famine," *IEEE Spectr.*, vol. 47, no. 10, pp. 26–31, Oct. 2010.
- [6] J. Reed *et al.*, "Spectrum access technologies: The past, the present, and the future," *Proc. IEEE*, vol. 100, no. Special Centennial Issue, pp. 1676–1684, May 2012.
- [7] B. Obama. (2010, Jun.) Presidential memorandum: Unleashing the wireless broadband revolution. [Online]. Available: <http://www.whitehouse.gov/the-press-office/presidential-memorandum-unleashing-wireless-broadband-revolution>
- [8] Radio spectrum policy programme. http://ec.europa.eu/information_society/policy/ecomm/radio_spectrum/rspp/index_en.htm.
- [9] R. B. Bacchus *et al.*, "Long-term, wide-band spectral monitoring in support of dynamic spectrum access networks at the IIT spectrum observatory," in *Proc. 3rd IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Chicago, IL, USA, 14–17 Oct. 2008.
- [10] R. Tandra *et al.*, "What is a spectrum hole and what does it take to recognize one?" *Proc. IEEE*, vol. 97, no. 5, pp. 824–848, May 2009.
- [11] FCC, "In the matter of unlicensed operation in the TV broadcast bands and additional spectrum for unlicensed devices below 900 MHz and in the 3 GHz band," FCC, Tech. Rep., Nov. 2008.
- [12] I. F. Akyildiz *et al.*, "Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey," *Comput. Netw.*, vol. 50, no. 13, pp. 2127–2159, Sep. 2006.
- [13] IEEE, "IEEE standard definitions and concepts for dynamic spectrum access: Terminology relating to emerging wireless networks, system functionality, and spectrum management (IEEE Std 1900.1-2008)," Sep. 2008.
- [14] ITU, "Definitions of software defined radio (SDR) and cognitive radio system (CRS) (SM.2152)," ITU, Tech. Rep., Sep. 2009.
- [15] T. Yücek *et al.*, "A survey of spectrum sensing algorithms for cognitive radio applications," *Commun. Surveys Tuts.*, vol. 11, no. 1, pp. 116–130, Jan./Mar. 2009.
- [16] I. Mitola, J. *et al.*, "Cognitive radio: making software radios more personal," *IEEE Pers. Commun.*, vol. 6, no. 4, pp. 13–18, Aug. 1999.
- [17] A. M. Wyglinski *et al.*, *Cognitive Radio Communications and Networks: Principles and Practice*, A. M. Wyglinski *et al.*, Eds. Elsevier Inc., 2010.
- [18] M. Fitch *et al.*, "Wireless service provision in TV white space with cognitive radio technology: A telecom operator's perspective and experience," *IEEE Commun. Mag.*, vol. 49, no. 3, pp. 64–73, Mar. 2011.
- [19] V. Gungor *et al.*, "Cognitive radio networks for smart grid applications: A promising technology to overcome spectrum inefficiency," *IEEE Veh. Technol. Mag.*, vol. 7, no. 2, pp. 41–46, Jun. 2012.
- [20] Y. Zhang *et al.*, "Cognitive machine-to-machine communications: visions and potentials for the smart grid," *IEEE Netw.*, vol. 26, no. 3, pp. 6–13, May/June 2012.
- [21] M. J. Bentum *et al.*, "The coexistence of cognitive radio and radio astronomy," in *16th Annual Symposium of the IEEE/CVT*, Louvain-la-Neuve, Belgium, 19 Nov. 2009.
- [22] S. Haykin, "Cognitive radio: brain-empowered wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 2, pp. 201–220, Feb. 2005.
- [23] FCC, "In the matter of fostering innovation and investment in the wireless communications market and a national broadband plan for our future," FCC, Tech. Rep., Nov. 2009.

- [24] P. Marshall, "Dynamic spectrum access as a mechanism for transition to interference tolerant systems," in *Proc. 4th IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Singapore, 6–9 Apr. 2010.
- [25] F. Granelli *et al.*, "Standardization and research in cognitive and dynamic spectrum access networks: IEEE SCC41 efforts and other activities," *IEEE Commun. Mag.*, vol. 48, no. 1, pp. 71–79, Jan. 2010.
- [26] S. Filin *et al.*, "International standardization of cognitive radio systems," *IEEE Commun. Mag.*, vol. 49, no. 3, pp. 82–89, Mar. 2011.
- [27] K. Moessner *et al.*, "Spectrum sensing for cognitive radio systems: technical aspects and standardization activities of the IEEE P1900.6 working group," *IEEE Wireless Commun. Mag.*, vol. 18, no. 1, pp. 30–37, Feb. 2011.
- [28] M. Murrioni *et al.*, "IEEE 1900.6: spectrum sensing interfaces and data structures for dynamic spectrum access and other advanced radio communication systems standard: technical aspects and future outlook," *IEEE Commun. Mag.*, vol. 49, no. 12, pp. 118–127, Dec. 2011.
- [29] M. Marcus, "Spectrum policy for radio spectrum access," *Proc. IEEE*, vol. 100, no. Special Centennial Issue, pp. 1685–1691, May 2012.
- [30] S. K. Jones *et al.*, "Evaluation of the performance of prototype TV-band white space devices phase II," Office of Engineering and Technology, Federal Communications Commission, Tech. Rep., Oct. 2008.
