

Development of a Real-Time Transient Analyser for the SKA

by

Antheun Botha

*Thesis presented in partial fulfilment of the requirements for the degree
Master of Engineering (Research) in the Faculty of Engineering at
Stellenbosch University*



Supervisor: Professor Howard Charles Reader
Department of Electrical & Electronic Engineering

Crtki2016

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Date: 18 November 2013

Abstract

The extension of the Karoo Array Telescope (KAT), MeerKAT, will be the most sensitive radio telescope in the southern hemisphere until it is superseded by the completion of the Square Kilometre Array (SKA). These instruments are to be constructed in the Karoo which is an area in South-Africa that is protected against Radio Frequency Interference (RFI) by the Astronomy Geographical Advantage (AGA) act. However, the telescope is also vulnerable to self-generated interference and specialised measurement systems are required to monitor RFI levels. The development of a Real Time Transient Analyser (RATTY) is described and two Experimental Development Models (XDM) are compared. The first uses a mixing philosophy, and the second direct-sampling. The selection of these models was based on the evaluation of several analogue Front-End (FE) designs. A stripline-filter design process is presented along with the results obtained for custom filters developed for the FEs. Several analyses were compared to measurements performed with one of the devices and good agreement was shown between the system characteristics. Issues regarding the Spurious Free Dynamic Range (SFDR) of the FE designs were identified in the process and measurement-corrected simulations used to predict the achievable ranges. The outcome of the XDM comparison promotes the continued development of a direct-sampling strategy to fulfil the short-term goals of the project. A static calibration procedure is demonstrated for the mixing system and implemented to account for different FE configurations. An overview of the digital and software components of the RATTY system is given and Electromagnetic Compatibility (EMC) principles are applied during the construction of both systems.

Opsomming

Die finale fase van die Karoo Reeks Teleskoop (KAT), MeerKAT, sal die sensitiefste radio teleskoop in die suidelike halfmond wees. Dit sal egter oortref word deur die vierkante kilometer reeks, wat die sensitiefste radio teleskoop ter wêreld sal word. Beide instrumente sal vatbaar wees vir radiofrekwensie steurings en sal opgerig word in 'n wet-beskermd omgewing. Die teleskope is ook vatbaar vir radiofrekwensie steurings wat deur interne stelsels opgewek kan word. Dus word gespesialiseerde meetapparate nodig om die betrokke area en substelsels van die teleskope te monitor. Die ontwikkeling van 'n meetinstrument vir die ontleding van kort-durasie tydseine (RATTY) word beskryf en twee eksperimentele ontwikkelings modelle word vergelyk. Hierdie modelle is gebaseer op die verfyning van voorafgaande ontwerpe vir die analoog substelsel van die instrument en hierdie proses word verduidelik. Die eerste model volg 'n menger strategie waar die tweede model direkte-monstering implementeer. 'n Dubbel-laag, mikrostrookfilter ontwerpproses word beskryf en die gemete resultate vir die ontwikkelde filters word bespreek. Verskeie ontledings is aangewend en vergelyk met die gemete resultate van die stelsels. Hieruit word bevredigende ooreenkomste getref. Die beperkings van die modelle, weens interne distorsie, word geïdentifiseer in dié proses en verdere skattings word gemaak d.m.v. simulاسies. Die eksperimentele modelle word vergelyk en die voorkeur van 'n direkte-monsterings stelsel word gemotiveer. Die digitale en sagteware komponente word oorsigtelik behandel. Tydens die konstruksie van die modelle word die toepassing van elektromagnetiese verenigbaarheids beginsels verduidelik. Laastens word 'n eenvoudige kalibrasie toegepas op die menger stelsel en 'n toepassing daarvan word behandel.

Acknowledgements

I would like to thank the following people for their help and support:

To Prof. Howard Reader, thank you for the guidance over the past two years. You lead by example and it was both inspiring and rewarding to be under your witty supervision.

Thanks to Rob Anderson, Prof. Petri Meyer, Wessel Croukamp, Paul van der Merwe, Joely Andraimbeloson, Nardus Matthysen and Necmi Tezel of SU. Without your assistance or theoretical and practical inputs this project would not have been possible.

Thanks to the following persons from the SKA-Africa: Jason Manley, Sias Malan, Henno Kriel, Richard Lord, Adriaan Peens-Hough, Renier Sieberts, Wesley New, Franciose Kapp, et al.. Your hard work formed the foundation of this project and your willing guidance was invaluable.

To Emeritus Prof. PW v.d. Walt, thank you for your insightful design inputs and shared wisdom.

To my family, thank you for the undying support and love.

Lastly, I would like to specially thank my mother and father for their love and devotion. You paved the way and provided me with the enriching opportunities that led me here.

Contents

Declaration	i
Abstract	ii
Acknowledgments	iii
Nomenclature	vii
List of figures	ix
List of tables	xii
1 Introduction	1
1.1 The Square Kilometre Array	1
1.2 The ReAl Time Transient AnALYser	2
1.3 Contents overview	2
2 Literature Study	4
2.1 Signal processing fundamentals.....	4
2.1.1 Digitisation.....	4
2.1.2 Frequency mixing	5
2.1.3 The Fourier, Discrete and Fast Fourier Transforms	6
2.1.4 Polyphase Filter Banks.....	7
2.1.5 Accumulation	7
2.2 Analogue system characterisation.....	8
2.2.1 Analogue noise characterisation	8
2.2.2 Analogue gain	9
2.2.3 Analogue distortion, compression, dynamic ranges and headroom.....	9
2.3 Analogue-to-digital converters	11
2.4 Stripline filter design.....	12
2.5 Electromagnetic compatibility	15
2.6 RATTY 1	16
3 RATTY 2 front-end developments	18

3.1	Practical considerations for the analogue subsystem	18
3.1.1	Analogue functionality	18
3.1.2	ADC specifications	19
3.1.3	MeerKAT requirements	20
3.1.4	Analogue component selection criteria	21
3.2	Front-end design exploration	21
3.2.1	RATTY 2.1: Single-mixer topology	22
3.2.2	RATTY 2.2: Dual-mixing FE topology	24
3.2.3	RATTY 2.3: Hybrid FE topology	26
3.2.4	RATTY SU: Refined single-mixing FE topology.....	27
3.2.5	RATTY SKA: Direct-sampling FE topology.....	29
3.3	Analogue filter development.....	30
3.3.1	RATTY SU filter designs	31
3.3.2	RATTY SKA filter designs.....	38
3.4	Front-end analysis	40
3.4.1	Component measurements	40
3.4.2	Spread-sheet analysis	41
3.4.3	Spurious analysis.....	45
3.4.4	Simulator analysis	46
3.4.5	Dynamic range specification.....	48
4	RATTY 2 system configuration, construction and integration	52
4.1	Digital subsystem description	52
4.2	Software description	54
4.2.1	Source code and configuration file	55
4.2.2	Operating scripts	55
4.3	FE layout, construction and testing.....	57
4.3.1	FE layout and construction.....	58
4.3.2	FE power supply unit and microcontroller board.....	60
4.3.3	FE reconfiguration	61

4.4	FE tests.....	62
4.4.1	FE gain measurement.....	62
4.4.2	FE SFDR measurement.....	63
4.5	Experimental development model comparisons	70
5	RATTY SU calibration, EMC testing and application	75
5.1	Calibration.....	75
5.2	EMC testing	77
5.3	Reverberation-chamber application	79
6	Conclusions and recommendations	85
	Bibliography	87
A	FE designs	89
A.1	RATTY 2.1: Single-mixer FE topology	89
A.2:	RATTY 2.2: Dual-mixing FE topology	91
A.3:	RATTY SU: Refined single-mixing FE topology	92
A.4:	MATLAB script for mixing spurious analysis.....	93
A.5:	FE Component lists.....	96
B	RATTY SU mixing plan for improved spurious response	98
C	Stripline filter designs	99
C.1	Transformation of a lumped-element to stripline filter.....	99
C.2	Supplementary figures to RATTY SU filter designs.....	101
C.3.	RATTY SU path 3 filters and measurements.	103
C.4.	Supplementary figures to RATTY SKA filter designs	107
D	Component measurements and analyses	108
D.1:	RATTY SU synthesiser.....	108
D.2:	RATTY SU mixer characterisation.....	109
D.3:	Supplementary documentation of FE analyses.....	111
E	RATTY script commands and help	113
F	Mixing experimental development model baseband SFDR	114
G	RATTY design documents	116

Nomenclature

Abbreviations

ADC	Analogue-to-Digital Converter
AWR	Applied Wave Research Incorporated
BPF	Band Pass Filter
CAD	Computer Aided Design
CAM	Control and monitoring
CM	Common Mode
CPU	Central Processing Unit
CST	Computer Simulation Technology™ Electromagnetic Simulation Software
dB	Decibel
dBc	Decibels relative to Carrier
DDR3	Double Data Rate type three
DFT	Discrete Fourier Transform
DSP	Digital Signal Processing
DUT	Device Under Test
E-field	Electric-field
EM	Electromagnetic
EMC	Electromagnetic Compatibility
ENOB	Effective Number of Bits
FD	Frequency Domain
FE	Front End
FPGA	Field-Programmable Gate Array
GbE	Gigabit Ethernet
GSa/s	Giga-Samples per Second
FFT	Fast Fourier Transform
FPGA	Field Programmable Gate Array
GHz	GigaHertz
HDF5	Hierarchical Data Format 5
IC	Integrated Circuit
IF	Intermediate Frequency
IMD	Inter-Modulation Distortion
IP3	Third-order Intercept Point
IP1dB	Input Power at 1 dB Compression Point
KAT	Karoo Array Telescope
KATCP	KAT Communications Protocol
LNA	Low Noise Amplifier
LO	Local Oscillator
LPDA	Logarithmic-Periodic Dipole Array
LPF	Low Pass Filter
MeerKAT	KAT telescope (MeerKAT)
MHz	MegaHertz

MWO	Microwave Office©
NF	Noise Figure
NPR	Noise performance Ratio
OIP2	Output Second Order Intercept Point
OIP3	Output Third Order Intercept Point
OP1dB	Output Power at 1dB compression point
OOP	Object-orientated Programming
P1dB	Power at 1 dB compression
PEC	Perfect Electric Conductor
PG	Pulse generator
PLR	Power Loss Ratio
POWER	Performance optimization with enhanced RISC
PowerPC	Power Performance Computing
RAM	Random Access Memory
RATTY	ReAl Time Transient AnalYser
RF	Radio Frequency
RFI	Radio Frequency Interference
RISC	Reduced Instruction Set Computing
ROACH	Reconfigurable Open Architecture Computing Hardware
R&S	Rohde and Schwarz
S-parameters	Scattering Parameters
SA	Spectrum Analyser
SFDR	Spurious Free Dynamic Range
SG	Signal Generator
SKA	Square Kilometre Array
SINAD	Signal-to-noise and Distortion Ratio
SMA	Sub-Miniature Version A
SNR	Signal-to-Noise Ratio
SQNR	Signal-to-Quantisation-Noise Ratio
SRAM	Static Random Access Memory
SU	Stellenbosch University
TD	Time Domain
TEM	Transverse Electromagnetic

Symbols

ϵ_R	Relative permittivity
F_s	Sampling Frequency
T_s	Sampling Period
T_{Sys}	System Temperature
T_N	Noise Temperature
Z_0	Characteristic Impedance
~	Average value

List of figures

2.1:	Schematic depiction of sampling in TD (above) and FD (below).....	5
2.2:	Plot of the simulated effects of accumulation on noise spectra.	8
2.3:	Example of a sparsely populated spectrum with harmonic and IMD.....	11
2.4:	A representation of a simple stripline filter structure.	13
2.5:	Schematic of a normalised, generic LPF circuit (Left) and its PLR (Right) response showing loss against frequency.	13
2.6:	A graphical representation of the first Kuroda transform.	15
2.7:	Illustration of CM-current flow on a coaxial cable with L-plate structure.	16
3.1:	A condensed representation of the RATTY 2 layout showing three subsystems.....	18
3.2:	Schematic illustration of sampling in different Nyquist zones.	19
3.3:	RATTY 2.1 design schematic.	22
3.4:	Schematic of an ideal BPF 1 response convolved with minimum LO.	23
3.5:	RATTY 2.2 FE schematic featuring baseband and dual-mixing paths.	24
3.6:	Lower sideband up-conversion with LO1 and second harmonic.....	25
3.7:	Upper side-band down-conversion showing LO 2 and related harmonics.	25
3.8:	A single-input, hybrid FE design schematic.	26
3.9:	An illustration of filter roll-off margins showing an aliased response.	27
3.10:	Schematic of a single-input layout for the RATTY SU design.	28
3.11:	Individual and combined passband characteristics of LPF 1 and 2.....	28
3.12:	Schematic of a single-input layout for the RATTY SKA design.	30
3.13:	Schematics of relevant system components (left) and ideal system responses (right).	31
3.14:	Schematic showing transformation of a normalised PLR from an ω to a θ plane.	32
3.15:	Schematics showing the derivation of a practical stripline filter.....	33
3.16:	Responses of lumped element and stripline filters as ideal models.....	33
3.17:	Responses of ideal and substrate dependent stripline filter models with adaptations.	34
3.18:	Comparison of measured and simulated BPF 1 forward transmission coefficient.....	36

3.19:	Simulated and measured responses for Add LPF 1	37
3.20:	Simulated response of the filter chain in signal path 2, RATTY SU.....	37
3.21:	Simulated and measured results of the RATTY SKA BPF 1 forward transmission coefficient.....	38
3.22:	Simulated and measured results of the RATTY SKA Add LPF 1 forward transmission coefficient.....	39
3.23:	BPF 1 and Add LPF 1 simulated combined response for RATTY SKA.....	39
3.24:	Spread-sheet calculations of the NF and gain parameters for path 2 in the mixing FE	42
3.25:	Spread-sheet calculations of the T_{sys} and OP1dB parameters for path 2 in the mixing FE.	43
3.26:	Spread-sheet calculations of the NF and gain parameters for path 1 of the mixing FE.....	43
3.27:	Plots of mixing spurious in the RATTY SU path 2.	46
3.28:	Simulated signal path responses of the RATTY SU FE.	47
3.29:	An illustration of SFDR definitions.....	48
3.30:	Schematic depiction of out-of-band suppression in the mixing path.....	51
4.1:	Block diagram of RATTY 2 experimental development model development.	52
4.2:	RATTY digital subsystem with key components outlined.	54
4.3:	Flow diagram representation of the measurement software routine.	57
4.4:	Stacked components with interconnecting SMA cables.	59
4.5:	Linear voltage regulators and twisted supply wires.....	59
4.6:	Photograph of DC-supply and SMA interfaces between the analogue and digital enclosures.	60
4.7:	Photograph of the optical isolation of FE-control lines.	61
4.8:	Reconfigured section of the mixing FE.	61
4.9:	Comparisons of measured and simulated signal-path gains.	62
4.10:	IMD measurements of the mixing path.	63
4.11:	Comparison between mixed IMD from arbitrarily spaced and factor-ten spaced inputs.	64
4.12:	Varying 3 rd order IMD levels and broadband FE outputs.....	64
4.13:	Worst-case IMD levels of the RATTY SU mixing path.....	68
4.14:	The RATTY measurement arrangement in the reverberation chamber.....	70
4.15:	Single-tone measurements with the second signal path of the direct-sampling system.	71

4.16:	Single-tone measurements with the second signal path of the mixing system.	72
4.17:	Low frequency RFI in the mixing path of RATTY SU.	72
4.18	Difference plot of RATTY SKA single-tone measurements.	73
4.19	Difference plot of RATTY SU single-tone measurements.	73
5.1:	RATTY SU baseband and mixing-path gain functions.	76
5.2:	Different RF-cable configurations and magnetic choke application.	77
5.3:	Difference plots of self-induced RFI on the digital-enclosure RF cable.	78
5.4:	Measurements of terminated digital enclosure.	78
5.5:	The SU reverberation chamber	80
5.6:	Direct pulse measurements with a sampling oscilloscope, RATTY SU baseband and mixing path.	81
5.7:	Spectra calculated from direct-pulse measurements with a sampling oscilloscope and RATTY.	82
5.8:	Simplified RATTY SU schematic illustrating positions of level-setting components.	83
5.9:	Reverberation chamber pulse measurements with a pulse train and single pulse.....	84
A.1:	A dual-input schematic of the RATTY 2.1 design.	89
A.2:	Ideal FD representation of path 2 signal processing in RATTY 2.1.....	90
A.3:	Schematic of RATTY 2.2 with single-valued gain and dynamic range calculation.....	91
A.4:	Schematic of RATTY SU with single-valued gain and dynamic range calculation.....	92
A.5:	Ideal FD representation of two stages of the FE signal processing in RATTY SU.....	92
A.6:	Worst-case mixing-spurious analysis for RATTY SU path 2.....	93
B.1:	Analyses of mixer spurious for improved SFDR.....	96
C.1:	Ideal stripline models (MWO) of a scaled, 4 th order Butterworth filter.....	98
C.2:	Comparison of BPF 1 responses with and without bends, cross sections and input strips.	99
C.3:	MWO and CST models of the final BPF 1 design and a flat pin SMA connector model.	99
C.4:	Comparison of BPF 1 responses with and without CST adaptive mesh refinement.	100
C.5:	Comparison of measured and simulated BPF 1 reflection coefficients.	100
C.6:	Comparison of measured and simulated Add LPF 1 reflection coefficients.	101
C.7:	Measured and simulated BPF 2 characteristics	102

C.8:	The MWO and CST models of BPF 2.....	103
C.9:	Measured and simulated Add LPF 2 characteristics.....	104
C.10:	Simulated response of the filter chain in signal path 3, RATTY SU.....	104
C.11:	Simulated response of the filter chain in signal path 3, RATTY SU.....	104
C.12:	Simulated and measured results of the BPF 1 reflection coefficient.	105
C.12:	Simulated and measured results of the Add LPF 1 reflection coefficient.....	105
D.1:	SA measurements of 1.5 GHz synthesiser outputs with and without filtering	106
D.2:	Mixer RF-port, reflection-coefficient measurements at different LO frequencies.	107
D.3:	Schematic showing the mixer-test configuration.....	108
D.4:	Measured and calibrated conversion loss curves for mixing signal paths 1 and 2.	108
D.5:	Mixer spurious measurements for RATTY SU path 2 range.....	109
D.6:	Spread-sheet calculations for the T_{sys} and OP1dB parameters of the RATTY SU, signal path 1.	109
F.1:	Worst-case SFDR performance of the baseband path.	113

List of tables

2.1:	The key specifications of RATTY 1.....	17
3.1:	Peak spurious products calculated for RATTY 2.2.	26
3.2:	Passband frequencies for the filters of the hybrid FE design.....	27
3.3:	Filtering specifications for RATTY SU.....	29
3.4:	Filtering and sampling specifications for RATTY SKA.	30
3.5:	Initial BPF 1 specifications.....	32
3.4:	Calculation of the BPF 1 scaling factor.	32
3.5:	Properties of the substrate used for BPF 1.....	35
3.6:	Measured BPF 1 characteristics.....	36
3.7:	Specifications for RATTY SU Add LPF 1 and HPF 1.....	37
3.8:	Comparison of spread-sheet parameter calculations for RATTY SU paths 1 and 2.	44
3.9:	Spread-sheet calculations of system parameters for RATTY SU path 1	

with varying attenuator settings.....	44
3.10: Comparison of spread-sheet and simulator parameter calculations for RATTY SU path 2.	47
3.11: Minimum FE SFDR calculations for RATTY SU paths 1 and 2.	49
3.12: RATTY SU path 1 and 2 SFDRs according ADC range definition.	50
3.13: Estimations of the maximum sensitivities for the RATTY SU signal paths.	50
4.1: Important data fields common or specific to TD and FD HDF5 files.	57
4.2: Measured output levels used for calculating intercept points and SFDRs of the mixing path.....	66
4.3: Measurement and simulation based calculations of key intercept points and SFDRs.	66
4.4: Comparisons of measured and simulated OIPs and SFDRs for the mixing path.	67
4.5: Baseband comparisons of measured and simulated parameters.	68
C.1: Initial BPF 2 specifications.....	101
C.2: Measured BPF 2 ranges.	102
D.1: Valon synthesiser specifications.....	106
D.2: Table of system parameters for RATTY SU path 2 at several step-attenuator settings.	110
D.3: Comparison of spread-sheet and simulator parameter calculations for RATTY SU path 1.....	110
E.1: List of terminal commands for RATTY control.....	111
F.1: Measured output levels used for calculating intercept points and SFDRs of the RATTY SU baseband path.	112
F.2: Measurement and simulation based calculations of key intercept points and SFDRs for the RATTY SU baseband.	112
F.3: Comparisons of measured and simulated OIPs and SFDRs for the mixing path.	112

Introduction

1.1 The Square Kilometre Array

The Square Kilometre Array (SKA) is a prestigious, international project which aims at developing the world's largest radio telescope. This will require significant technological developments to be successful and will provide the field of radio astronomy with an instrument of unprecedented scientific capability. The main objectives are to achieve sensitivities and survey speeds well beyond the specifications of contemporary radio telescopes. These features present challenges which impact on each other; compromises must be made to retain feasible goals for the project.

The full SKA will be divided into low and mid-frequency, as well as approximately 2500 high frequency telescopes. Southern Africa was awarded the dominant portion of the telescope project. As the name suggests, the telescope will provide a collective aperture of 1 000 000 m^2 . Although there will be remote stations spread throughout Africa, the majority of the telescope elements will be concentrated at the core site. The farthest stations are estimated to be up to 3000 km away from the core, which is projected to span only 15 – 20 km. This centre will be positioned at a site that was carefully selected and was the main motivation for constructing a significant portion of the project in South Africa. The site is located in a sparsely populated area in the Northern Cape Province, also known as the Karoo. In addition to excellent climatic conditions for radio observations, the site was thoroughly examined and is considered to have low Radio Frequency Interference (RFI) characteristics. The RFI quietness is crucial and will allow sensitive surveys to be conducted without contamination from unwanted sources.

A distinction is made between internal and external sources of RFI. The former is related to self-generated emissions from the telescope systems and the latter is interference caused by radiation from systems outside of the telescope. The sensitivity of the instrument will make it extremely vulnerable to either type of RFI. Indeed, RFI threatens the success of the project and it is thus being given significant prominence in all aspects of system design. The core site is protected against external RFI by the Astronomy Geographical Advantage Act, which places stringent limits on the allowable levels of Radio Frequency (RF) emissions inside a large area surrounding the location. To avoid self-generated interference, the subsystems of the telescope must undergo rigorous testing.

The SKA will be constructed in several phases estimated to commence in 2018. Meanwhile, a South-African telescope is under development. This is an expansion of the Karoo Array Telescope (KAT), known as MeerKAT. The latter is set to be the most sensitive radio telescope in the southern hemisphere until the SKA is completed in 2024. The development of this telescope stimulated a wide field of related technological innovation, which forms part of, and contributes to, the SKA project. MeerKAT provides the background and motivation for the work described in this thesis [31].

1.2 The ReAl Time Transient Analyser

The ReAl Time Transient Analyser (RATTY) is a development from the MeerKAT digital backend team and is based on Reconfigurable Open Architecture Computing Hardware (ROACH). The system utilises the processing capabilities of the ROACH technology to perform broadband, RF measurements. Ultimately, the goal is to contribute to both the on-site monitoring of RFI and emission-verification requirements of the MeerKAT and SKA projects. These are long-term aspirations and this research describes intermediate developments in the exploration of such possibilities. In attempting to achieve the stringent requirements of these large projects, a versatile measurement system will be presented which can also be used in other applications.

A number of natural explorations were pursued to develop an improved design leading to two Experimental Development Models (XDM). These systems were mainly differentiated by the analogue Front End (FE) designs and sampling strategies that were employed. The two XDMs were pursued by SKA and Stellenbosch University (SU) teams, respectively. It is worthwhile documenting the unsuccessful paths in the development for at least two reasons: 1) the findings motivate the selection of the XDM designs and 2) it is worth alerting others to the shortcomings of some approaches. The constructed systems were subjected to several evaluations; one of the designs will be recommended for future developments.

1.3 Contents overview

The second chapter provides the necessary theoretical background for the work in later sections. This includes an overview of fundamental signal-analysis techniques, analogue-system characterisation, Analogue-to-Digital Converter (ADC) characteristics and Electromagnetic (EM) Compatibility (EMC) concepts. The principles of stripline-filter designs are also described as a reference to the development of analogue filters in Chapter 3. Lastly, a summary of the RATTY 1 instrument provides an introduction to the system layout.

Chapter 3 focusses on the design of an RF FE and begins with an explanation of the basic functionality of this subsystem. The incorporation of the ADC specifications is explained and the requirements of the MeerKAT project outlined. This is followed by the criteria considered in selecting connectorised, analogue components. The different FE designs are then evaluated and the selection of two strategies for the comparison of XDMs motivated. A section was added to describe the development of several custom filters for the FEs. The chapter concludes with an analysis of the designs. This involves component measurements with the evaluation of the system performance using several simulation techniques.

A description of the digital and software subsection on the RATTY system can be found in Chapter 4. The important features of the system construction are described with the focus on the implementation of EMC principles. A set of FE tests were performed and a comparison is made between the simulated and measurement gain and SFDR characteristics. This is followed by a comparison of the XDMs, which is concluded with the selection of the favoured design.

The fifth chapter describes a static calibration procedure and EMC testing that were performed on one of the XDMs. An application of this device is then described, namely the instrument in a reverberation chamber.

Chapter 6 concludes the study making several recommendations for future improvements. The appendices that follow contain supplementary figures, tables and other additions that are referenced throughout the thesis.

Chapter 2

Literature study

2.1 Signal processing fundamentals

2.1.1 Digitisation

An analogue signal contains an infinite amount of information, taking on any value within an infinitesimally fine resolution at any time. To constrain the information for storage on a computer, analogue signals must be discretised. This is termed digitisation and involves two steps, namely, sampling and quantisation.

Sampling involves discretising a signal in time and the ideal case is represented in this domain as the multiplication of a signal by an impulse train with a repetition rate equal to the Sampling Period (T_s). The result is an impulse train with sequential amplitudes equal to that of the analogue signal at the sampling instance. This signal can now be parameterised in discrete time (n), as opposed to continuous time (t). Since a discrete sequence has a finite amount of samples, the sampled spectrum is also multiplied by a windowing function, the default shape being a rectangular window.

The implication of sampling in the Frequency Domain (FD) is the convolution of the input spectrum and a uniform impulse train with a period that is the inverse of the sampling period, i.e., the Sampling Frequency (F_s). This gives rise to the Nyquist criterion, which states that a sampled signal's spectrum must be limited to half of F_s to avoid signal degradation.

A signal is filtered prior to sampling to limit its bandwidth and fulfil the Nyquist criterion. However, ideal LPF's do not exist and time-limited signals technically have infinite bandwidth, so aliasing is unavoidable. This is the process whereby a signal's spectrum extends beyond half of F_s , effectively being mirrored around this point and distorting the lower part of its spectrum. Hence, the quality of the filtering applied to a signal impacts significantly on the effects of sampling. The sharper the roll-off of the filter, and the higher the suppression of out-of-band responses, the less the effects of aliasing. Figure 2.1 depicts sampling in the TD and FD with the effect of aliasing. Note that the spectrum is periodic and that only the content up to half of F_s is unique.

The sampled spectrum is also convolved with the Fourier transform of the windowing function, which leads to distortion in the FD. Since the Fourier transformed window typically displays sidelobes, unlike the ideal impulse function, this leads to spurious products and is referred to as spectral leakage. As a result, single components can cause non-zero values in the spectrum around them. By selecting a tapered window with lower sidelobe levels, a better choice in windowing function can limit this unwanted effect. However, the compromise is wider main lobe widths, which degrade the resolution [5].

Quantisation is the second step in signal digitisation, which involves the discretisation of a signal's amplitude. By limiting the values the amplitude can take in its discrete representation, it can be cast to a

binary value. The possible input range is divided into intervals and the value of the input signal is equated to the quantisation interval in which it resides.

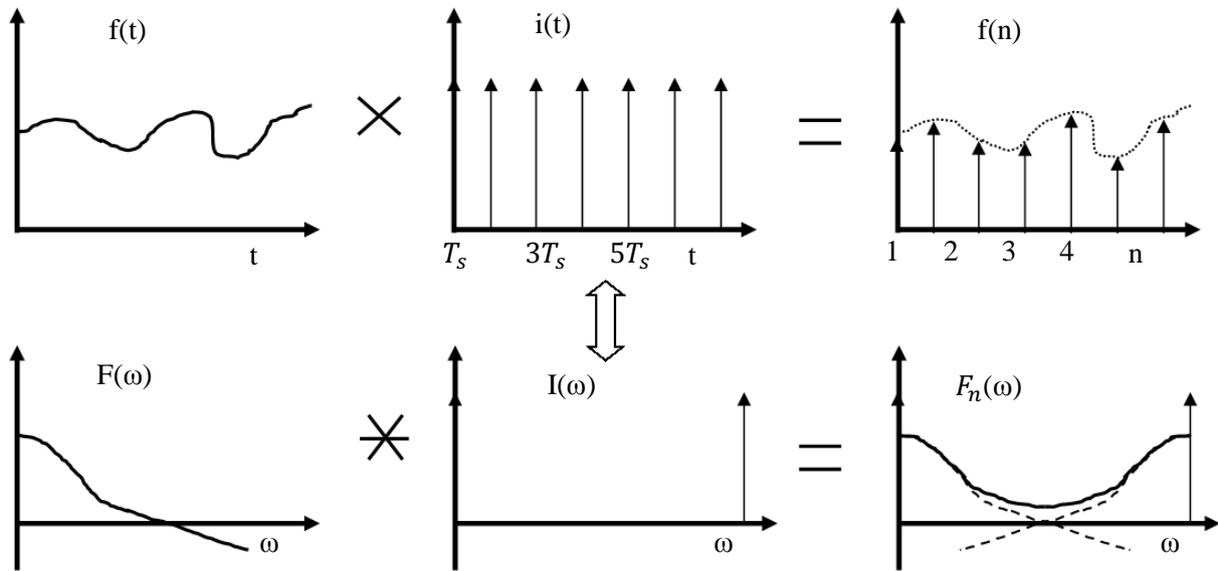


Figure 2.1: Schematic depiction of sampling in the TD (above) and FD (below) (after [18]). Note: 1) the Fourier Transformation between the two domains, 2) aliasing in the sampled spectrum and 3) the effect of windowing.

Since the digitised signal value is no longer exact, this introduces quantisation errors and noise, which limits the achievable Signal-to-Quantisation-Noise Ratio (SQNR). The noise is minimised by increasing the amount of intervals used, reducing the average error that will be made. Since the value needs to be in a binary form, there are always 2^N intervals, where N is the number of bits used to represent the value. The relation between N and the SQNR for sinusoidal inputs is given by equation 2.1, as presented in Proakis [5]. This serves as the standard for all real signals.

$$SQNR(dB) = 1.76 + 6.02 * N \quad (2.1)$$

2.1.2 Frequency Mixing

Frequency mixing is the process whereby a signal is modulated by a carrier in order to shift its spectrum to a different frequency. The ideal modulation is represented as the convolution of the input spectrum with that of a carrier sinusoid's, which implies that the signals are multiplied in the TD.

The component used to achieve this modulation is a three terminal device called a mixer. The modulating sinusoid input is typically referred to as the Local Oscillator (LO) and the remaining ports as the RF and Intermediate Frequency (IF) signals, respectively. The nomenclature for the latter signals typically used and followed here is defined for an RF receiver system. The RF is the signal at the higher frequency, since receivers are designed to shift radiated signals downward in frequency, i.e. down-conversion. This enables the use of detection electronics operating at the lower, IF frequencies. The circuitry preceding and subsequent to the mixer are typically also referred to as the RF and IF sections, respectively. If used for up-conversion, the mixer ports are reversed; the IF is used as input and the RF as output.

Mixing is implemented practically with nonlinear switching circuits driven by the LO. The result is the input spectrum convolved with multiple harmonics of the carrier. However, only the convolution of the input and LO fundamental component is desired and the rest are considered spurious products. Due to the non-ideal nature of the switching electronics, the RF and LO inputs modulate switch behaviour and this results in RF to IF and LO to IF signal leakage.

The distortion is typically summarised in spurious tables that predict the level of unwanted products relative to that of the desired output component. These tables list the mean value of products produced across a mixer's operational bandwidth and spurious levels may be worse than predicted for specific implementations. Due to the broadband nature of spurious products, the carrier frequency and bandwidth of the modulated signal must be wisely chosen to avoid contaminating it with the unwanted products. Since mixer spurious is non-linear, the levels of products are dependent on the signal power and are reduced at lower levels [4].

Mixers also exhibit conversion loss, noise and Inter-Modulation Distortion (IMD), due to the use of nonlinear components and resistive losses. An important factor in mixer performance is the LO power level, since the aforementioned non-ideal characteristics are affected by this. Typically, the higher the LO level is, the closer the switching is to the ideal process. However, a high power LO signal can be problematic if it leaks into the signal path and is not sufficiently suppressed. There are various types of mixers with different characteristics and a good choice is dependent on the application.

2.1.3 The Fourier, Discrete and Fast Fourier Transforms

The Fourier transform is a mathematical function that enables both the conversion of TD information to the FD and the inverse. It is classically derived from the Fourier series expansion of periodic signals. By time gating and expanding any real, aperiodic signal to a periodic repetition of itself, the Fourier series representation of such a signal can be determined. However, by allowing the period of such an expansion to strive to infinity, the spectral information becomes infinitesimally tightly packed and the series expansion transforms into a continuous function. This provides the frequency content of the input signal at all frequencies as an analytical solution and equation 2.2 defines this transform, as presented in Lathi [17].

$$X(\omega) = \int_0^{\infty} x(t)e^{-j\omega t} dt \quad (2.2)$$

where $x(t)$ is a TD signal, $X(\omega)$ is its FD transform and ω is the continuous variable of frequency (rad/s).

The Discrete Fourier Transform (DFT) is derived from the Fourier transform and calculates the spectral components of a discrete TD signal as a finite series. This can be regarded as sampling in the FD and leads to the periodic repetition of the TD samples, due to the convolution with an impulse train. If N is the length of the time sequence fed into a DFT algorithm, there are only $N/2$ unique samples in the resulting FD sequence. The DFT is defined by equation 2.3 [17].

$$X(k) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi kn/N} \quad k = 0, 1, \dots, N-1 \quad (2.3)$$

where k is the discrete FD variable, n is the discrete TD variable and N is the length of.

Efficient implementations of the DFT have been developed and dubbed Fast Fourier Transforms (FFT). These exploit the symmetric and periodic properties of the DFT by sequentially rearranging the input sequence into two dimensional matrices. These matrices reduce in size until they are factored into their smallest prime dimensions, which leads to significant improvements in the computational complexity of the DFT. A particularly efficient and common implementation is possible when the TD sequence has a length that is a power of two, which is referred to as Radix-2 algorithms. For this reason, TD data is often extended to a radix-2 length with the addition of zeroes, termed zero padding. Although this does not increase the effective resolution of the resulting spectrum, i.e. how well closely spaced spectral components can be resolved, it provides a more densely sampled spectrum and a better representation in this domain [5].

2.1.4 Polyphase Filter Banks

A polyphase filter is a digital filter implementation that acts on discrete data by adding delays and applying TD filter coefficients to a shifted input sequence. Similarly, polyphase filter banks consist of multiple polyphase structures in parallel with distinct phase responses, hence the terminology [5]. Polyphase filter banks are employed to limit spectral leakage and scalloping loss, by applying an optimised TD windowing filter to input samples before employing a DFT. Scalloping loss is caused by a curved main lobe in the FD transform of a windowing function, which leads to losses when the spectrum is sampled by the DFT, since samples are not necessarily taken at the centre of the lobe.

A polyphase filter bank method that is commonly employed for use in spectrometers splits the input data stream of S samples into subsets. The number of subsets (T) is referred to as the taps of the polyphase filter bank and this gives rise to the structure of the filter with $S = T * N$. A window function is multiplied with the data stream within N sub-filters, each with T coefficients, by entering the data to these parallel filters through a rotating commutator delaying the samples systematically. The filter coefficients are arranged to match the commutator operation and the parallel output streams are then entered to an N -point DFT. In this process, the T subsets are optimally windowed, added together and a weighted spectrum is yielded [11].

2.1.5 Accumulation

In all natural systems there is thermal noise present - a random signal caused by the thermal agitation of electrons. It is characterised as white noise, since its probability density spectrum is Gaussian at RF frequencies. In a measurement system there are multiple sources of thermal noise, usually defined by the ambient measurement and system noise temperatures, with the resulting noise power assumed to be constant throughout the bandwidth. The level of noise power that is measured is referred to as the noise floor of the system and characterised by equation 2.2 for a matched system, as presented in Pozar [1].

$$N_{sys} = kT_N B \quad (2.4)$$

Where N_{sys} = system noise power, k = Boltzmann's constant, T_N = Noise Temperature and B = processing bandwidth.

When consecutive spectra from a system are averaged over time, the randomness of the resulting noise floor is decreased. This does not reduce the noise floor level, but it enables the detection of small signals that are otherwise obscured by the random fluctuations in the noise floor. Rather, averaging enables a better approximation of the noise floor, since the power density spectrum of the noise is then calculated with more data. The process of adding and averaging spectra is referred to as accumulation and figure 2.2 shows simulated results that demonstrate its effects. This involved generating random data sequences with normal distribution properties and comparing the spectrum of a single sequence to averaged spectra.

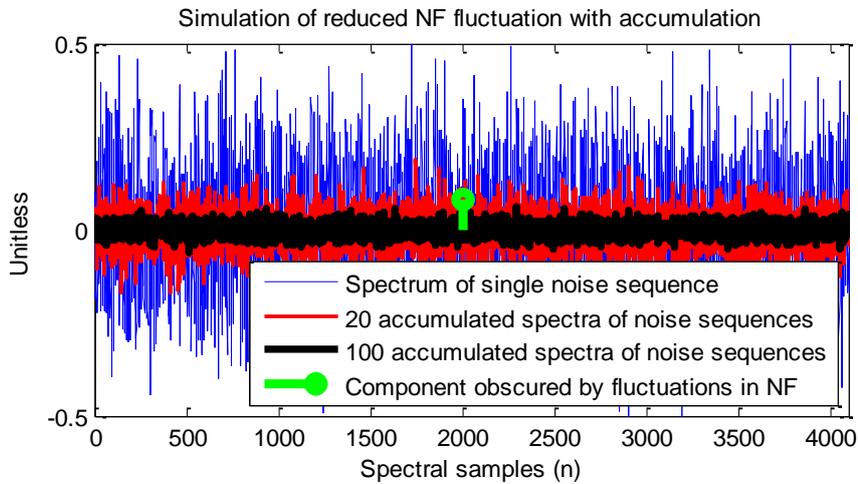


Figure 2.2: Plot of the simulated effects of accumulation on noise spectra. Note that the fluctuations in the noise floor are reduced and that the component at $n = 2000$ becomes visible as the number of accumulated spectra increases.

2.2 Analogue system characterisation [1]

There are several important parameters that are used to characterise an analogue, RF-receiver system. The noise floor, noise temperature (T_N), noise figure (NF), gain, input 1 Decibel (dB) compression point (IP1dB) and headroom of an analogue system are described here.

2.2.1 Analogue noise characterisation

In the conventional definition, system noise is referred to the input and modelled as a matched load at the calculated system temperature (T_{sys}). Furthermore, a field probe will present a measurement temperature (T_{meas}) as input to the system. Thus, the system is modelled as being noiseless with the sum of the self-generated and measured noise powers as an input. The noise power is calculated from the T_N using equation 2.4, with the NF and T_N related through equation 2.5, as presented in Pozar [1]. Note that the NF is characterised by the ratio of T_N to the ambient temperature of 290 K.

$$NF = 10 \log_{10} \left(1 + \frac{T_N}{T_0} \right), \quad (2.5)$$

where $T_N = T_{Sys} + T_{Meas}$ and $T_0 = 290$ K, the ambient temperature.