- [31] FCC, "In the matter of unlicensed operation in the TV broadcast bands and additional spectrum for unlicensed devices below 900 MHz and in the 3 GHz band," FCC, Tech. Rep., Sep. 2010.
- [32] Office of Communications, "Statement on licence-exempting cognitive devices using interleaved spectrum," Office of Communications, Statement, Jul. 2009.
- [33] C. Stevenson *et al.*, "IEEE 802.22: The first cognitive radio wireless regional area network standard," *IEEE Commun. Mag.*, vol. 47, no. 1, pp. 130–138, Jan. 2009.
- [34] E. Axell *et al.*, "Spectrum sensing for cognitive radio: State-of-the-art and recent advances," *IEEE Signal Process. Mag.*, vol. 29, no. 3, pp. 101–116, May 2012.
- [35] D. B. Čabrić, "Cognitive radios: System design perspective," Ph.D. dissertation, University of California, Berkeley, 2007.
- [36] S. Shellhammer *et al.* (2006, Jul.) Performance of the power detector with noise uncertainty. PowerPoint Presentation. [Online]. Available: http://www.ieee802.org/22/Meeting_documents/2006_July/22-06-0134-00-0000_Performance-of-the-power-detector-with-Noise-Uncertainty.ppt
- [37] A. Wagstaff *et al.*, "Man-made noise measurement programme," Mass Consultants Ltd, Tech. Rep., Sep. 2003.
- [38] ITU, "Man-made NF measurements," ITU, Tech. Rep., 2009.
- [39] ITU, "Man-made noise measurements in the HF range (SM.2155)," ITU, Tech. Rep., Sep. 2009.
- [40] T. Erpek *et al.*, "DSA operational parameters with wireless microphones," in *Proc. 4th IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Singapore, 6–9 Apr. 2010.
- [41] Microsoft Research. (2012, Nov.) WhiteFiService. Tower data of Mar. 26, 2010. [Online]. Available: <http://whitespaces.msresearch.us>
- [42] Shared Spectrum Company, "Spectrum occupancy measurements: Chicago, Illinois, November 16–18, 2005," Shared Spectrum Company, Tech. Rep., 2005.
- [43] Shared Spectrum Company, "Spectrum occupancy measurements location 4 of 6: Republican National Convention, New York City, New York, August 30, 2004 – September 3, 2004, revision 2," Shared Spectrum Company, Tech. Rep., 2005.
- [44] Shared Spectrum Company, "Spectrum occupancy measurements location 5 of 6: National Radio Astronomy Observatory (NRAO), Green Bank, West Virginia, October 10 – 11, 2004, revision 3," Shared Spectrum Company, Tech. Rep., 2005.
- [45] V. J. Arkesteijn *et al.*, "Jitter requirements of the sampling clock in software radio receivers," *IEEE Trans. Circuits Syst. II*, vol. 53, no. 2, pp. 90–94, Feb. 2006.
- [46] A. Tasic *et al.*, "Design of multistandard adaptive voltage-controlled oscillators," *IEEE Trans. Microw. Theory Tech.*, vol. 53, no. 2, pp. 556–563, Feb. 2005.
- [47] P. Ruiippo *et al.*, "An UMTS and GSM low phase noise inductively tuned LC VCO," *IEEE Microw. Wireless Compon. Lett.*, vol. 20, no. 3, pp. 163–165, Mar. 2010.
- [48] B. Murmann. (2011, Mar.) ADC performance survey 1997–2011. [Online]. Available: <http://www.stanford.edu/~murmman/adcsurvey.html>
- [49] B. Murmann, "A/D converter trends: Power dissipation, scaling and digitally assisted architectures," in *Proc. Custom Integrated Circuits Conf. (CICC)*, San Jose, CA, USA, 21–24 Sep. 2008, pp. 105–112.
- [50] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge University Press, 2004.
- [51] C. Rauscher, *Fundamentals of Spectrum Analysis*, 1st ed. Rohde & Schwarz, 2001.
- [52] X. Gao *et al.*, "Jitter analysis and a benchmarking figure-of-merit for phase-locked loops," *IEEE Trans. Circuits Syst. II*, vol. 56, no. 2, pp. 117–121, Feb. 2009.

- [53] M. C. M. Soer *et al.*, "A 0.2-to-2.0GHz 65nm CMOS receiver without LNA achieving >11dBm IIP3 and <6.5 dB NF," in *Proc. IEEE Int. Solid-State Circuits Conf. - Dig. Tech. Papers (ISSCC)*, San Francisco, CA, USA, 8–12 Feb. 2009, pp. 222–223, 223a.
- [54] Z. Ru *et al.*, "Digitally enhanced software-defined radio receiver robust to out-of-band interference," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3359–3375, Dec. 2009.
- [55] J. Greenberg *et al.*, "A 40MHz-to-1GHz fully integrated multistandard silicon tuner in 80nm CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf. - Dig. Tech. Papers (ISSCC)*, San Francisco, CA, USA, 19–23 Feb. 2012, pp. 162–164.
- [56] D. Murphy *et al.*, "A blocker-tolerant wideband noise-cancelling receiver with a 2dB noise figure," in *Proc. IEEE Int. Solid-State Circuits Conf. - Dig. Tech. Papers (ISSCC)*, San Francisco, CA, USA, 19–23 Feb. 2012, pp. 74–76.