When components are cascaded, the resulting T_{Sys} of the combined response is determined by the characteristics and order of the elements. Equation 2.4 shows this interdependence and indicates that the first component in the chain can have a significant influence on the T_{Sys} .

$$T_{Sys} = T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + \dots \quad (2.6)$$

where T_1, T_2, T_3, \dots and G_1, G_2, \dots are the noise temperatures and gains of consecutive components in a system.

Since spectra are mostly produced with FFT implementations, it is important to consider the effects this has on the noise floor. An FFT performs a sweep-like action, whereby the spectrum is essentially sampled at the FFT resolution. This acts as a resolution bandwidth (RBW) filter and changes the level of the noise floor due to the bandwidth dependency described by equation 2.2. Therefore, a distinction can be made between the analogue and digital noise floors and they are related by equation 2.4, as presented in [19].

$$\text{Noise Floor}_{FFT} = \text{Noise Floor}_{Anlg} - 10 \log_{10}(M/2) \quad (2.7)$$

where Noise Floor_{FFT} = the digital processed noise floor, $\text{Noise Floor}_{Anlg}$ the analogue noise floor and M is the length of the FFT.

2.2.2 Analogue gain

The gain of an analogue component is defined through scattering parameters (S-parameters) in equation 2.8, as presented in Pozar [1]. The equation is divided into three terms; the middle term is the maximum realisable gain with the first and last terms correcting for impedances mismatches at the input and output ports, respectively. The maximum gain is achieved when the section is perfectly matched, although this is mostly not the case and the actual gain is dependent on the cascade of components. However, the maximum gain can be assumed, if the parts are well matched. This may be the case when amplifiers are used in cascade with well-matched passive elements, since amplifiers typically dominate the response due to their good reverse isolation (S_{12}) characteristics. If the maximum gain is assumed, the total gain of a cascaded system is simply the sum of the individual gains of the parts. The accuracy of a simplified analysis should be evaluated and compared to more precise calculations, a measured response or both.

$$G = \frac{1 - |\Gamma_S|^2}{|1 - \Gamma_{In} \Gamma_S|^2} \cdot |S_{21}|^2 \cdot \frac{1 - |\Gamma_L|^2}{|1 - S_{22} \Gamma_L|^2} \quad (2.8)$$

$$\text{with } \Gamma_{In} = S_{11} + \frac{S_{12} S_{21} \Gamma_L}{1 - S_{22} \Gamma_L}$$

where G = realised gain, Γ_S = source reflection coefficient, Γ_L = load reflection coefficient, Γ_{In} = adjusted input reflection coefficient, S_{11} , S_{22} , S_{12} and S_{21} are scattering parameters of the gain section.

2.2.3 Analogue distortion, compression, dynamic ranges and headroom

In general, no component has a purely linear response, i.e. the desired output is only linear for a range of input power levels and it will contain unwanted products. Furthermore, the unwanted products include harmonic and IMD spurious components. The harmonic elements repeat at fixed intervals of the input frequency, whereas IMD is the interaction of input components and occurs at the sum or difference of input frequencies and their harmonics. The levels of distortion components are typically specified relative to that of the desired response, i.e. Decibels relative to carrier (dBc). These levels are dependent on the signal power, since a distortion product increases in magnitude exponentially, raised to the power of the product order.

The upper limit of linear response is typically the Power at 1 dB compression (P1dB) point. This point is reached when the output of a device is 1 dB lower than expected for a linear response and it can be defined for input (IP1dB) or output (OP1dB) power levels. Another figure of merit is the Third-order Intercept Point (IP3), which is the hypothetical point where third order IMD products would be equal in magnitude to the desired output if a component were to remain in a linear state at such a power level. The P1dB figure degrades when a component is placed in a cascaded system and this must be calculated throughout a layout. Equation 2.9 is used to adjust the OP1dB of a component to incorporate the effect of the part preceding it. This is adapted from the IP3 adjustment presented in [1], which assumes that the relationship between these figures remains constant.

$$P1dB_2 = -10 * \log_{10} \left(\frac{1}{p1dB_2} + \frac{1}{p1dB_1 * G_2} \right) \quad (2.9)$$

where $P1dB_2$ = adjusted OP1dB of the second component (dB), $p1dB_2$ = initial, linear OP1dB of the second component, G_2 = the component gain and $p1dB_1$ = linear OP1dB of the preceding component.

The dynamic range of a device is the maximum power range that results in a linear response. The dynamic range lower limit is the noise floor and the upper limit is the P1dB, which is also the definition of headroom. There is also a Spurious Free Dynamic Range (SFDR) figure and it is the key parameter used to specify the measurement range of a device. The SFDR definition mostly used in this document has an upper limit bound by the P1dB and the lower limit is set by the highest spurious component in the output, hence the terminology. Any spectral components within the SFDR are present at the input and are not false products as a result of non-linear effects. Equation 2.10 is a typical maximum SFDR definition based on the Output IP3 (OIP3) and the noise floor, as presented in Pozar [4]. However, this equation is for the maximum SFDR relative to third order IMD products and it is only used as a reference figure in this document.

$$SFDR (dB) = \frac{2}{3} * (OIP3 - N_o) \quad (2.10)$$

where N_o = output noise power level.

Furthermore, if a spectrum is sparsely populated, spurious products can typically be identified by the frequency relation to the higher level components present. This allows a measurement range to be interpreted beyond the SFDR, since actual, lower-level signals can be distinguished from false products. With an FFT lowering the NF, the observed digital dynamic range is greater than the analogue dynamic range figure. The Signal to Noise Ratio (SNR) is Figure 2.3 shows a sparsely populated spectrum with IMD and harmonic distortion elements. The Signal to Noise Ratio (SNR) is the difference between the signal level and the noise floor, which is typically less than the SFDR.

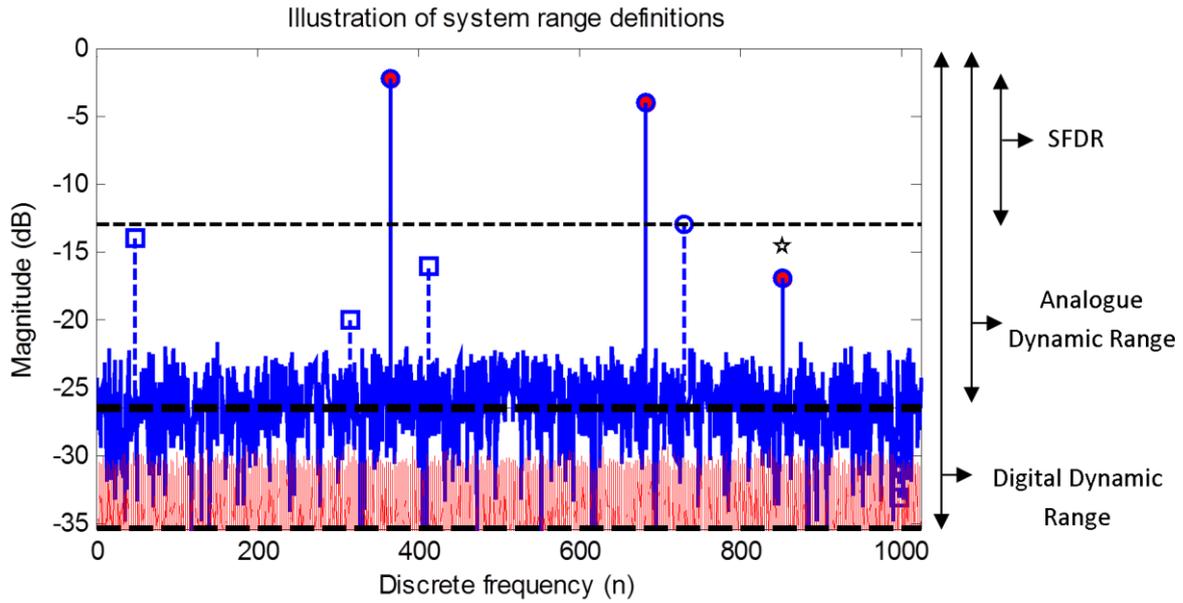


Figure 2.3: Example of a sparsely populated spectrum with harmonic and IMD. The IMD is shown with square and harmonic distortion with round markers, respectively, featuring dashed stem lines. Note that the analogue and digital NF levels are shown with horizontal, dashed lines and that the digital level is lower. Also, the digital and analogue dynamic ranges are indicated alongside the SFDR, assuming that the P1dB is at the 0dB level. A low, actual component is indicated with a star and could be distinguished from the other components, although it is below the SFDR.

2.3 Analogue-to-digital converters

There are several non-ideal effects present and dealt with in sampling theory, as described in a Section 2.1. However, an ADC typically consists of several stages of complex electronics and practical sampling presents challenges that lead to additional unwanted distortion in the output of ADCs. In addition to IMD, an ADC also causes harmonic distortion of input signals. These spurious products occur in a sampled spectrum as components aliased around F_s and are observed at frequencies described by equation 2.4, as presented in [19]. Since the spectrum is periodic, each spectral repetition falls within a range called a Nyquist zone spanning half of F_s . Only the information in the first zone is retained and is called the baseband range.

$$f_{spur} = |\pm K * F_s \pm n * f_{In}| \quad (2.10)$$

where f_{spur} = spurious frequency, f_{In} = input frequency and $K, n = 1, 2, 3, \dots$

There are several parameters that are used to describe ADC performance. These include the SFDR and IMD, which are identical to the descriptions provided in section 2.2. The remaining performance parameters of importance are defined according to the descriptions provided in the datasheet of an ADC part that was used in the systems that are described in later sections [7]. This includes the Signal to Noise Ratio (SNR), which excludes the power of first 25 spurious harmonics to provide the ratio of the signal level relative to the ADC noise floor. There is also a Signal-to-noise and Distortion Ratio (SINAD) that includes the power of all the unwanted components in the ratio. The Effective Number of Bits (ENOB) is calculated from the SINAD figure and provides a measure of the realised dynamic performance of a device in bits. Furthermore, the

Noise Power Ratio (NPR) is a test of the ADC performance for broadband input signals. The importance of these characteristics will be made clear within the context that they are used.

The assumption that quantisation errors are uniformly distributed, which is used for the calculation of the quantisation noise level, is not valid for simple, low-level signals. Therefore, the ratio of low level signals to the resulting spurious products can be significantly less than the SFDR realised for large input signals that span a bigger range of the ADC. Typically, this affects signals below 10 quantisation intervals in magnitude most significantly. Since the quantisation intervals are significant relative to low amplitudes, small signals are coarsely digitised and result in step-like outputs. This produces high harmonic distortion levels and is similar in effect to a saturated signal.

The spurious performance and resolution of the ADC can be improved by adding a controlled amount of broadband noise to the input, which is called dithering. In combining the small signal with noise and applying averaging, the quantisation errors are randomised for small signals and the harmonic distortion can be significantly reduced. The cost of the process is a decrease in SNR and conversion efficiency, due to the added noise and averaging, respectively [20]. There are also alternative dithering techniques and the optimal amount of dithering power depends on the application. Furthermore, the method can be useful across the ADC range and also improves harmonic distortion caused by inherent nonlinearities for high-level inputs [28].

A measure of the broadband performance of an ADC is indicated by the NPR. This is tested using a notch-filtered, wideband noise source as input to the ADC. The notch suppression must typically be greater than the SFDR to allow ADC products to be observed. The NPR is then measured as the ratio of the wideband, input power level to the narrower, in-notch power level. The NPR may be less than the single-tone SFDR and could influence the measurements of broadband signals [7].

An ADC is a sensitive and complex device, which is vulnerable to destructive effects that can materialise through several mechanisms. Although the source of the damage and the process may vary, exceeding any absolute maximum ratings of an ADC can impair the device by degrading performance parameters or destroying it. Furthermore, some damaging effects, such as electrostatic discharge, can be cumulative and worsen the behaviour of a part over extended periods. Therefore, care must be taken to protect an ADC from damage and most parts have protecting features. A device may require a settling time prior to an over-voltage condition, to regain the specified sensitivity [25].

In order to increase the sampling frequency, multiple ADCs circuits can be used in an interleaved unit with each circuit sampling out-of-phase to the others. However, these components typically have worse spurious responses than single-part, monolithic ADCs.

2.4 Stripline Filter Design [1]

A stripline is a Transverse Electromagnetic (TEM) transmission line structure similar to a microstrip line. In the simplest case, the stripline circuit is developed on one substrate layer and covered with another that has been stripped of copper on one side, effectively sandwiching the circuit within two layers of dielectric and earth plates, as illustrated by Figure 2.4.

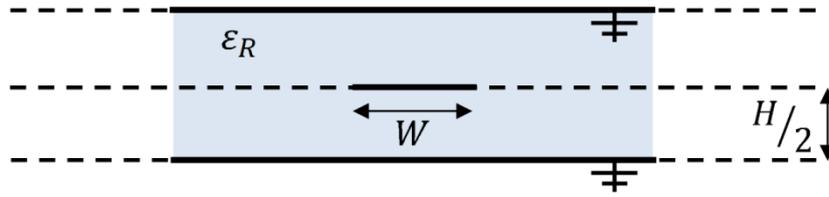


Figure 2.4: A representation of a simple stripline-filter structure. Note that the outer plates are earthed, the substrate (shaded) has a relative permittivity of ϵ_R , the total height is H and the line width is W .

Some advantages of using striplines in filter designs, as opposed to microstrip, are better bandwidth performance and improved symmetry in Band Pass Filter (BPF) roll off. In addition, striplines are less dispersive due to the field being wholly contained within the substrate (full TEM behaviour). This is often described as having a unity filling factor and is similar to the properties of wave propagation in coaxial transmission lines. There are also disadvantages to using stripline: a more complex production method is required and the width of traces in stripline circuits is lower than for equivalent microstrip traces. This can lead to problems related to increased design constraints and lower accuracy due to the production process of the circuit [3].

A stripline filter is traditionally conceived through a stepwise process. Although there are multiple methods, a preferred procedure is presented here. The first step is the derivation of a lumped-element Low Pass Filter (LPF) with the desired characteristics using the insertion loss method. This involves defining the filter response in terms of its input reflection coefficient, which in turn is derived from its input impedance. The reflection coefficient is included in the Power Loss Ratio (PLR), which provides an insightful equation for visualising a filter's response and is shown in figure 2.5. Taking into account the causal properties of passive networks, it can be shown that the PLR is an even function in ω and a particular form can then be assigned to it. Equations 2.11 and 2.12 define the reflection coefficient and PLR, respectively, as presented in Pozar [1].

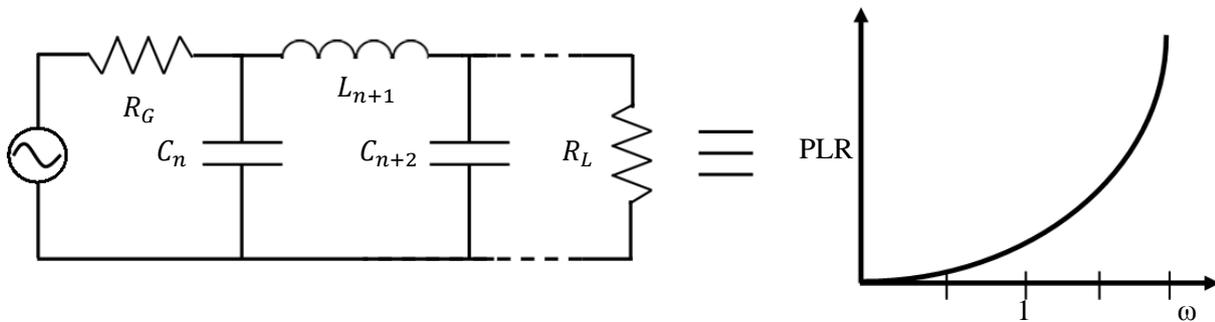


Figure 2.5: Schematic of a normalised, generic LPF circuit (Left) and its PLR (Right) response showing loss against frequency.

$$\Gamma(\omega) = \frac{Z_{in} - Z_0}{Z_{in} + Z_0} \quad (2.11)$$

where Z_0 = characteristic impedance, Z_{in} = input impedance and Γ is the reflection coefficient

$$\text{PLR} = \frac{\text{Power from source}}{\text{Power at load}} = \frac{1}{1 - |\Gamma(\omega)|^2} = 1 + \frac{M(\omega^2)}{N(\omega^2)} \quad (2.12)$$

where M and N are even polynomials in ω .

Since the PLR is related to the structure of the filter through the input impedance, the lumped-element values are determined by expanding the PLR to include them and solving for the constraints introduced by the form

of the function chosen and its evenness. This has been done for various functions, which realise particular filter characteristics, and the values for normalised circuit elements corresponding to these are tabulated.

A lumped-element filter is transformed into a distributed equivalent through the Richard's transform, which maps from the ω plane to the θ plane. This enables the transformation of inductors to series shorted and capacitors to parallel open TEM-lines. Furthermore, the characteristic impedance is adjusted to match the element values and the length of the lines to the desired frequency response, as described by equations 2.13 to 2.15 [1]. Note that the transformation is only equivalent at the design frequency. For a given substrate geometry, relative permittivity (ϵ_R) and loss tangent, the characteristic impedance of a trace is logarithmically related to its width, i.e. the impedance increases with a decrease in width and vice versa.

$$S = j\omega = j\tan(\beta l) = j\tan(\theta), l = \frac{\lambda}{2} \text{ at } \omega_c \quad (2.13)$$

$$jX_L = j\omega L = jL\tan(\beta l) = jL\tan(\theta), l = \frac{\lambda}{2} \text{ at } \omega_c \quad (2.14)$$

$$jB_C = j\omega C = jC\tan(\beta l) = jC\tan(\theta), l = \frac{\lambda}{2} \text{ at } \omega_c \quad (2.15)$$

where β = lossless propagation constant, ω_c = cut-off frequency, L = lumped inductance, C = lumped capacitance and l = line length.

The distributed equivalent will be significantly different to that of the lumped-element structure, as their frequency dependent behaviour is governed by different equations. An ideal lumped element reacts according to complex impedance characteristics, but the distributed structure will display unique transmission line properties as determined by a combination of the substrate and line variables. Mapping the lumped-element filter response to its distributed equivalent reveals the periodic nature of the latter which can be utilised in design. The transformed LPF characteristics display repeating passbands that provide BPF operation at all frequencies that are harmonically related to the baseband response.

To clarify, the relative bandwidth of the distributed filter is related to that of the cut-off of the lumped-element filter, but the frequency response of the distributed filter depends on the electrical length chosen for the lines. Thus, a distributed LPF or BPF can be realised from a lumped-element design and the bandwidth controlled by scaling the impedances. Equations 2.16 and 2.17 demonstrate the scaling of inductors and capacitors, respectively, as presented in Pozar [1].

$$L' = \frac{R * L}{\omega_c} \quad (2.16)$$

$$C' = \frac{C}{R * \omega_c} \quad (2.17)$$

where L' = adjusted inductance and C' = adjusted capacitance.

Lastly, the distributed filter must be changed to a practical topology, since series TEM-lines are not practically possible in stripline structures and parallel lines require transmission line sections that separate them. This is accomplished by adding unit elements to the design and using the Kuroda identities to transform series structures to parallel equivalents. A unit element is a section of TEM-line with a characteristic impedance matching the design, therefore only adding delay without changing the magnitude response. The Kuroda identities are a group of transforms that provide parallel equivalents to series structures and the use of the first identity is depicted in figure 2.6.

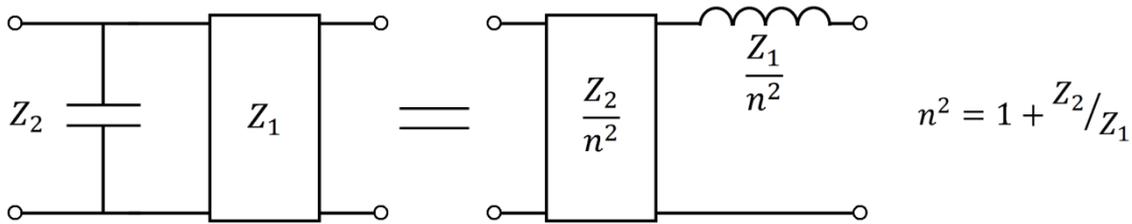


Figure 2.6: A graphical representation of the first Kuroda transform (after [1]). Note that the relevant equations are formulated for impedance values.

Computer simulations greatly simplify the design process by verifying that distributed equivalent circuits respond desirably and also enabling the designer to utilise the transmission line properties of striplines to achieve better filter performances. This is helped by the use of powerful optimisation algorithms that are built into design software. By allowing the user to enter the required filter performance, variables are trained for an optimal response with the given material properties whilst remaining within TEM propagation and production constraints. However, such optimisation tools are limited to closed form simulations of TEM structures and verification of the performance of such an analysis should be conducted in a full EM solver.

2.5 Electromagnetic compatibility [23]

EMC studies are concerned with the interaction between electronic systems, their environments and, in particular, any RFI that may result from such behaviour. The interference is typically called coupling and propagates as conducted electric current at low frequencies (< 30 MHz) or radiated EM fields (> 30 MHz) at high frequencies [23]. Note that this is defined relative to the source of the interference and that radiated energy can also induce currents. The classification of radiated coupling is dependent on the field component causing the interference, which is inductive for magnetic and capacitive for electric field coupling. This terminology is related to the models used for analysing the process at relatively low frequencies. Nonetheless, the interaction can be very complex at shorter wavelengths and interference propagated through the ambient medium can simply be referred to as radiated. It requires in-depth investigations to characterise these processes across a broad bandwidth, even for systems with simple geometries.

Any interference is related to the geometries of the systems involved, i.e. the presence of current paths and conductive apertures. In this context, a current path is any closed, galvanically-connected loop that will allow the flow of electrons if excited in some way. There are several simple and practical measures that can be implemented to reduce the likelihood of EMC related RFI. Firstly, separating systems can reduce issues, as EM-field strength and density are spatially dependent. Magnetic-field coupling can be reduced by twisting conductor pairs and minimising the coupling area, resulting in the cancellation of the induced voltages in subsequent loops. Furthermore, configuring current paths to be close to a metallic ground-reference structure also reduces the coupling area and perturbs the magnetic field, further reducing interference related to this component. The use of individual earth-return conductors avoids resistive coupling via common ground paths. The electric-field interference mostly requires the application of conductive shielding for reflecting the field. This can take on many forms of varying complexity and effectiveness.

These measures are important, although they may only be effective over a portion of the bandwidth of concern in this document. Note that these preventative actions serve a dual purpose, since reducing the RFI susceptibility of a layout also lowers the level of interference it produces. The RFI susceptibility of a

structure is typically provided by the shielding effectiveness. This parameter is used to characterise the ability of an enclosure to reduce interference and is typically parameterised for a range of frequencies. Another important definition is for Common-Mode (CM) current, which is the net current and is common to both conductors in a dual-conductor configuration. Figure 2.7 is an illustration of CM-current flow on a coaxial cable and the use of an L-plate structure to constrain the CM-current loops, which is commonly implemented as part of EMC mitigation.

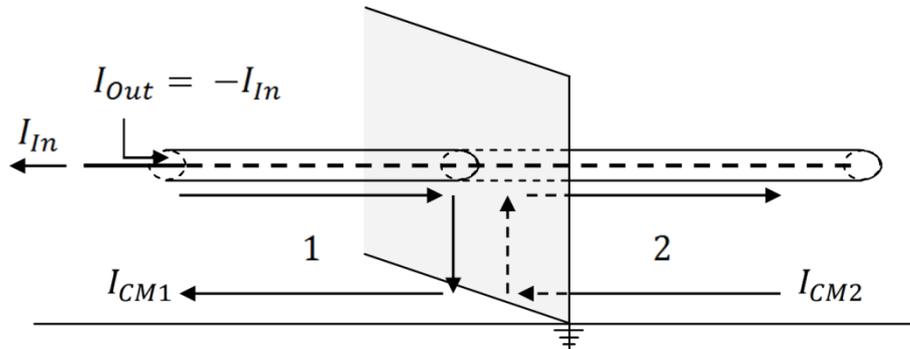


Figure 2.7: Illustration of CM-current flow on a coaxial cable with L-plate structure. Note that the plate limits the current loops. I_{In} is the current on the inner conductor, I_{Out} is the current on the inside of the outer conductor, I_{CM1} and I_{CM2} are the CM currents flowing on the outside of the outer conductors. The latter currents are induced and confined to areas 1 and 2, respectively. Note that the currents inside the coaxial cable are equal in magnitude, but that the presence of I_{CM1} results in a net current on the coaxial structure.

To summarise, EMC issues can present complex problems and require in-depth analyses at high frequencies. This makes measurements invaluable to EMC investigations and most self-induced interference can be mitigated once unwanted currents are identified. Lastly, measurement configurations can also be subject to or introduce RFI into a system and the configuration of measurement paths must be carefully considered to avoid additional problems [21]. The EMC principles described here were used as guiding principles throughout developments described in later sections.

2.6 RATTY 1

The first RATTY system consists of three sections: 1) an analogue FE, 2) a digital unit and 3) a software package. It can output TD or FD data by functioning as either a transient analyser or an RF spectrometer, respectively. The digital and software parts are distinguished from each other, since they are physically separated units. This is also true of the analogue and digital stages, though the analogue parts are housed in the digital enclosure. A basic block-diagram of the RATTY 2 system is shown in figure 3.1, which is identical to the RATTY 1 layout.

The analogue FE is simplistic, consisting of two amplifiers and a 31.5 dB, adjustable attenuator. The second amplifier follows after the attenuator in the signal path and both are integrated parts on the ADC integrated circuit (IC). This combination enables the gain to be adjusted, thereby increasing the range of input signals processable with the ADC, as explained in section 3.1.1. An LPF limits the bandwidth to 828 MHz to avoid aliasing at the ADC sampling rate of 1.8 Giga-samples per second (GSa/s). There is a spectral margin retained here for the filter roll-off to reach an out-of-band rejection ratio greater than the ADC SNR. The

input is guided via a semi-rigid, coaxial cable from the back panel of the enclosure to its front, where the available space is used to mount the analogue components.

The digital unit implements a ROACH 1 platform with an 8-bit, 1.8 GSa/s ADC as input to it. A ROACH 1 board consists of three main subsections: 1) the main Field-Programmable Gate Array (FPGA), 2) a PPC* Central Processing Unit (CPU) and 3) a monitoring FPGA, subsequently referred to as the management subsystem. The FPGA performs all the signal processing and data acquisition in the digital system, whilst the CPU facilitates the interfacing between the FPGA and a controlling computer. The latter is used to run the software package and the intercommunication uses an Ethernet connection. An SD card was used as the storage capacity for the FPGA programming and the CPU also performs this task upon initialisation. The management system monitors the power supply unit's outputs and gauges the temperature of several core parts. An automated power-down procedure is executed if any of these signals reach hazardous levels. It is noteworthy to mention that the ADC used for RATTY 1 is an interleaved component, which results in more spurious products in digitised spectra than present in single ADC sampling.

The samples of the digitised input signal are transmitted to the FPGA via a high speed connector and buffered in internal block RAM. This buffering procedure can be controlled during TD measurements by specifying a trigger level, enabling oscilloscope-like functionality for capturing transients. During TD operation, the unprocessed ADC samples are read from the FPGA. The FD output is an average, accumulated spectrum generated with a 2^{15} -point, polyphase filter bank implemented on the FPGA. This realises a spectral resolution of 54.8 kHz.

A python-based software package was developed for controlling RATTY via a standard computer with a Linux based operating system. The programming comprises three main scripts that initialise the setup of RATTY and controls either the TD or FD functionality, respectively. Several other scripts and libraries developed for the KAT are also utilised for communication protocols and general interfacing with ROACH systems. The software is used for storing and monitoring the data during measurements. For TD measurements, a generic python FFT provides spectral output of transient events. This and the FD spectra are calibrated for the FE configuration and system response in software. All three subsystems are described in more detail in section 4.3, since the original layout was adapted and implemented in the RATTY 2 system.

TD record length	8.192 μ s
Spectrograph resolution	54.8 kHz
ADC SFDR	~40 dBc
Adjustable gain	~31.5 dB
Adjustable range (ADC SFDR + adjustable gain)	~80 dB

Table 2.1: The key specifications of RATTY 1.

A RATTY 1 system has been deployed and used as an on-sight monitoring system at the SKA core site.

* Performance Optimization With Enhanced RISC (POWER) – performance computing

Chapter 3

RATTY 2 front-end development and analysis

3.1 Practical considerations for the analogue subsystem

The development of RATTY 2 is based on its predecessor, RATTY 1. Three subsections can be distinguished and are shown in figure 3.1. The analogue Front End (FE) is the first subsystem of the RATTY whole that an input signal will enter after being captured by either an electric-field (E-field) or magnetic-field probe. Due to the nature of system noise and sampling requirements, analogue signal conditioning prior to digitisation is paramount to the development of RATTY. The desired performance characteristics include maximised bandwidth, sensitivity, dynamic range and reconfigurability. These goals compete for priority, since choices made to improve one property may impose limitations on the performance of the others. Before proceeding to the design, it is instructive to speak about some practical considerations.

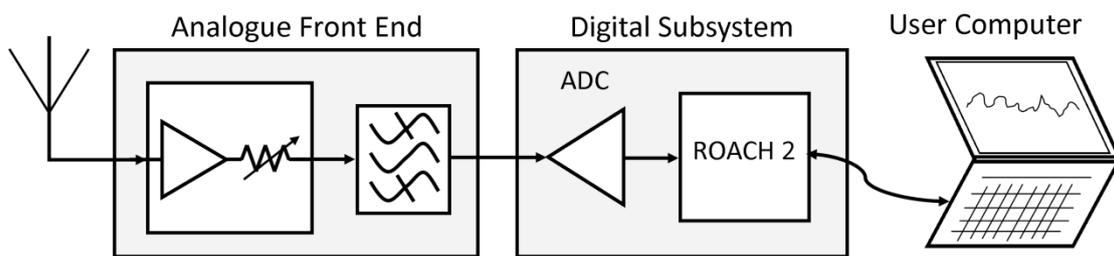


Figure 3.1: A condensed representation of the RATTY 2 layout showing the three subsystems. Note the directions in which data flows, which is bidirectional for the control from the user computer.

3.1.1 Analogue functionality

There are various strategies that can be implemented for signal conditioning, although two processes can be highlighted and must be performed: level setting and frequency conditioning. These form functional blocks in the FE and are designed to meet the characteristics of the ADC. This is evident from the system layout shown in figure 3.1, as the ADC follows directly after the FE. Furthermore, the requirements of the MeerKAT project provide performance goals that the design must pursue.

Level setting is required to amplify low level signals and maintain the appropriate input power into the ADC. This determines the instantaneous sensitivity and adjustable dynamic range performance of RATTY. The amplification control is realised through a combination of amplifiers in the FE, which yield a maximum gain, and adjustable attenuators. Thus, the full gain can be decreased by increasing the attenuation. This functionality enables the setting of maximum and minimum detectable input levels during measurements, hence the term adjustable dynamic range. The minimum detectable level can be increased beyond this limit if the ADC output is accumulated and averaged.

Frequency conditioning involves limiting the bandwidth of an ADC input to comply with the sampling criterion. This could also include frequency mixing, which is required if the input frequency range falls across an aliasing axis, i.e. a harmonic of half of F_s , or is out of the ADC's input bandwidth. Although the sampling theorem is derived classically for baseband sampling, i.e. the input's spectral content is limited to below half of F_s , the requirements are more general than this. A signal can be digitised without the occurrence of aliasing, if its spectral content is limited to any of the Nyquist zones. When a baseband spectrum is sampled, its aliased mirror image in the second Nyquist zone is reversed, due to the convolution of the baseband spectrum and the impulse train periodic in F_s . Likewise, if an even numbered Nyquist zone is sampled, the aliased version of the input spectrum in the baseband will be reversed. This is shown in figure 3.2.

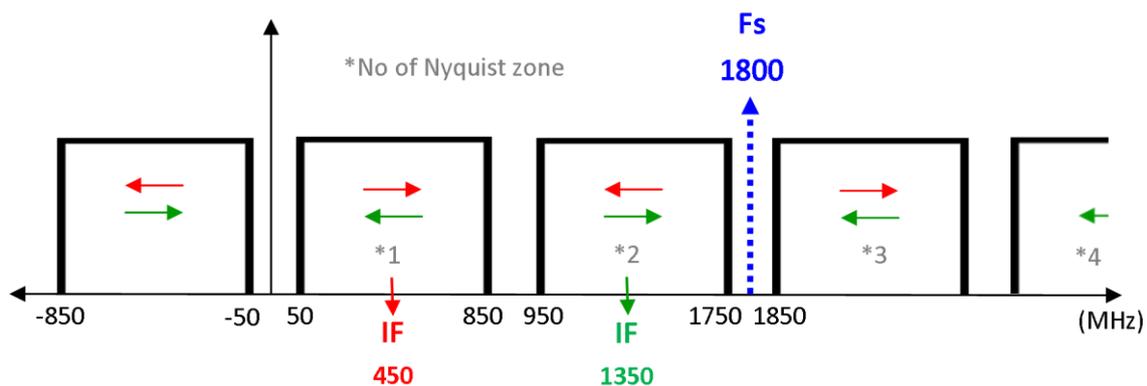


Figure 3.2: Schematic illustration of sampling in different Nyquist zones. Note that the orientations of sampled spectra are shown for odd (red) and even (green) numbered Nyquist zones.

An inverted spectrum in the baseband, where points decrease in frequency along the horizontal axis, is only a processing issue and can be easily corrected. Since there is no mixing involved, a sampling strategy that utilises several Nyquist zones is referred to as a direct-sampling strategy. This gives rise to the possibility of two different strategies for frequency conditioning: 1) Frequency mixing or 2) Direct sampling.

3.1.2 ADC specifications [7]

Since the FE design is guided by the ADC specifications, an introduction to the RATTY 2 ADC is required. Its key characteristics include the sampling rate, dynamic range, input power levels and input impedance performance. These characteristics are translated into design constraints and guidelines. For example, the upper limit of the sampling rate determines the maximum nominal bandwidth that can be processed by the ADC, in accordance with the Nyquist criterion. Also, the SQNR sets a requirement for the out-of-band suppression that the anti-aliasing filters must provide; the analogue input power levels determine how much gain is required to achieve a desired sensitivity and the useful bandwidth of the input impedance determines the upper limit for input spectra.

A component that was central at the start of the RATTY 2 development was a 10-bit, monolithic ADC claiming a 2.2 GSa/s maximum sampling rate. However, a current IC layout with this part has only achieved a rate of 1.8 GSa/s. The maximum nominal bandwidth in the design was thus set at 900 MHz, which affects the planning of the frequency conditioning strategy. Sampling in the ADC's baseband requires only that the input spectrum is limited to below the nominal bandwidth by appropriate filtering. However, sampling the

higher spectral ranges necessitates employing either: 1) frequency mixing, or 2) isolating specific bands within the Nyquist zones for direct sampling using filters.

The 10-bit sample resolution provides a measure of the realisable SQNR. Although the relevant SQNR is dependent on the input power density distribution, single-tone characteristics are used as the standard for characterising ADCs (2.1). Following this standard, the theoretical SQNR limit is ~ 62 dB, although the realised SFDR is only 52 dB for the worst case. This SFDR figure could only be realised with a monolithic ADC.

Although cannot be completely avoided due to the causal limits of realisable filters, a specification can be set for the suppression filters must have beyond their passbands, i.e. out-of-band suppression. This goal is derived from the SQNR figure, seeing that aliasing below this level is not detectable without accumulation. The out-of-band suppression sought in the RATTY design was 60 dB, which approximated the theoretical SQNR.

The maximum input power of the ADC is limited to -2 dBm (250 mV peak sinewave). Subtracting the worst-case SFDR of the ADC from this upper limit provides the minimum effective range of the ADC. By amplifying the input, weaker signals can be shifted into this range, i.e. -2 to -54 dBm. This sets the target for the level setting sections to provide enough gain to achieve a desired sensitivity.

The ADC has a 3dB nominal input bandwidth of 3.3 GHz, which allows the direct sampling of inputs to this limit. However, its response is usable to 3.6 GHz, enabling the use of the first four Nyquist zones at maximum F_s . This could be seen as a limit on the input bandwidth achievable in a direct-sampling system, whilst a mixing system bandwidth is constrained by other factors, such as the range of the LO source or amplifiers.

3.1.3 MeerKAT requirements

An RFI measurement system document was compiled during the early stages of the design as a goal-setting guide for the project. It can be found in Appendix G.1 [15]. This document presents the required measurement sensitivity needed for the verification of MeerKAT subsystem emissions, which was calculated according to industry standards. Several RATTY parameters could only be estimated at this stage, which include the channel bandwidth and input antenna gain, and these determine the system noise temperature. The document also recommends the correlation of two independent, simultaneous measurements for a further reduction in the measurement NF. This was later considered to be an excessive endeavour, since the NF is dominated by the ambient temperature and the improvement with correlation would not be warranted.

The requirements were very demanding and the predicted performance characteristics for RATTY were ambitious. Nonetheless, the document set a long term goal for the sensitivity of the device at approximately -118 dBm, which was used as a rough projection of the required gain for the FE. Since the total gain needed was in excess of 110 dB, this was not considered for the initial designs described in the subsequent sections. The philosophy behind this decision was that an active antenna would be used and that the system would be reconfigured to add gain if extremely sensitive measurements were attempted. Due to unforeseeable developments, several other requirements were also not realised. This includes the self-generated RFI, bandwidth and remote operation goals. Thus, the document serves as a reference for future developments and the 70 MHz – 3 GHz bandwidth was considered in the selection of components to facilitate this.

3.1.4 Analogue component selection criteria

The FE is required to deliver a stringent performance criteria and this is inevitably constrained by practical realities. With a design environment unknown to both teams and project timelines introduced, an early decision was made to procure connectorised components for the FE construction. This simplifies design reconfiguration after an experimental model is built, as opposed to redesigning and reconstructing a PCB implementation with surface mounted parts. Also, a connectorised layout allows the use of current probes around interconnecting cables, which is advantageous for RFI debugging. Conversely, the connectorised approach limits the available selection of parts and the layout is typically bigger. Subminiature Version A (SMA) connectors and cables are used throughout the analogue system.

The flatness of a component's forward transducer gain is important, since the effect of cascading several parts with variation across their operational ranges causes a complicated system response. This reduces the overall SFDR, because different spectral ranges are not equally amplified. The spurious products from higher input levels with more gain could obscure small, real components undergoing less amplification. Additionally, the sensitivity of the system varies along with its gain.

The input impedance characteristics of all components were evaluated as an important criterion, to avoid standing-wave effects in the transmission lines connecting components. For this reason, 3 dB attenuation pads are used to buffer adjacent components that are not well matched across the operational bandwidth .

Some components were already available and well suited for use in RATTY. Hence, several design choices were steered towards incorporating these parts, which include filters used in MeerKAT systems and a Valon controllable synthesiser. The latter was selected as an LO and ADC clock-signal source. The selection of the remaining components formed a key part of the investigation.

3.2 Front-end design exploration

With the framework set according to the ADC specifications and project requirements, a goal-driven design of the FE commenced. From the preceding sections it should be apparent that several signal paths are necessary in the FE, each dedicated to a spectral range and collectively providing continuous coverage. In order to make a clear distinction between the digital and analogue subsystems, the term paths will be used to refer to the analogue section, as opposed to the term channels used for the frequency bins in the Digital Signal Processing (DSP).

Several core designs were explored to test the possibilities of different strategies within performance and component constraints. These developments are documented chronologically and concepts are introduced when relevant to the context. Each model was evaluated expediently, with the anticipation of making rapid progress and to refine the strategies towards realistic implementations. The spectral coverage, complexity, availability of components and spurious rejection of the proposed topologies were considered. The outcome provided insights for the final two system configurations. Note that a set of components was compiled for each strategy and used for the assessments.