- [57] S. Youssef *et al.*, "Active feedback technique for RF channel selection in front-end receivers," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 3130–3144, Dec. 2012.
- [58] C. Andrews *et al.*, "A passive mixer-first receiver with digitally controlled and widely tunable RF interface," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2696–2708, Dec. 2010.
- [59] J. A. Weldon *et al.*, "A 1.75-GHz highly integrated narrow-band CMOS transmitter with harmonic-rejection mixers," *IEEE J. Solid-State Circuits*, vol. 36, no. 12, pp. 2003–2015, Dec. 2001.
- [60] H.-K. Cha *et al.*, "A CMOS wideband RF front-end with mismatch calibrated harmonic rejection mixer for terrestrial digital TV tuner applications," *IEEE Trans. Microw. Theory Tech.*, vol. 58, no. 8, pp. 2143–2151, Aug. 2010.
- [61] H. Zhang *et al.*, "A harmonic-rejection mixer with improved design algorithm for broadband TV tuners," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Montréal, Canada, 17–19 Jun. 2012, pp. 163–166.
- [62] J. Ryu *et al.*, "Double quadrature harmonic rejection architecture insensitive to gain and phase mismatch for analog/digital TV tuner IC," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Baltimore, MD, USA, 5–7 Jun. 2011, pp. 1–4.
- [63] B. Razavi, "Cognitive radio design challenges and techniques," *IEEE J. Solid-State Circuits*, vol. 45, no. 8, pp. 1542–1553, Aug. 2010.
- [64] S. Park *et al.*, "Spectrum broadening due to phase noise interaction in cognitive radio systems," *IEEE Commun. Lett.*, vol. 14, no. 10, pp. 891–893, Oct. 2010.
- [65] X. Gao *et al.*, "Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector," *IEEE J. Solid-State Circuits*, vol. 45, no. 9, pp. 1809–1821, Sep. 2010.
- [66] M. Kitsunezuka *et al.*, "A 30-MHz-2.4-GHz CMOS receiver with integrated RF filter and dynamic-range-scalable energy detector for cognitive radio systems," *IEEE J. Solid-State Circuits*, vol. 47, no. 5, pp. 1084–1093, May 2012.
- [67] J. Park *et al.*, "A fully integrated UHF-band CMOS receiver with Multi-Resolution Spectrum Sensing (MRSS) functionality for IEEE 802.22 cognitive radio applications," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 258–268, Jan. 2009.
- [68] W. Namgoong, "A channelized digital ultrawideband receiver," *IEEE Trans. Wireless Commun.*, vol. 2, no. 3, pp. 502–510, May 2003.
- [69] B. Sadhu *et al.*, "A 5GS/s 12.2pJ/conv. analog charge-domain FFT for a software defined radio receiver front-end in 65nm CMOS," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Montréal, Canada, 17–19 Jun. 2012, pp. 39–42.
- [70] F.-K. Wang *et al.*, "A CMOS spectrum sensor using injection locking of two voltage-controlled oscillators for cognitive radio system," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Baltimore, MD, USA, 5–7 Jun. 2011, pp. 1–4.
- [71] S. Mandal *et al.*, "A bio-inspired active radio-frequency silicon cochlea," *IEEE J. Solid-State Circuits*, vol. 44, no. 6, pp. 1814–1828, Jun. 2009.
- [72] K. A. Jenkins *et al.*, "Integrated CMOS spectrum analyzer for on-chip diagnostics using digital autocorrelation of coarsely quantized signals," US Patent 7 218 091, May 15, 2007.
- [73] F. K. Bowers *et al.*, "Quantization noise of correlation spectrometers," *Astron. Astrophys. Suppl. A*, vol. 15, pp. 373–380, Jun. 1974.
- [74] D. T. Lin *et al.*, "A 600MHz to 3.4GHz flexible spectrum-sensing receiver with spectrum-adaptive reconfigurable DT filtering," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Montréal, Canada, 17–19 Jun. 2012, pp. 269–272.
- [75] P. F. Marshall, "Dynamic spectrum management of front end linearity and dynamic range," in *Proc. 3rd IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Chicago, IL, USA, 14–17 Oct. 2008.
- [76] D. H. Mahrof *et al.*, "On the effect of spectral location of interferers on linearity requirements for wideband cognitive radio receivers," in *Proc. 4th IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Singapore, 6–9 Apr. 2010.

- [77] H. Urkowitz, "Energy detection of unknown deterministic signals," *Proc. IEEE*, vol. 55, no. 4, pp. 523–531, Apr. 1967.
- [78] F. F. Digham *et al.*, "On the energy detection of unknown signals over fading channels," *IEEE Trans. Commun.*, vol. 55, no. 1, pp. 21–24, Jan. 2007.
- [79] S. Atapattu *et al.*, "Analysis of area under the ROC curve of energy detection," *IEEE Trans. Wireless Commun.*, vol. 9, no. 3, pp. 1216–1225, Mar. 2010.
- [80] A. Sonnenschein *et al.*, "Radiometric detection of spread-spectrum signals in noise of uncertain power," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 28, no. 3, pp. 654–660, Jul. 1992.
- [81] R. Tandra *et al.*, "SNR walls for signal detection," *IEEE J. Sel. Topics Signal Process.*, vol. 2, no. 1, pp. 4–17, Feb. 2008.