Although the final FE design is to be implemented in a parallel system with dual analogue sections, each providing an input to one of two ADC's, the layouts are presented for a single ADC. A dual input system

forms part of future work and will make measurements with dual polarisation or the combined bandwidth of two signal paths possible.

3.2.1 RATTY 2.1: Single-mixer FE topology

An initial design proposal for the FE features a dual-path layout featuring a baseband and single mixing strategy, as shown in figure 3.3.* The Low Noise Amplifier (LNA) is the first component in the signal path and contributes most significantly to the noise performance of the system. In addition to a low noise figure, this component should also have high gain and a well matched input impedance, both constant across the system bandwidth. The combination of a low noise figure and high gain minimizes the amount of system noise, as predicted by eq. 2.5. A good input match limits power loss resulting from reflection and dominates the input characteristics of the FE. However, signal degradation occurring prior to the LNA, such as noise introduced by a lossy cable, will dominate a system's performance, which emphasises the importance of the LNA placement in a measurement chain.

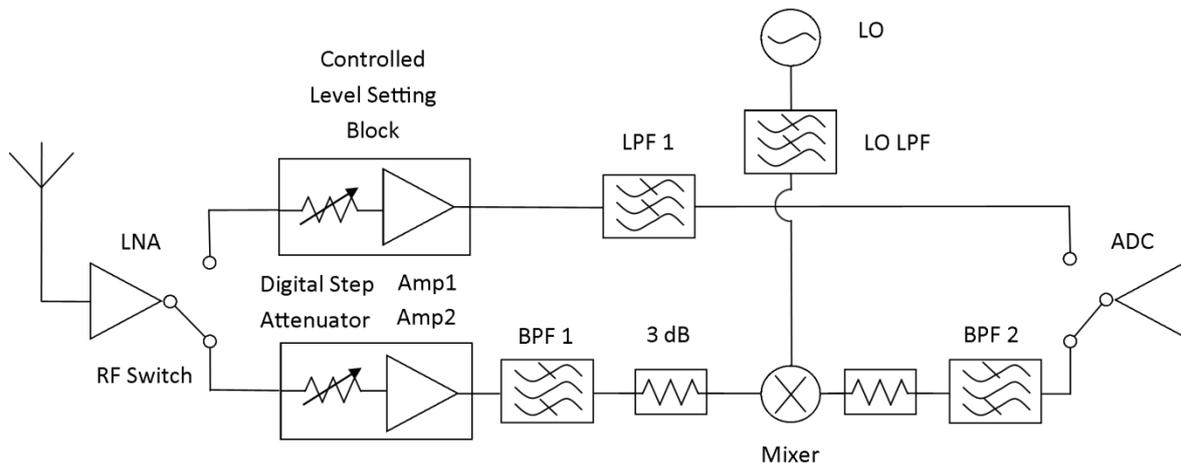


Figure 3.3: RATTY 2.1 design schematic.

A resistively terminated switch follows the LNA and enables the selection of the desired frequency path. The isolation provided between inactive and active ports is very important in limiting signal leakage, since the switches are essentially used to select filters needed to provide the out-of-band rejection. As switches are passive, they can be reversed and used to form multiple signal paths in pairs. The combined isolation provided by a pair should not lead to leakage above the filter suppression (60 dB).

The maximum input power of the system is matched to the upper limit of the LNA, since LNAs typically cannot deliver high output power, maintain a large bandwidth and have good noise performance. By placing attenuators before the additional amplifiers in the level-setting blocks, these can be kept from compression. However, if a mixer with low power handling capabilities is selected, it could be the first component to compress.

A quick analysis of the FE gain and input limits was done according to single valued gain or attenuation levels, ignoring any mismatches or cascading effects. These values are extracted from the datasheets of selected components and a guesstimate made where the component loss is unknown. In order to maintain

* RATTY 2.1 strategy proposed by Jason Manley, SKA-Africa.

equal dynamic range and sensitivity across the system bandwidth, amplifiers are selected to provide roughly equal amounts of gain for all paths.

The LNA places a limit on the bandwidth of the design with the lower-frequency cut-off of the design set to 50 MHz. However, the LNA upper limit of 3 GHz will only affect future developments.

The baseband path LPF 1 is a MeerKAT filter with a 0 – 828 MHz passband and >45 dB out-of-band rejection from 1 GHz to 1.9 GHz. It requires additional filtering to fulfil the 60 dB specified, particularly to suppress the upper frequency response of LPF 1, which provides little rejection. This emphasises that the rejection of out-of-band signals is obtained with the combined response of multiple components. Hence, components do not have to perform strictly according to the specifications, but the system response must satisfy the performance needs.

The mixing path down-converts the remainder of the system bandwidth into the second Nyquist zone of the ADC, using a controlled synthesiser as LO and setting it to several frequencies. Each of the LO settings represents an additional signal path and the system's range could be covered with three settings. Due to the specific requirements of the design, the BPF's are designed and produced in-house. BPF 1 is required to suppress content outside the 800 – 3000 MHz range, avoiding spectral leakage above 3 GHz into the second Nyquist zone and limiting the range of spurious products. A schematic depicting an ideal BPF 1 response convolved with the minimum LO frequency (2.45 GHz) during lower-sideband down-conversion and the first two Nyquist zones is shown in figure 3.4.

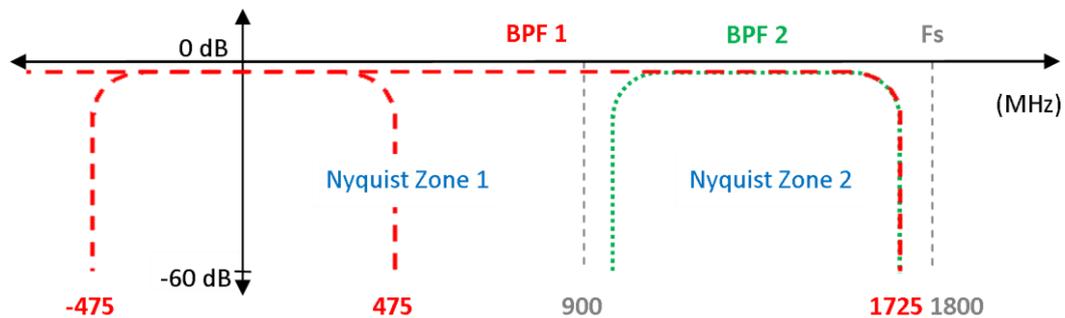


Figure 3.4: Schematic of an ideal BPF 1 response convolved with the minimum LO (2.45 GHz). This shows lower-sideband down-conversion and sampling of 800 – 1550 MHz. Note the BPF2 response and ‘folding’ of the mixed spectrum beyond 0 Hz.

In support of the first, BPF 2 provides the necessary anti-aliasing filtering for digitisation. To ensure that a design matches the criteria, a sufficient analysis must be performed to derive satisfactory results. The LO LPF is added to suppress harmonics of the modulating signal, which could leak into the signal path and increase the level of spurious products.

Although this design presents an attractively simple layout, it has several flaws that cannot be remedied. Since the mixer RF and IF bands overlap, the RF-to-IF isolation requirement for this part would be > 60 dB, which is not feasible. There are also spurious products present in the signal paths, which can be shown with an FD analysis. Furthermore, the bandwidth of BPF 1 may be difficult to realise whilst maintaining the required rejection ratio and filter roll-off. Appendix A.1 shows a schematic of the layout with single-valued gain and dynamic range analyses according to proposed components. A graphical representation of the FD signal processing is also included and shows selective filtering with varying LO frequencies.

3.2.2 RATTY 2.2: Dual-mixing FE topology

A FE design incorporating a baseband and dual-mixing strategy is depicted in figure 3.5. The baseband path is similar that of the first design, with additional gain added for increased sensitivity and LPF 1 moved to the end of the FE chain. In the mixing path, an input spectrum is first up-converted to a higher frequency and then down-converted to the ADC's baseband for sampling, referred to as stage 1 and 2 mixing, respectively. Moving the LPF 1 to the back of the FE utilises the part as an anti-aliasing filter for both paths. However, the design is vulnerable to spurious interference and the choices regarding the LO frequencies are important.

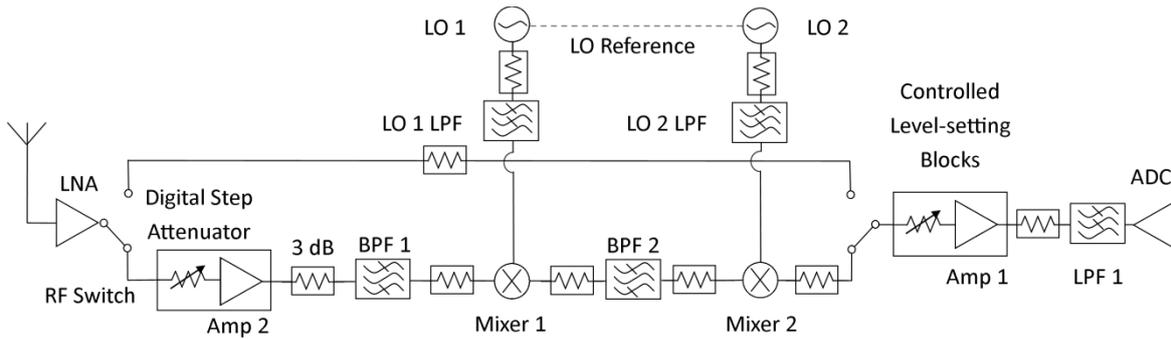


Figure 3.5: RATTY 2.2 FE schematic featuring baseband and dual-mixing paths. Note the buffering of mixer and filter ports.

The mixing plans were developed by assessing the LO and direct signal leakage as fundamental problems and then evaluating the spurious responses. This allows the identification of frequency ranges that avoid contamination from the LO, which is a high power signal, relative to low input levels. Although LO leakage is limited by mixer isolation properties, there are no components that perform sufficiently to maintain the desired SFDR and frequencies must be selected with care.

Using the lower sideband of the up-converted spectrum, harmonic leakage of the LO 1 source into the RF band is avoided.* If an RF band is selected at a range above the input frequencies and BPF 1 provides sufficient rejection, there is no problematic leakage into the mixed band. Figure 3.6 is a normalised, FD representation of the lower-sideband up-conversion in stage 1 mixing for the second path.

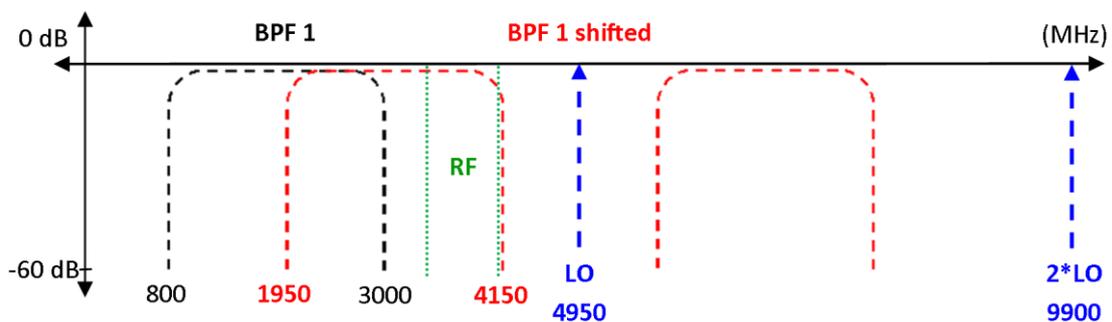


Figure 3.6: Lower sideband up-conversion with LO1 and second harmonic. The ideal BPF 1 response (black), mixed spectrum (red) and RF (green) ranges are shown. Note the first two LO harmonics (blue) for sampling the 800 – 1550 MHz band with the RF centred at 3.75 GHz.

*Design input from Adriaan Peens-Hough, SKA-Africa.

The LO 1 products that result from up-conversion must be sufficiently suppressed prior to stage 2 mixing, to prevent these from featuring in the output spectrum as aliased products. Such aliasing is limited by the combined response of mixer 2, amplifier 1, LPF 1, and BPF 2. A commercially available Valon synthesiser can be used as LO 2 source, if upper-sideband down-conversion is employed with mixer 2. The synthesiser also provides a reference signal to lock the two LO's for improved stability. A normalised, FD depiction of the down-conversion process showing the down-converted spectrum shifted into the negative baseband frequency range and 2nd LO harmonic is illustrated in figure 3.7.

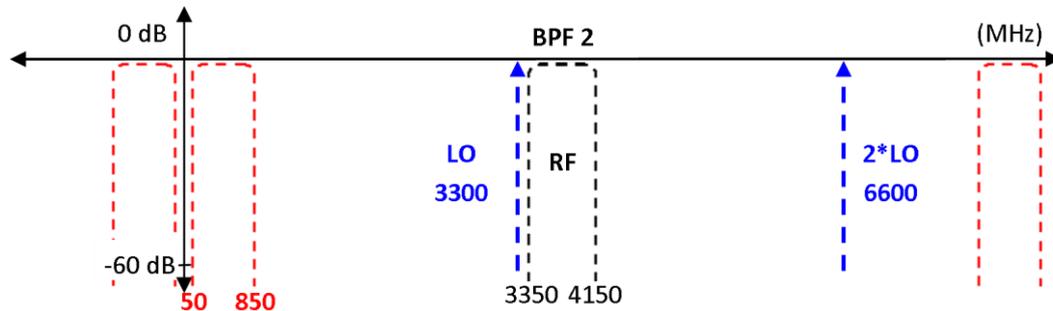


Figure 3.7: Upper side-band down-conversion showing LO 2 and related harmonics. Note the output band of stage 1 (black) and stage 2 (red) mixing with LO harmonics (blue).

A more comprehensive FD analysis of the design reveals that there are mixer related spurious products in the RF and baseband spectra, as a result of up-conversion and down-conversion, respectively. At the maximum input power, i.e. the LNA IP1dB, some products will be within the SFDR specification. Note that this is representative of the worst case of spurious contamination and that the attenuation of components separating the LNA and mixer are included. Using the spurious table of a proposed mixer component and adjusting it for higher input power, the levels of unwanted products within the sampled band are calculated. The highest spurious levels are listed in table 3.1.

The level-setting in this design featured more gain and adjustable attenuation than present in the first layout. This is for increased sensitivity and dynamic range with amp 2 added to balance the gains of the paths. An analysis based on single-valued component gains and compression points indicates that the maximum input level to this layout is again limited by the output capabilities of the selected LNA. Appendix A2 includes a schematic with a gain and power range assessment.

This design indicates that mixer spurious products are problematic when broadband signals are being processed. However, if the input is up-converted to an RF frequency high enough for the bandwidth to be only a fraction of the centre frequency, the RF window can be devoid of the highest spurious components. This would require finding several components operating at frequencies well beyond the range of the proposed parts, which are readily available.

Moving the level-setting block in mixer path to a post-mixing position would reduce the bandwidth requirements of these components to the range of BPF 2. Also, the likelihood of spurious issues is lessened, due to the reduced bandwidth of the signals being amplified.

Lower sideband up-conversion spurious products in output spectrum	
Spurious product	Level (dBc)
$1*LO - 2*RF$	60
$1*LO - 3*RF$	67
$2*LO - 3*RF$	72
Upper sideband down-conversion spurious products in output spectrum	
Spurious product	Level (dBc)
$2*LO - 2*RF$	57
$3*LO - 3*RF$	65

Table 3.1: Peak spurious products calculated for RATTY 2.2. This is according to the performance of a proposed mixer at 0dBm input power.

3.2.3 RATTY 2.3: Hybrid FE topology

A hybrid design presents another solution for the RATTY FE, with a combination of mixing and direct-sampling strategies that provide continuous spectral coverage when used in succession.* The system features several signal paths that are selected with multi-port switches and an adjustable ADC clock source. A schematic of a single-input layout of this system is shown in figure 3.8.

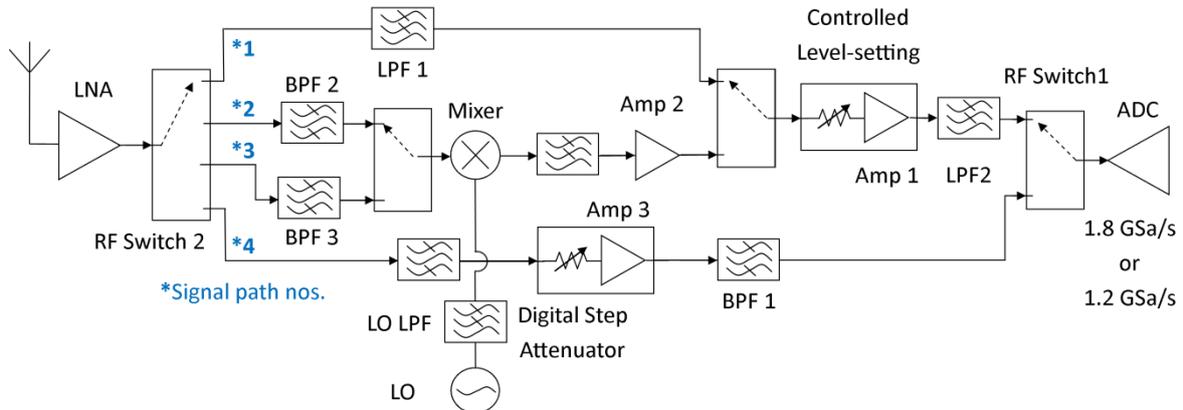


Figure 3.8: A single-input, hybrid FE design schematic. Note the dual sampling frequencies and omission of buffering attenuation pads.

The first three signal paths represent a single mixer topology in a revised form. The path at the top of the schematic covers the baseband, with the level-setting at the back of the signal chain. This enables sharing the anti-aliasing LPF 2 between the paths. The mixing path uses lower-sideband down-conversion for shifting the ranges of BPF 2 and 3 into the baseband as path 2 and 3, respectively. This scheme prevents LO contamination, but does allow several spurious products to fall within the output range. The passband ranges for the filters are listed in table 3.2.

*Design input from Emeritus Prof. P.W. van der Walt of our dept.

Filter	Passband range (MHz)
LPF 1	0 - 828
LPF 2	0 - 900
BPF 1	975 - 1725
BPF 2	1725 - 2475
BPF 3	700 - 1100

Table 3.2: Passband frequencies for the filters of the hybrid FE design.

By limiting the path bandwidth to 750 MHz or less, a minimum of 150 MHz is left as leeway for the filter roll-off to reach 60 dB out-of-band rejection. This limit in roll-off is determined to be a sensible figure, to pursue filter designs without necessitating the use of complex structures and materials with high order filters. The available roll-off margins are calculated assuming symmetric filter responses that are shifted into the centre of the Nyquist zone with a nominal bandwidth of 900 MHz. With a filter bandwidth of 750 MHz, a 75 MHz margin is left between the filter cut-off and Nyquist boundary. However, since aliased spectra are mirrored around the Nyquist zones, the margin for roll-off is doubled, as illustrated by figure 3.9. A LPF 2 is added to sufficiently reduce the LO, upper-sideband and other conversion products before strengthening the signal with Amplifier 2.

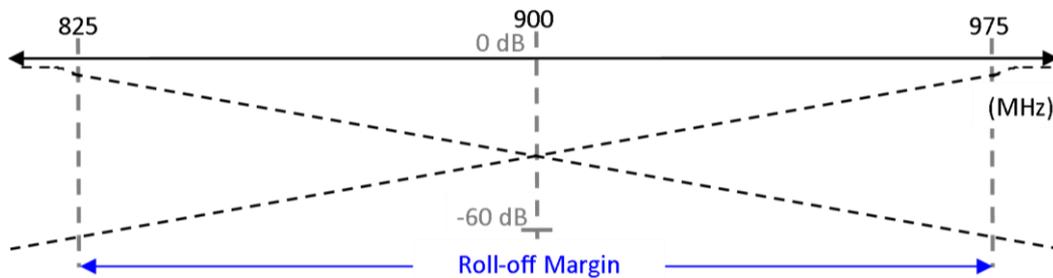


Figure 3.9: An illustration of filter roll-off margins showing an aliased response.

There remains a frequency range, 828 – 975 MHz, which is not covered by the first three paths. This band is not mixed, since the direct leakage of this spectral content is in the ADC's baseband at the maximum F_s . Therefore, this range is sampled directly in the second Nyquist zone at a lowered F_s with BPF 3 applied. By lowering the sampling rate, the second zone is shifted down until centred over the guard band, which decreases the bandwidth. This completes the spectral coverage of the design up to the upper limit of BPF 2. The direct-sampling path has a dedicated level-setting block covering its bandwidth. Note that the spectra of paths 2 to 4 will be reversed with regards to frequency in the baseband; this results from lower-sideband down-conversion or aliasing in the mixing and direct-sampling paths, respectively. The output must then be flipped and shifted to its original range in software.

This complex configuration with several switches and dual sampling rates must be controlled by the RATTY digital and software systems. The available control is limited by digital hardware and depends on the control implementation, since several strategies are possible. In this sense, it is best to implement the simplest solution. The design is also subject to mixer spurious interference in paths 2 and 3.

3.2.4 RATTY SU: Refined single-mixing FE topology

As a result of there being two strategies to follow for the FE designs, two approaches were pursued and developed into Experimental Development Models (XDM). Each XDM plan was undertaken by one of the

project teams and named accordingly: 1) RATTY SU and 2) RATTY SKA, respectively. These designs are also referred to as the mixing and direct-sampling systems, due to the approaches followed in the frequency conditioning. It is worth noting that only the FEs require distinct hardware, although the software must be uniquely configured to implement the FE control and data processing.

The RATTY SU incorporates a mixing strategy with down-conversion in upper frequency paths enabling baseband sampling. A schematic of a single-input layout for the design is shown in figure 3.10.

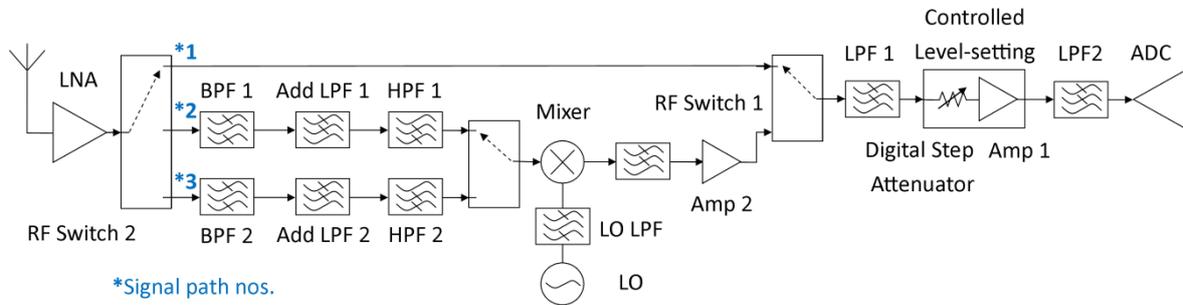


Figure 3.10: Schematic of a single-input layout for the RATTY SU design. Note that buffering attenuation pads are omitted.

The layout consists of a switched filter bank to select paths and a level-setting common to all settings, with the exception of amplifier 2 for gain balance. Both baseband filters now aid in anti-aliasing, as well as limiting the bandwidth before and after the final amplification. These also suppress up-converted mixer bands. Figure 3.11 shows the measured individual and simulated, combined passband responses of LPF 1 and the part selected for LPF 2.

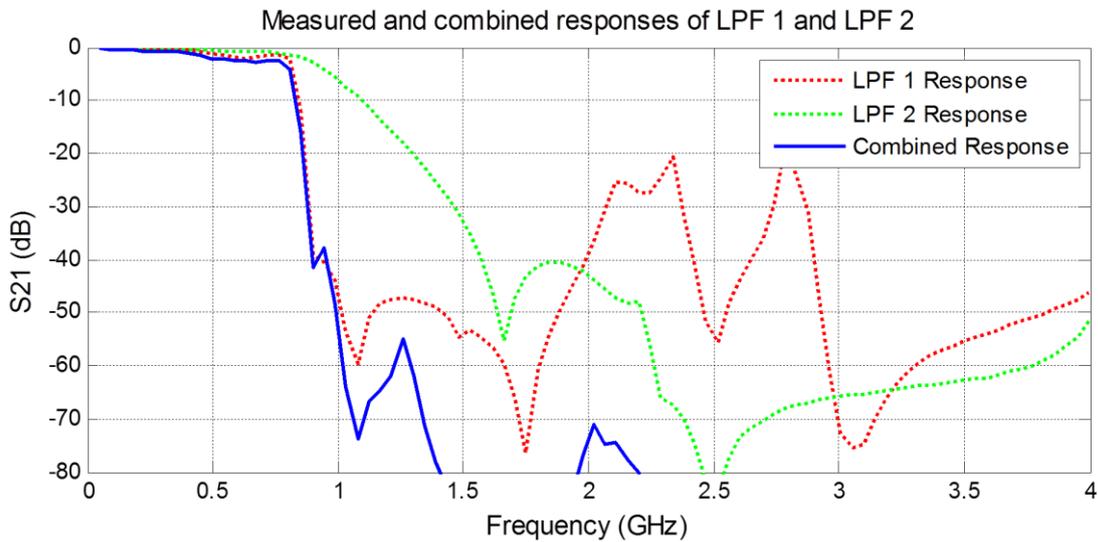


Figure 3.11: Individual and combined passband characteristics of LPF 1 and 2. Note that the individual responses were measured and that the combined response was simulated from this data.

There are LPF and HPF additions to paths 2 and 3, since the custom BPFs were developed as stripline implementations with inherent harmonic responses that require additional suppression outside the passbands.* Each filter bank now contains a chain of parts and the design goals for these are listed in table

*Choice of filter type and first design initiated by Prof. Petrie Meyer in our dept.

3.3. Note that the none-flat responses of other components also contribute to the system's out-of-band performance, but this is ignored for the filter specifications.

Signal path and filters	Passband range	Ranges with 60 dB out-of-band rejection (MHZ)
Path 1 LPF 1 + LPF 2	50 – 828 MHz	above 970 MHz
Path 2 BPF 1 + Add LPF 1 + HPF 1	800 – 1400 MHz	below 650 MHz above 1550 MHz
Path 3 BPF 2 + Add LPF 2 + HPF 2	1400 – 2150 MHz	below 1250 MHz above 2300 MHz

Table 3.3: Filtering specifications for RATTY SU.

In path 2, there is direct spectral leakage present in the baseband, since BPF 1 roll-off occurs below 800 MHz, leaving the filter response within the SFDR. Notwithstanding the leakage, this allows the input spectrum around the Nyquist boundary to be shifted and sampled by sacrificing a portion of the usable bandwidth. The detailed design constraints are documented in section 3.3.

The maximum input frequency to the FE in figure 3.10 is only 2150 MHz, contrary to the goal set at 3 GHz. This is based on a decision to develop the XDM models as proof of principle devices. Increasing the input range to the desired maximum would require the addition of another filter chain and tunable LO frequency.

An active mixer with an integrated LO square-wave amplifier was selected. Unlike most parts that require an LO drive > 7 dBm, this enables the use of the Valon synthesiser. The synthesiser can output 8 dBm, but it must be filtered and buffered at the mixer port; the resulting loss making it unsuitable for use with other components. With good, broadband matching, isolation and conversion loss characteristics, the mixer is also fit for developing a system with a larger bandwidth .

Several mixing spurious products fall within the output band. A layout free of the highest interference products can be realised, though it implies a significant reduction in the monitoring bandwidth and the addition of several signal paths. The frequency plan of such a design is demonstrated in appendix B. A FD representation of analogue signal processing of the RATTY SU XDM is included as appendix A.4, as well as a component list, single-valued gain and compression analysis. A detailed analysis of the system is documented in section 3.4.

3.2.5 RATTY SKA: Direct-sampling FE topology

A direct-sampling XDM was designed by the SKA-Africa team and is dubbed RATTY SKA for reference in this document. The design utilises multiple sampling frequencies in order to shift the Nyquist zones, which enables continuous bandwidth coverage without applying any mixing. It consists of a switched filter-bank and level-setting components, which is illustrated in figure 3.12 as a single input layout.

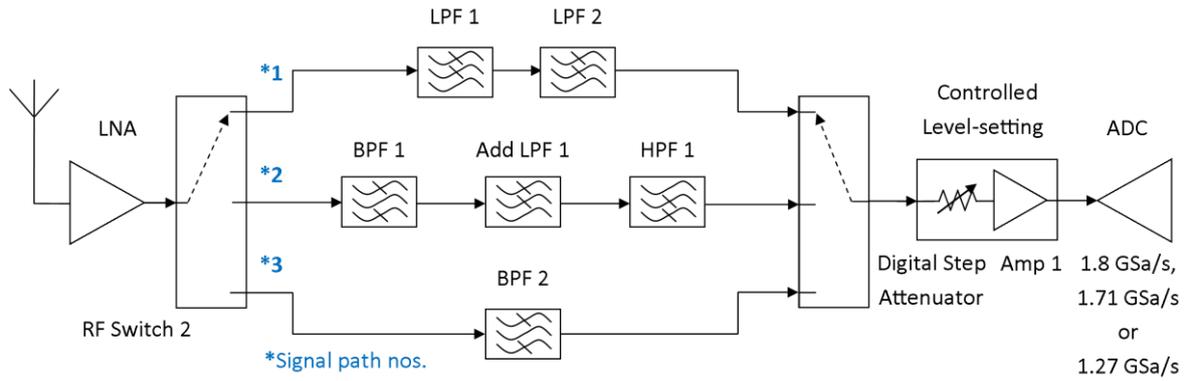


Figure 3.12: Schematic of a single-input layout for the RATTY SKA design. Note the multiple sampling frequencies and omission of buffering attenuation pads.

The path 1 baseband LPF 1 and 2 components are identical to the parts depicted in figure 3.10. Path 3 is dedicated to sampling in the second Nyquist zone and contains BPF 2, which is another MeerKAT filter that is optimised for use at an ADC clock frequency of 1.71 GHz. In order to sample the gap in spectral coverage left as a guard band at the higher sampling rates, the ADC clock frequency is lowered. For the proposed rate of 1.2 GHz, the second Nyquist window is used and the path 2 filters are designed accordingly. These form part of the stripline filter designs described in section 3.3. Table 3.4 contains the filtering requirements for the respective signal paths.

Signal path and filters	Passband range (ADC sampling rate)	Ranges with 60 dB out-of-band rejection
Path 1: LPF 1 + LPF 2	0 – 828 MHz (1.8 GSa/s)	above 970 (MHz)
Path 2: BPF 1 + Add LPF 1 + HPF 1	700 - 1100 MHz (1.2 GSa/s)	below 500 MHz above 1300 MHz
Path 3: BPF 2	900 - 1670 MHz (1.71 GSa/s)	below 812 MHz above 1758 MHz

Table 3.4: Filtering and sampling specifications for RATTY SKA.

The level setting strategy for the RATTY SKA layout is similar to the systems previously described. Although a detailed analysis of the direct-sampling XDM is not presented in this document, the assessments of the baseband path in the RATTY SU model are applicable to the RATTY SKA design and reference is made to it throughout the subsequent chapters. A component list is included in appendix A.

3.3 Analogue filter development

The main filters of the RATTY SU and some in RATTY SKA FEs are developed as stripline implementations. The decision is based on stripline properties, the complexity of the design and the production capabilities available. These processes are described in detail for the first filter with the focus on the requirements and results of the subsequent parts. The software used for the closed-form optimisation is [9] and the EM field simulations were run with [10].

3.3.1 RATTY SU filter designs

3.3.1.1 BPF 1

The first and detailed filter design is based on BPF 1 in the RATTY SU XDM (Figure 3.7). The goal is to realise the required out-of-band rejection with the minimum roll-off margin, thus maximising the bandwidth. This involves developing a normalised, sufficient-order, lumped-element filter and with a desired shape, which is then transformed into a stripline structure.* The distributed filter is then optimised using a software package that performs a goal driven analysis on the stripline configuration and alters its variables to provide and improved response.

Preceding the design process is the derivation of specifications for the filter application. BPF 1 must limit the input bandwidth in a manner that allows the sampling of the spectrum around the aliasing frequency at half the F_s . Figure 3.13 shows the relevant section of the system layout with its normalised FD response. The design constraints are derived for an ideal filter with linear roll-off. Note that the lower-sideband down-converted BPF 1 response (red) is adjusted for conversion loss and the RF input (blue) is attenuated by the mixer isolation, which increases the usable bandwidth.

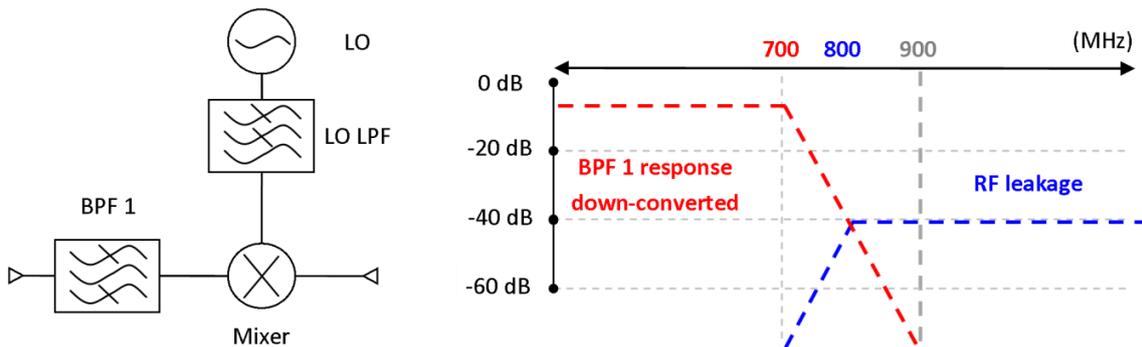


Figure 3.13: Schematics of relevant system components (left) and ideal system responses (right). The latter is used to define BPF 1 constraints of RATTY-SU path 2.

An 8 dB conversion loss and 40 dB minimum RF-to-IF isolation are assumed, as this is presented in the mixer datasheet. If the desired suppression is achieved with a roll-off of 150 MHz, the RF leakage should be 72 dB below the maximum power level in the mixer output. This leaves a margin of several dB; a sensible safeguard, since the datasheet predictions reflect the results of specific tests that vary with applications. The initial specifications of BPF 1 are summarised in table 3.3.

The element values of the generic, lumped-element filter (Figure 2.12) are chosen according to a Butterworth expansion, which translates to a PLR featuring a maximally-flat response. Since BPF 1 is the main filter in path 2, it plays a significant role in shaping the passband and must cause the minimum amount of variation in this range. The form of this PLR is shown in equation 3.8, as presented in Pozar [1].

*The design was initiated by Prof. Petrie Meyer of our department to accelerate progress. An alternative design is documented here to show understanding of the process.

Passband	800 - 1450 MHz
Down-converted passband	700 - 50 MHz
60 dB out-of-band suppression	Below 650 MHz and above 1600 MHz
Roll-off rate	60 dB / 150 MHz

Table 3.5: Initial BPF 1 specifications.

$$PLR = 1 + k^2 \left(\frac{\omega}{\omega_c} \right)^{2N} \quad (3.8)$$

where ω_c = cut-off frequency of the filter, N is its order and k is a constant.

A flat PLR predicts a roll-off of $20 \times N$ dB per decade, if defined for 3 dB corner frequencies, i.e. $k = 1$. This suggests that the required order is very high. However, since stripline filters typically provide better roll-off characteristics than single element implementations, a filter order of 4 was chosen. Figure 3.14 illustrates the transformation of the PLR from a normalised ω to a θ plane through equation 3.9.

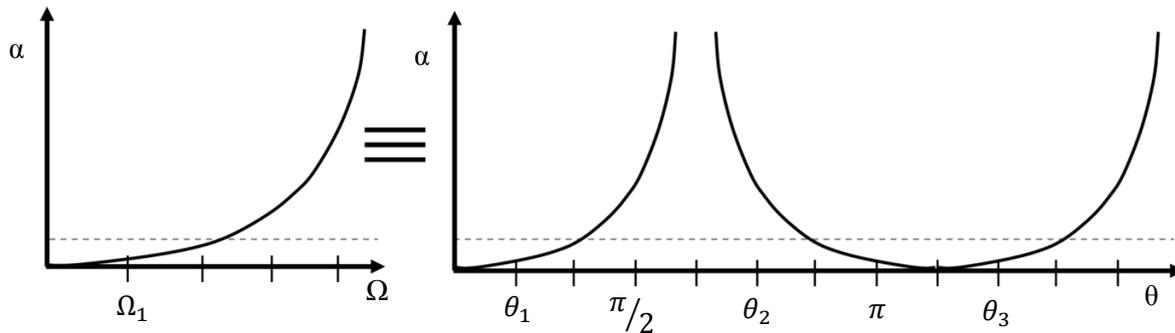


Figure 3.14: Schematic showing transformation of a normalised PLR from an ω to a θ plane. Note the BPF response shown on the right and the dashed line indicating the 3 dB loss margin (after [3]).

The range in figure 3.11 between θ_2 and θ_3 will be used for BPF 1. These variables are used to specify the bandwidth, by exploiting the repetitive nature inherent to the transform. This gives rise to equations 3.9 to 3.11 [2] with table 3.4 showing the calculation of the scaling factor used to adjust the bandwidth. The scaled element values are then implemented as the characteristic impedances of stripline sections.

$$\theta_1 = \tan^{-1}(\Omega_1) \quad (3.9)$$

$$\theta_2 = \pi - \theta_1 / \theta_3 = \pi + \theta_1 \quad (3.10)$$

$$\text{Bandwidth Ratio} = 2\theta_1/\pi \quad (3.11)$$

where Ω_1 and θ_{1-3} = corner frequencies of lumped-element and distributed responses, respectively.

Bandwidth ratio:	Scaling factor:
$650 / 1125 = 0.533$	$\Omega_1 = 1.28$
$\theta_1 = 0.9076$	

Table 3.4: Calculation of the BPF 1 scaling factor.

A practical stripline implementation is derived by adding the necessary unit elements to the generic structure and applying the Kuroda transform. In essence, series shorted lines, which are the distributed equivalents of inductors, are transformed into parallel stub lines. Secondly, single shunt lines are replaced with dual, parallel sections and all impedances adjusted to the 50 ohm system. Since capacitance adds in parallel configurations, and the admittance of lines are equated to it, the required impedance is doubled via parallel

sections. This results in decreased line widths and avoids very broad cross sections where parallel stubs are connected. If parallel lines are too broad, non-TEM behaviour may occur, which typically falls beyond the constraints of closed-form approximations. However, too narrow a line is also problematic and a compromise must be made. The initial and stripline filter structures are shown in figure 3.15. A detailed graphical representation of this transformation, the equations resulting from it, and the impedance values, are documented in appendix C.1.

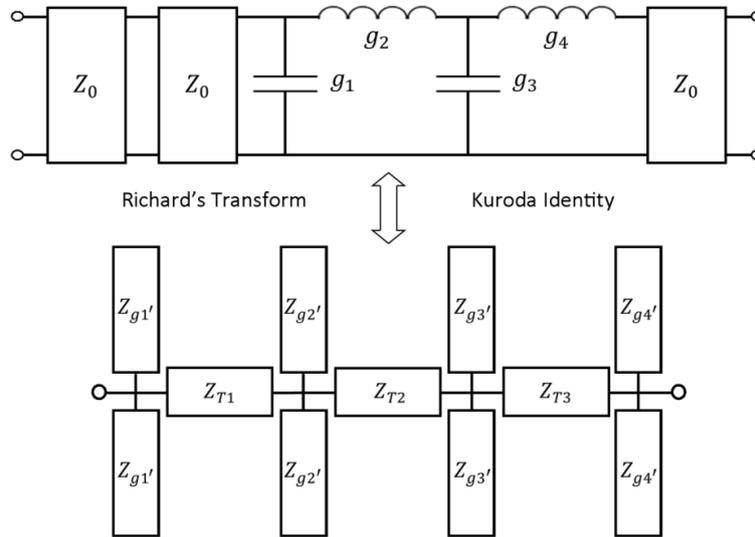


Figure 3.15: Schematics showing the derivation of a practical stripline filter. Note that the stripline impedances (below) are scaled element values of the initial filter (above) and that the lines are $\lambda/2$ long at the design frequency.