- [82] J. E. Salt *et al.*, "Performance prediction for energy detection of unknown signals," *IEEE Trans. Veh. Technol.*, vol. 57, no. 6, pp. 3900–3904, Nov. 2008.
- [83] ITU, "Radio noise measurements (SM.2055)," ITU, Tech. Rep., 2006.
- [84] R. Tandra *et al.*, "Fundamental limits on detection in low SNR under noise uncertainty," in *Proc. Int. Wireless Netw., Comm. and Mobile Comp. Conf.*, vol. 1, Maui, HI, USA, 13–16 Jun. 2005, pp. 464–469.
- [85] P. De *et al.*, "Blind spectrum sensing algorithms for cognitive radio networks," *IEEE Trans. Veh. Technol.*, vol. 57, no. 5, pp. 2834–2842, Sep. 2008.
- [86] M. Naraghi-Pour *et al.*, "Autocorrelation-based spectrum sensing for cognitive radios," *IEEE Trans. Veh. Technol.*, vol. 59, no. 2, pp. 718–733, Feb. 2010.
- [87] Y. Zeng *et al.*, "Eigenvalue-based spectrum sensing algorithms for cognitive radio," *IEEE Trans. Commun.*, vol. 57, no. 6, pp. 1784–1793, Jun. 2009.
- [88] C. P. Domizioli *et al.*, "Noise correlation in compact diversity receivers," *IEEE Trans. Commun.*, vol. 58, no. 5, pp. 1426–1436, May 2010.
- [89] P. Wang *et al.*, "Multiantenna-assisted spectrum sensing for cognitive radio," *IEEE Trans. Veh. Technol.*, vol. 59, no. 4, pp. 1791–1800, May 2010.
- [90] K. Milne, "Theoretical performance of a complex cross-correlator with Gaussian signals," *IEE Proc. Radar Signal Process.*, vol. 140, no. 1, pp. 81–88, Feb. 1993.
- [91] J. Briaire *et al.*, "Uncertainty in Gaussian noise generalized for cross-correlation spectra," *J. Appl. Phys.*, vol. 84, no. 8, pp. 4370–4374, Oct. 1998.
- [92] C. Ciofi *et al.*, "A new method for high-sensitivity noise measurements," *IEEE Trans. Instrum. Meas.*, vol. 51, no. 4, pp. 656–659, Aug. 2002.
- [93] M. Heskamp *et al.*, "Sub-noise primary user detection by cross-correlation," in *Proc. Int. Conf. on Communications (ICC)*, Dresden, Germany, 14–18 Jun. 2009.
- [94] M. Sampietro *et al.*, "Spectrum analyzer with noise reduction by cross-correlation technique on two channels," *Rev. Sci. Instrum.*, vol. 70, no. 5, pp. 2520–2525, May 1999.
- [95] B. P. Kibble, "Measurements for the future—a complete SI from physical constants," *IEEE Trans. Instrum. Meas.*, vol. 61, no. 8, pp. 2125–2131, Aug. 2012.
- [96] A. Thompson *et al.*, *Interferometry and Synthesis in Radio Astronomy*, 4th ed. Krieger Publishing Company, 1998.
- [97] M. J. Sousa, "Limitations of crosscorrelation radiometer performance," in *Proc. IEEE Military Communications Conference (MILCOM)*, 1992, pp. 816–821.
- [98] N. A. Moseley *et al.*, "A spectrum sensing technique for cognitive radios in the presence of harmonic images," in *Proc. 3rd IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Chicago, IL, USA, 14–17 Oct. 2008.
- [99] J. Green, P., "The output signal-to-noise ratio of correlation detectors," *IRE Trans. Inf. Theory*, vol. 3, no. 1, pp. 10–18, 1957.
- [100] G. Roe *et al.*, "Probability density functions for correlators with noisy reference signals," *IRE Trans. Inf. Theory*, vol. 7, no. 1, pp. 13–18, Jan. 1961.
- [101] D. Cooper, "The probability density function for the output of a correlator with band-pass input waveforms," *IEEE Trans. Inf. Theory*, vol. 11, no. 2, pp. 190–195, Apr. 1965.
- [102] G. Jenkins *et al.*, *Spectral Analysis and its Applications*. Holden-Day, 1968.
- [103] S. Niranjanayan *et al.*, "On the integrated cross-noise term in correlation detectors," *IEEE Trans. Commun.*, vol. 57, no. 11, pp. 3244–3248, Nov. 2009.
- [104] A. M. Yaglom, *Correlation Theory of Stationary and Related Random Functions*, ser. Springer Series in Statistics. Springer-Verlag, 1987, vol. I: Basic Results.
- [105] B. Farhang-Boroujeny, "Filter bank spectrum sensing for cognitive radios," *IEEE Trans. Signal Process.*, vol. 56, no. 5, pp. 1801–1811, May 2008.
- [106] J. D. Bunton, "SKA correlator advances," *Exp. Astron.*, vol. 17, no. 1–3, pp. 251–259, Jun. 2004.
- [107] K. C. Rovers, "Front-end research for a low-cost spectrum analyser," Master's thesis, University of Twente, Jun. 2006.

- [108] M. A. T. Sanduleanu *et al.*, *Power Trade-Offs and Low-Power in Analog CMOS ICs*, ser. The Springer International Series in Engineering and Computer Science. Springer Netherlands, 2003, vol. 662, ch. Power Considerations in Sub-Micron Digital CMOS, pp. 9–29.
- [109] Y. Chen *et al.*, “A 2.4-Gsample/s DVFS FFT processor for MIMO OFDM communication systems,” *IEEE J. Solid-State Circuits*, vol. 43, no. 5, pp. 1260–1273, 2008.
- [110] T.-H. Yu *et al.*, “A 7.4-mW 200-MS/s wideband spectrum sensing digital baseband processor for cognitive radios,” *IEEE J. Solid-State Circuits*, vol. 47, no. 9, pp. 2235–2245, Sep. 2012.
- [111] F. Crupi *et al.*, “Enhanced sensitivity cross-correlation method for voltage noise measurements,” *IEEE Trans. Instrum. Meas.*, vol. 55, no. 4, pp. 1143–1147, Aug. 2006.
- [112] E. Rubiola *et al.*, “Very high frequency and microwave interferometric phase and amplitude noise measurements,” *Rev. Sci. Instrum.*, vol. 70, no. 1, Part 1, pp. 220–225, Jan. 1999.
- [113] E. Rubiola *et al.*, “Improved interferometric method to measure near-carrier AM and PM noise,” *IEEE Trans. Instrum. Meas.*, vol. 48, no. 2, pp. 642–646, Apr. 1999.
- [114] E. Salik *et al.*, “Dual photonic-delay line cross correlation method for phase noise measurement,” in *Proc. IEEE Int. Frequency Control Symp. and Exposition*, 23–27 Aug. 2004, pp. 303–306.
- [115] E. Rubiola *et al.*, “The effect of AM noise on correlation phase-noise measurements,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 5, pp. 926–932, May 2007.
- [116] H. Pan *et al.*, “Spectral spurs due to quantization in Nyquist ADCs,” *IEEE Trans. Circuits Syst. I*, vol. 51, no. 8, pp. 1422–1439, Aug. 2004.
- [117] B. Le *et al.*, “Analog-to-digital converters: A review of the past, present and future,” *IEEE Signal Process. Mag.*, pp. 69–77, Nov. 2005.
- [118] C. Andrews *et al.*, “A passive-mixer-first receiver with baseband-controlled RF impedance matching, <6dB NF, and >27dBm wideband IIP3,” in *Proc. IEEE Int. Solid-State Circuits Conf. - Dig. Tech. Papers (ISSCC)*, San Francisco, CA, USA, 7–11 Feb. 2010, pp. 46–47.
- [119] B. W. Cook *et al.*, “Low-power 2.4-GHz transceiver with passive RX front-end and 400-mV supply,” *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2757–2766, Dec. 2006.
- [120] M. C. M. Soer *et al.*, “Unified frequency-domain analysis of switched-series-passive mixers and samplers,” *IEEE Trans. Circuits Syst. I*, vol. 57, no. 10, pp. 2618–2631, Oct. 2010.
- [121] C. Andrews *et al.*, “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, Dec. 2010.
- [122] E. Rubiola *et al.*, “Correlation-based noise measurements below the thermal noise floor,” in *Proc. Joint Meeting of the European Frequency and Time Forum and the IEEE Int. Frequency Control Symp.*, vol. 2, Besançon, France, 13–16 Apr. 1999, pp. 1125–1128.
- [123] A. A. Abidi *et al.*, “De-embedding the noise figure of differential amplifiers,” *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 882–885, Jun. 1999.
- [124] M. Kitsunezuka *et al.*, “A 30MHz–2.4GHz CMOS receiver with integrated RF filter and dynamic-range-scalable energy detector for cognitive radio,” in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Baltimore, MD, USA, 5–7 Jun. 2011, pp. 1–4.
- [125] S. Pollin *et al.*, “Digital and analog solution for low-power multi-band sensing,” in *Proc. 4th IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Singapore, 6–9 Apr. 2010.
- [126] M. Ingels *et al.*, “A 5mm² 40nm LP CMOS 0.1-to-3GHz multistandard transceiver,” in *Proc. IEEE Int. Solid-State Circuits Conf. - Dig. Tech. Papers (ISSCC)*, San Francisco, CA, USA, 7–11 Feb. 2010, pp. 458–459.
- [127] W. Sansen, “Distortion in elementary transistor circuits,” *IEEE Trans. Circuits Syst. II*, vol. 46, no. 3, pp. 315–325, Mar. 1999.
- [128] F. Gatta *et al.*, “An embedded 65 nm CMOS baseband IQ 48 MHz–1 GHz dual tuner for DOCSIS 3.0,” *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3511–3525, Dec. 2009.
- [129] J. Xiao *et al.*, “A high dynamic range CMOS variable gain amplifier for mobile DTV tuner,” *IEEE J. Solid-State Circuits*, vol. 42, no. 2, pp. 292–301, Feb. 2007.
- [130] S. Lerstaveesin *et al.*, “A 48–860 MHz CMOS low-IF direct-conversion DTV tuner,” *IEEE J. Solid-State Circuits*, vol. 43, no. 9, pp. 2013–2024, Sep. 2008.