The first simulations incorporate ideal transmission line models for the stripline layouts, which are combined with lumped elements for testing the initial configuration. Figure 3.16 shows the responses of the lumped element filter, its stripline equivalent and the transformed structure.

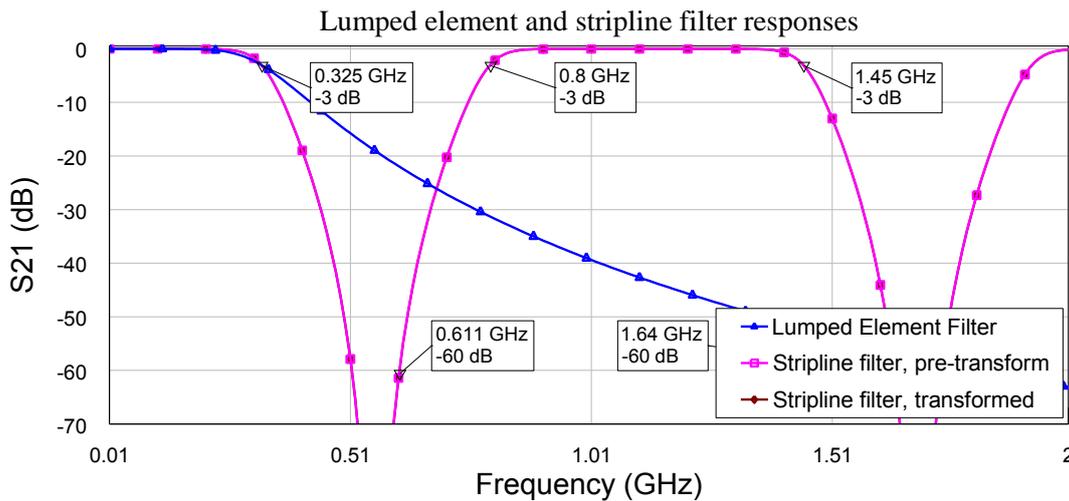


Figure 3.16: Responses of lumped element and stripline filters as ideal models. Note that the stripline filter responses are identical and true to the Kuroda identity's uniformity. Also, the stripline implementations have better roll-off, flatter passbands and periodic responses. The equivalence of the Richard's transform is shown at 325 MHz where the graphs cross.

The model is adapted for implementation on a real substrate. The ϵ_R and loss tangent are assumed to be invariant with frequency and are entered as single values for the calculation of the required width and length of lines. These calculations are performed using an integrated function of the optimisation software. By varying the width, the impedance is adjusted and the half-wavelength is matched to the BPF centre frequency. It is advantageous to minimise the number of variables used in optimisation, since this quickens the process. If this 4th order filter is symmetrical, the variables are reduced from 14 to 8. In spite of the difference between the impedances, the $Z_{g1'}$, $Z_{g2'}$ and Z_{T1} values are similar to $Z_{g3'}$, $Z_{g4'}$ and Z_{T3} , respectively. By setting the comparable values to their geometric average, a symmetric layout is derived.

The substrate properties used during the first simulations are similar to those of readily available substrates in the department. Due to the available photolithographic production process, the impedance of the stubs at the edges of the layout is not achievable. Therefore, a 0.2 mm limit is set as the minimum width. Figure 3.17 presents four simulated responses. These include the exact filter, a symmetric implementation of it, a model of the latter incorporating the substrate and this arrangement with limited widths. The simulations incorporating the substrate feature line models that approximate losses and the end effects of open lines.

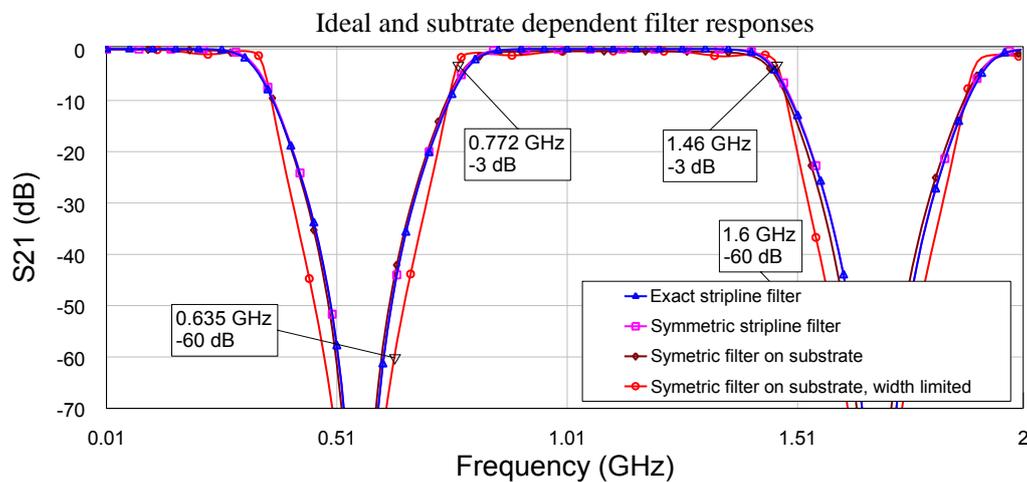


Figure 3.17: Responses of ideal and substrate-dependent filter models with adaptations. Note that the passband and 60 dB rejection ranges of the width-limited model are shown.

Figure 3.17 suggests that the filter response is not unduly sensitive to geometry changes and that a symmetric design will perform well, even if the trace widths are limited. Although the roll-off requirements are met, the width-limited result has a rippled passband with a small shift in the centre frequency. To improve this, an optimisation algorithm is selected.

The desired out-of-band rejection and a low passband reflection coefficient are set as the goals to the optimisation process, which allows the filter parameters to be trained towards a desired shape. The key variables are then identified and constrained. This keeps the structure within the design limits and enables control of the structure size. A rule of thumb is used to keep lines separated by a minimum distance of 5 linewidths and the constraints are configured accordingly [3]. This separation can be increased if coupling is shown to degrade the overall filter performance. During optimisation, the filter response will be analysed for systematic changes to the variables.

The response of the software between subsequent iterations depends on the type of algorithm selected. For example, a random algorithm makes arbitrary changes to the structure and calculates if the resulting performance is closer to the goals. If the result is favourable, this change is kept and another modification is tested. The number of iterations is specified and goals may have different levels of priority. When the

number of executions has been run, the software will resort to the solution that came closest to satisfying the objectives. At this stage, the dielectric substrate properties are entered. Table 3.5 records the values of the substrate used.

ϵ_R :	3.355
Loss tangent:	0.0035
Height (Single layer):	3.048 mm (1.524 mm)
Conductor thickness:	30 μm

Table 3.5: Properties of the substrate used for BPF 1.

The model used during optimisation includes approximations of bends, cross-sections where lines meet and unit elements that are added for the attachment of connectors. The bends are used to make the structure as small as possible. These additions take into account changes in the length of lines due to their intersections. Figure C.2 shows the influence of these additions. The optimised layout is different to the initial arrangement in several ways, with both widths and lengths changed for better results. However, the bandwidth is reduced in order to achieve the required roll-off characteristics.

The final configuration is analysed with an EM-field simulator for improved accuracy and substantiation of results prior to construction. A parameterised model is configured to simplify geometric changes, starting with two blocks of dielectric material assigned with the substrate properties. The lines are modelled as Perfect Electric Conductor (PEC) material with no thickness, since this dimension is too small to mesh. A flat-pin SMA connector is used at the filter ports, since these can be soldered onto the input strips. A Computer Aided Design (CAD) drawing of the filter and a detailed image of the SMA port from the Computer Simulation Technology (CST) software are shown in Figure C.3.

This structure is enclosed with PEC boundaries, as the built filters are also enclosed by conductive material. By grouping lines in the model with similar orientations and widths, the mesh around these can be manually configured to ensure that the fields are sampled sufficiently. An automatic adaptive mesh refining function is used to improve the analysis to a set error margin. Figure C.4 shows the comparison of the analysis before and after refinement, indicating that the manual setup performs well.

The photo-etching process uses negative masks and 2:1 scale prints of the design. Once the circuits are developed, they are cut to the correct size. A small groove is then filed away, leaving space for the flat pin of the connector. All surfaces are cleaned mechanically with alcohol. The pins are subsequently soldered into position, which requires some care, as all parts are still separate and the connectors must be attached perpendicularly. A sheet of glue cut to the filter shape is then placed in between the layers of substrate. Using an aluminium plate to apply equal pressure across the structure, the filter is clamped and baked for several minutes. The glue has a negligible effect on the filter response, but is essential for binding the top and bottom substrates together, ensuring that the height remains uniform across the assembly.

The completed filter is characterised with a calibrated measurement on a VNA. Figure 3.18 compares the BPF 1 forward transmission coefficients predicted by simulations and measured. The corresponding reflection coefficients are plotted in figure C.5 and the measured ranges are listed in table 3.6.

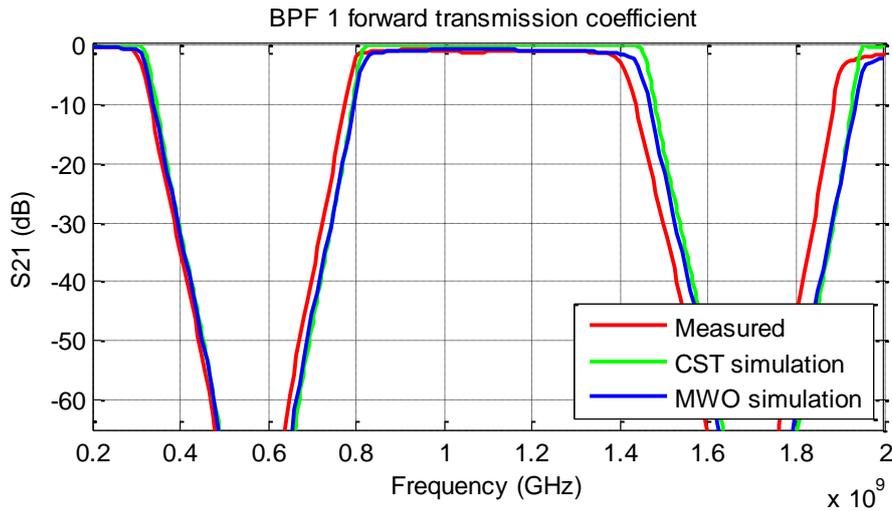


Figure 3.18: Comparison of measured and simulated BPF 1 forward transmission coefficient.

Passband	790 - 1400 MHz
60 dB out-of-band suppression	Below 644 MHz and above 1590 MHz

Table 3.6: Measured BPF 1 characteristics.

The measured data indicates that the passband was shifted downward in frequency and that it is less symmetric than predicted. The main reason for this is a difference between the actual and simulated values of ϵ_R . If the ϵ_R value in the simulation is increased to result in the same centre frequency, a more accurate value can be used during optimisation. The agreement between the predicted and achieved behaviour will then be improved. However, as the simulations implement single valued properties of the dielectric, the variation with frequency of the actual parameter also contributes to the difference between simulated and measured parameters. Furthermore, the limited accuracy of the production process may also contribute to the deviations. Notwithstanding these points, the final filter is functional for path 2 and its refinement is only necessary for future work. The system was characterised with this component and the reflection coefficient is satisfactory (below -10 dB) throughout the passband.

3.3.1.2 Additional signal path 2 filters

There are two additional filters required for the path 2 filter chain (Figure 3.7). The first, Add LPF 1, suppresses the upper periodic response of BPF 1 and the second, HPF 1, the lower. The LPF is developed as a stripline implementation and the required out-of-band rejection ranges for both filters are listed in table 3.7. Figure 3.19 shows the simulated and measured forward transmission gains of an optimised Add LPF 1 design.

Add LPF 1: 60 dB out-of-band rejection	Above 1.9 GHz
HPF 1: 60 dB out-of-band rejection	350 MHz

Table 3.7: Specifications for RATTY SU Add LPF 1 and HPF 1. Note that the passband requirements are to coincide with that of BPF 1.

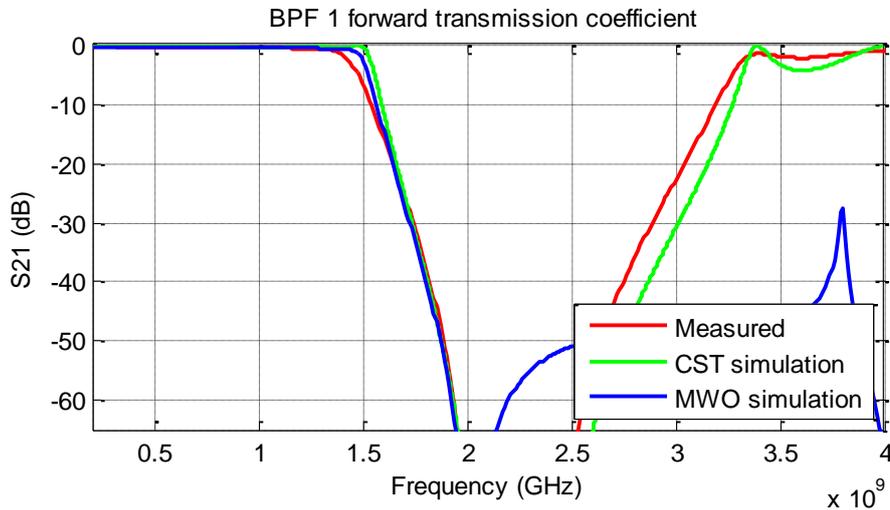


Figure 3.19: Simulated and measured responses for Add LPF 1.

Since the same substrate was used for the construction of Add LPF 1 and BPF 1, the downward shift of the measured response in figure 3.19 is expected. Though the filter chain will provide the required suppression to avoid mixing related signal leakage, it will not achieve this for frequencies above 2.5 GHz. Figure 3.20 shows the simulated path 2 filter chain response from measured component parameters. A well-suited HPF was purchased and three attenuation pads are added for improved matching.

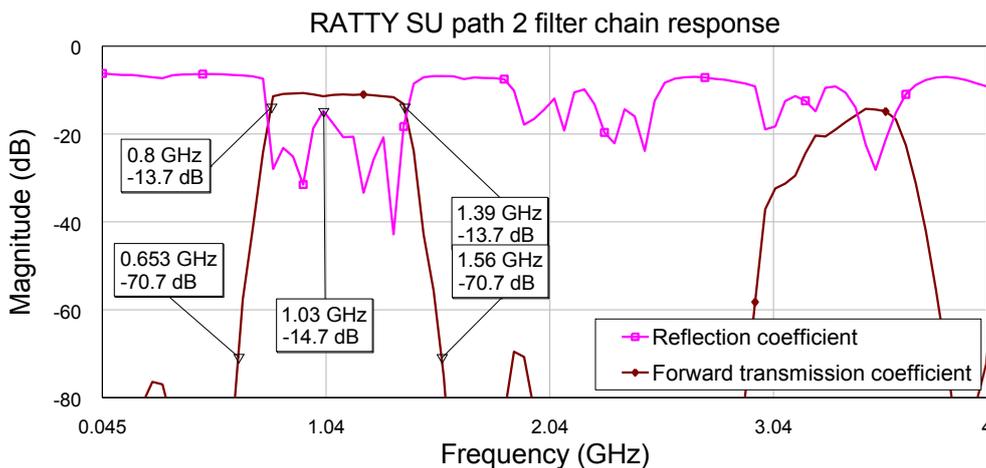


Figure 3.20: Simulated response of the filter chain in signal path 2, RATTY SU. Note that the limits of the passband and out-of-band ranges are indicated.

The lack of suppression above 3 GHz, which is apparent in figure 3.20, is compensated for by the anti-aliasing filters. With this considered, all the principle specifications are met by the combined response of the three filters. However, the bandwidth is 8.3% lower than in the intended design.

There were also two filters developed for signal path 3 and the process is described in appendix C.3. An upward frequency shift of a 120 MHz occurred in the passband of this filter chain. If the system is to maintain continuous bandwidth coverage, the design must be adjusted for the actual ϵ_R value and the filters reproduced. However, the project timelines did not allow further activity here and this is regarded as future work, since the second signal path serves as a proof of principle layout. Nonetheless, the third channel was constructed and tested. Reference is made to it where appropriate.

3.3.2 RATTY SKA filter designs

The filter requirements for the second signal path of the RATTY SKA FE are listed in table 3.2, with the BPF 1 and Add LPF 1 developed as stripline implementations at SU. The BPF 1 functions as the defining component to the band shape, filtering out the second Nyquist zone at a sampling rate of 1.2 GHz. Both stripline layouts are optimised 4th order filters with tapered input lines and the results for BPF 1 are shown in figure 3.21.

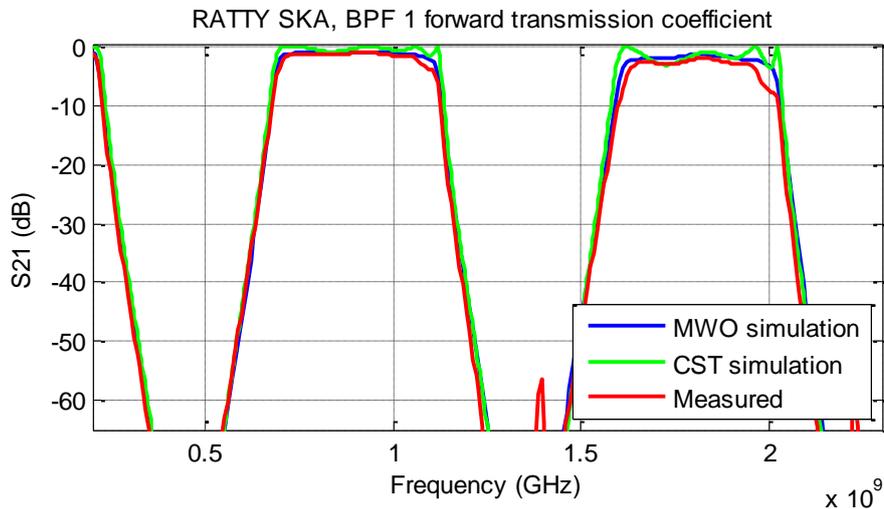


Figure 3.21: Simulated and measured results of the RATTY SKA BPF 1 forward transmission coefficient.

The built BPF 1 performed very closely to its simulated models, since a corrected ϵ_R value was used. This is calculated using a frequency correction factor and an ϵ_R value recommended for design purposes. Both values are specified in the substrate datasheet. The roll-off ratio achieved is better than its requirement, since it is given higher priority during optimisation. This is possible with the reduced bandwidth ratio, compared to those of the RATTY SU filter designs. The reflection coefficient also remains good across the passband, with the exception of the last several MHz (Figure C.11). The requirements of Add LPF 1 are derived to support the BPF, i.e. a 60 dB rejection ratio above 1.6 GHz and a minimised reflection coefficient within the passband. The measured and simulated forward transmission coefficients are presented in figure 3.22 and the reflection coefficients in figure C.12.

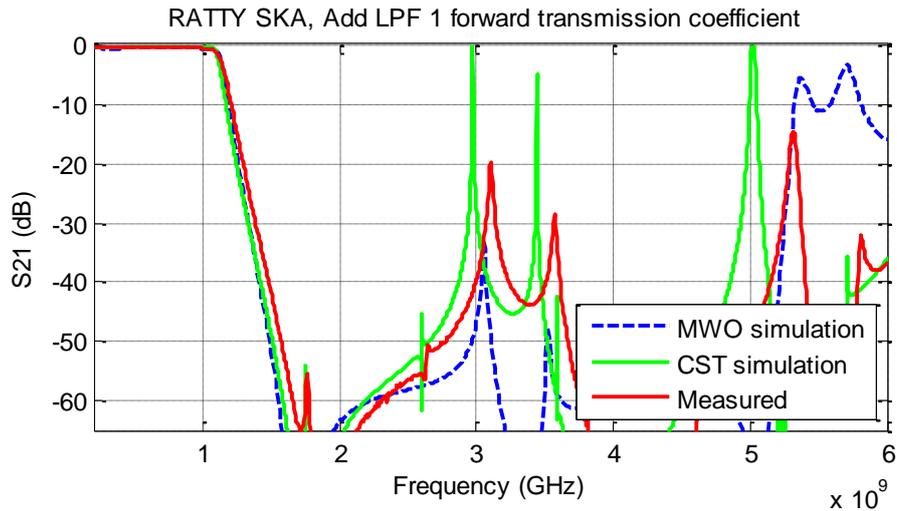


Figure 3.22: Simulated and measured results of the RATTY SKA Add LPF 1 forward transmission coefficient.

The LPF design also benefitted from the ϵ_R adjustment and nearly matched its simulated predictions. However, the variation of the actual ϵ_R between the centre frequency and the start of the maximum rejection range resulted in a diminished roll-off rate. This could also have been affected by manufacturing inaccuracies. The lowered slope of the roll-off, combined with its repeating passbands and peaks in the response, leads to the combined response of the path 2 chain not fulfilling the suppression requirements above 1.5 GHz. This is shown in figure 3.23, which is a simulated plot of the measured filter properties combined.

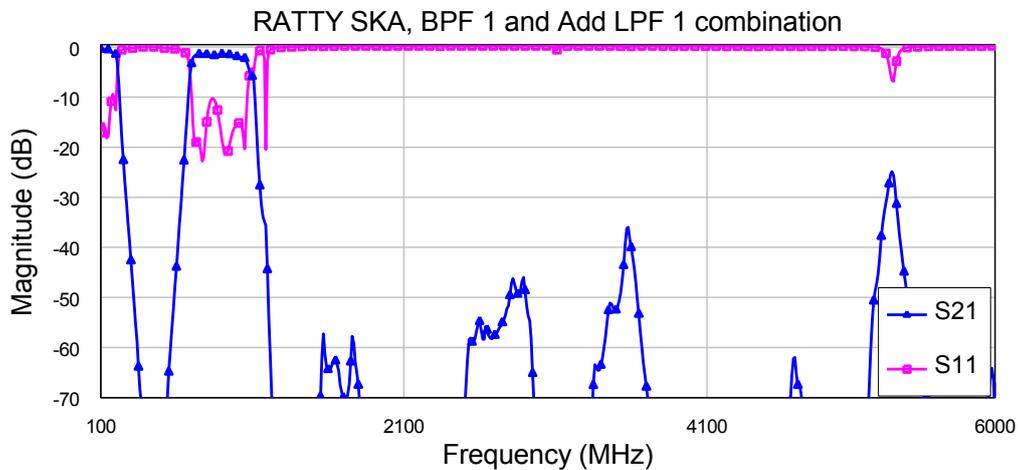


Figure 3.23: BPF 1 and Add LPF 1 simulated combined response for RATTY SKA.

As is apparent in figure 3.23, there is additional LPF suppression required to achieve the specified rejection ratio and prevent out-of-band contents from aliasing into the output band. Also, a suitable HPF is needed for the lower frequency response and both the latter issue was addressed with a readily-available component. The combined filters will pass an output band that satisfies the sampling criteria and provide 60 dB of rejection.

A recommendation for future designs of stripline BPFs is to optimise the BPF and LPF combinations simultaneously, which will enable the reduction of periodic responses. Additionally, an optimised layout can be tested for sensitivity to small variations in the ϵ_R ; this would test to see if small fluctuations would violate

the requirements. However, the attenuation caused by several filters in cascade is unwanted and different filter implementations should be considered.

3.4 Front-end analysis

This section documents the analyses of the mixing and direct-sampling XDMs. Since components were characterised with measurements, the first section is a description of this process. Several simulations are then described and used to calculate system parameters, along with the interpretation of the results. Although reference is made to the RATTY SKA design, the analysis of the mixing design is mainly presented.

There were three analyses applied to the FE designs: 1) a spread-sheet calculator*, 2) a commercially available simulator [9] and 3) a spurious assessment script (Appendix A.4). The goal of the investigation was to calculate several system parameters and to aid the design process, by providing insight into the effects that changes in component parameters or the system configuration have. The parameters that were evaluated include the SFDR, headroom, gain, P1dB, NF and T_{sys} .

3.4.1 Component measurements

Measurements were taken to study the broadband responses of all components. This allows an accurate analysis of the system. To remain within the project time constraints, components were purchased at an early phase of the design. With the lead times of parts varying by several weeks, the parts were tested upon delivery within this window period.

An Agilent 8510C VNA was principally used for the measurements and calibrated with a 3.5 mm, precision connector calibration set. The VNA was switched on approximately 30 minutes prior to the measurements to allow the electronics to reach stable operating temperatures and to maintain the calibration accuracy. A measurement range of 45 MHz – 18 GHz was used with 801 measurement points and an averaging factor of 16. Although the frequency resolution of ~22.4 MHz is too coarse to characterise the cut-off features of the stripline filters exactly, the precision is sufficient for broadband system analysis. The wide measurement range was also selected to characterise several of the components that have operational ranges greater than the system bandwidth .

The active components were powered from a bench power supply unit for several minutes and measured at operating temperatures. The VNA power settings were configured to keep amplifiers within their linear ranges and to restrain input-power levels to several dB below the 17 dBm limit. A separate power supply unit was used for controlling switches and the combined port parameters provided both insertion loss and isolation levels. For the attenuators, all the possible settings were characterised and the responses are also used for calibration purposes. In general, the measurements closely followed datasheet predictions.

The synthesiser used as the LO source in the mixing XDM was tested to see that it provides enough power to the mixer and that harmonic components are sufficiently suppressed by the LO LPF. Initial SA measurements

* The initial spread-sheet calculator was compiled by Adriaan Peens-Hough of SKA-Africa; this was adapted for further analyses.

indicated that the highest harmonic was in excess of 30 dB below the fundamental product at the path 2 LO frequency of 1.5 GHz. However, 10 MHz harmonics were observed throughout the ADC baseband. This broadband, self-generated RFI is related to the 10 MHz reference source of the synthesiser. To limit interference from this, an HPF was added as an LO filter. The selection of the LO filters, synthesiser measurements and the interference assessment is documented in appendix D.1.

Although the reflection coefficients of the mixer ports were measured with a VNA, the conversion loss, spurious response and RF-to-IF isolation were tested on a trusted SA using Signal Generators (SG) as inputs. These tests were conducted in a configuration similar to the mixing XDM layout with the synthesiser and accompanying filters used as the LO source. Also, the input was filtered with the appropriate filter chains and the mixer ports were buffered with 3 dB attenuators. Since conversion loss and spurious levels are typically dependent on the specific application, the tests provide a good indication of the expected mixer performance in the design. By injecting the input via the filter chains, the SG harmonics are kept below 70 dBc of the fundamental component, limiting interference with several, frequency-related spurious products.

At an SG level of 7 dBm, the resulting mixer-input power was ~ -10 dBm, matching the level at which the spurious tables are specified. It is also a good approximation of the highest spurious levels that will occur during system operation; a maximum, mixer-input level of -6 dBm is expected at the LNA OP1dB. No mixing spurious was observed within 60 dBc throughout the path 2 range and this agrees with the datasheet predictions. Furthermore, the conversion loss was found to be -8 dB at worst. The conversion loss measurement also demonstrated the RF-to-IF isolation, which was 25 dB in the worst case. This is 15 dB less than predicted by the datasheet and has implications on the out-of-band rejection of the channel (Section 3.3.1.1). The mixer measurements are documented in more detail in appendix D.2.

There are several units of some passive components in the designs and not all these parts were measured. Since there was good agreement between the measured responses of identical components, the averaged parameters were used in simulation.

3.4.2 Spread-sheet analysis

A spread-sheet calculator was used from an early stage in the mixing XDM design to estimate the system parameters at single-frequency points. In particular, it was used to evaluate the gain of the level-setting amplifiers. It performs a systematic implementation of the RF-receiver equations without any mismatch corrections (Section 2.2). This allows the state of the signal path to be followed through the design and problem areas to be identified. Initially, the component specifications were set according to datasheet figures or guesstimates, which were later replaced with measured values for the final evaluations. In this refined calculation, the component values correspond to the testing frequencies and are changed to the output frequencies in the post-mixer stages. Note that the step attenuator remains at a 0 dB setting for all the assessments, unless stated otherwise.

During the gain assessment, several amplifiers were considered for the mixing system. The goal was to realise the maximum P1dB and SFDR, whilst minimising the NF. This commenced with a trial and error approach that was guided towards a conclusion by the calculator results. The Amp 2 part was selected to balance the gains of the signal paths (Figure 3.10). However, the addition of LPFs and HPFs to the filter chains was not considered at this time, with the filter designs still in a conceptual phase. This led to an imbalance between the path gains and some degradation in the NF of this XDM. Notwithstanding the

unanticipated losses, the design was built with several attenuator buffers for improved matching. Figure 3.24 shows the calculated NF and gain parameters of the implemented path 2 design with measured component specifications. Note that the simplified circuit diagrams shown at the bottom of these figures indicate the relative position in the particular signal path. In order to maintain the legibility of the diagrams, only the key components are shown. However, all the parts are accounted for on the horizontal axes and reference is made to the relevant, full schematics in the captions of the figures.

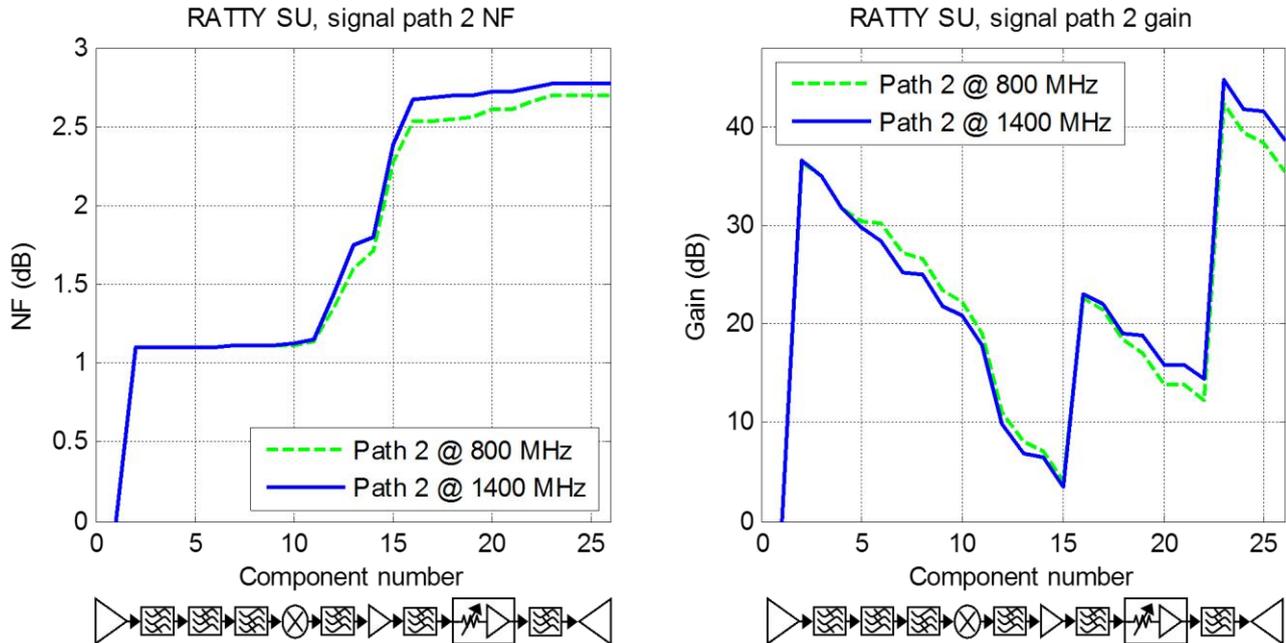


Figure 3.24: Spread-sheet calculations of the NF (left) and gain (right) parameters for path 2 in the mixing FE. Note that the diagrams at the bottom of the figures are simplified and based on Figure A.4 (Figure 3.10), which can be referred to for clarification of the component number on the horizontal axis. These diagrams do not show 3 dB attenuators or switches. Also, the input frequencies are given.

The results in figure 3.24 are calculated at the cut-off frequencies of the second signal path. The 800 and 1400 MHz points down-convert to 700 and 100 MHz, respectively. Due to these points falling within the limits of some component frequency-ranges, they represent the worst-case performances, i.e. with attenuation at a maximum and amplification at a minimum.

It is apparent that the NF degradation is most significant in the post-mixer stage, where the gain reaches a minimum. Since there is less pre-mixer loss at this frequency, the NF is better at 800 MHz than at 1400 MHz. However, the frequency inversion of down-conversion combined with the increased attenuation of the anti-aliasing LPFs at higher frequencies affects the gain oppositely. This results in a gain decrease of 3 dB lower than at 1400 MHz.

Figure 3.25 shows the path 2 T_{sys} and represents the noise performance trend in K. Furthermore, the OP1dB only reaches the 10 dBm limit of the LNA with the other components remaining within linear operating ranges. These OP1dB values do not include the ADC compression level and indicate the output capabilities of the FE.

An identical spread-sheet analysis was performed for the mixing path 1 configuration. The results obtained for the NF and gain are shown in figure 3.26. A figure showing the T_{sys} and OP1dB results for path 1 is included in appendix D.3.

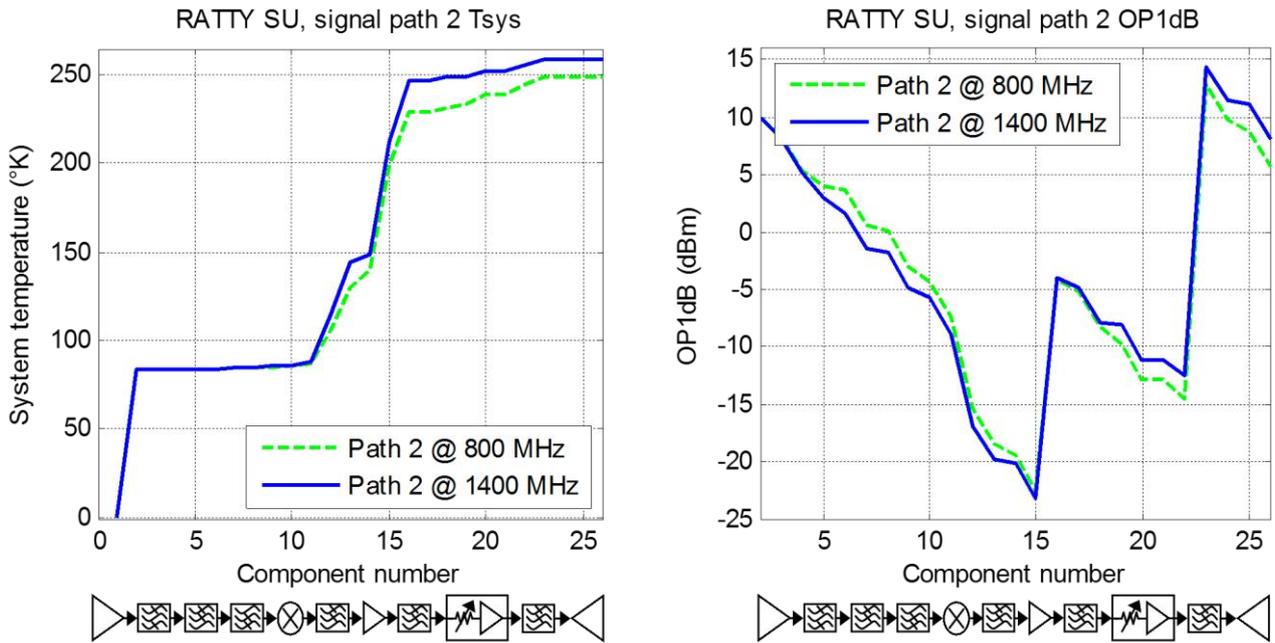


Figure 3.25: Spread-sheet calculations of the T_{sys} (left) and OP1dB (right) parameters for path 2 in the mixing FE. Note that the diagrams at the bottom of the figures are simplified and based on Figure A.4 (Figure 3.10), which can be referred to for clarification of the component number on the horizontal axis.

These diagrams do not show 3 dB attenuators or switches. Also, the input frequencies are given.

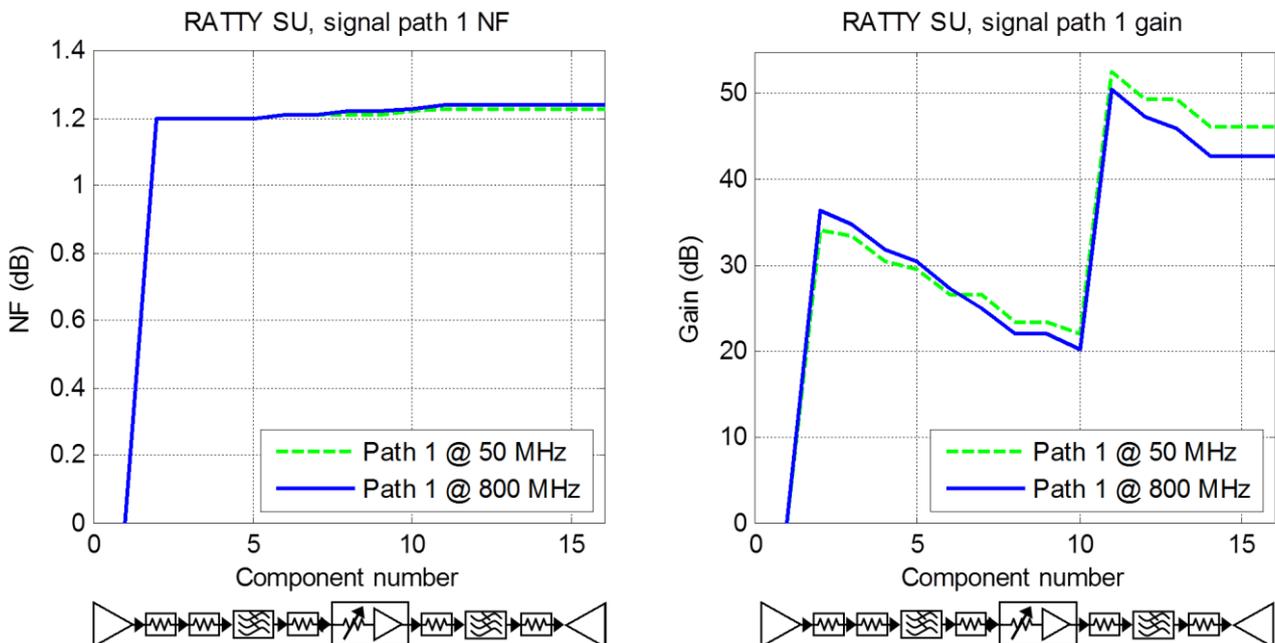


Figure 3.26: Spread-sheet calculations of the NF (left) and gain (right) parameters for path 1 of the mixing FE. Note that the diagrams at the bottom of the figures are simplified and based on Figure A.4 (Figure 3.10), which can be referred to for clarification of the component number on the horizontal axis. These diagrams do not show switches.

The NF of the baseband path remains close to that of the LNA (1.2 dB). Due to the gain remaining above 20 dB, there is only a small amount of degradation in the noise performance ($T_{sys} = 96 \text{ }^{\circ}K$). The $> 3 \text{ dB}$ difference in gain between these frequency points is due to the anti-aliasing LPF attenuation.

It is apparent from a comparison between the signal path performances that the best noise performance does not necessarily occur at the maximum gain. However, the shape of the noise floor will be more dependent and reflect the gain functions, since the noise figure variations have less influence in this regard.

A comparison can now be made between the baseband and mixing paths of the XDM design. Table 3.8 contains system parameters of either path. This evaluation is done at frequencies that fall within the best operating ranges of the system.

RATTY SU, Path 1 (100 MHz)		RATTY SU, Path 2 (1100 MHz)	
NF	1.2 dB	NF	2.5 dB
T_{Sys}	94 °K	T_{Sys}	227 °K
Gain	47.4 dB	Gain	38.6 dB
Gain variation	4.5 dB	Gain variation	3.3 dB
OP1dB	10.4 dBm	OP1dB	8 dBm
Input noise power	-78 dBm	Input noise power	-76.7 dBm

Table 3.8: Comparison of spread-sheet parameter calculations for RATTY SU paths 1 and 2.