- [131] Y. Kim *et al.*, “A multi-band multi-mode CMOS direct-conversion DVB-H tuner,” in *Proc. IEEE Int. Solid-State Circuits Conf. - Dig. Tech. Papers (ISSCC)*, San Francisco, CA, USA, 5–9 Feb. 2006, pp. 2504–2513.
- [132] M. Koutani *et al.*, “A digitally controlled variable-gain low-noise amplifier with strong immunity to interferers,” *IEEE J. Solid-State Circuits*, vol. 42, no. 11, pp. 2395–2403, Nov. 2007.
- [133] H. Dogan *et al.*, “Analysis and design of RF CMOS attenuators,” *IEEE J. Solid-State Circuits*, vol. 43, no. 10, pp. 2269–2283, Oct. 2008.
- [134] A. Youssef *et al.*, “Digitally-controlled RF passive attenuator in 65 nm CMOS for mobile TV tuner ICs,” in *Proc. IEEE Int. Circuits and Syst. Symp. (ISCAS)*, Paris, France, 30 May – 2 Jun. 2010, pp. 1999–2002.

- [135] H. Dogan *et al.*, "Intermodulation distortion in CMOS attenuators and switches," *IEEE J. Solid-State Circuits*, vol. 42, no. 3, pp. 529–539, Mar. 2007.
- [136] Y.-Y. Huang *et al.*, "Highly linear RF CMOS variable attenuators with adaptive body biasing," *IEEE J. Solid-State Circuits*, vol. 46, no. 5, pp. 1023–1033, May 2011.
- [137] M. Granger-Jones *et al.*, "A broadband high dynamic range voltage controlled attenuator MMIC with IIP3 > +47dBm over entire 30dB analog control range," in *Proc. IEEE MTT-S Int. Microwave Symp. Digest (MTT)*, Baltimore, MD, USA, 5–10 Jun. 2011, pp. 1–4.
- [138] A. J. Scholten *et al.*, "The new CMC standard compact MOS model PSP: Advantages for RF applications," *IEEE J. Solid-State Circuits*, vol. 44, no. 5, pp. 1415–1424, May 2009.
- [139] X. Li *et al.*, "Benchmark tests for MOSFET compact models with application to the PSP model," *IEEE Trans. Electron Devices*, vol. 56, no. 2, pp. 243–251, Feb. 2009.
- [140] P. Bendix *et al.*, "RF distortion analysis with compact MOSFET models," in *Proc. Custom Integrated Circuits Conf. (CICC)*, Orlando, FL, USA, 3–6 Oct. 2004, pp. 9–12.
- [141] C. C. McAndrew, "Validation of MOSFET model source–drain symmetry," *IEEE Trans. Electron Devices*, vol. 53, no. 9, pp. 2202–2206, Sep. 2006.
- [142] B.-H. Ku *et al.*, "6-bit CMOS digital attenuators with low phase variations for X-band phased-array systems," *IEEE Trans. Microw. Theory Tech.*, vol. 58, no. 7, pp. 1651–1663, Jul. 2010.
- [143] F. Bruccoleri *et al.*, "Wide-band CMOS low-noise amplifier exploiting thermal noise canceling," *IEEE J. Solid-State Circuits*, vol. 39, no. 2, pp. 275–282, Feb. 2004.
- [144] A. B. J. Kokkeler *et al.*, "A correlating receiver for OFDM at low SNR," in *IEEE Vehicular Technology Conference (VTC)-Spring*, Budapest, Hungary, 15–18 May 2011, pp. 1–5.
- [145] A. J. v. Heusden, "Automatic gain control ADC based on signal statistics for a cognitive radio cross-correlation spectrum analyzer," Master's thesis, University of Twente, Enschede, The Netherlands, Aug. 2011.
- [146] P. L. Bicker, "Analysis and design of an ADC for a spectrum analyzer," Master's thesis, University of Twente, Enschede, The Netherlands, Sep. 2011.
- [147] A. R. Smeenge, "Improving cross-correlation spectrum sensing using two antennas," Master's thesis, University of Twente, Enschede, The Netherlands, Oct. 2010.
- [148] P. J. Prins, "On the effect of antenna coupling on spectrum sensing using a cross-correlation spectrum analyser with two antennas," Master's thesis, University of Twente, Enschede, The Netherlands, Jan. 2012.
- [149] G. Scandurra *et al.*, "Multi-channel cross-correlation for increasing sensitivity in voltage noise measurements," in *IEEE Int. Instrumentation and Measurement Technology Conf. (I2MTC)*, Graz, Austria, 13–16 May 2012, pp. 1524–1528.
- [150] N. M. Blachman, "The intermodulation and distortion due to quantization of sinusoids," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 6, pp. 1417–1426, Dec. 1985.
- [151] M. Abramowitz *et al.*, *Handbook of Mathematical Functions With Formulas, Graphs and Mathematical Tables*. Dover Publications, Jun. 1964.
- [152] M. F. Wagdy, "Effect of various dither forms on quantization errors of ideal A/D converters," *IEEE Trans. Instrum. Meas.*, vol. 38, no. 4, pp. 850–855, Aug. 1989.
- [153] M. F. Wagdy *et al.*, "Validity of uniform quantization error model for sinusoidal signals without and with dither," *IEEE Trans. Instrum. Meas.*, vol. 38, no. 3, pp. 718–722, Jun. 1989.