The gain variation is the difference between the lowest and highest figures. Furthermore, the input noise power includes 300 K that is added as the noise captured by an antenna. It is apparent that the path 2 response is poorer in most respects, except for the gain variation. The latter figure is better, due to the reduced bandwidth of path 2. The variation will be worse for a path-3 analysis with a 700 MHz bandwidth, since the anti-aliasing LPFs will shape such response more drastically. The difference in the path gains will affect the respective measurement ranges and SFDRs, since both are dependent on the noise floor.

All the results documented thus far are for the lowest step attenuator setting. The FE OP1dB values show that the maximum input power level of the ADC can be exceeded by several dB, which requires a systematic configuration of the attenuator settings during measurements. An important aspect of path 1 is that the LNA is not the first component to enter compression. The final amplifier will respond non-linearly first at full gain in the baseband, which can be avoided by using the attenuator. Table 3.9 lists the system parameters calculated at several attenuator settings for path 1, since the baseband is more complexly affected by the change. An equivalent table for path 2 is included as appendix D.3.

RATTY SU, Path 1 parameters (100 MHz)				
Attenuator setting	Gain (dB)	NF (dB)	OP1dB (dBm)	Input noise power (dBm)
0 dB	47.4	1.2	10.4	-78
-10 dB	37.4	1.4	8.1	-77.8
-20 dB	27.4	2.9	1.2	-76.3
-31 dB	16.4	9.6	-9.4	-69.7

Table 3.9: Spread-sheet calculations of system parameters for RATTY SU path 1 with varying attenuator settings.

The decrease in the OP1dB in table 3.9 does not coincide with the increase in attenuation, which indicates the compression of the final amplifier. Once the path 1 attenuation increases beyond ~15 dB, the compression is limited by the LNA. Thus, the maximum input level for mixing system is ~-27 dBm. If a

baseband measurement is performed at this level, it would require an attenuator setting of -23 dB to remain below the ADC limit of -2 dBm. Also, note that the degradation of the NF will not lead to an increase in the noise floor level of measurements. In fact, the analogue noise floor will be reduced as a result of the lowered gain. This illustrates the usefulness of the attenuators in reducing the total gain of the FE.

The signal paths in the direct-sampling FE are similar in operation to the baseband path of the mixing XDM. However, there are two significant differences between the systems: 1) the SKA XDM was constructed with three consecutive level-setting sections at the output and 2) there is no anti-aliasing filters preceding the ADC (Appendix A.5). The first dissimilarity arises from the pursuit of the goals set in the requirement document (Section 3.1.2), which proposed a very high gain for the systems. This increases the instantaneous sensitivity and adjustable dynamic range, although the noise floor is only reduced by ~1.2 dB relative to the second RATTY SU path. Secondly, the lack of anti-aliasing filters implies that the out-of-band rejection is solely dependent on the filter-bank components and the out-of-range responses of the level-setting components. Although the initial filters do not satisfy the design criteria in this regard, improvements are considered to be future work.

The losses caused by impedance mismatches and connecting cables are ignored in the spread-sheet calculations. It is conceivable that these errors can lead to gain variations within stages of the design, which require an increase in the analysis frequency resolution to be observed. Nonetheless, the spread-sheet calculator is a powerful tool for gaining insight into and improving a design.

3.4.3 Spurious analysis

The system simulator does not provide an active-mixer model and a spurious assessment script (Appendix A.4) was compiled to incorporate the characteristics of the RATTY SU mixer. Although the mixer spurious is analysed separately, the result can be used to determine the superimposed SFDR for the system. The analysis presented here is for the worst-case mixing spurious products of path 2, which is explained in section 3.4.2. A simulated response of the FE subsection preceding the mixer incorporates measured component characteristics and is used as broadband input in the script. This provides an accurate representation of the bandshape that is input to mixer. Figure 3.27 is a plot of the script output showing the baseband range and the highest spurious components.

The graph on the left of figure 3.27 reveals that there are several spurious components that exist in the path 2 frequency range. These products are within 60 dBc of the passband levels, therefore within the design SFDR. However, the highest of these spurious are $|n*LO \pm 1*RF|$ products. This indicates a lack of filter out-of-band rejection at higher frequencies, which is an issue related to the filter-designs rather than mixing-spurious. To illustrate this, the graph on the right of figure 3.27 shows the spurious response without the filter leakage. In this case, the path range can be seen to exhibit nearly 60 dB of dynamic range clear of mixer interference. Only two products remain within the SFDR: 1) direct signal leakage, due to the reduced RF-to-IF isolation, and 2) filter leakage in the centre of the band. Both these interference issues can also be resolved with improved filter designs. Thus, the spurious analysis suggests that mixing spurious is not an issue for real-time measurements, since a mixing SFDR of 60 dBc is expected.

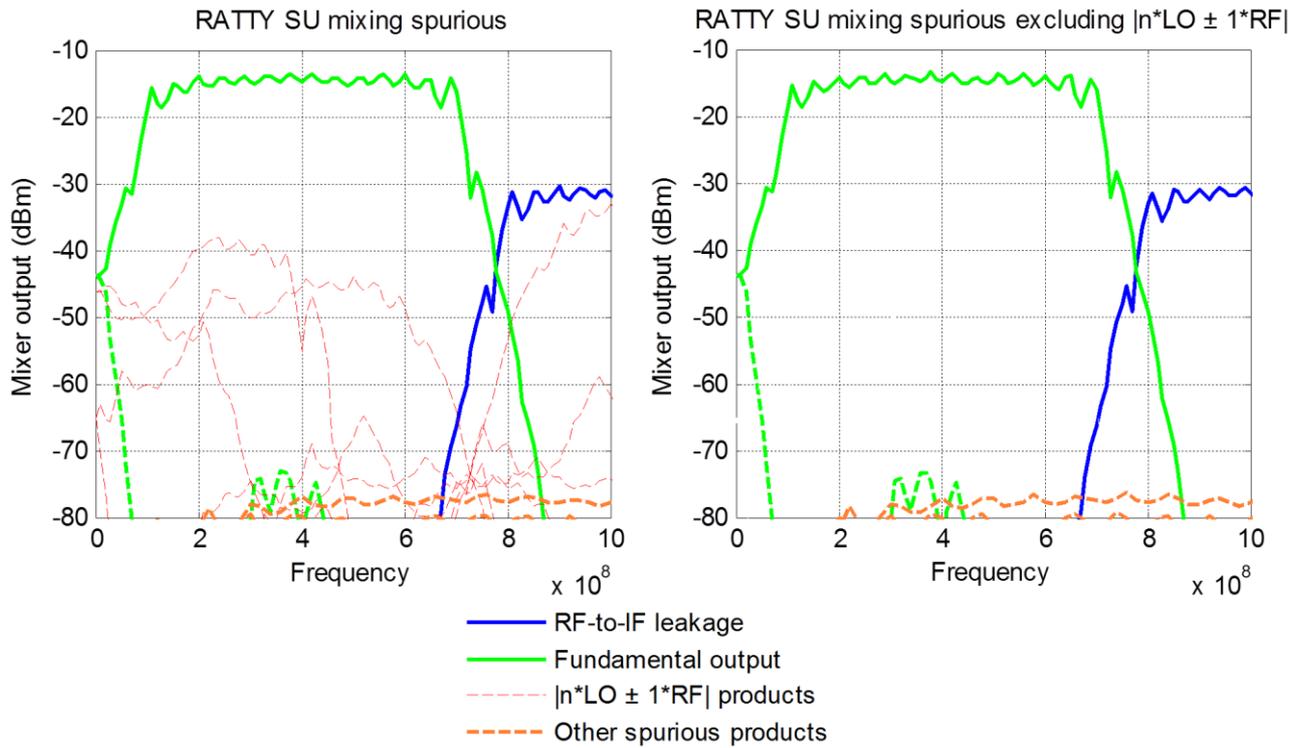


Figure 3.27: Plots of mixing spurious in the RATTY SU path 2. Note that the figure on the right excludes all $|n \cdot LO \pm 1 \cdot RF|$ products.

The rippled response of the filter chain is caused by the increase in resolution through interpolation in the simulator. The responses outside of the baseband are suppressed by the subsequent anti-aliasing LPFs and are not considered. In summary, the predicted mixing SFDR is better than the ADC range. However, the mixing SFDR must be considered if integration is used to increase the dynamic range.

3.4.4 Simulator analysis

The Applied Wave Research (AWR) simulator provides a good platform for FE analyses that incorporate measured component parameters. This solver includes the effect of reflection coefficients in gain, compression and noise performance calculations [9]. Figure 3.28 shows the simulated path responses of the RATTY SU FE.

Figure 3.28 demonstrates the bandwidth coverage in the mixing system and the imbalance in the path gains. The baseband response roll-off does not precisely satisfy the design specification and aliasing will occur within 60 dBc in the 760 – 800 MHz range. Due to the reduced bandwidth of path 2, the only aliasing issue in this path is the lack of suppression at 1.8 GHz. Furthermore, the gain at the lower end of this band is clearly affected by the anti-aliasing-filter attenuation.

Note that mixing-path bandshape is less flat than in the baseband, due to the complexity of the layout. This is an important characteristic to consider, since the simplicity of a direct sampling configuration makes it similar to the baseband layout and the bandshape is easier to control. The path 3 response is also shown in figure 3.28 and the frequency shift in the stripline filters is apparent, leaving a gap in the bandwidth coverage from 1.4 to 1.5 GHz.

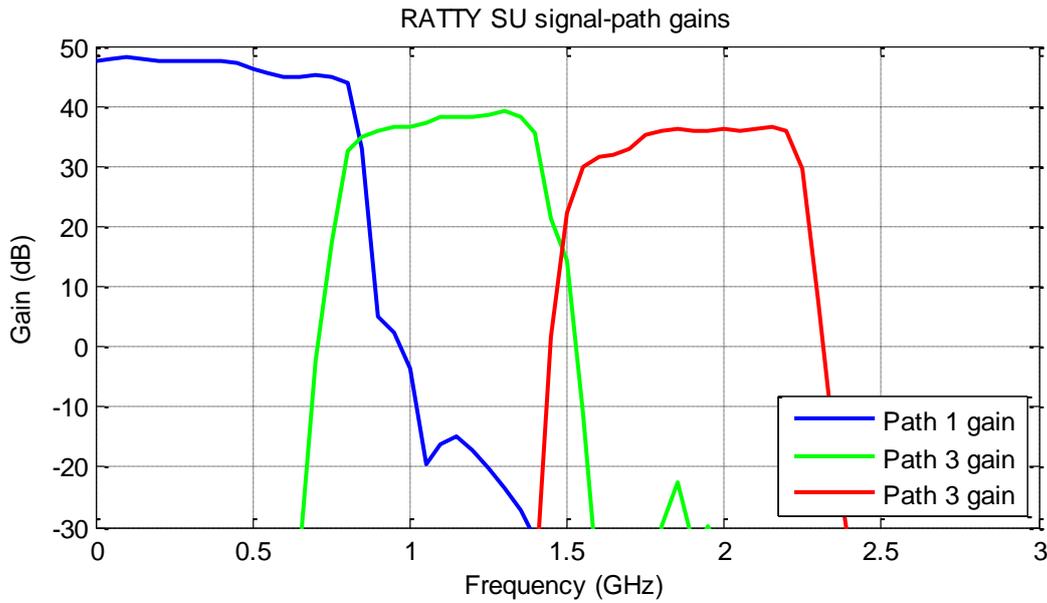


Figure 3.28: Simulated signal path responses of the RATTY SU FE. Note that the simulated gains are based on measured component parameters.

As the simulated results follow the parameter trends presented in section 3.3.2, comparisons between the analyses are only evaluated at single-frequency points. Table 3.10 lists the spread-sheet and simulator calculated parameters for path 2 of the mixing system. These values show the largest discrepancies between the assessments and the equivalent table for path 1 is included in appendix D.3.

RATTY SU, Path 2				
Parameters	AWR simulator		Spread-sheet calculations	
	Worst-case	Best	Worst-case	Best
NF	4.23	2.65	2.77	2.5
T_{Sys} ($^{\circ}\text{K}$)	477	243	258	227
Gain (dB)	33.1	39.1	35.3	38.6
Input noise power (dBm)	-73.4		-76.4	-76.7
Maximum Gain variation (dB)	6		3.3	
Maximum FE OP1dB (dBm)	6.9		8	

Table 3.10: Comparison of spread-sheet and simulator parameter calculations for RATTY SU path 2.

The best NF calculated with the simulator is higher than the spread-sheet prediction. This is likely as a result of the mixer model implemented in the simulator, since this component has been shown to have a significant effect on the noise characteristics. Furthermore, the worst-case, simulator NF is very poor and deviates even more from the standard formulation implemented spread-sheet. However, the minimum simulator gain is > 2 dB lower than the spread-sheet calculation and this is caused by the interpolation at the cut-off frequency of the BPF. Due to the rapid variation in attenuation around this point, the interpolation is less accurate and several deviations from the simpler analyses occur. This also accounts for the worsened gain variation, input noise power and T_{Sys} values. Note that only the maximum simulated input noise power has been listed. Notwithstanding the differences between the assessments, it is apparent that the best figures are nearly in agreement. The equivalent baseband comparison is in good agreement and reflects a direct sampling system.

There will be differences between the measured and the actual responses of components that are integrated into the system. These will mainly be caused by variations in operating temperatures and supply voltages, which can only be approximated in the measurement setup. Furthermore, the interconnecting-cable and connector losses are not included in the analyses and will lead to some additional deviations in the built FE.

3.4.5 Dynamic range specification

With several FE parameters and levels calculated, the theoretical dynamic ranges of the measurement system can be specified. These ranges are described in section 2.2.4 and the figures that are of utmost importance are the headroom and SFDR. The interpretable range of a measurement depends on the smaller one of these specifications. If the headroom is greater than the SFDR, the latter poses a limitation on the useful dynamic range. Similarly, if the headroom is smaller than the SFDR, the magnitude of an input signal can only increase linearly until the P1dB limit is reached and the SFDR has no influence in the real-time spectrum.

To re-emphasize, the P1dB and noise floor levels are used in the SFDR calculation. A sensible SFDR must be defined for measurements and the range must be intuitively interpreted in a graphical representation. Figure 3.29 is an illustration of several SFDR definitions.

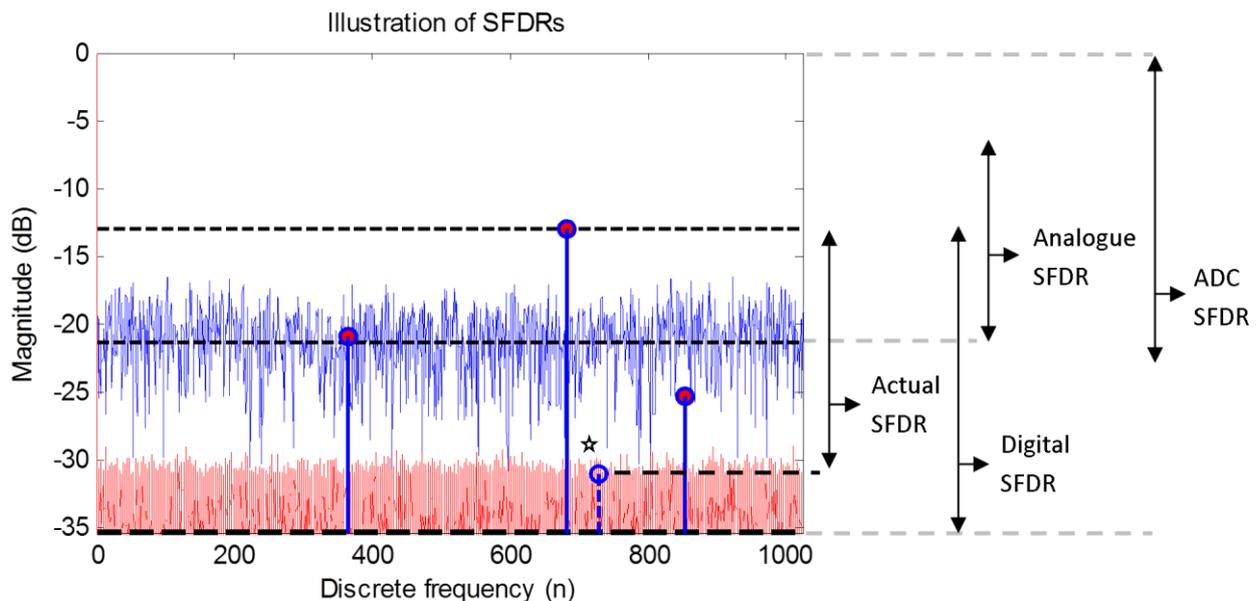


Figure 3.29: An illustration of SFDR definitions (After [19]). Note that the actual SFDR is limited by the presence of the spurious product indicated with a star and that the digital noise floor (red) is lower than the analogue noise floor (blue). The ranges are indicated with the P1dB and ADC maximum input level assumed to be at 0 dB.

Since the digital noise floor is lower than the analogue level and the same P1dB figure is used for both SFDR calculations (eq.2.10), the digital SFDR will always be greater than the analogue limit. However, note that the digital range is defined at a lower power level. This is typically below the specified ADC SFDR, which is defined with the lower limit provided relative to the maximum input level. Since the goal is to amplify signals into the ADC range, the FE SFDR definition must also be calculated at these levels. The minimum ADC SFDR is -2 to -54 dBm, although the design has incorporated a 60 dB figure in the preceding sections and -2 to -60 dBm is used in this analysis.

The maximum FE SFDR equation is adapted for an output power level that results in third order IMD products nearly matching a -60 dBm ADC noise level. This is used as a reference and the realised SFDR of the system may be reduced, due to non-linear effects other than third order IMD. The alternative SFDR definition is shown in eq. 3.12. Note that the headroom will be greater than the SFDR if the latter remains above 20 dB, according to this definition.

$$SFDR (dB) = \frac{2}{3} * (OIP3 - N_{ADC}) \quad (3.12)$$

$$OIP3 = OP1dB + 10$$

where N_{ADC} = ADC noise level.

The adaptation of the OP1dB to the OIP3 in eq.3.12 is a general rule that is typically used to characterise practical components. The 10 dB difference accounted for here is conservative and below suggested values [1]. This analysis is focussed on specifying a minimum FE SFDR, as this figure has broadband implications. If a minimum is set, this level of spurious can occur anywhere in the spectrum and is dependent on the input spectral components. It is apparent from (3.12) that the minimum range will coincide with the lowest OIP3 value. Table 3.11 lists several SFDR values that are calculated for the lowest OIP3 values from the preceding spread-sheet section. This is done at several attenuation settings for both signal paths.

RATTY SU, Path 1 minimum SFDR (800 MHz)			
Attenuator setting	Gain (dB)	OIP3 (dBm)	SFDR
0 dB	42.9	18.6	52.4
-10 dB	37.4	14.7	49.8
-20 dB	27.4	6.3	44.2
-31 dB	16.4	-4.5	37
RATTY SU, Path 2 minimum SFDR (800 MHz)			
Attenuator setting	Gain (dB)	OIP3 (dBm)	SFDR
0 dB	38.6	16.2	50.8
-10 dB	28.6	8.3	45.5
-20 dB	18.6	-1.4	39.1
-31 dB	7.6	-12.4	31.7

Table 3.11: Minimum FE SFDR calculations for RATTY SU paths 1 and 2.

The results in table 3.11 suggest that there is a SFDR bottleneck in the design. Although the FE SFDR realised at the maximum gain nearly matches the minimum ADC limit, an increase in attenuation leads to a reduction in both the OIP3 and SFDR figures. There are two competing goals: 1) lowering the OP1dB to protect the ADC and 2) increasing the output capabilities to improve the SFDR. A solution to this issue could be an output amplifier in the FE with an improved difference in P1dB and OIP3 parameters. This combination could both limit power to the ADC and provide a greater SFDR, although the OIP3 adaptation used in these calculations is very typical and these characteristics are not readily available.

The values of table 3.1 must be considered as limitations on the SFDR characteristics of the initial RATTY SU design and are also relevant to the direct-sampling FE. The implication of this is that the SFDR of the devices will decrease when the attenuation is used to reduce the gain for higher power measurements. Note

that these calculations are relative to the ADC noise floor. If input levels are above the upper limit of the defined ranges, the FE SFDR will be further reduced and must be recalculated. Table 3.12 lists the SFDRs for both mixing-system paths, calculated at the maximum ADC input level, which is identical to the ADC range definition.

RATTY SU minimum SFDR (800 MHz)					
Path	Gain (dB)	OIP3 (dBm)	SFDR (dB)	Upper limit (dBm)	Lower limit (dBm)
1	42.9	18.6	41.2	-2	-43.2
2	38.6	16.2	36.4	-2	-38.4

Table 3.12: RATTY SU path 1 and 2 SFDRs according ADC range definition.

The real-time sensitivity of the system is limited by the maximum ADC input and analogue noise floor levels. If the maximum FE amplification delivers a noise floor several dB below the maximum ADC input level, an FFT can reduce the measurement noise floor to a level approximately equal to the lower limit of the ADC SFDR. This also allows the optimum range of the ADC to be used during sensitive measurements. The reduction is related to a length-dependent FFT factor (eq. 2.7). Note that the digital noise floor is the measurement noise floor, if the analogue noise floor is higher than the ADC lower limit. If a 2^{16} -point FFT is implemented, the output noise floor will be reduced by ~ 45 dB. This represents the polyphase filter bank resolution of the spectrometer configuration of RATTY 2 that is documented in section 4.1. Furthermore, all the preceding input-noise-floor calculations have considered the 2.95 GHz input bandwidth. These values are reduced by an additional 5.2 dB factor for the ADC bandwidth of 800 MHz, which is included here. Table 3.13 lists the specifications of a RATTY system that can realise the maximum sensitivity with the current ADC unit.

RATTY SU specifications for maximum sensitivity.					
Path	Gain (dB)	FE input noise power (dBm)	ADC input noise power (dBm)	System noise floor (dBm)	Maximum sensitivity (dBm)
1	76	-78	-7	-57.3	-128.3
2	74.1	-76.7	-7	-57.3	-126.4

Table 3.13: Estimations of the maximum sensitivities for the RATTY SU signal paths.

Table 3.13 suggests that gain increases of 23.5 and 31 dB are required for RATTY SU paths 1 and 2, respectively. Hence, an additional level-setting stage in the mixing FE should enable the maximum sensitivity to be realised. Assuming that the maximum input level remains at -27 dBm, the total adjustable dynamic range of the system would be ~ 100 dB. However, the achievable sensitivity may be affected by the FE or ADC SFDR and the instantaneous SFDR would significantly reduce for high power measurements.

There is an oversight in the out-of-band rejection ratio requirements of the filters, as this was determined to match the ADC SNR relative to the peak in-range power. To maintain this ratio at any linear power level, the filters must be capable of suppressing a maximum input level to 60 dB below the upper limit of a selected range. In the extreme case, if a measurement is conducted at the highest sensitivity with an out-of-range signal at the maximum input power, the required suppression is the adjustable dynamic range of the system. This ~ 100 dB requirement is stringent and presents a great challenge for future filter developments and switch selection, as the latter components do not meet the criteria either. However, this is only required of ranges that are mixed directly into the output range, since component responses will attenuate signals in other ranges. This is illustrated by figure 3.30 for the path 2 of RATTY SU.

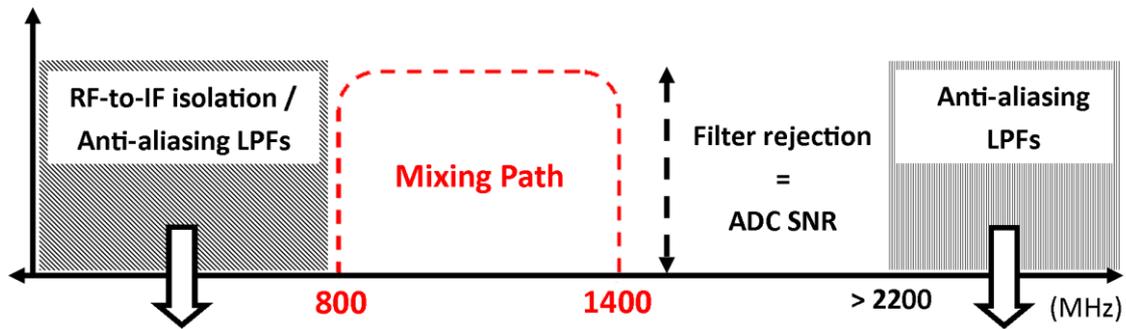


Figure 3.30: Schematic depiction of out-of-band suppression in the mixing path. Note that the 1600 to 2200 MHz band is a fundamental mixing product within the output range and is mainly suppressed by the filter chain. In the down-converted spectrum, signals above 2200 MHz and below 800 MHz will be also be suppressed by the anti-aliasing filters, with the direct leakage from the latter range lowered by the RF-to-IF isolation.

The out-of-band suppression of components will not be sufficient to maintain the more stringent rejection requirements and this affects both XDMs. For the current systems, the issue can be evaluated by observing the maximum signal levels in all paths, particularly for measurements in unknown environments.

Chapter 4

RATTY 2 system configuration, construction and integration

A block diagram representation of RATTY 2 operation is shown in figure 4.1. Although there is commonality between the first and the second generation RATTY devices, the latter RATTY features several significant improvements on the first, in addition to the FE designs described in chapter 3. This chapter documents the configuration of the remaining subsystems, practical aspects in the development of the analogue parts and the integration of the unit.

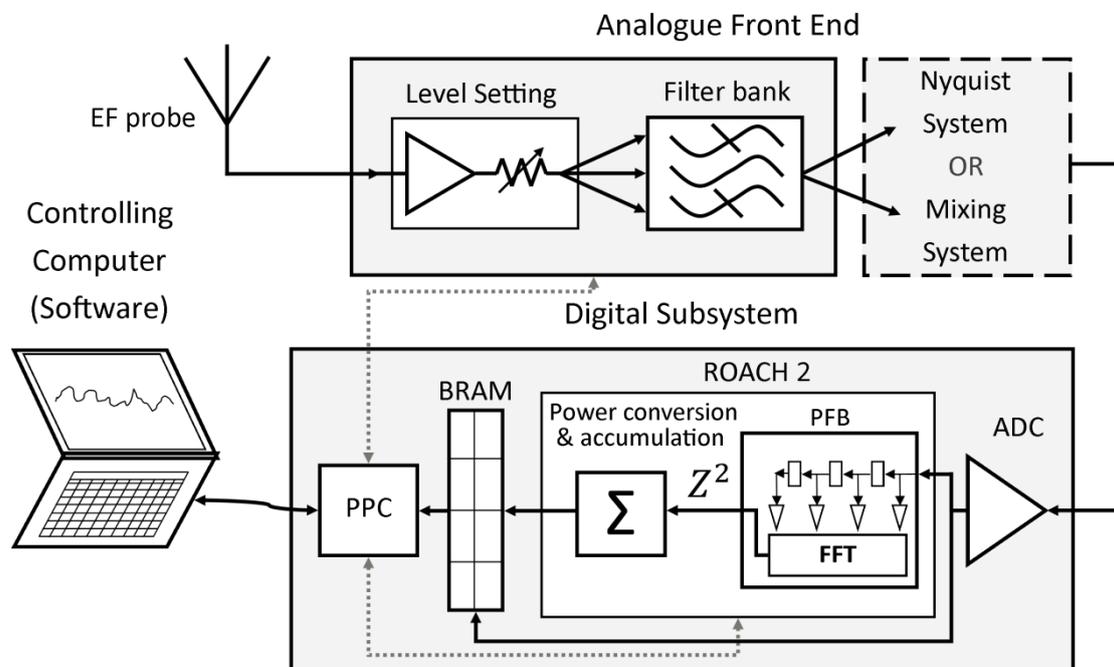


Figure 4.1: Block diagram of RATTY 2 XDM development. Note the two possible FE configurations and the dual signal paths in the digital subsystem for FD or TD operation, respectively.

4.1 Digital system description [12]

RATTY 2* was developed with the second generation ROACH board (ROACH-2) as its corner stone. The digital unit consists of two main subsystems: 1) an FPGA and a 2) PPC† CPU. The PPC has dedicated FLASH memory containing a bootloader, an operating system (Linux), and a root filesystem (Debian-based) containing tcpborphserver3. The ROACH uses open-source software exclusively and is capable of standalone operation. On power-up, the bootloader starts the OS, which mounts the filesystem and starts the

* The design of the digital subsystem for the project was undertaken by the SKA-SA digital backend team.

† Performance optimization with enhanced RISC Performance Computing

tcpborphserver3 network server. Tcpborphserver3 provides the software infrastructure for remote interaction with the device over Ethernet. It implements the Karoo Array Telescope Control Protocol (KATCP) for transferring Ethernet data between the FPGA and an external computer [32].

An application on a ROACH FPGA is designed with the MATLAB[®], Simulink[®], System Generator and EDK (MSSGE) toolflow. This enables a high level design process for the generation of FPGA code. A graphical interface is used to plan a device schematically using Simulink[®]. There are existing libraries of functional blocks for interfacing with ROACH hardware and for performing FPGA-based Digital Signal Processing (DSP). The output of the compile process is a .bof file, which is uploaded from the controlling computer and programmed onto the FPGA via the CPU [11].

The RATTY 2 device is configured for two modes of operation: 1) TD transient capturing and 2) FD spectrometer functionality. The TD output from the ROACH is a buffered array of unprocessed ADC samples with an oscilloscope-like behaviour. A trigger monitor can be set to evaluate the incoming data for fluctuations and allows the device to record selectively important events. With a portion of the data preceding a triggered event still retained in the buffer, the input signal leading up to the event is also captured. Thus, the state of the instrument leading up to an event, i.e. the quiescent state, can also be analysed. This is valuable for investigating self-generated RFI. It also ensures that subtler behaviour related to a triggered occurrence may be recorded, providing more information on the transient event as a whole.

In this early version of RATTY 2, the TD data is stored in the on-board FPGA black-RAM and this limits the realisable time record length. With several FPGA configurations compiled with different buffer dimensions, the recording length can be varied in discrete steps between 4 and 144 μ s by reprogramming the device. There are also larger Static RAM (SRAM) and DDR 3* Dynamic RAM (DRAM) capacities available as peripheral components to the FPGA. These could be utilised in future revisions to enable recording for durations in excess of a second.

A spectrometer is implemented on the FPGA for FD measurements and the process involves several sequential DSP operations. The buffered input samples are passed to an overlapping, 8-tap FIR filter. A 2^{16} point FFT is then applied to the filter output, together forming a PFB. With the total, nominal bandwidth of 900 MHz, a spectral resolution of ~ 27.5 kHz is realised. The voltage spectrum is then squared to convert to a power measurement and the latter is accumulated. This operation is described as real-time spectral averaging, since the calculation is done for successive input sequences and all the digitised data has a weighted influence in the output spectrum. Should the integration period be increased, the effect of averaging becomes more severe and the influence of an isolated transient event is less notable or vice versa.

The bit-width and numeric representation of the data is carefully controlled throughout the digital FD path to avoid that any corruption occurs, due to saturation or over-flow. The spectral output is in the form of 64-bit integers, which provides 384 dB of accumulated digital dynamic range. This ensures that there are no spurious components in the output as a result of the FPGA processing. A calibration function is applied to the FD output in the software, which described in the next section. The block diagram representation in Figure 4.1 shows the TD and FD paths in the digital subsystem block. Figure 4.2 shows the digital enclosure with the key components outlined and numbered.

* Double Data Rate type three

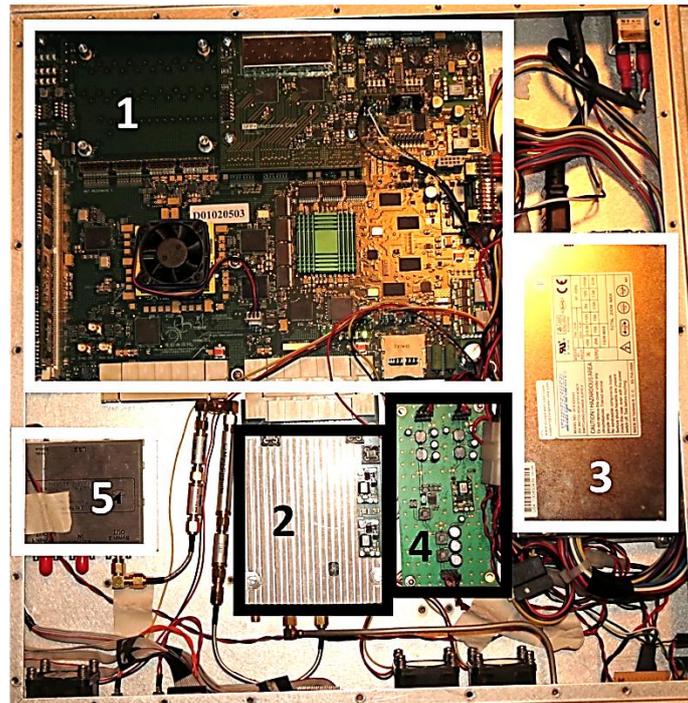


Figure 4.2: RATTY digital subsystem with key components outlined. The numbered parts are: 1) the ROACH-2 board, 2) ADC, 3) mains power supply unit and 4) FE power supply unit and 5) ADC clock source.

The CPU currently limits both TD and FD functionality, since the interface with the controlling computer can only facilitate transferring the output at a minimum rate approximately 1Hz. This constrains the analysis of transients, since the likelihood missing events becomes higher. Secondly, the accumulation time of FD operation must be kept at 1 s for the spectrometer to maintain real-time operation, which masks the responses of singular events. There are also 1 gigabit Ethernet (GbE) and 10GbE ports directly connected to the FPGA, that could be utilised in revised designs for higher FD and TD output rates by avoiding data dispatch by the PPC.

The front panel of the ROACH board features two mechanical switches used for powering on the devices and resetting it for a reboot procedure, respectively. There are also three sets of light emitting diodes (LED) that are used as visual indicators of the power supply, CPU and FPGA states. The switches each have an LED, with the power LED flickering to indicate when a supply is connected and remaining lit after power on. A green LED turns on when the CPU is idle or flashes when the processor is busy and a red LED flashes if an error exists.

4.2 Software description

The original RATTY software suite was adapted for controlling the second generation device.[†] It comprises three sets of files: 1) source code, 2) a default textfile, and 3) operating scripts. The operating scripts are run from a Linux terminal on the controlling computer. These are used to interact with the digital and analogue subsystems. The remaining files in the RATTY suite contain the additional functions and definitions needed

[†] The core RATTY 1 & 2 software was written by Mr Jason Manley of the SKA-Africa.

to support the operation. In addition, this software implements several SKA libraries that provide general functions in support of ROACH interaction. In this section, the key software features of the RATTY code are described with a high-level overview of the structure.

All the scripts were written in Python: a common, high-level programming language that is open source. Since Python facilitates object-orientated programming, the architecture of the software is shaped to incorporate this paradigm. It allows multiple data structures and functions to be attributed to a single name, i.e. an object. This is achieved by describing a class and assigning the object to this definition. In the RATTY software, the source code defines classes and functions used in the operating scripts, which avoids having large and cumbersome sets of code for every mode of operation [13]. By abstracting sections of the code, it is also easier to maintain generic definitions within each part and adapt the set to incorporate changes.

4.2.1 Source code and configuration file

A Control And Monitoring (CAM) script contains the definition of a class that is used for TD and FD operation. The two remaining source files are also imported in the cam code to perform calibration and configuration, respectively. This forms a link between the source codes and allows access to all the peripheral functions they contain. As the name suggests, the functions attributed to the cam class are central to interaction with the other subsystems. The key actions performed with these functions are briefly described in the section 4.2.2. The cam file starts with the initialisation of an object, which is executed by default when the class is assigned to it. This involves creating several attributes that represent the parameters of the analogue and digital subsystems configurations. Here the attributes are set with processes defined in the configuration file, which uses a default textfile as input for assigning some of the values. This allows several parameters to be changed manually, including the specification of a bof file and the directories of calibration files. There are also attributes related to the ADC and FPGA FFT, which are used for calibration.

Note that two different initialisations are referred to in this chapter: 1) digital and 2) software. To recapitulate, digital initialisation is the independent booting procedure that is executed when the ROACH is powered on.

The calibration source code contains a class definition that is used for the accurate representation and storage of data. It contains two main functions that support TD and FD data processing, respectively. Both functions correct spectra for ADC bandpass characteristics, the ADC voltage scaling factor and the FE configuration. In addition to this calibration, the FD function also compensate for the on-FPGA accumulation and the overall digital algorithm gain. A TD spectrum is obtained via an FFT algorithm that is imported from a Python library. There are several supporting functions in the class that perform interpolation, data smoothing, the extraction of calibration functions from files and the conversion of units, e.g. converting dBm values to dB μ V. Lastly, response of an input field probe can also be compensated, by uploading its transfer function to the calibration file directory, which contains all the correction files.

4.2.2 Operating scripts

Once the ROACH board has power up, the operating scripts are run sequentially: 1st) system initialisation, 2nd) FE configuration and 3rd) TD or FD measurements. Each of the files contains an options-generating function that enables numeric values to be passed to the script as user input from the terminal. A help feature presents a list with the available input options and a description of each possibility. In addition to these prompts, there is error logging implemented throughout the software package and this provides information on the origin of faults. The terminal commands and an example of one of the help menus are presented in appendix E.

The initialising script (*r2_init.py*) can be run once the ROACH board has been powered on and the CPU LED remains lit, indicating that its booting process was successful. It establishes communication with the ROACH, checks the state of the digital system, programs the FPGA and sets the values of some of the digital parameters in succession. These actions are performed by several cam functions and once the FPGA is programmed, this step can be skipped and the other commands can be continuously iterated.

A KATCP-serial script was developed for serial communication between the controlling computer and a microcontroller board in the FE, travelling via the CPU.* This forms part of the underlying SKA software and is used to configure the FE hardware, i.e. adjust the attenuator levels and select a signal path. It is called from within the FE configuration script (*r2_init_serial.py*), which accepts input arguments for the attenuation level of each attenuator and the selection of a desired measurement range. These requests are then translated by the KATCP script to the required pin assignment on the microcontroller board. This provides a means to reconfigure the pin layout of the microcontroller for a specific control implementation.

The two remaining operational scripts are used to access either TD (*r2_time.py*) or FD (*r2_spectrum.py*) data, respectively. They have similar architectures and can be used in one of two modes of operation. The first mode is for performing measurements, which writes the output data of the ROACH board to a file. The second mode is used to access and review files containing previously captured data. To avoid confusion, these actions are consequently referred to as measuring and reviewing, respectively. The two key functions that are performed by both scripts are plotting and file writing. Figure D.1 shows examples of both TD and FD plots. The file format used for data storage is Hierarchical Data Format 5 (HDF5), which is a dynamic file structure that has no size limitations and supports multiple data object types. This file format has widespread applications and there are libraries available for interaction with it on most common platforms, such as Python or MATLAB [14].

When a measurement is started, a cam object is created and a connection is established with the ROACH board. An HDF5-file object is initiated and configured to have the necessary data sets with attributes. Data is then read from the FPGA memory using a cam function and is temporarily stored in the assigned data fields of the object. This is followed by the initialisation of a figure window and the execution of a recursive plotting function (*DrawDataCallBack*). This function iteratively calls the reading function to obtain the next ROACH-data bundle, passes it to the calibration function, appends another data set to the HDF5 file and updates all the plots. It allows the system to run without continued user input, overwriting the object data fields for every ROACH board output and it is terminated when the process is killed in the terminal. There are distinct cam and calibration functions used for the TD and FD operations, respectively. Furthermore, the main data fields written to the HDF5 files are unique in either script, although many attributes are identical,

* The serial control was developed by the SKA-Africa digital backend team.

and the key fields listed in table 4.1. Nonetheless, the basic software routine is similar for measuring in either domain and is illustrated as a flow diagram in figure 4.3.

Data fields common to FD and TD HDF5 files:	
Timestamps, ADC overrange, ADC shutdown, ADC power level	
Data fields specific to FD HDF5 files:	Data fields specific to TD HDF5 files:
FPGA FFT spectra	ADC data
FFT overrange	Trigger levels

Table 4.1: Important data fields common or specific to TD and FD HDF5 files.

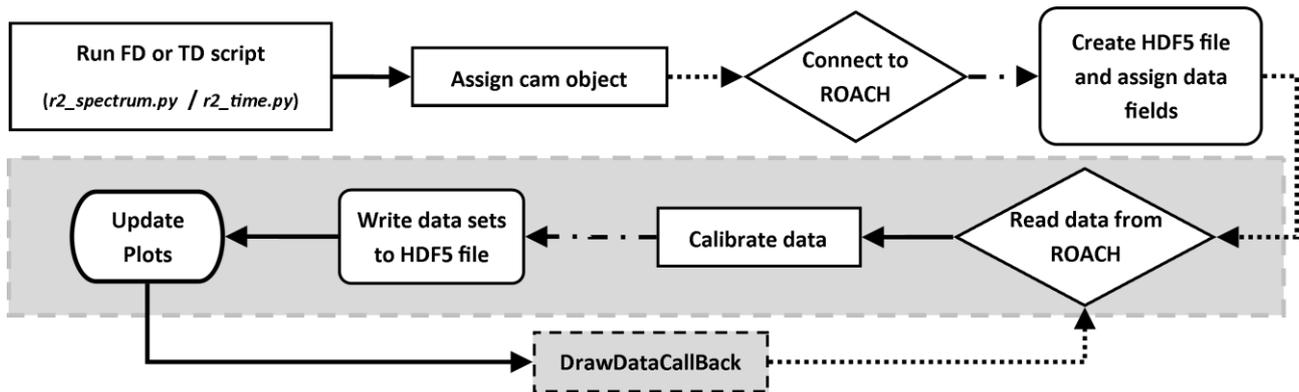


Figure 4.3: Flow-diagram representation of the measurement software routine. Note that the interactions between the software scripts, ROACH board and HDF5 files are graphically distinct.

In the reviewing mode, an HDF5 file is read by specifying the file option and its name. This action is identical to the measurement routine, except that the interaction with the ROACH board is replaced with the reading function and the file writing is bypassed. There are attributes that log both the ADC and FPGA status, are written to file and alert the user of fault states. These include ADC overrange, which results in data clipping, and FFT overflow. The user is warned of errors with text warnings and a red title box in the figure window. The help output in appendix E.1 lists several additional software features and there is a wide scope for future development in this regard. Lastly, several edits were made to the software package and these changes are referred to in subsequent sections.

4.3 FE layout, construction and testing

A preliminary design feature was the use of separate enclosures for housing the digital and analogue subsystems. The apparent lack of space in the digital enclosure necessitated this decision, although separating the high-power digital electronics from the sensitive analogue parts was also a preventative EMC measure. The following subsections document the design and practical issues regarding the construction of the mixing FE, although these are generally also relevant to the direct sampling system. Lastly, an initial test of the FE is documented in conclusion of this material.

4.3.1 FE layout and construction

An aluminium plate was used as a base for the construction of the mixing FE. The dimensions of the plate were based on the geometry of a standard, 1.3 unit ROACH enclosure, which is the structure intended to house a final system [11]. These enclosures are alodined to ensure corrosion-free and good electrical-conducting surfaces for earthing mechanisms. Furthermore, the cases are commonly used in SKA applications and enable a system consisting of several separate units to be placed into available rack mounts. To form a structure similar to the enclosure sidewall, an extra surface was added to the plate and bent perpendicularly to the base. This forms an L-plate structure used to mount both SMA and DC-supply connections from the digital system.

There are two main reasons for using the plate as the XDM base.† Firstly, the mixing system is elaborate and requires a significant portion of an enclosure surface area for mounting all the components. The space restriction encumbers the construction of the FE inside the case, since the flanges hinder soldering tasks or the tightening of bolts and screws close to the side walls. Secondly, the aluminium is inexpensive and error corrections or reconfigurations in an open layout are easily performed. In contrast, only the necessary holes are to be drilled into the coated enclosure and an initial plate could be used as a template for the final configuration. The open arrangement also better facilitates current-probe measurements of CM currents on interconnecting cables, which were conducted during the EMC investigations.

A dual-input FE layout was attempted that would enable housing the subsection in a single enclosure, as this was considered relevant for the comparisons of the XDMs. Using cutouts of component outlines, a multi-layered FE design was conceived. A key restriction considered in the placing of parts was the minimum bending radius of the semi-rigid, SMA cables used for connecting components. This restriction was avoided in several cases by using right-angle connectors and the initial, single-input layout was limited to approximately half of the available surface area.

All the active components were bolted onto the base plate, which also functions as the heat sink and the ground reference. The passive parts were connected in fitting positions in-between the fixed parts, conserving space and retaining bending room for the cables. This avoids placing excessive stress on the cables or SMA connectors and enables the removal of components. The stacked components were cushioned and stabilised with sponge-like tape at essential points. This also separates some active and passive parts, allowing the former to dissipate heat without direct contact to the latter. Figure 4.4 is a photograph of the stacked components.

The construction of the FE was governed by EMC principles. In addition to isolating the noisy digital components, this involved minimising all galvanic loops in the layout; a general strategy aimed at minimising magnetic coupling to the signal path. There are two DC-voltage levels generated in the digital enclosure and supplied to the analogue subsection. Each of these lines is connected into a star configuration from where active components are connected in parallel. A linear voltage regulator was fitted in the proximity of the supply-connections to provide specific supply levels. A more detailed account of the power supply unit configuration is documented in section 4.3.2. Figure 4.5 is a photograph showing several voltage regulators and twisted wire pairs in the FE.

† Mr Wessel Croukamp provided valuable practical input regarding the mixing XDM layout and construction, particularly for the assembling of semi-rigid cable connections.

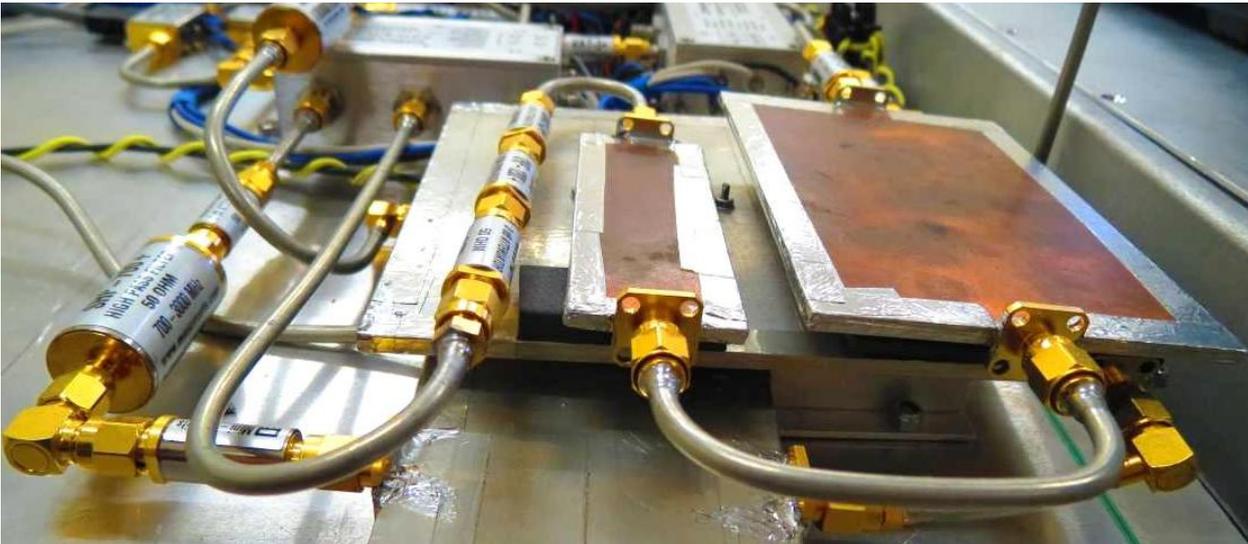


Figure 4.4: Stacked components with interconnecting SMA cables. Note that there are three layers in the forefront: 1) path 3 filter chain (top), 2) path 2 filter chain (middle) and 3) LO synthesiser (bottom). Also, the stacked sections are separated with sponge tape and there are right-angled connectors in the path 2 filter chain to reduce the layout area.



Figure 4.5: Linear voltage regulators and twisted supply wires. Note the proximity of several supply-wire pairs and SMA cables of the signal-path.

All supply-wire pairs were twisted and laid out close to the ground plane. This was also done for dual switch-control conductors, although the attenuator cables are shielded and it is impractical to twist the 4-way-switch lines. These were simply grouped together. A semi-rigid SMA cable with a solid outer sheath was used to manufacture the connecting cables. These cables provide good shielding and the connectors are expected to be the most vulnerable section to RFI, particularly at high frequencies [22]. Furthermore, the LO synthesiser contains digital circuitry and is placed with the analogue parts. This is contrary to the initial philosophy, but the addition of SMA connections between the enclosures is avoided and the selected module is housed inside a shielded enclosure, which reduces the risk of interference.

There is a trade-off inherent to the construction of a compact FE. Although short connecting cables present minimised geometries in the signal path, the space between supply lines is reduced. Either of these aspects could be more influential to interference levels and worsen the EMC performance of the system. This is a complex issue and could be assessed as part of future work. As suggested in section 2.5, the EMC measures

applied to the open conductors may only be effective for the lower portion of the system bandwidth. This emphasises the importance of using shielded enclosures and minimising any RFI close to the FE.

4.3.2 FE power supply unit and microcontroller board

It was important to keep the galvanic connections between digital and analogue enclosure to a minimum and limit current paths that could introduce RFI issues. However, the selected components required positive and negative voltage-supplies at several levels. This necessitated the use of switch-mode voltage converters for adapting the 12 V available from the standard ROACH-2 supply. From this, two output levels were generated that enable the use of linear regulators to provide specific voltages to individual components. Thus, a power supply unit featuring an 18 V and a -10 V output was developed.‡ There are LPFs on the input and outputs of the unit that suppress spectral products that are produced by the switching operations of the power supplies. The 900 kHz and several harmonics of this switching frequency are relatively low and are also attenuated by the LPF characteristics of the stability capacitors in the analogue section.

Of primary concern was the leakage of RFI related to the digital electronics, since previous tests on ROACH-2 boards had indicated emissions across the system bandwidth [11]. This could induce interference on the output conductors of the power supply unit subsequent to the integrated filters. Therefore, the live wires of the DC-supplies are passed through LPFs with broadband suppression, using feed-through capacitors in the sidewalls of both enclosures. The return conductors are terminated in earth studs to retain the reference potential to the ground plane. All the voltage-supply wiring in the digital enclosure was completed with the application of EMC measures previously described and with current fuses for fault protection. A single analogue power supply unit is capable of powering a dual-input FE system with the unit rated for 18 W and the peak power consumption of a single-chain, mixing layout is ~6.5 W. Figure 4.6 is a photograph of the L-plate interface of the mixing FE.

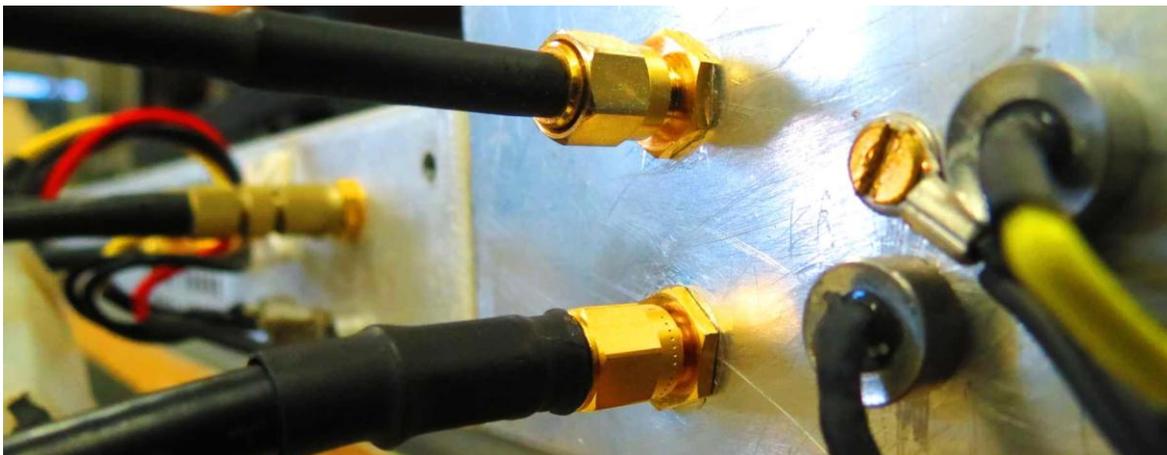


Figure 4.6: Photograph of DC-supply and SMA interfaces between the analogue (right) and digital (left) enclosures. Note the feed-through capacitors and earth stud on the far right.

The inclusion of the microcontroller board in the analogue enclosure was essential to enable FE configuration via the software interface.§ This introduced a digital component into the analogue subsection

‡ The PSU development was undertaken by Mr Renier Sieberts of SKA-Afrika.

§ The FE control board and related software was SKA-Africa developments.

with the potential of causing RFI. To avoid conducted interference, an optical connection was implemented between the ROACH CPU and the microcontroller. Thus, the board also features an optical transducer section. Although a customised IC was developed for this purpose, it was not available at the time of the system integration and a preliminary board was implemented [24]. The 16 MHz clock frequency generated on the initial module was of particular concern as a potential interference source. Figure 4.7 is a photograph of the control interface showing the optical fibres and microcontroller.

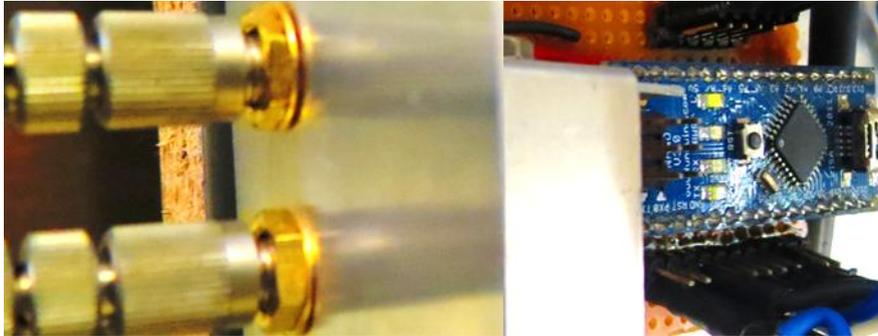


Figure 4.7: Photograph of the optical isolation of FE-control lines. Note the optical fibre connections (left), the metallic L-plate (middle) and the preliminary microcontroller board (right).

The pin layout was configured in the software script and matched to the hardware outputs. Since only 20 output pins are available on either of these microcontrollers, only a single-input FE can be independently controlled with one IC. A future improvement for a mixing system will require the development of serial control of the LO synthesiser, which can only be manually configured in the XDM. This is needed if more signal paths are added, requiring adjustable LO frequencies.

4.3.3 FE reconfiguration

After assessing the initial component and FE testing, the mixing XDM was modified. A broadband level-setting section was added to the FE with a step-attenuator and an amplifier following successively after the LNA. The modified section of the FE is shown in figure 4.8.

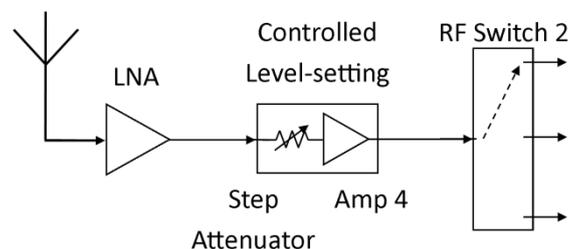


Figure 4.8: Reconfigured section of the mixing FE.

A readily-available amplifier with ~ 21 dB gain and an ~ 11 dBm OP1dB was selected. The 2.4 dB noise figure of the part does not have a notable effect on this parameter and was assessed through simulation. The input power remains limited by the LNA with the gain and adjustable dynamic ranges of all signal paths increased. The added amplifier will not reach compression for attenuator settings above 20 dB. Moreover, a signal can be amplified prior to mixing and this increases the headroom relative to the LO interference (Appendix D.1). The change was mainly implemented for the reasons mentioned above and the

characteristics of the revised system are detailed in subsequent sections. Further reference to the mixing FE includes these changes, with the exception of section 4.4.1.

4.4 FE tests

4.4.1 FE gain measurement

A preliminary FE test was undertaken prior to the integration of the subsystems or the FE reconfiguration described in section 4.3.3. The two main voltages and the control levels were supplied with bench power supply units, since neither the microcontroller nor the analogue power supply unit were available yet. All these measurements were taken with the attenuation at a minimum. A manual measurement of the path gains was performed at 40 MHz and 20 MHz resolutions in the passbands and in the filter-roll-off ranges, respectively. This was done with a SG tone as input and a SA measuring the FE output. Figure 4.9 shows comparisons of the simulated and measured gains of the first two signal paths. Note that the SG input was not calibrated and the gain functions are also not corrected for losses of the cables used to connect the equipment.

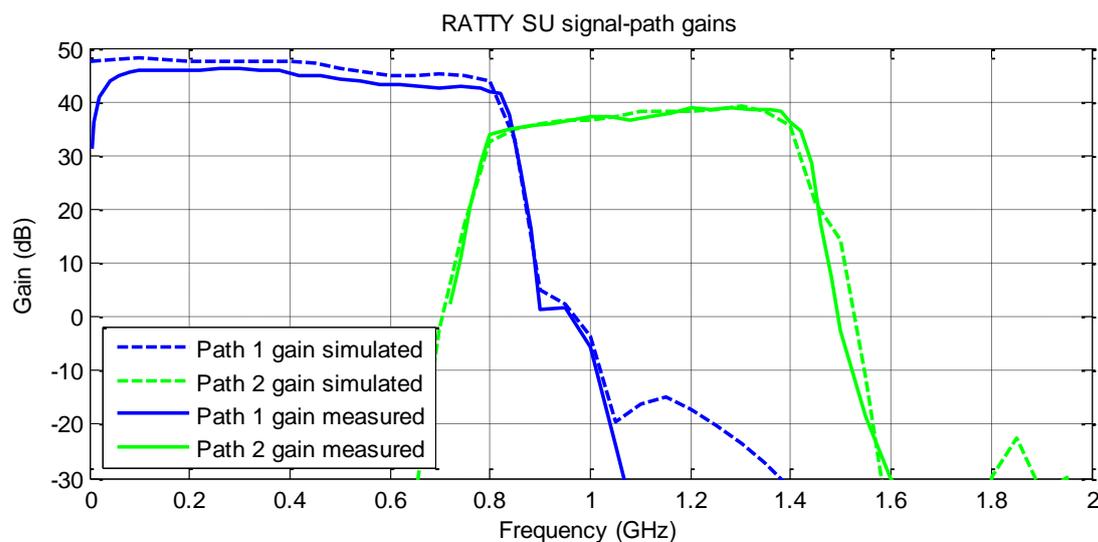


Figure 4.9: Comparisons of measured and simulated signal-path gains.

There is good agreement between the measured and simulated FE gains of both paths, although the path 1 result shows a ~ 2 dB difference. There are several factors that lead to discrepancies in the comparison, in addition to the measurement equipment that is not accounted for. The linear regulators of the active components do not provide the exact supply levels at which components were characterised and simulated. Also, a worst-case conversion loss value was used for the analyses with the measured parameter indicating a 1.1 dB variance for the path 2 bandwidth and the interconnecting-cable losses are not simulated. Thus, the path 1 evaluation contains fewer unknowns and is more insightful. In this case, the lower, measured gain gives an indication of the losses that result from the interconnecting cables of the FE. The difference at the lowest frequencies results from the lack of measured data in the simulation, which is only valid from 45 MHz.

4.4.2 FE SFDR measurement

A set of tests were performed to assess the actual SFDR performance of the mixing FE. The theoretical predictions of SFDR are based on a standard mathematical model with several assumptions. This only provides an expected range and may be inaccurate. Furthermore, the calculations are based solely on third order products and there are several harmonic and IMD products that may occur within the broad bandwidth ranges of the system.

For the measurements, a power splitter was implemented to combine the outputs of two SGs for IMD tests. The individual components also produce harmonic spurious and the superimposed input enables the characterisation of both distortions. It can be shown that there are constant differences between the two distortion products [4]. To ensure that only self-generated spurious elements were evaluated, the harmonic contents of the SGs were limited with the application of an LPF suitable to relevant signal-path range. Prior to each measurement, the output levels of the SGs were matched according to an SA readout. This device was also used for taking the measurements.

Since spurious products rise rapidly with increases in power, the maximum false outputs occur within the upper output limits of the FE. Figure 4.10 shows two mixing path measurements at different power levels with the main spurious elements identified.

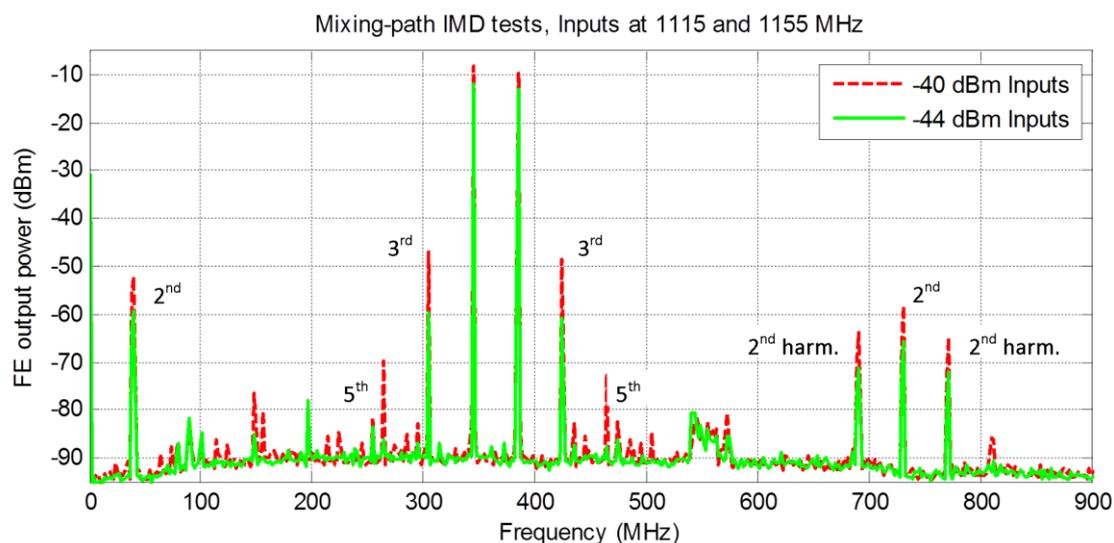


Figure 4.10: IMD measurements of the mixing path. Note that the orders of the highest spurious products are labelled and that the harmonic (harm.) elements are indicated. Also, there are some RFI elements present in the graphs.

The false products are identified according to frequency, although the amount of increase a tone exhibits also indicates the order of the spurious. In the logarithmic domain, the exponential growth of a spurious product is transformed to linear increases with a coefficient equal to the order [27]. For example, for the 4 dB power increase showed, the second and third order distortions increased by ~8 and ~12 dB, respectively. A crucial characteristic is the strong presence of the second order components, which are greater than the third order products at the lower power level. Moreover, these products will be the highest spurious if the power is reduced even further, due to the slower rate decrease. Thus, the standard SFDR calculations may not provide good approximations for all signal levels, since the second-order non-linear effects may be the highest

generated products in this broadband system and must be taken into account. Note that the harmonic products are ~ 6 dB below the related IMD levels.

Furthermore, there are mixing spurious products within the presented range. These components coincide in frequency with IMD products, due to the particular choice in input frequencies. This is not a general rule and mixer-related spurious does not overlap with IF IMD points for arbitrarily spaced inputs. The relation is also true for RF IMD elements, which are down-converted into the baseband. Figure 4.11 Shows a comparison between a graph from figure 4.10 and an FE output with one input frequency randomly selected.

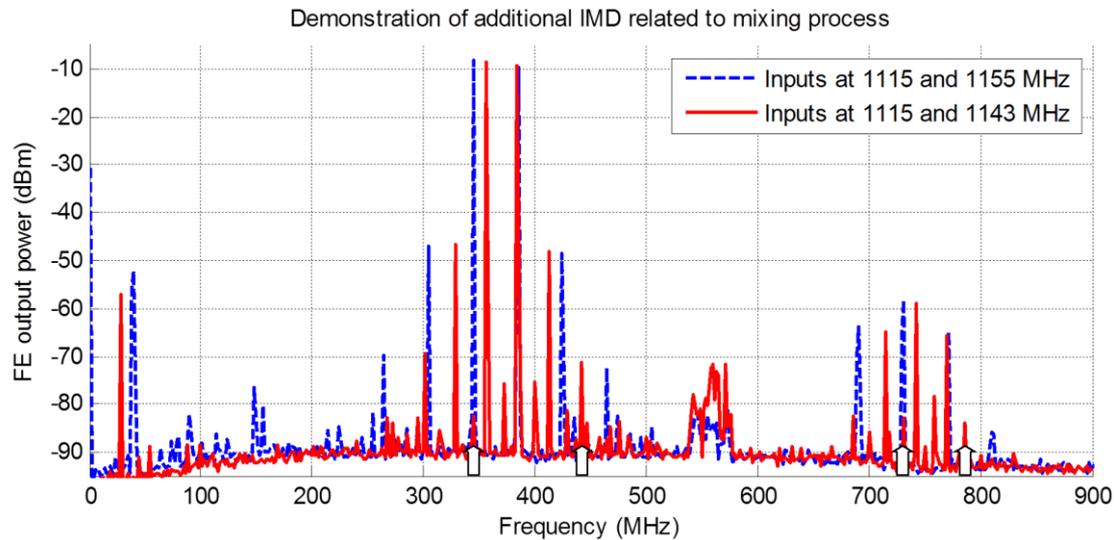


Figure 4.11: Comparison between mixed IMD from arbitrarily spaced and factor-ten spaced inputs. Note that additional spurious elements are indicated with arrows.

Figure 4.11 demonstrates that a mixing strategy results in additional false products; a combination of mixer spurious and RF IMD products. This is a weakness of this strategy and has implications if integration is attempted to improve the system sensitivity, since these elements will lead to low-level contamination of spectral channels. It is apparent from this plot that the most substantial IMD products are caused by the IF section and this is confirmed by the close agreement between the levels of the main distortion elements in the figure. Nonetheless, coinciding IMD components represent the maximum distortion magnitudes and the input frequencies used in the remainder of this section were chosen to represent this condition. Figure 4.12 shows two graphs; an example of variation in IMD levels and broadband mixing path outputs.

As shown on the right-hand side of figure 4.12, the broadband outputs contain no notable false components beyond the baseband. This is the combined effect of decreasing out-of-band gain and anti-aliasing filtering. Since the latter feature is not present in the direct-sampling FE, there may be distortion elements present that could alias into output spectra. Furthermore, the plot on the left in figure 4.12 demonstrates how the attenuation configuration affects the IMD products, in this case the third order distortion. For the 26 dB settings, the expected three-fold variation in the IMD can be seen relative to the unity increase in signal power. This relationship is not maintained for the alternate, 30 dB arrangement and this can be explained by a more detailed assessment.

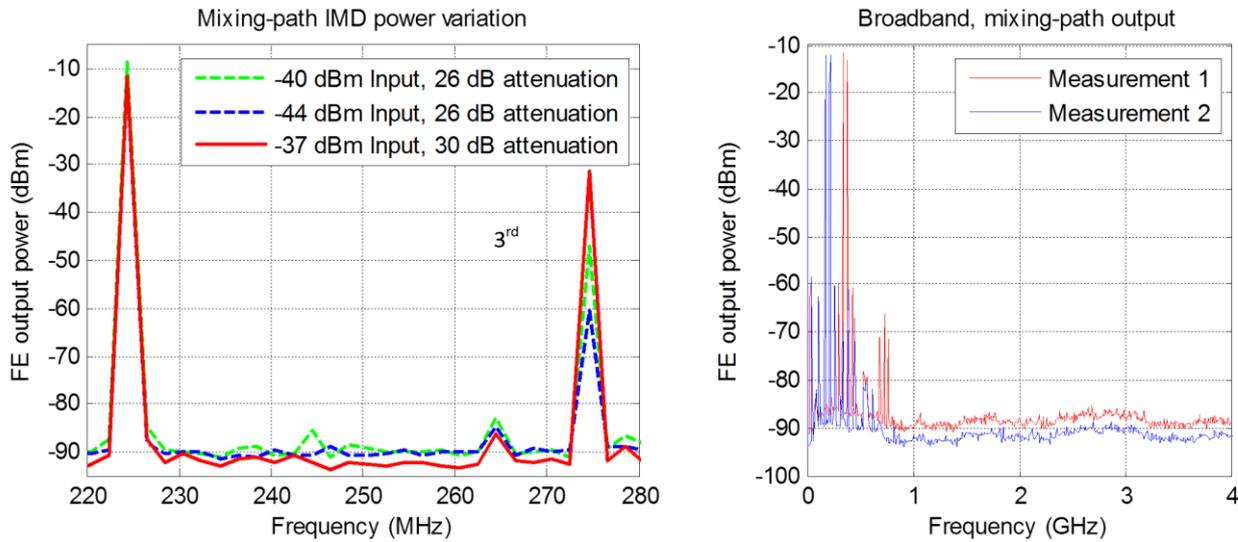


Figure 4.12: Varying 3rd order IMD levels (left) and broadband FE outputs (right). Note that the 3rd order distortion components are shown at three different input levels with different attenuation settings. The 26 dB attenuation setting was realised with 13 dB at both the pre (RF) and post (IF) mixer positions, respectively.

The 30 dB was divided with 20 and 10 dB step-attenuator settings at the RF and IF stages, respectively.

A general form of linear relationship that is the basis of equation 2.10 is presented in equation 4.1, as shown in [26]. This enables the calculation of intercept points for any order of spurious, based on measured data. The SFDR calculations of chapter 3 are based on datasheet predictions of P1dB values and the OIP3 values are approximated from this. This leads to a lack in the accuracy of SFDR estimations and a measurement was performed to provide measured data of the system intercept points. Furthermore, a second-order IMD SFDR calculation can also be performed by using the relationships in equation 4.1 to calculate the Output Second Order Intercept Point (OIP2). This is vital, since this product is significant and there is no data provided for this parameter in the component datasheets. Although there are general rules for approximating the OIP2 with the P1dB value, these are typically not as accurate as the OIP3 approximation.

$$OIP_N = \frac{(SFDR_N)}{(N-1)} + P_N \quad (4.1)$$

$$\text{with } SFDR_N = P_{Out} - P_N$$

where OIP_N = Nth order intercept point and $SFDR_N$ = SFDR between the fundamental output power (P_{Out}) and the Nth order IMD power (P_N).

A measurement procedure presented in [26] was followed and involved several sequential steps. In addition to filtering and matching the input power levels, the IMD characteristics were carefully assessed to ensure that the system was still well within the linear range. This required setting the input power level and attenuation configuration to preserve a 10 dB minimum margin from component compression points. The margin ensures that measurable distortion levels are generated and allows accurate characterisation of the intercept points. However, the levels were not calibrated for the output cable loss and this introduced some errors. Tables 4.1 and 4.2 present measured results with calculated intercept points and SFDRs for the mixing path. Appendix F contains the equivalent tables and interpretation for the baseband path.

RF and IF attenuation, respectively (dB)	Fundamental (dBm)	2 nd IMD (dBm)	3 rd IMD (dBm)
0 0	-2.4	-49.2	-49
0 3	-5.5	-54.3	-55.3
3 0	-5.4	-55.9	-57.5
3 3	-8.4	-60.7	-63.9

Table 4.2: Measured output levels used for calculating intercept points and SFDRs of the mixing path. (Input power = -60 dBm)

RF and IF attenuation, respectively (dB)	Measurement OIP2 (dBm)	Measurement $SFDR_2$ (dB)	Simulated OIP2 (dBm)	Simulated $SFDR_2$ (dB)
0 0	44.4	46.8	44.4	46.8
0 3	43.3	48.8	43.7	49.2
3 0	45.1	50.5	44.4	49.8
3 3	43.9	52.3	43.7	52.1
RF and IF attenuation, respectively (dB)	Measurement OIP3 (dBm)	Measurement $SFDR_3$ (dB)	Simulated OIP3 (dBm)	Simulated $SFDR_3$ (dB)
0 0	20.9	46.6	20.9	46.7
0 3	19.4	49.8	20.2	51.5
3 0	20.65	52.1	20.9	52.6
3 3	19.35	55.5	20.2	57.2

Table 4.3: Measurement and simulation based calculations of key intercept points and SFDRs.

The tables above show that the reduction of distortion levels was less effective for the IF than RF attenuation changes, which is related to the position and configuration of the level-setting components. Since distortion levels depend on signal-power and OIP levels, an attenuation change is effective if the power can be reduced without notably lowering the intercept point. This was the case for the RF attenuation variations demonstrated in these IMD measurements. However, RF level-setting adjustments also become less effective beyond 22 dB of attenuation, when the OIPs are lowered as a result. Thus, there is a limiting mechanism inherent to the efficiency by which distortion can be decreased and the characteristics of this process must be evaluated for specific configurations.

The spread-sheet calculator was used for predicting the OIPs and SFDRs, by implementing equation 4.1 and using the measured data as inputs. The results of the first measurement configuration in table 4.2 were used to determine modification factors for adapting the simulator P1dB to the OIPs. This resulted in excellent agreement between the measured and simulated figures for the first settings. However, the simulated results can be significantly affected by measurement errors with the corrections based only on a single observation. Nonetheless, the calculated parameters closely follow the trends of the measured figures for the settings presented in table 4.3. For the mixing path test, ~11.5 and ~35 dB was added to the P1dB as approximations of the OIP3 and OIP2 figures, respectively. The first value indicates that the 10 dB correction implemented in chapter 3 is a relatively good general rule for OIP3s. Similar results are presented for the baseband path in appendix F. These also show good agreement between the calculated and measured values for 3 dB variations in attenuation. Note that the calculations are only valid for the linear operation of the FE.

The baseband adaptations applied to the P1dB for the OIP2 and OIP3 figures were 17.2 and 10 dB, respectively. In comparison to the mixing path values, these also indicate that the OIP3 figure is reasonably predictable and that the OIP2 displays more variation. The IMD calculations are sensitive to changes in the OIPs of key components, particularly the amplifiers that limit the output of the system. Furthermore, variations in attenuator settings change the influence that particular components have on the generation of distortion products. Hence, a single adaptation based on one measurement does not describe and correct for this complex interaction between the FE components and leads to inaccuracies. This configuration dependency also explains the difference between the OIP2 adaptation values of the baseband and mixing paths. An improved analysis requires the characterisation of individual parameters, especially for the amplifiers, which should be considered in future developments.

To demonstrate these issues, the spread-sheet predictions were compared to several other measurement configurations with varying input powers. The measured and simulated parameters of the mixing-path are listed in table 4.4 and an equivalent assessment of the baseband path is presented in appendix F.

RF and IF attenuation, respectively (dB) [Input power (dBm)]	Measurement		Simulated	
	OIP2 (dBm)	SFDR ₂ (dB)	OIP2 (dBm)	SFDR ₂ (dB)
19 25 [-33]	25.3	44.3	25.7	44.7
25 31 [-33]	25.2	55.2	13.8	43.8
13 13 [-40]	36	44.1	37.1	45.2
RF and IF attenuation, respectively (dB) [Input power (dBm)]	Measurement		Simulated	
	OIP3 (dBm)	SFDR ₃ (dB)	OIP3 (dBm)	SFDR ₃ (dB)
19 25 [-33]	-2.15	33.7	0.3	38.5
25 31 [-33]	-8.55	42.9	-9.72	40.6
13 13 [-40]	11.3	38.8	13.6	43.5

Table 4.4: Comparisons of measured and simulated OIPs and SFDRs for the mixing path. Note the varying input powers and attenuation settings.

Table 4.4 shows that calculated OIPs and SFDRs approximate the actual parameters reasonably well. However, the SFDR₂ and OIP2 levels are less accurate for the second measurement, since the latter figure is not well characterised for RF attenuation settings above 22 dB. Nonetheless, the simulation provides useful estimates of the generated distortion levels. All of the baseband results demonstrate good coherence between simulations and measurements, since the additional measurements have configurations with similar characteristics to the settings for which the OIP adaptations were calculated. Table 4.5 lists two second-order comparisons that demonstrate that the baseband simulation is also imprecise for certain settings.

RF and IF attenuation, respectively (dB) [Input power (dBm)]	Measurement		Simulated	
	OIP2 (dBm)	$SFDR_2$ (dB)	OIP2 (dBm)	$SFDR_2$ (dB)
0 20 [-54]	17.8	27.8	20.3	30.3
0 0 [-54]	39.5	43.7	27.7	31.9

Table 4.5: Baseband comparisons of measured and simulated parameters. Note that the second comparison deviates in a similar way to the inaccurate calculations for mixing-path shown in table 4.4.

Table 4.5 shows that the second-order calculations are also inaccurate for low attenuation settings. An OIP2 adaptation of ~ 29 dB is required to match the parameters at zero attenuation, which is similar to the mixing-path value. It is interesting to note that the third order predictions are less dependent on precise component values and show less deviation for configuration changes. Figure 4.13 is a plot of the worst-case SFDR for the mixing-path, which is measured at near-maximum input power. An equivalent figure for the baseband is included in appendix F.

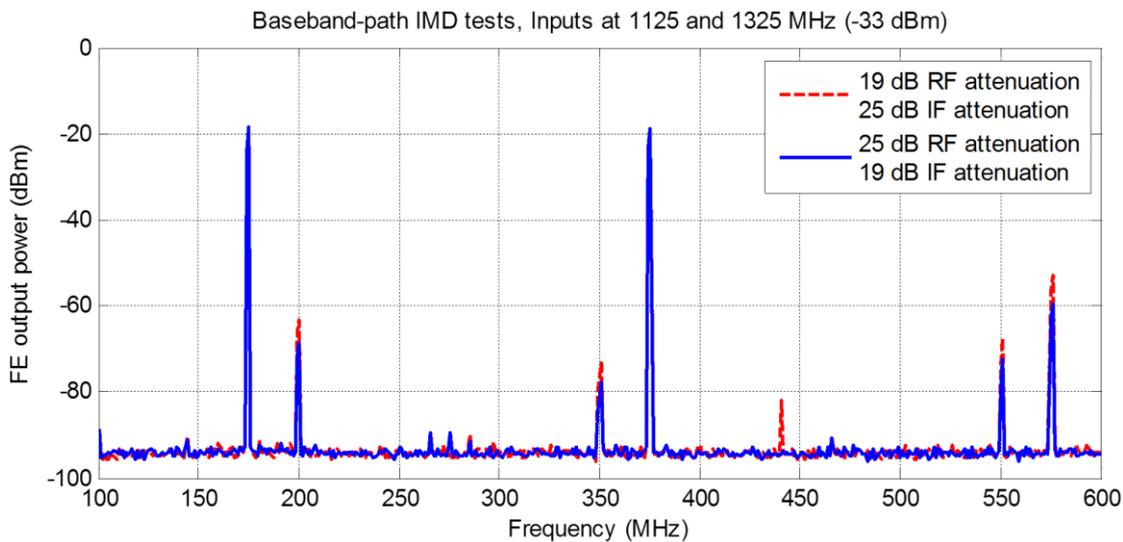


Figure 4.13: Worst-case IMD levels of the RATTY SU mixing path. Note that the second attenuation configuration only results in a slight improvement in the SFDR.