- [154] M. S. Keshner, "1/f noise," *Proc. IEEE*, vol. 70, no. 3, pp. 212–218, Mar. 1982.
- [155] A. B. J. Kokkeler *et al.*, "Modeling correlation of quantized noise and periodic signals," *IEEE Signal Process. Lett.*, vol. 11, no. 10, pp. 802–805, Oct. 2004.
- [156] J. H. Shirley *et al.*, Eds., *Encyclopedia of Planetary Sciences*. Springer, 1997.
- [157] C. A. Balanis, *Antenna Theory*, 3rd ed. Wiley, 2005.
- [158] D. Childers *et al.*, "Space-time cross-correlation functions for antenna array elements in a noise field," *IEEE Trans. Inf. Theory*, vol. 11, no. 2, pp. 182–190, Apr. 1965.
- [159] L. Fang *et al.*, "New method of performance analysis for diversity reception with correlated Rayleigh-fading signals," *IEEE Trans. Veh. Technol.*, vol. 49, no. 5, pp. 1807–1812, 2000.

LIST OF PUBLICATIONS

Peer-reviewed

- [MOA:1] **M. S. Oude Alink**, A. B. J. Kokkeler, E. Klumperink, K. C. Rovers, G. Smit, and B. Nauta, "Spurious-free dynamic range of a uniform quantizer," *IEEE Trans. Circuits Syst. II*, vol. 56, no. 6, pp. 434–438, Jun. 2009.
- [MOA:2] **M. S. Oude Alink**, E. A. M. Klumperink, M. C. M. Soer, A. B. J. Kokkeler, and B. Nauta, "A 50MHz-to-1.5GHz cross-correlation CMOS spectrum analyzer for cognitive radio with 89dB SFDR in 1MHz RBW," in *Proc. 4th IEEE Symp. on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Singapore, 6–9 Apr. 2010.
- [MOA:3] **M. S. Oude Alink**, A. R. Smeenge, A. B. J. Kokkeler, E. A. M. Klumperink, G. J. M. Smit, and B. Nauta, "Exploring the use of two antennas for crosscorrelation spectrum sensing," in *IEEE Vehicular Technology Conference (VTC)-Fall*, San Francisco, CA, USA, 5–8 Sep. 2011.
- [MOA:4] **M. S. Oude Alink**, A. B. J. Kokkeler, E. A. M. Klumperink, G. J. M. Smit, and B. Nauta, "Lowering the SNR wall for energy detection using cross-correlation," *IEEE Trans. Veh. Technol.*, vol. 60, no. 8, pp. 3748–3757, Oct. 2011.
- [MOA:5] W. Cheng, **M. S. Oude Alink**, A. J. Annema, G. J. M. Wienk, and B. Nauta, "A wideband IM3 cancellation technique for CMOS attenuators," in *Proc. IEEE Int. Solid-State Circuits Conf. - Dig. Tech. Papers (ISSCC)*, San Francisco, CA, USA, 19–23 Feb. 2012, pp. 78–79.
- [MOA:6] **M. S. Oude Alink**, E. A. M. Klumperink, A. B. J. Kokkeler, M. C. M. Soer, G. J. M. Smit, and B. Nauta, "A CMOS-compatible spectrum analyzer for cognitive radio exploiting crosscorrelation to improve linearity and noise performance," *IEEE Trans. Circuits Syst. I*, vol. 59, no. 3, pp. 479–492, Mar. 2012.
- [MOA:7] **M. S. Oude Alink**, E. A. M. Klumperink, A. B. J. Kokkeler, W. Cheng, Z. Ru, A. Ghaffari, G. J. M. Wienk, and B. Nauta, "A CMOS spectrum analyzer frontend for cognitive radio achieving +25dBm IIP3 and -169 dBm/Hz DANL," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, Montréal, Canada, 17–19 Jun. 2012, pp. 35–38.
- [MOA:8] **M. S. Oude Alink**, E. A. M. Klumperink, A. B. J. Kokkeler, G. J. M. Smit, and B. Nauta, *Cognitive Communications: Distributed Artificial Intelligence (DAI), Regulatory Policy and Economics, Implementation*. Wiley, Aug. 2012, ch. 15: CMOS RF Transceiver Considerations for Dynamic Spectrum Access, pp. 417–464.
- [MOA:9] **M. S. Oude Alink**, A. B. J. Kokkeler, E. A. M. Klumperink, Z. Ru, W. Cheng, and B. Nauta, "Improving harmonic rejection for spectrum sensing using crosscorrelation," in *Proc. European Solid-State Circuits Conf. (ESSCIRC)*, Bordeaux, France, 17–21 Sep. 2012, pp. 361–364.
- [MOA:10] W. Cheng, **M. S. Oude Alink**, A. J. Annema, J. A. Croon, and B. Nauta, "RF circuit linearity optimization using a general weak nonlinearity model," *IEEE Trans. Circuits Syst. I*, vol. 59, no. 10, pp. 2340–2353, Oct. 2012.
- [MOA:11] W. Cheng, **M. S. Oude Alink**, A. J. Annema, G. J. M. Wienk, and B. Nauta, "A wideband IM3 cancellation technique for CMOS Π - and T-attenuators," *IEEE J. Solid-State Circuits*, vol. 48, no. 2, pp. 358–368, Feb. 2013.