The solid-line graph in figure 4.13 represents the best SFDR performance the FE mixing path can deliver at the maximum input level and with highest distortion component matching the theoretical ADC lower limit of -60 dBm. For the complete system, the realised FE SFDR for this measurement will be slightly worse than the ~ 41 dBc demonstrated here, since the attenuation must be reduced to match actual ADC limit. Due to the higher gain in the baseband, the SFDR limit is ~ 23 dB for near-maximum input power.

The current FE SFDR remains limited by the OIPs and both these figures are more sensitive to changes in the IF than RF attenuation, as a result. However, the RF attenuation has a greater influence on the noise figure of FE than the IF stage and the optimal configuration depends on the application. Nonetheless, the maximum FE SFDR will be realised at the minimum attenuation. If this is calculated to match the minimum ADC range (-2 to -54 dBm), the SFDR limit is 48.4 dBc for the mixing path and 46.8 dBc for the baseband path. For these figures, the second-order distortion is the highest. If appropriate attenuation settings are implemented for increasing power levels, the realisable SFDR will gradually decrease towards the worst-case values specified in the previous paragraph. In this process, the third order distortion will also increase beyond the second-order level. However, the selection of attenuator settings is nontrivial and implementing and optimisation algorithm to perform this task should be considered for future improvements.

In conclusion, this evaluation indicates that the FE poses the SFDR limit for the mixing system and can at best come within several dB of matching the minimum ADC range. However, this statement is based on the combined range of both subsystems. For FE outputs below the upper limit of SFDRs defined in this manner, the ADC will pose the limit on the measurement range. This is because the ADC SFDR reduces for lower power signals, whilst the FE range increases

The behaviour of the baseband channel gives an indication of the operational characteristics that is expected of the direct-sampling FE. The OP1dB of the direct-sampling system is slightly higher than the mixing FE output limit, and a similar SFDR is anticipated. This is also dependent on the characteristics of the output amplifier and the consecutive level-setting stages in this system will exhibit similar interactions to those of the mixing system. However, the SFDR characteristics of the final amplifier can be retained if the attenuation of the preceding level-setting blocks does not reduce the OP1dB. A detailed characterisation of the final amplifier is required for more detailed analysis of this model.

4.5 XDM comparisons

A comparison of the developed XDMs was required to determine which design would satisfy the goals of the project best. The evaluation involved a set of elementary measurements that could be used to assess several key parameters including the bandshape, spurious performance and RFI quietness. The intent was to observe the basic characteristics of the systems to motivate the use of one of the strategies for future developments.

The tests involved performing two consecutive, single-tone measurements with each FE. This methodology enabled the generation of difference plots by performing spectral subtraction. A difference plot reveals only input related products, assuming that all other spectral components are constant features of the particular XDM and equal for both tone measurements. An SG was used as input to the systems and an LPF was added for harmonic suppression. With high spectral resolution and dynamic range, the spectrometer operation was favoured for this comparison. A 15 minute measurement period was implemented for integration, providing spectra with reduced noise-floor variations and accurate representations of low-power signals. The accumulation and subtraction were performed in the post-processing of the HDF5 files.

These measurements were performed in a reverberation chamber at SU, which was used as a shielded room. The reverberation chamber features a filtered power source and broadband shielding (> 40 dB), which suppress conducted and radiated environmental interferences, respectively. Since these shielding features were implemented, any measured interference signals are assumed to have been caused by one of the RATTY subsystems. Also, the SG was positioned outside the room and connected via shielded SMA cabling. Figure 4.14 shows the arrangement inside the reverberation chamber.

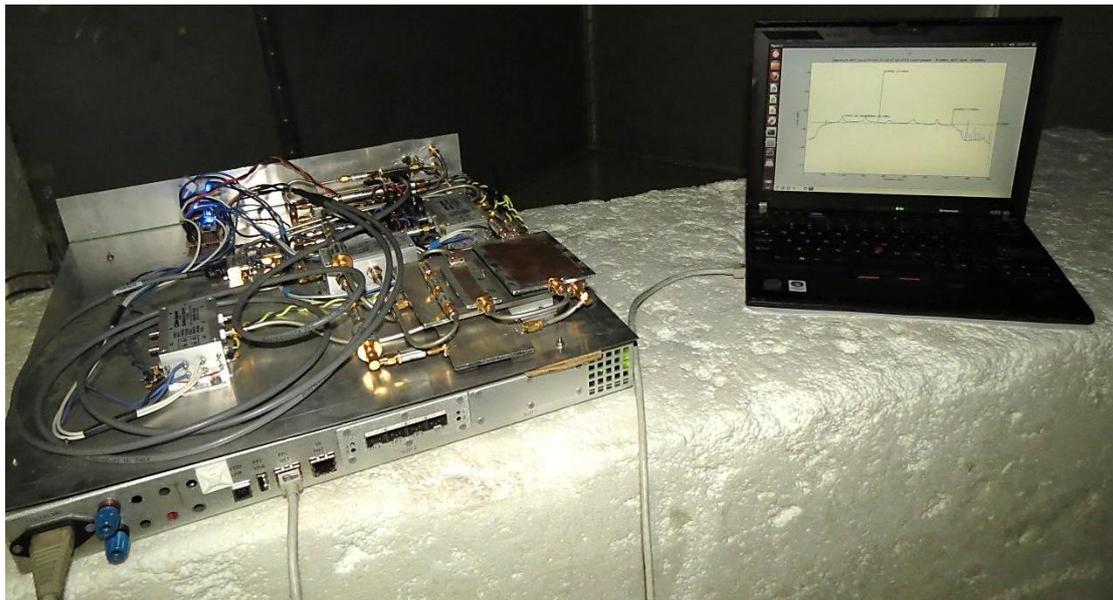


Figure 4.14: The RATTY measurement arrangement in the reverberation chamber. Note that the FE is placed on top of the digital sub-system (left) and that the laptop is battery powered. The SG tones were injected on the far side of L-plate, which is not visible.

The tone frequencies were selected to be non-harmonically related during the subsequent measurements, to avoid coinciding ADC-spurious products. However, for the 1050 to 1400 MHz input range of the mixing path, all $|n \cdot \text{LO} - n \cdot \text{RF}|$ mixing-spurious products will overlap with the n^{th} ADC harmonics. This is not the case for inputs from 800 to 1050 MHz, since the corresponding 450 to 700 MHz output frequencies will produce aliased ADC harmonics and the mixed products mentioned above will be suppressed by the anti-aliasing filtering. This will not have any significant effect, as mixer has been shown to produce low-power distortion that will not add notably to the ADC spurious.

There are also FE spurious components that coincide with ADC products in the direct-sampling design. For the digitisation of the second Nyquist zone, the n^{th} harmonic of ADC distortion occurs at the same frequency as the aliased n^{th} harmonic of the FE, i.e. $|n \cdot F_S - n \cdot F_{FE}|$. This is also true of the baseband paths of either system. A spread-sheet calculation is useful for estimating the spurious levels of the FE, since the power levels of the latter are similar and could be higher than the false products of the ADC. Thus, a calculation can indicate which source of distortion is dominant for a particular configuration and signal power. There are frequencies produce several superimposed false components at one point, e.g. 720 MHz for 1.8 GSa/s sampling. However, this was not assessed and it assumed that the SFDR will not decrease notably, since the magnitude of higher order products reduce rapidly.

The system configurations were set to result in similar noise floor values and output magnitudes, although the difference in the gains required that a higher input power be used for the RATTY SU measurement. Figures 4.15 and 4.16 present plots of single-tone outputs of the direct-sampling and mixing systems, respectively.

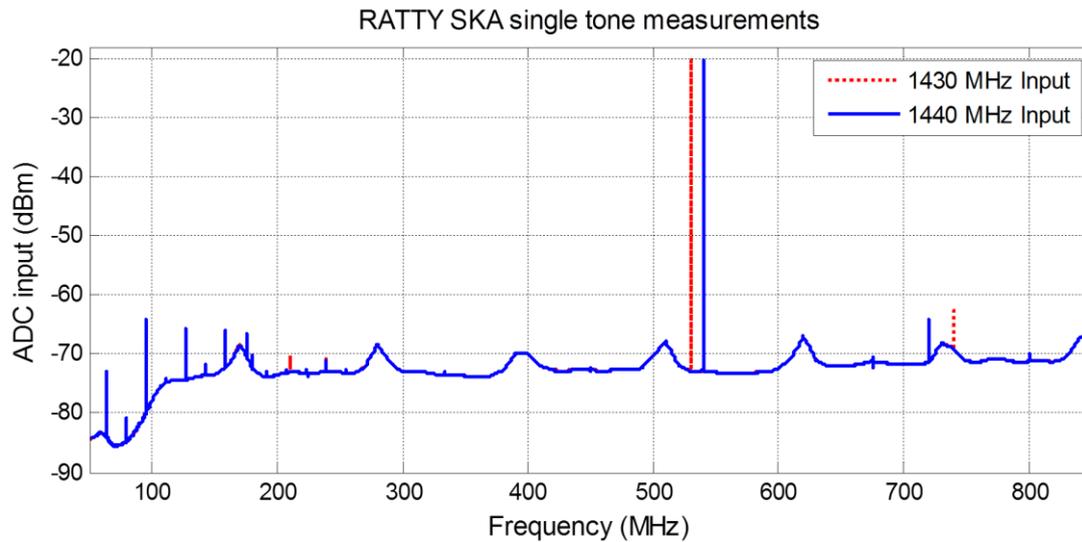


Figure 4.15: Single-tone measurements with the second signal path of the direct-sampling system. Note the SFDR of ~ 43 dB, the aliased output frequencies and the RFI tones on the left of the plot.

There are several important features in figure 4.15 that must be clarified. Firstly, these measurements were conducted with the second signal path and the range presented here is the ADC baseband, thus the tones are at aliased frequencies. The lobes in the noise floor spaced at ~ 120 MHz are caused by a gain function that is implemented to calibrate the spectrometer data, which is discussed in more detail in section 5.2. The input noise floor can also be calculated by considering the effect of the polyphase filter bank. If a ~ 44.6 dB reduction in the noise power is assumed, the input noise floor is estimated at -28 dBm. At this level, the analogue noise will have some dithering properties, although the ADC typically requires higher noise power for a notable improvement in the SFDR [28]. The integration yields a smooth noise floor and reveals the bandshape of this signal path. Note that aliased upper corner frequency (1670 MHz) is distinguishable, although the lower cut-off (900 MHz) is not clear. The lower range is obscured by aliasing, since the measurement was performed at $F_s = 1.8$ GSa/s and not the intended 1.71 GSa/s. This change was not essential to test the operation of the FE and the bandshape exhibits 3 dB of gain variation for the input range of 1000 to 1670 MHz. The ADC calibration is disregarded in this assessment, as it is common to both systems.

An average SFDR of 43.1 dB is demonstrated in the graph and the highest false components are the second order harmonics. These are related to the sampling process, as the second-order FE harmonics that also occur at these points are well above the operational ranges of several output components and duly suppressed. Note that this does not represent the best ADC performance, which requires a higher input level. It is apparent that the spurious products also appear at aliased points in the spectrum. Although it is trivial to correct the spectrum to display the actual input frequencies, the baseband range reveals the nature of the RFI components in the system. There are several harmonics elements to left of the of figure 4.15 that were present at equal levels for both measurements. These components are all 16 MHz harmonics with the most prominent levels at 64, 96, 128 and 160 MHz. It was concluded that this RFI is caused by the 16 MHz clock signal of the FE control board.

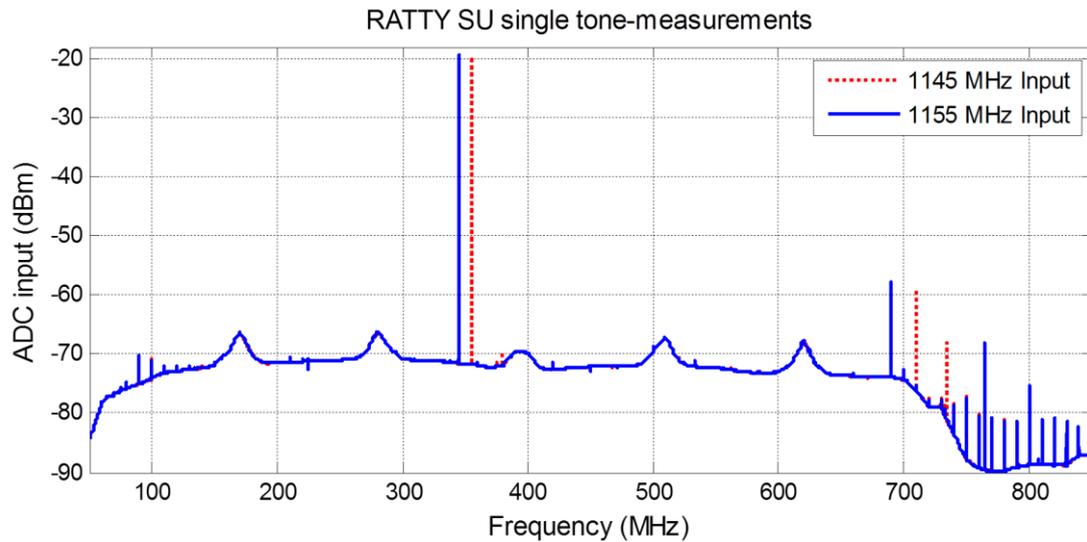


Figure 4.16: Single-tone measurements with the second signal path of the mixing system. Note the SFDR of ~ 39 dB, the down-converted output frequencies and RFI in the plot.

The plots in figure 4.16 also show the ADC calibration function and the input noise floor is estimated at -29 dBm for a ~ 43.4 dB reduction. The bandwidth (800 to 1400 MHz) and bandshape of the mixing path are distinct and the latter displays roughly 3.6 dB of variation. There is a downward slope towards the corner frequency at 700 MHz and this is caused by the anti-aliasing LPFs. Nonetheless, these components are essential for suppressing out-of-range distortion components, although the benefit of this is not demonstrated here. In this case, the average SFDR was 38.9 dB, which is 4.1 dB less than the range of the direct-sampling system. A spread-sheet calculation estimates a 60 dB FE range for this configuration and output power. Thus, this difference is related to the variation of the ADC performance and the main false are related to sampling. Figure 4.17 shows a higher resolution graph of the mixing-path, low-frequency RFI.

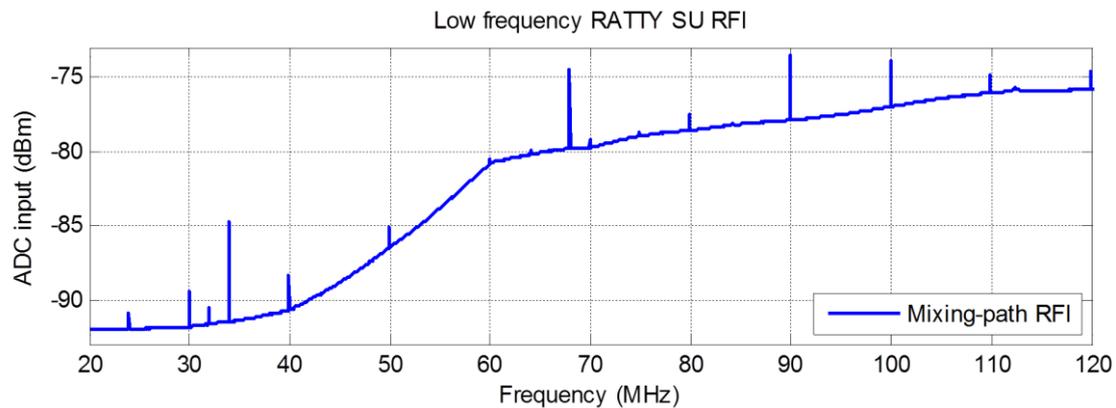


Figure 4.17: Low frequency RFI in the mixing path of RATTY SU. Note that the clock-signal interference of the synthesiser and control-board (32 and 64 MHz) can be seen as 10 and 16 MHz harmonics, respectively.

The self-generated interference in the RATTY SU system is clearly different to the issues exhibited by the direct-sampling system. Although some harmonics of control-board clock can be visible in figure 4.17, the main concern for the mixing measurements is leakage of the 10 MHz reference signal of the LO synthesiser. This occurs as broadband interference and is particularly problematic above 700 MHz, where the LO HPF is ineffective. As suggested in appendix D.1, this problem must be resolved with narrow BPF applied to the LO. Figures 4.18 and 4.19 show the difference plots for the SKA and SU XDMs, respectively.

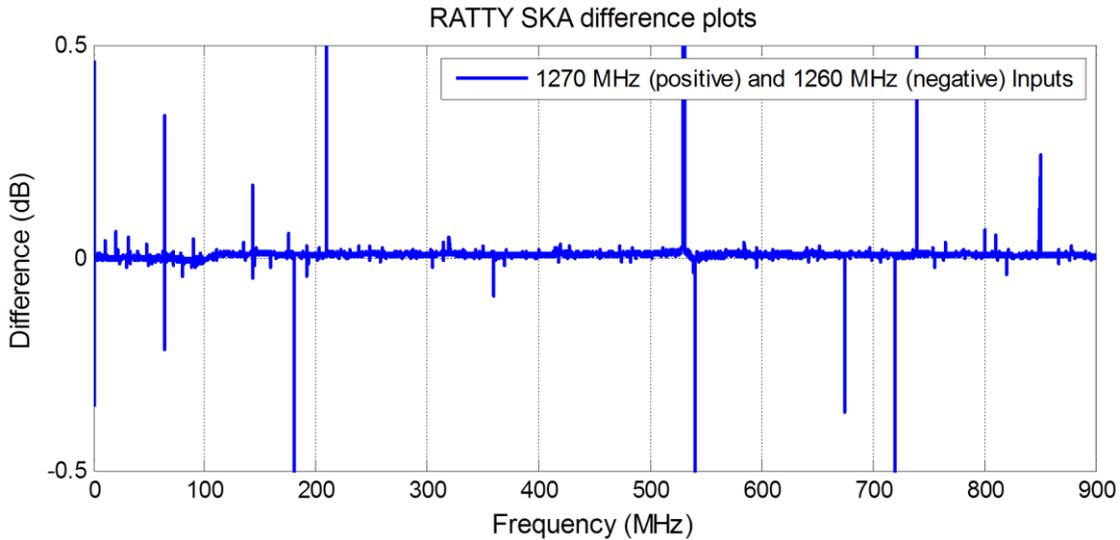


Figure 4.18 Difference plot of RATTY SKA single-tone measurements.

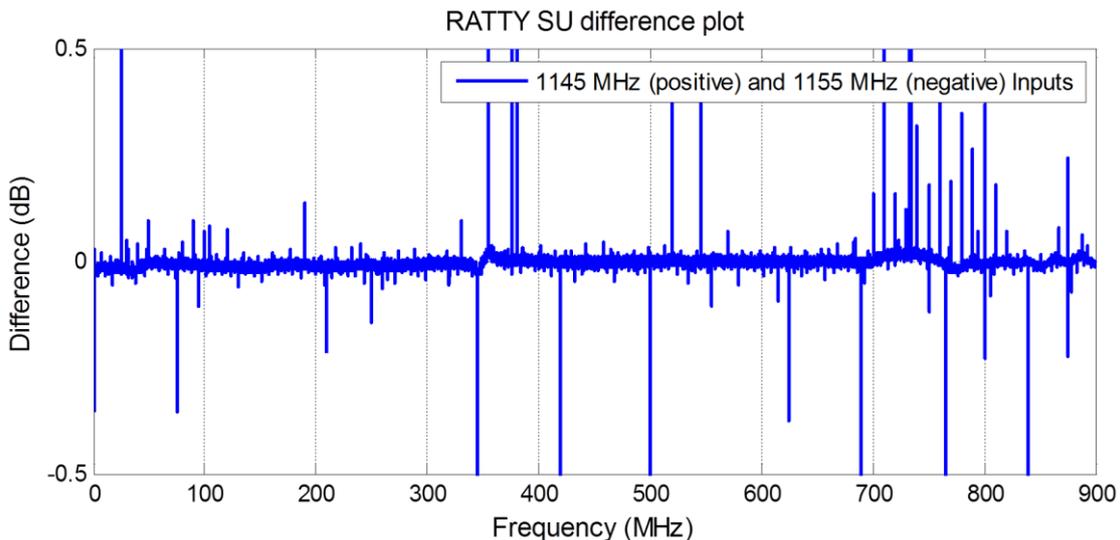


Figure 4.19 Difference plot of RATTY SU single-tone measurements.

The difference plots indicate a higher the noise floor of the mixing system than for the direct-sampling FE. This demonstrates the difference in the T_{Sys} of the measurements and that the mixing path requires longer integration periods to improve the sensitivity equivalently to the direct-sampling path. Note this is only an indication of the difference between the systems and not an accurate comparison. Furthermore, the RATTY US system contains more spikes in the difference plots that also contribute to the noisiness. This is caused by variations in the levels of the LO reference RFI, which was confirmed upon closer inspection of the 10 MHz harmonics. These components occur throughout the baseband range; the level of the leaked power varied for consecutive measurements, although the configuration was not changed. If this mixing system were to be modified with additional amplification for improved sensitivity, the general RFI issues would be exacerbated. The control-board leakage would also increase to levels equivalent to those in the direct-sampling XDM.

Lastly, there are small perturbations in the noise floor at the base of the larger fundamental ADC outputs. This demonstrates the low spectral leakage that occurs with the use of a polyphase filter bank. However, any noise-floor variations, spurious or RFI features limit broadband sensitivity that can be gained with

integration. Moreover, these single-tone tests demonstrate that a sub-ADC SFDR measurement interpretation will require detailed analysis, particularly for the characterisation of arbitrary signals.

The lack of depth in this evaluation limits the complexity of a comparison between the XDMs. Nonetheless, there were several key observations that allowed a sensible decision to be made for the development of more refined systems. Firstly, there are several RFI issues that require further EMC mitigation and this is clearly a greater issue for the RATTY SU design. Either of the XDM models can benefit from refined filter designs and FE output stages. The FE SFDR for such systems could be identical. This was shown in section 4.4 with the distortion performances shown to be component dependent. However, there are additional mixer spurious components that could influence measurements with long integration periods.

In addition to the aforementioned issues, a mixing system is also more complex, power consuming and costly. The first of these differences is of particular importance, since the more complex system presents a greater challenge for further EMC improvements, has a greater risk of failure and requires more expertise in a production environment.

Nevertheless, a mixing design has offers the possibility for bandwidth development, because it is not limited by the ADC bandwidth. Furthermore, ADC performance is generally better at lower frequencies and mixing systems can utilise this for an entire system input range. There are more readily available components for IF frequencies, whereas a direct sampling system requires more specific parts. The bandwidth coverage of a mixing system is also more efficient than a direct-sampling strategy, since there are no guardbands that require additional signal paths for continuous spectral coverage.

In conclusion, a direct-sampling strategy suites the project requirements and resources best for short-term developments. An improved mixing design will likely have to be developed for a RATTY system requiring a much larger bandwidth, due to ADC limitations. The requirements for application may favour either one of the designs, although a direct-sampling strategy should be the first choice.

Chapter 5

RATTY SU calibration, EMC testing and application

5.1 Calibration

A static calibration procedure was followed for RATTY SU, which entails the characterisation of the signal-path gain functions. These set functions are implemented to adjust spectra in order to account for the system response and reflect the true measured spectrum. Similarly, linear factors can be derived from the calibration functions and used to scale time domain signals. Several methods were considered with broadband and single-tone sources available. Ultimately, a swept-tone approach was used and allowed the precision of the calibration to be tested against trusted measuring devices. The method involves performing tone measurements at a set frequency resolution across the bandwidth with RATTY and a reference device. The difference between the measurements is the gain function.

A Rohde and Schwarz (R&S) SA was used as the reference for the calibration and its precision was compared to a calibrated power meter. The comparison was made at several points throughout the 100 MHz to 1400 MHz range with a SG output level of -10 dBm. An average variation of 0.2 dB was observed with a maximum difference of 0.28 dB. This was considered to be acceptable, as inherent variation in the system response was unknown and improved accuracy may require a different calibration method. An R&S SML 03 SG was used and the harmonic performance of the device was assessed prior to the power-meter comparisons. There were no harmonics within 40 dBc of fundamental outputs and this would not amount to notable errors in the power-meter readings.

All the equipment was powered on for several minutes prior to taking measurements, to limit deviations caused by temperature variations. The measurement layout included two SMA cables that were not fixtures in the system. These cables have good phase stability and low-loss characteristics, with one used as the digital-to-analogue-subsystem connection and the other used to connect the SG. The latter cable was included in the reference measurements and did not affect the calibration, although the connection between the subsystems formed part of the gain function calculations. If a different cable is used in RATTY measurements, it must be characterised and the calibration corrected for this change.

A complete calibration procedure would involve performing measurements at all the possible attenuator settings, power levels and with a fine frequency resolution. This is not practical and would require performing several hundred measurement sets. Moreover, the range of the reference SA is limited and the device requires careful adjustment for accurate measurements at low power levels. Subsequently, a limited set of measurement were performed that expediently sampled the operational range of the RATTY SU system. This required that attenuation configurations be corrected with the measured responses of the components. The philosophy was to indicate the extent of errors that result from non-linearity and the corrections for attenuations settings. This was applied to the TD operation of RATTY SU, although a single calibration was also performed in the FD mode.

A 2^{16} -point FFT was applied to the TD data, to sample the peaks of the tone powers accurately, although the sweeping resolution was only 10 MHz. This was selected as a sensible spacing, which limited the amount of calibration points and retained enough accuracy to portray the shape of the gain functions.

The input power used for the calibration was adjusted to four levels, which represented high, mid and low level RATTY inputs. The minimum was restricted by the SA limit. Several attenuation configurations were applied and the highest settings were implemented at the maximum power levels, to avoid saturating the active components. The specific attenuators applied in the RATTY SU FE were not characterised for the calibration. A comparison of the responses of two characterised components showed maximum variations of > 0.1 dB within the system bandwidth range. The differences increased for higher settings and frequencies and this introduces errors in the calibration process. Furthermore, equation 2.8 was used along with the measured port characteristics of the components adjacent to the attenuators to correct the loss functions of the latter for mismatches.

The measured gain functions were corrected for the attenuation settings, in order to represent the maximum gain. This demonstrated the reversed calculation for static calibration and the correction accuracy was compared in the process. Figures 5.1 shows the gain functions calculated for the baseband and mixing paths.

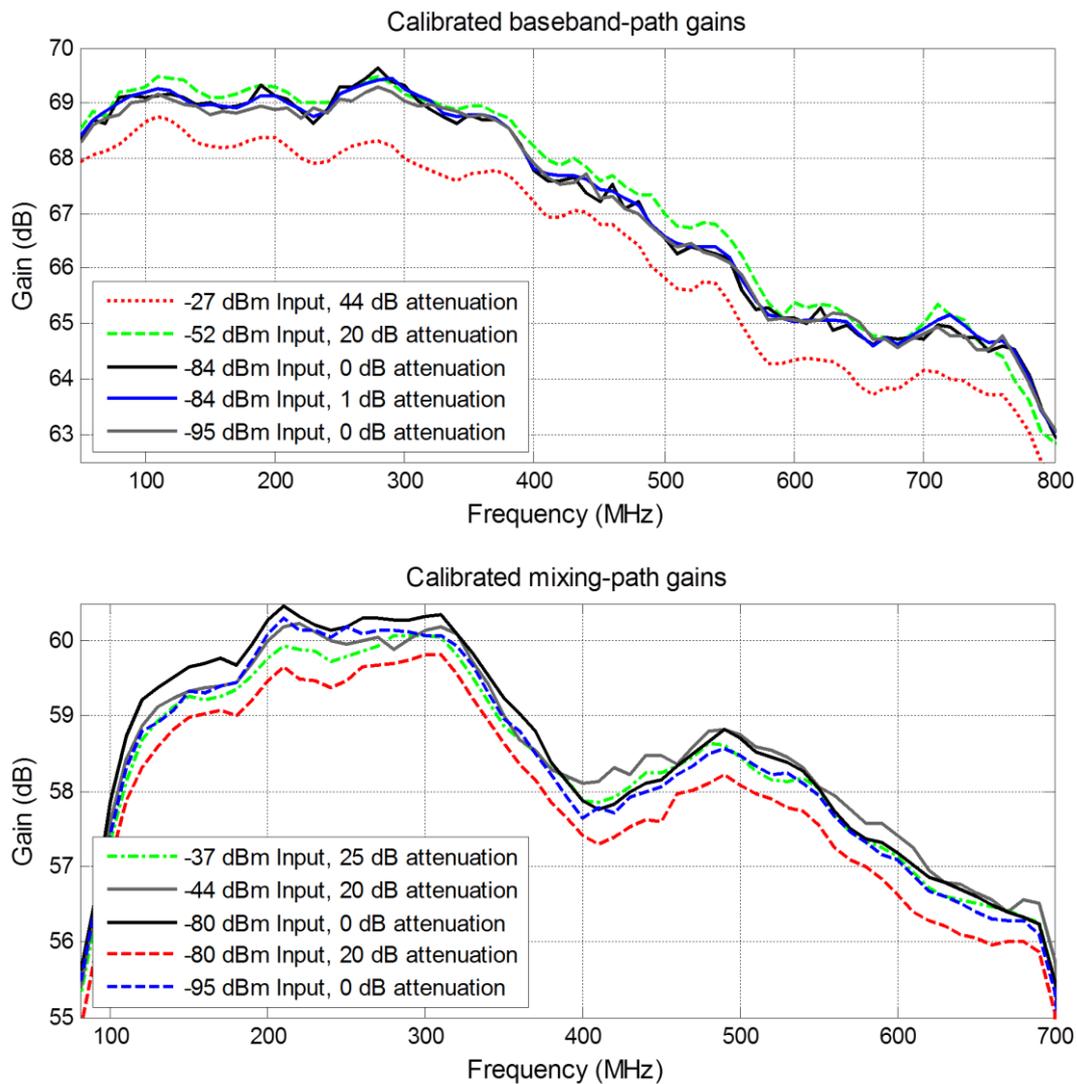


Figure 5.1: RATTY SU baseband (top) and mixing-path (bottom) gain functions.

The baseband calibration produced variations below 0.5 dB, with the exception of the maximum input power (-27 dBm) measurement which clearly indicated compression. A less accurate result was produced for the

mixing-path measurements, although no apparent non-linear behaviour was displayed. The two curves corresponding to -80 dBm deviated the most and this indicates a possibility of a measurement error, since this same reference data was used for both measurements. The difference between the results is also partially due to the attenuator correction errors, which can be corrected by characterising the integrated attenuators. Average calibration functions were calculated from the measurements that displayed good agreement and the accuracy of the calibration is assumed to be within 2 dB across the operational range. A more thorough assessment of the test space is recommended to provide better accuracy.

The calculated gain functions are applied in software and FE-configuration variables are recorded as attributes to the HDF5 files. There are several effects that may cause variations in the system response, including temperature or supply voltage changes. This may be particularly consequential in harsh environments with significant temperature variations and a long-term development for a monitoring system will be a self-calibration procedure.

5.2 EMC testing

The RFI issues that were identified during measurements were investigated through several tests. This involved measurements using both RATTY and a hand-held SA inside the REVERBERATION CHAMBER to avoid external interference. Difference plots were generated from RATTY measurements with the input of the digital subsystem terminated in a matched load. This configuration enabled the characterisation of self-generated interference and provided insight into the mechanisms by which this occurs. A 10 min integration period was applied in each test to accurately represent constant interference components and discrete changes were made to the configuration during subsequent tests.

The RF-cable, FE supply and system supply were adjusted to evaluate the influence of each and the first was found to have the most significant influence. In consecutive measurements, the SMA cable to the ADC was inductively filtered with increasing numbers of magnetic chokes. These chokes suppress CM currents induced on the outside of the signal path. The RF cable was also replaced with a shortened connection to the front of the enclosure. The second cable was a semi-rigid type with improved shielding, although a right-angle connector was present in this case. Figure 5.2 shows the different cable layouts and the application of a magnetic choke inside the digital enclosure.

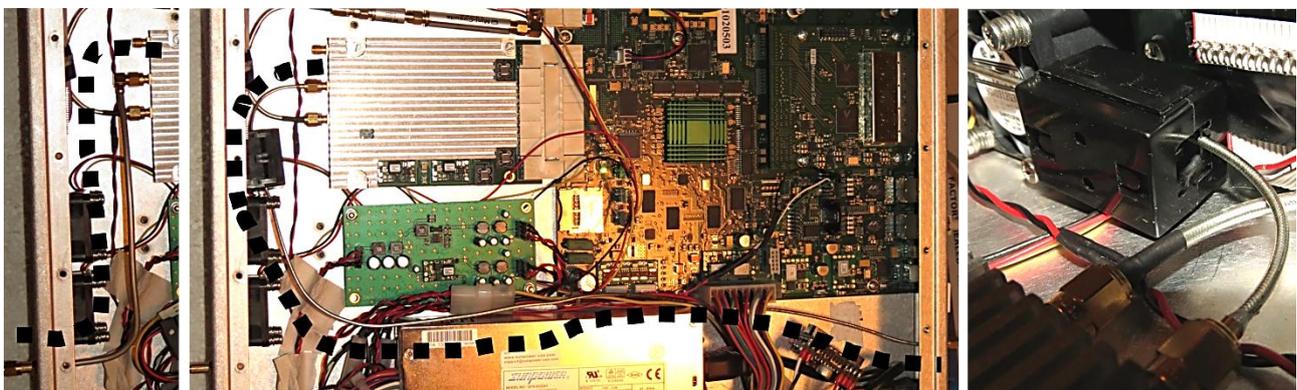


Figure 5.2: Different RF-cable configurations and magnetic choke application. A magnetic choke clamped around the RF cable is shown on the right with the ADC input visible in the foreground. The different cable layouts in the digital enclosure are indicated with broken lines.

It is clear from figure 5.2 that the second SMA cable (left) is shorter than the first (middle) and exits to the front of the digital enclosure. Furthermore, the second RF cable does not come within the same proximity of the digital and supply units as the first cable and possesses a smaller coupling aperture. The difference plots for the choke and cable changes are shown in figure 5.3.

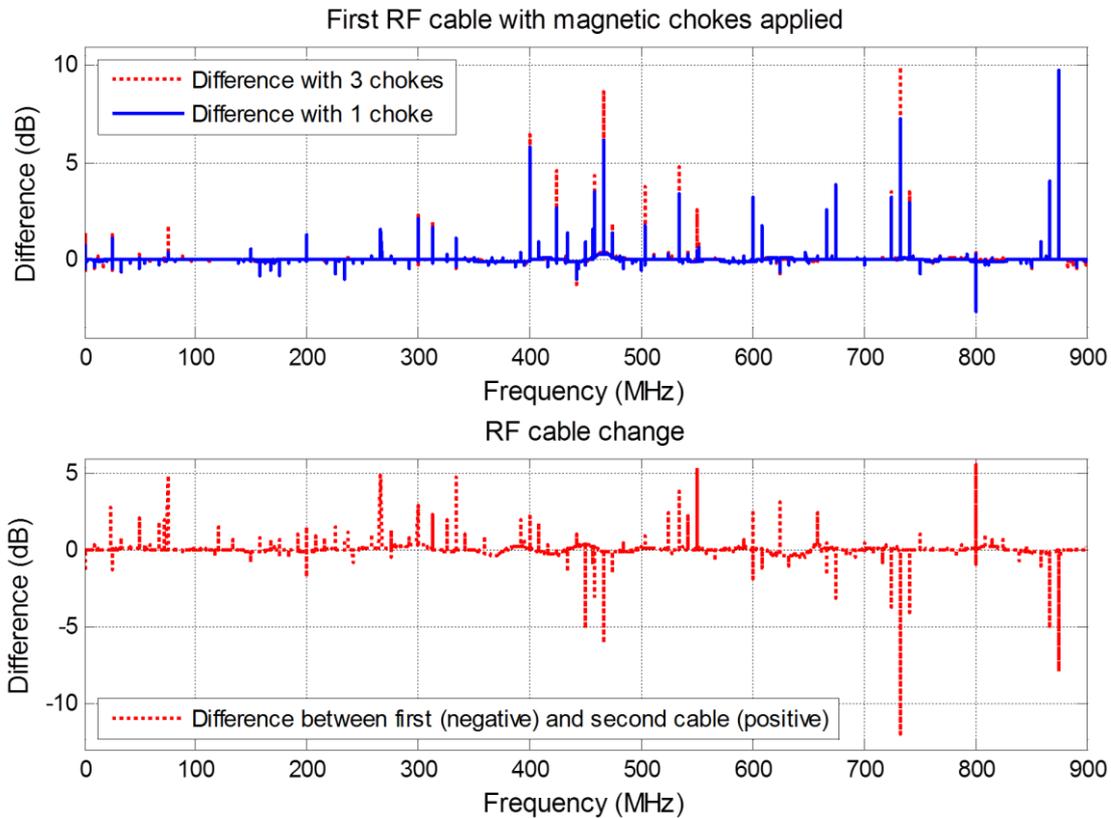


Figure 5.3: Difference plots of self-induced RFI on digital-enclosure RF cable. The effects of applying magnetic chokes are shown at the top and the result produced by the cable change is shown below.

The use of more chokes reduced the level of interference, which is reflected positively in the top graph of figure 5.3. Although the second cable layout reduced the highest interference components, there was an increase in several other products. This demonstrates the complex nature of the coupling, which likely leaks into the signal path either via the connectors or the ADC IC. Figure 5.4 shows the measured levels of interference for either cable.

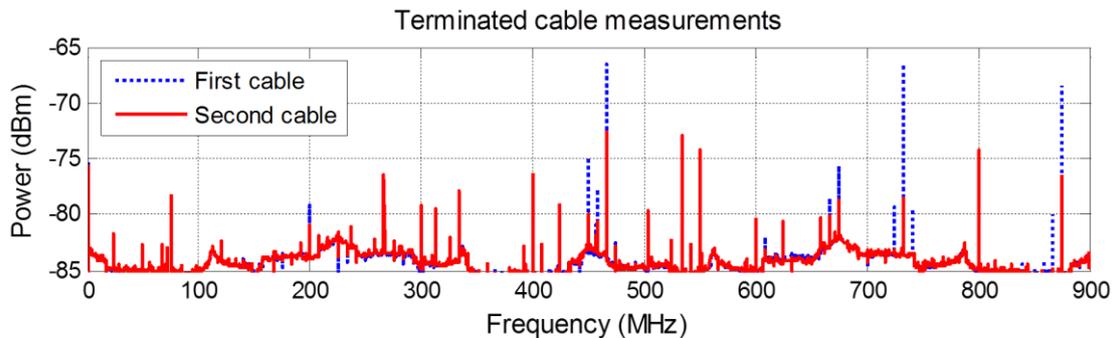


Figure 5.4: Measurements of terminated digital enclosure. Note that the second cable layout reduced the highest components, although several other products gained in strength.

Both these tests indicate that interference in the digital enclosure couples to the RF cable and can be suppressed with the application of CM-current rejection techniques. Similar tests were conducted with the FE

disconnected and the mains supply replaced with a battery source, although no significant changes were observed for these changes.

The hand-held SA measurements were mainly focussed on the FE EMC and implemented either CM-probes or sniffer antennas, which are E-field probes used for close range observations. The 14 MHz control-board clock signal and several harmonics thereof were measured on both supply and control conductors. This illuminates the presence of this interference in RATTY measurements. These signals were also present on the LO supply wires, although the 10 MHz reference of this unit was not observed. A different model of the synthesiser without a shielded enclosure was also evaluated and 10 MHz harmonics were present on the supply lines in this case. This indicated that the shielding of synthesiser in the mixing FE was effective and that this part was also protected against interference from the control-board signals.

A transient interference signal was also associated with the FE control and captured using the TD operation of the system. This pulse was related to the operation of the fibre-optic transceiver, since it coincided with the serial communication between the CPU and the control board. An initial pulse with a subsequent, larger pulse was observed and this was independent of the level-setting configuration. These transients enter the signal path through the final level-setting components, as the attenuation had no effect on the level of the pulses. Also, the two pulses were likely produced by the optical receiver and transmitter, respectively.