- [MOA:12] **M. S. Oude Alink**, E. A. M. Klumperink, A. B. J. Kokkeler, Z. Ru, W. Cheng, and B. Nauta, "Using crosscorrelation to mitigate analog/RF impairments for integrated spectrum analyzers," *IEEE Trans. Microw. Theory Tech. (invited paper)*, vol. 61, no. 3, pp. 1327–1337, Mar. 2013.
- [MOA:13] D. H. Mahrof, E. A. M. Klumperink, **M. S. Oude 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)*, Seattle, WA, USA, 2–4 Jun. 2013, pp. 601–604.
- [MOA:14] **M. S. Oude Alink**, A. B. J. Kokkeler, E. A. M. Klumperink, G. J. M. Smit, and B. Nauta, "Sensitive spectrum sensing with high SFDR using crosscorrelation energy detection," *IEEE J. Emerg. Sel. Topics Circuits Syst. (submitted for publication)*, vol. 3, no. 4, Dec. 2013.

- [MOA:15] N. Haider, **M. S. Oude Alink**, D. Caratelli, E. A. M. Klumperink, and A. G. Yarovoy, "Frequency-tunable antenna by varying the input impedance of the frontend CMOS chip," in *European Microwave Week (EuMW)*, Nuremberg, Germany, 6–11 Oct. 2013 (accepted for publication).
- [MOA:16] C. W. Korevaar, **M. S. Oude Alink**, P.-T. de Boer, A. B. J. Kokkeler, and G. J. M. Smit, "Correlations involving Hermite functions," *IEEE Trans. Signal Process.* (submitted for publication).

Other

- [MOA:17] **M. S. Oude Alink**, "Increasing the spurious-free dynamic range of an integrated spectrum analyzer," Master's thesis, University of Twente, Enschede, The Netherlands, Nov. 2008.
- [MOA:18] **M. S. Oude Alink**, D. H. Mahrof, A. Ghaffari, S. Subhan, E. A. M. Klumperink, A. B. J. Kokkeler, G. J. M. Smit, and B. Nauta, "AD-REM radio project overview," in *15th Annual CTIT Symposium - Reinventing ICT*, Enschede, The Netherlands, 16 Jun. 2009.
- [MOA:19] **M. S. Oude Alink**, "Cognitive radio," in *16th URSI Forum of the Benelux (invited talk)*, Brussels, Belgium, 18 May 2010.
- [MOA:20] **M. S. Oude Alink**, D. H. Mahrof, A. Ghaffari, S. Subhan, E. A. M. Klumperink, A. B. J. Kokkeler, G. J. M. Smit, and B. Nauta, "Cognitive radio – a revolution in wireless communications," in *STW Jaarcongres, Nieuwegein*, The Netherlands, 7 Oct. 2010.
- [MOA:21] **M. S. Oude Alink**, E. A. M. Klumperink, M. C. M. Soer, A. B. J. Kokkeler, G. J. M. Smit, and B. Nauta, "A crosscorrelation CMOS spectrum analyzer with improved SFDR for cognitive radio," in *20th Annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, Veldhoven, The Netherlands, 18–19 Nov. 2010.
- [MOA:22] **M. S. Oude Alink**, "Cognitive radio research at the University of Twente," in *17th Annual Symposium on Communications and Vehicular Technology in the Benelux (invited talk)*, Enschede, The Netherlands, 24 Nov. 2010.
- [MOA:23] **M. S. Oude Alink**, "Cognitive radio," *I/O Vivat*, vol. 26, no. 2, pp. 32–34, 2011.
- [MOA:24] **M. S. Oude Alink**, A. R. Smeenge, A. B. J. Kokkeler, E. A. M. Klumperink, G. J. M. Smit, and B. Nauta, "Crosscorrelation spectrum sensing exploiting two antennas," in *21th Annual Workshop on Circuits, Systems and Signal Processing (ICT.Open)*, Veldhoven, The Netherlands, 14–15 Nov. 2011.
- [MOA:25] **M. S. Oude Alink**, "A wideband IM3 cancellation technique for CMOS attenuators," in *RFIC and Technology Workshop (invited talk)*, Nijmegen, The Netherlands, 29 Feb. 2012.
- [MOA:26] **M. S. Oude Alink**, "Cross-correlation spectrum sensing - breaking the trade-off between linearity and noise performance," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC) (invited talk in tutorial 'RF Spectrum Sensing and Signal Feature Detection Circuits')*, Montréal, Canada, 17–19 Jun. 2012.
- [MOA:27] **M. S. Oude Alink**, A. B. J. Kokkeler, E. A. M. Klumperink, G. J. M. Smit, and B. Nauta, "UTSFIX: An IC for spectrum sensing to improve noise, linearity, and harmonic rejection," in *21st Annual Workshop on Circuits, Systems and Signal Processing (ICT.Open)*, Rotterdam, The Netherlands, 22–23 Oct. 2012.



Mark Oude Alink received his Master's degrees in Computer Science and Electrical Engineering with honors from the University of Twente in 2008. This thesis is the result of his Ph.D. research, conducted at the Integrated Circuit Design and Computer Architectures for Embedded Systems groups of the CTIT Research Institute at the same university.

His research interests include spectrum sensing, modulation, and analog and digital circuit-level and system-level design for software-defined and cognitive radios.