There are several short-term solutions to address the issues described above. An L-plate structure and high-precision SMA connectors can be implemented for the ADC connection. The plate can be earthed to the floor of the digital enclosure, which will provide CM-currents with a low-resistance path and divert related interference from the ADC. Furthermore, an improved control-board layout has been completed, featuring filtered outputs and power supply connections. This was not yet implemented upon completion of this document, but the new unit is expected to improve on the current RFI performance of the FE control.

The RFI issues described here are only the main issues present in the XDMs and more demanding EMC conduct may be required to harden the devices for any environment. Nevertheless, the precautionary measures employed in the construction of the devices (Section 4.3) are vital to avoid basic interference problems.

5.3 Reverberation-chamber application [29]

The RATTY SU system was applied in a reverberation chamber study conducted by Mr Joely Andriambelason. This research focusses on the shielding effectiveness of several enclosures and cable systems. The studies are important to the SKA-Africa project, which has acquired a reverberation chamber for conducted in-house measurements. The measurements are typically aimed at characterising EM emissions from telescope subsystems. Furthermore, there is general SKA interest in shielding-effectiveness studies, as shielding is required to avoiding self-generated interference throughout the project. A RATTY receiver is also intended for several other SKA RFE tests and this section provides a brief overview of some key aspects on the use of the device in the SU reverberation chamber.

A reverberation chamber is principally a resonant metallic cavity that is used for environmentally-independent EM measurements. The highly conductive walls of the room reflect EM waves when illuminated with radiation, which can propagate in any direction. If the field components of waves are in the same orientation at a given observation point, the net result is the superposition of the components. The process results in standing wave effects throughout the room. If a characteristic pattern is formed by these interactions, it is typically referred to as a mode and identified by the amount of discrete repetitions that

occur in the pattern. The possibility of exciting modes in a reverberation chamber depends on the size of the cavity and the wavelength. A wave distribution can only exist effectively in the room if the frequency is high enough and this introduces a lower limit on the operational range of a chamber [30].

Most reverberation chambers typically feature metallic stirrers, which are semi-randomly shaped reflectors inside the chamber. The rotation of these structures change the inner surface of the room and this alters the patterns of the existing modes. A reverberation chamber must be equipped with stirrers that enable the randomisation of modes across the operational range of the room. This is necessary to carry out typical reverberation chamber measurement procedures, which are concerned with RF emissions. Figure 5.5 shows the SU reverberation chamber.



Figure 5.5: The SU reverberation chamber. Note the metallic walls, DUTs and stirrers for rotation in perpendicular planes. There are also filtered-mains outlets inside the reverberation chamber, visible on the right wall.

When a radiating device under test (DUT) is placed inside the room, it will produce an arbitrary EM-field distribution. This is measured with an EM-field probe that is usually placed at a convenient position inside the reverberation chamber. However, the observed field could at worst be at a maximum or minimum, depending on the position of the probe. To ensure that the actual radiated power of the DUT is characterised, the stirrers are rotated to vary the field and an average, or peak-hold, measurement is carried out. A sufficient number of measurements must be taken with the stirrers rotated to an adequate amount of positions (~72) to produce a truly randomised distribution. This generally makes reverberation chamber measurements resource intensive. Alternatively, a DUT can be rotated inside the chamber to produce varying emission patterns without stirrers or the probe can be moved, although these techniques were not used.

There are several types of field probes that can be used inside the reverberation chamber and a broadband Logarithmic-Periodic Dipole Array (LPDA) antenna was mostly used as an electric-field (E-field) probe. A key aspect of a reverberation chamber is that the room also protects measurements from external RFI through its characteristic shielding effectiveness. This enables sensitive measurements in urban areas where high levels of RF emissions are ubiquitous. Interference from equipment is also excluded from measurements by keeping noisy devices outside the reverberation chamber.

The time required to complete a measurement procedure is highly dependent on the type of receiver that is used. For sweeping receivers, a measurement must be taken at every frequency and at all stirrer positions to enable an average power calculation at each wavelength. Thus, the required spectral resolution of a measurement influences the time linearly, as shown by equation 5.1.

$$T_{\text{Tot}} = N_f * P_s * T_{\text{Meas}} \quad (5.1)$$

where T_{Tot} = total time of reverberation chamber measurement procedure, N_f = number of frequencies, P_s = number of stirrer positions and T_{Meas} = the dwell time for individual frequencies.

If a TD receiver is used, however, each measurement provides a large bandwidth at a resolution that is determined by the number of samples used in the FFT. This can speed up the procedure and is the motivation for implementing a RATTY device for reverberation chamber studies. Furthermore, a broadband input can be used to perform these real-time measurements. A pulse generator (PG) capable of producing sufficiently short rise times is such a broadband source and these devices formed part of the study.

The shielding effectiveness testing procedure can now be described, using a PG as input and the RATTY SU for recording data. A direct pulse measurement is first performed, by connecting the PG to RATTY through all the cables involved. This calibrates everything except the reverberation chamber and the antennas. The pulse amplitude must be sufficient to produce a measurable signal inside the chamber after being attenuated by the shielding effectiveness of the DUT. This requires an external, adjustable attenuator to limit the pulse amplitude during the direct measurement, although the component is well characterised and is accommodated in the result. The reverberation chamber measurements are then performed and involve injecting the pulse into the DUT whilst measuring the emissions. To compensate for the effect of the room, the LPDA antenna is calibrated using a standard factor. Figure 5.6 shows plots of direct-pulse measurements with a sampling oscilloscope and both RATTY SU signal paths.

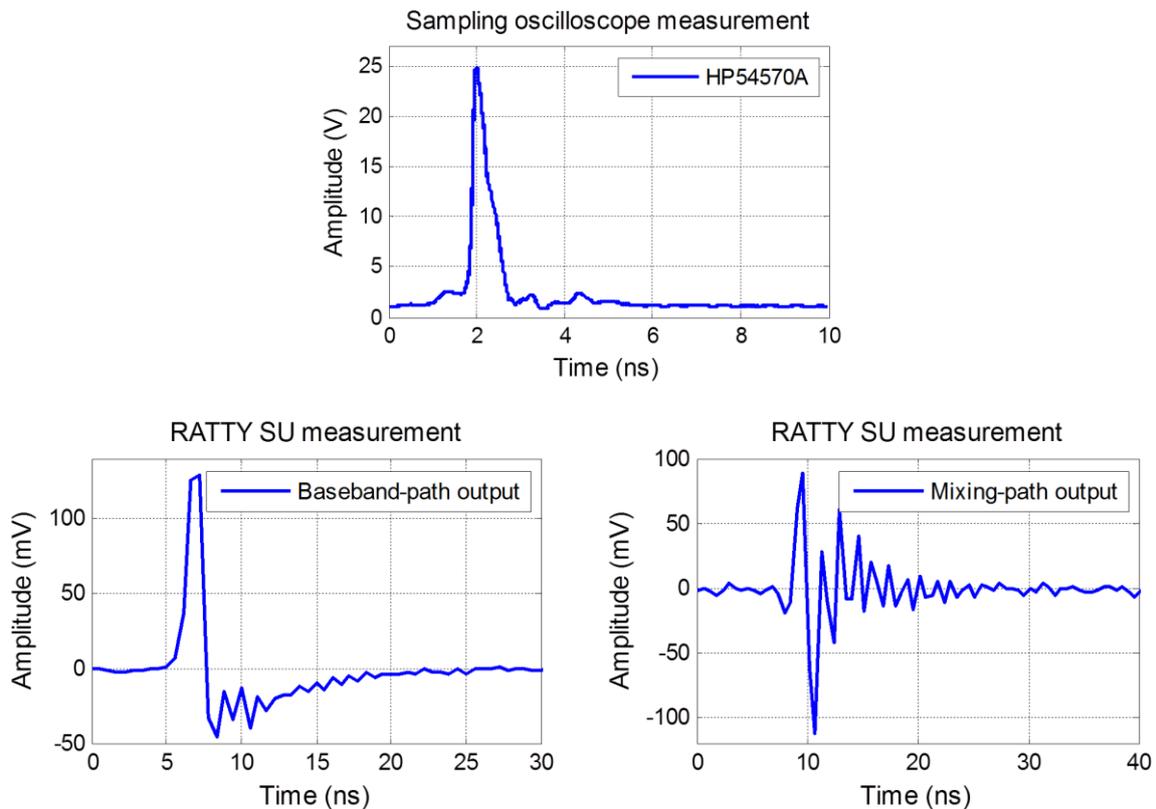


Figure 5.6: Direct pulse measurements with a sampling oscilloscope (top), RATTY SU baseband (bottom left) and mixing path (bottom right). Note that the oscilloscope measurement was corrected for the external attenuator loss and depicts the output capability of the PG.

It is apparent that the bandwidth-limited RATTY measurements distort the TD pulse. Nevertheless, the pulse produces sharp rising edges in both signal paths and this is used to trigger measurements in the TD operation. Also, it is clear that the pulse is too strong and must be attenuated to conduct the direct measurements. Figure 5.7 is a comparison of the pulse spectra.

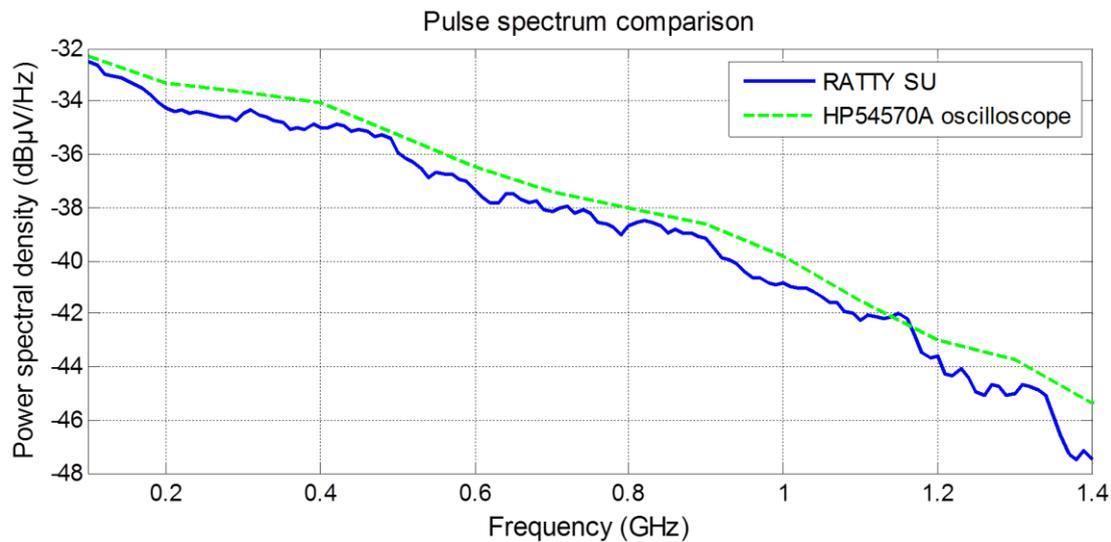


Figure 5.7: Spectra calculated from direct-pulse measurements with a sampling oscilloscope and RATTY. Note that the RATTY spectrum was calibrated for the signal-path gains and that this spectrum has a higher resolution than the oscilloscope calculation.

The differences between the spectra in figure 5.7 can be related to calibration errors and mismatches between the devices, although the low resolution on the oscilloscope spectrum limits the comparison. Nonetheless, the spectra display similar trends in the power densities. Note that the power spectral density (Decibels microvolts per Hertz) is independent of the measurement bandwidth. The RATTY spectra are in agreement where the path bandwidths join (800 MHz), indicating that the calibration is reasonably accurate. However, errors present in both direct and reflected measurements are inconsequential and cancel in shielding effectiveness calculations. A dynamic range requirement can be estimated from this spectrum. Firstly, the maximum variation in the input spectra is assumed to be below 10 dB across the signal-path bandwidth. If the gain difference is below 6 dB and any other measurement deviations is limited to 6 dB, a SFDR of ~26 dB should ensure that measurements are free of artefacts. This must be evaluated according spread-sheet calculations. The measurements should be conducted at the maximum level to avoid the influence of the noise floor, since the pulse spectrum is broadband and at relatively low power and repetition rates.

The reverberation chamber measurement procedure was automated and controlled from the computer managing the stirrer positions.* This was connected to the RATTY laptop, which was prompted to perform the TD measurements. The FE configuration procedure was added to the TD measurement script to reduce measurements to single commands and the laptop stored the data.

A manual setup was performed prior to initialising a complete procedure and was focussed on ensuring that the FE remains within the linear operating range. A single-point calculation was added to the TD-script to aid the manual FE configuration. This involved calculating the maximum, linear input voltage to the amplifiers, which are the components at risk of compression. The attenuator settings are evaluated to determine which amplifier in the signal path will compress first and this depends on the signal path. For baseband path measurements, attenuation above 15 dB protects the output amplifier from non-linear

* The automation was configured by Joely Andriambelosen.

behaviour. This is not the case for the mixing path, although the second attenuator suppresses the LO interference. Furthermore, the first level-setting amplifier can compress the signal for settings below ~ 22 dB of the preceding attenuator and this is valid in for both paths. Figure 5.8 is a simplified schematic of the RATTY SU system which illustrates the positions of the level-setting blocks.

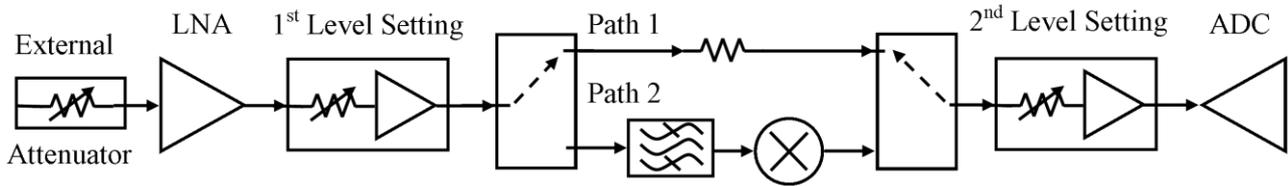


Figure 5.8: Simplified RATTY SU schematic illustrating positions of level-setting components. Note that the LNA can be prevented from compression by external attenuation.

Using the cases described above, along with the maximum path gains from the first position of compression, the linear input limit of the ADC is calculated. The input level is based on the amplitude of a sinusoid at IP1dB powers. This is not valid for a broadband pulse, since this signal requires uniform amplification across its bandwidth for the peak to increase according to a single gain value and the real gains are not constant. Thus, the actual pulse peaks are less than the calculated values. These figures were used as guidance, along with systematic adjustments of the attenuators to evaluate the linearity of pulse amplitudes.

The requirement of external attenuation depends on the DUT and PG. The level required must be determined experimentally with systematic adjustments. For every external-layout test, the 1st and 2nd level setting attenuators are initialised at 31 and 20 dB, respectively. The first attenuation can then be gradually lowered until the system triggers on pulses reliably. This is continued until the attenuation is at a minimum or the amplifiers approach non-linearity. Subsequently, the second attenuation is lowered to utilise the ADC range, unless the input to the digital-subsystem is already sufficient. This also ensures that the optimal noise performance of the FE is attained. Since changes in the stirrer positions are known to produce maximum variations of ± 3 dB, the setup limits pulse peaks to $\sim 50\%$ of the ADC range to avoid saturation during the reverberation chamber procedure.

The manual adjustment is first completed for the baseband path, since the spectral density of pulses is highest at low frequencies. Furthermore, the first level-setting configuration is retained for mixing-path measurements. This is required, as compression is known to produce broadband distortion and the amplitude of the mixed pulse does not clearly indicate non-linearity. However, the second attenuator can be set to any level for the second-path measurements. A SFDR simulation also can be performed to evaluate this parameter for a particular configuration, although the required range is not high and is realised for most FE settings. Figure 5.9 shows reverberation chamber pulse measurements.

This section has described the application of the RATTY SU system and some key features for TD operation were highlighted. The post-processing of the TD HDF5 files is based on a MATLAB platform and falls beyond the scope of this document. However, the integration of all the software and hardware parts is a future development that is well suited to the open-source environment of RATTY. Although an in-depth analysis on the speed-up of measurements has not been conducted, the initial reverberation chamber measurements have been shortened by 60% in comparison to an available VNA method. This does not account for the increase in resolution of the TD measurements, which would reflect a more significant improvement according to the definition of equation 5.1 and determined by the sample length used for FFT calculations.

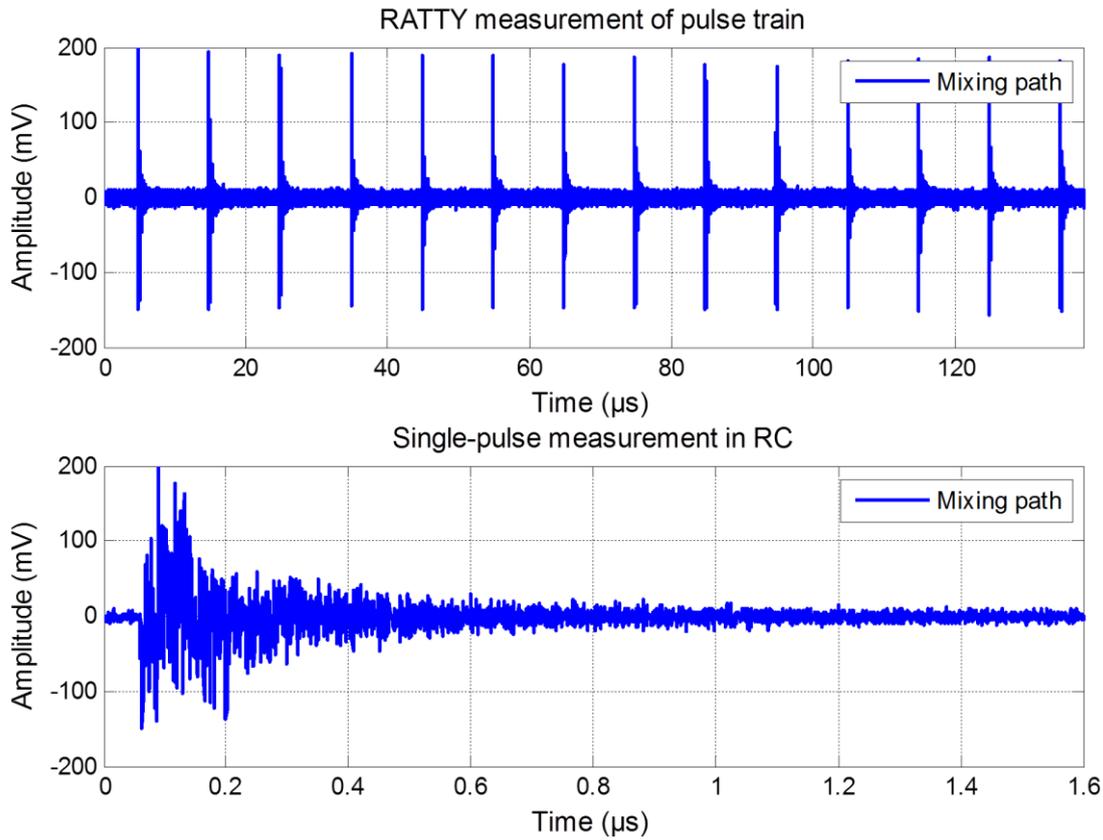


Figure 5.9: Reverberation chamber pulse measurements with a pulse train (above) and single pulse (below). Note these are mixing-path measurements and that the pulse exhibits multiple reflections inside the chamber that exponentially decrease in amplitude.

There are no further results presented for the shielding effectiveness research, although preliminary results have shown good correlation with theoretical predictions [29]. Despite there being several aspects of the RATTY system that require further development, the device has provided a powerful alternative to standard reverberation chamber measurement procedures.

Chapter 6

Conclusions and recommendations

The development of an RF measurement system was documented with the principle goals of realising high DR, sensitivity and a wide bandwidth. These characteristics were motivated by requirements of the MeerKAT project, although this system can be utilised in multiple applications.

Several FE designs were evaluated and each presented a unique frequency-conditioning plan. The large bandwidth introduced spurious interference in the mixing strategies, which could not be avoided with the implementation of readily-available, connectorised components. Ultimately, two XDMs with simplified layouts were pursued to demonstrate different strategies for the analogue subsystem: 1) mixing and 2) direct sampling.

Both models required customised filtering and several stripline components were developed for this purpose. The bandwidth and out-of-band rejection ratio proved to be challenging to realise and additional LPFs were required to suppress the harmonic behaviour of the stripline structures. There were differences between the dielectric-substrate properties of simulations and the real materials, which caused significant deviations in the responses of several filters. This was improved in subsequent designs, although the large bandwidth still exhibited slight dielectric-substrate-related or manufacturing-related variations. Nevertheless, the initial filters were effective for the XDM systems, although the components were lacking in suppression at higher frequencies. It is recommended that the filter developments are revised through integrated optimisation of the LPF and BPF designs. Alternative methods could also be reconsidered to limit the size of the components and improve on the stripline performance.

The mixing XDM was subjected to several analyses with comparisons made between the performances of the upper and lower signal paths. For the direct-sampling model, key references were made to the evaluation of the similar, baseband path in the mixing system. There were important issues identified during this process, including the effects of gain variations and SFDR limitations. The component selection for the SU XDM resulted in an imbalance between the mixing and baseband path gains, which resulted in an inferior noise-figure calculated for the former. However, the effect of the T_{sys} on the noise power is limited, since measurement noise is expected to be dominated by observation temperatures. In consideration of this, gain variation is more significant with regard to the measurement ranges of the paths. Furthermore, the gain functions shape the noise floors and reduce the SFDRs.

A comparison was made between a spread-sheet calculator and a commercial simulator package. There was good agreement between the calculated parameters, with the exception of differences occurring due to simulation errors that were expounded. The capacity of the commercial simulator to perform broadband analyses was demonstrated by projecting gain functions based on measured component characteristics.

A self-compiled, mixing-spurious calculator was used to incorporate the characteristics of the active mixer of the SU FE. The results indicated that the lack a continuous suppression in the stripline-filter responses was problematic and could result in significant interference levels if out-of-band signals were present in measurements. However, this is related to filter issues and unavoidable mixer products were shown not to implicate the real-time SFDR (60 dB).

After measuring the characteristics of the LO synthesiser and improving the output filtering for this unit, an additional level-setting stage was implemented. This enabled an increase in the headroom above the broadband interference caused by reference signal of the LO source. Nevertheless, these were short-term improvements and it is recommended that the issue is addressed with notch filters. The reconfigured system also featured improved adjustable DR.

The simulated SFDRs were compared to a practical assessment of this parameter. This was done for the mixing FE with references made to the SKA XDM. An initial finding was that second order IMD components were of significant importance to the broadband system and the spread-sheet calculations were adapted to include this. Furthermore, there were additional IMD components observed and related to the mixing process. This is a weakness of this strategy that is consequential for increased-SFDR measurements with extended integration periods. A systematic test of the dynamic FE behaviour for reconfigured attenuation indicated that subsequent level-setting sections are effective at preserving the SFDR. However, this is only possible if the OIPs of the FE remain unaffected and makes the behaviour of the reconfigurable FEs more complex. Further comparisons of simulated and measured figures indicated that the accurate predictions of the IMD levels require the characterisation of the key component OIPs, especially for second order IMD. It was concluded that the FEs limit the SFDR performance of the systems and that this is worse for higher input powers. It is recommended that component OIPs are characterised in future systems, to enable accurate calculations of the SFDR.

A comparison between the XDMs showed that both devices produced self-generated interference, although the mixing design exhibited more issues in this regard. Thus, it was decided that a direct-sampling strategy would be followed for short-term developments of RATTY systems. However, the benefits of a mixing FE may prove worthwhile for some applications and it is recommended that long-term developments further explore both strategies. This will be important for the production of RATTY systems with significantly improved BWs.

A static calibration procedure was applied to the mixing system. This involved calculating gain functions at several input power levels and for different attenuator configurations. The testing strategy also indicated the accuracy of the calibration, although the result showed variations up to 0.8 dB. Improving on this requires a more accurate and thorough test of the operational range of the FE and the implementation of measured responses for the step attenuators.

The EMC issues of the mixing system were investigated through several measurements. In addition to the LO contamination, the clock signal of the control-board was detected on most supply and control cables in the FE. Furthermore, RFI related to the digital subsystem was shown to be present on the RF cable in this enclosure. A transient pulse was also characterised and related to the FE control. A refined control-board configuration with filtered ports will be implemented and should improve issues related to this unit.

The use of the mixing system in a reverberation chamber study was described with several results shown. The preliminary results of these shielding effectiveness measurements have been promising and this research will be continued. An improved RATTY system will make reverberation chamber measurements more expedient and less labour intensive. This is also encouraging to the prospect of applying RATTY systems within more fields of research.

In conclusion, the implementation of dynamic, RF receivers on ROACH 2 platforms was described and demonstrated. Several limitations were identified in the process and recommendations are now cleared for improving the performance of the initial designs. Following the completion of these refinements, two direct-sampling receivers will be constructed for use in the SKA-Africa: 1) as an enhanced on-site monitoring system and 2) as a reverberation-chamber receiver. There remains significant scope for the development of these systems in SKA-Africa applications and internationally.

Bibliography

- [1] David M. Pozar, *Microwave and RF Design of Wireless Systems*, New York, John Wiley & Sons, 2001.
- [2] P. Meyer. (2013). High frequency Techniques 414 - Classical Design of TEM-filters. Stellenbosch University. RSA. [Online], Available: http://courses.ee.sun.ac.za/HFTegniek_414/
- [3] (2011, July). Stripline. P-N Designs Inc.. USA. [Online]. Available: <http://www.microwaves101.com/encyclopedia/stripline.cfm>
- [4] David M. Pozar, *Microwave Engineering*, New York, John Wiley & Sons, 2012.
- [5] John G. Proakis, Dimitris G. Manolakis, *Digital Signal Processing*, 4th ed., Pearson Prentice Hall, New Jersey, 2007.
- [6] F. Marki, C. Marki. (2010). Mixer Basics Primer. Marki Microwave. USA. [Online]. Available: www.markimicrowave.com/menus/appnotes/mixer_basics_primer.pdf
- [7] (2009). 10-bit 2.2 Gsps ADC. e2v. USA. [Online]. Available: <http://symmetron.ru/suppliers/e2v/files/pdf/e2v/AT84AS008.pdf>
- [8] (2012, Jul.). TWO-TONE-TERMINATOR MIXER/LO-AMPLIFIER T3A3-08. Marki Microwave. USA. [Online], Available: <http://www.markimicrowave.com/Assets/datasheets/T3A3-08.pdf>
- [9] *Applied Wave Research (AWR) Incorporated's Microwave Office*© (MWO) 2010
- [10] *Computer Simulation Technology™ (CST) Studio Suite*® 2012
- [11] (2009). Collaboration for Astronomy Signal Processing and Electronics Research (CASPER). CASPER Group. [Online]. Available: https://casper.berkeley.edu/wiki/Main_Page
- [12] A.R. Botha, H.C. Reader, J. Manley, et al., "Dynamic RFI measurement systems on a ROACH-2 platform," *Electromagnetics in Advanced Applications (ICEAA)*, 2013 International Conference on.
- [13] (2013, April 02). Classes. Python Software Foundation. [Online]. Available: <http://docspy2zh.readthedocs.org/en/latest/tutorial/classes.html>
- [14] (2011, May 16). What is HDF5?. The HDF5 Group. [Online]. Available: <http://www.hdfgroup.org/HDF5/whatishdf5.html>
- [15] R. Lord, A. Peens-Hough. (2012, July). RFI Measurement System - Requirements Specification for Ratty-2. SKA-Africa. RSA. [Online]. Available: <https://drive.google.com/folderview?docId=1S3yb9yCfqndVbR2sPQadFChphbH83ZLLIqoHBNYo cQg&id=0B5HUkgt4UqvCMXViN3hOa09pdDg#>
- [16] Sias Malan, Jason Manley, Richard Lord, et al.. (2013, Sept.). SKA_SA_RATTY2_design_doc. [Online]. Available: <https://drive.google.com/folderview?docId=1S3yb9yCfqndVbR2sPQadFChphbH83ZLLIqoHBNYo cQg&id=0B5HUkgt4UqvCMXViN3hOa09pdDg#>
- [17] B.P. Lathi, *Modern Digital and Analogue Communication Systems*, 3rd ed., New York: Oxford Univ. Press, 1998

- [18] E. O. Brigham, "The Fast Fourier Transform", New Jersey, Prentice Hall, 1974, ISBN 0-13-307496-X
- [19] W. Kester. (2009). MT-003 tutorial [Online]. Analog devices. USA. [Online] Available: <http://www.analog.com/static/imported-files/tutorials/MT-003.pdf>
- [20] L. Melkonian. (1992, Feb.). AN-804 Improving A/D Converter Performance Using Dither. Texas Instruments. USA. [Online]. Available: <http://www.ti.com.cn/cn/lit/an/snoa232/snoa232.pdf>
- [21] TEW Stuart, HC Reader, "Apparently Simple Low Frequency Measurements: An EMC View", in South African Symposium Communications and processing, Cape Town, 1998, pp. 319 – 322
- [22] FM Tesche, MV Ianoz, et al., *EMC analysis methods and computational models*, New York: Wiley, 1997.
- [23] (1995). Electromagnetic Compatibility. Schneider Electric. [Online]. Available: http://www.engineering.schneider-electric.dk/Attachments/ia/instal/electromagnetic_compatibility_install_guide.pdf
- [24] Arduino Nano. Arduino. [Online]. Available: <http://arduino.cc/en/Main/ArduinoBoardNano>
- [25] W. Kester. (2004, Mar.). ANALOG-DIGITAL CONVERSION. Analog Devices Inc.. USA. [Online]. Available: http://www.analog.com/library/analogdialogue/archives/39-06/data_conversion_handbook.html
- [26] C. Henn. (1994, Apr.). INTERMODULATION DISTORTION (IMD) (Application Bulletin). Burr-Brown International. [Online]. Available: <http://www.ti.com/lit/an/sboa077/sboa077.pdf>
- [27] (2000, Sept.). Intermodulation Distortion (IMD) Measurements. Anritsu Microwave Division. [Online]. Available: <http://downloadfile.anritsu.com/RefFiles/en-US/Services-Support/Downloads/Application-Notes/Application-Note/11410-00257a.pdf>
- [28] (2007). Dithering in Analog-to-digital Conversion. e2v. USA. [Online]. Available: <http://www.e2v.com/e2v/assets/File/documents/broadband-data-converters/doc0869B.pdf>
- [29] J.A. Andriambelason, H.C. Reader, A.R. Botha, "Time-domain nested-enclosure technique for shielding effectiveness characterisation," in *Electromagnetics in Advanced Applications (ICEAA)*, 2013 International Conference on.
- [30] B. S. Guru, H.R. Hiziriğlu, "Introduction", in *Electromagnetic Field Theory Fundamentals*, 2nd ed. New York: Cambridge Univ. Press, 2005, ch. 10, sec. 10.1, pp. 502-505.
- [31] (2013). SKA Africa. SKA South Africa. RSA. [Online]. Available: <http://www.ska.ac.za/pubdocs/index.php>
- [32] S. Rajan, "Automated Gateway Discovery Using Open Firmware", M.S. Thesis, UCT, Cape Town, RSA, 2013

Appendix A: FE designs

This appendix documents several FE design configurations, presenting schematics of system layouts with gain analyses, signal processing diagrams, component lists and spurious analyses of several of the concepts. Appendices A.1-A.3 present mixing FE designs and appendix A.4 is the MATLAB script used for spurious analyses of the mixing systems. Appendix A.5 contains the component lists for the mixing and direct-sampling front ends.

Appendix A.1: RATTY 2.1: Single-mixer FE topology

The RATTY 2.1 configuration consists of a baseband path (0 – 800 MHz) and a broad-band, down-converting path (800 – 3000 MHz).

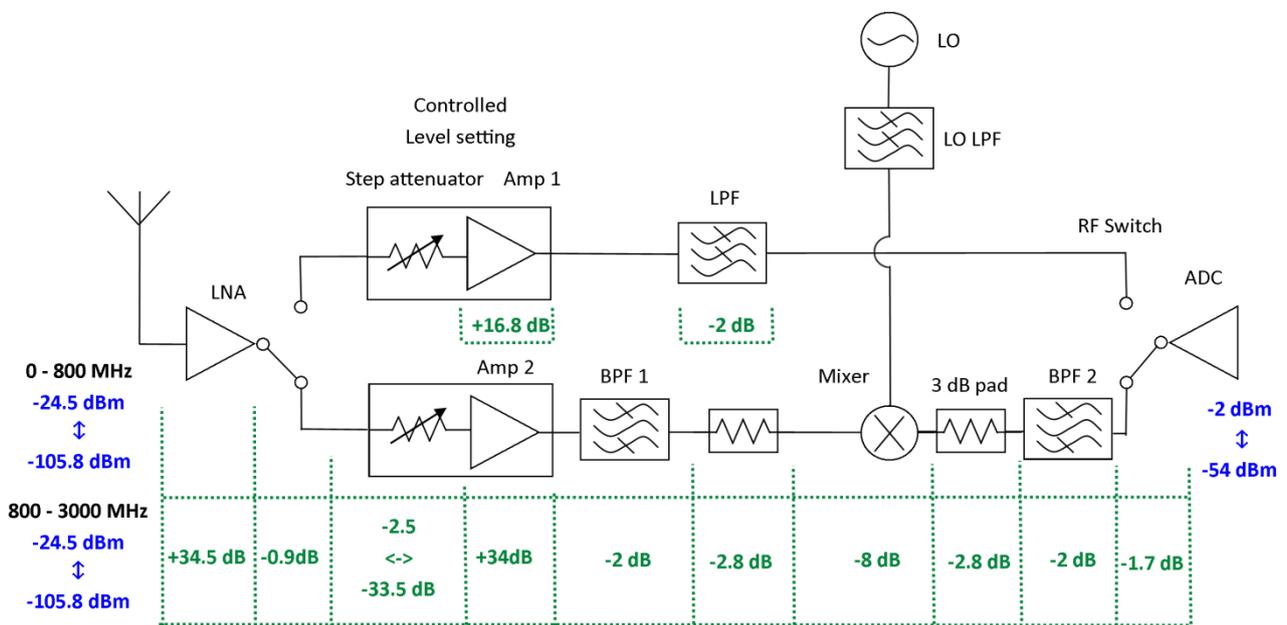
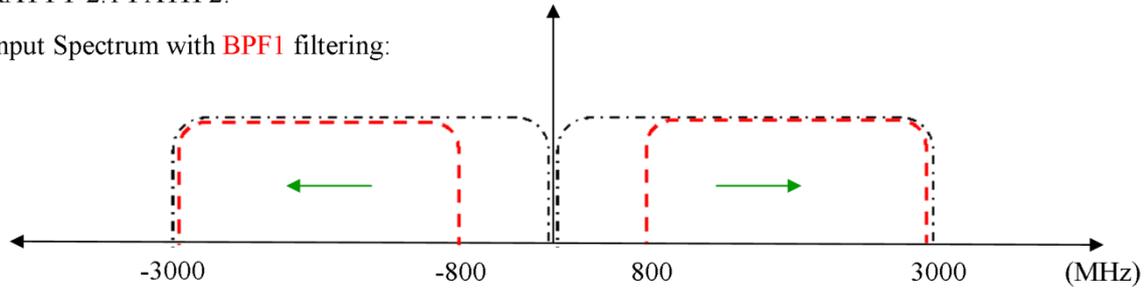


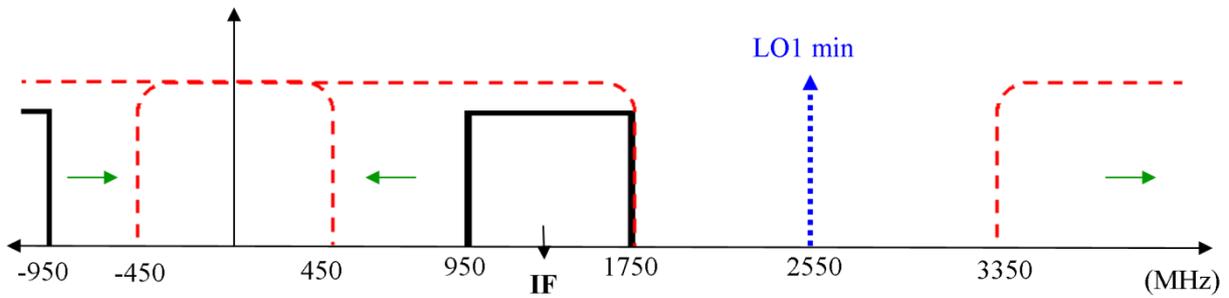
Figure A.1: A dual-input schematic of the RATTY 2.1 design. The figures enclosed in dashed lines are single-valued gain approximations, whilst the ranges are estimates of the input and output power ranges. Note that the component are either based on selected components or guesstimates. The range on the right is the minimum ADC range and the maximum input power is limited by the LNA IP1dB.

RATTY 2.1 PATH 2:

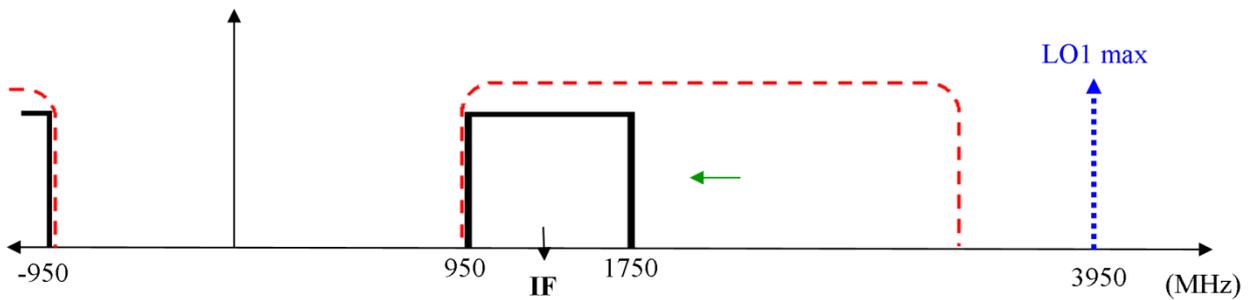
Input Spectrum with BPF1 filtering:



Convolved with LO1 min and IF filtering:



Convolved with LO1 max and IF filtering:



Sampled Spectrum:

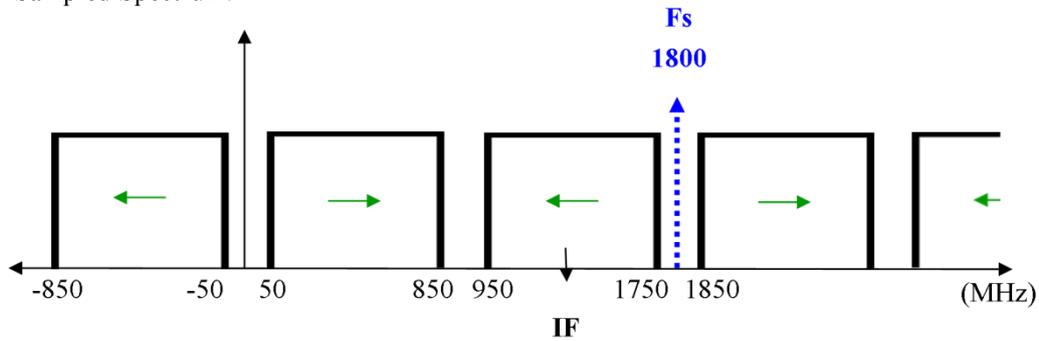


Figure A.2: Ideal FD representation of path 2 signal processing in RATTY 2.1. Note that the filter responses are drawn according to 3 dB cut-off specifications and that the green arrows indicate direction of increasing frequency in measured spectra.

Appendix A.2: RATTY 2.2: Dual-mixing FE topology

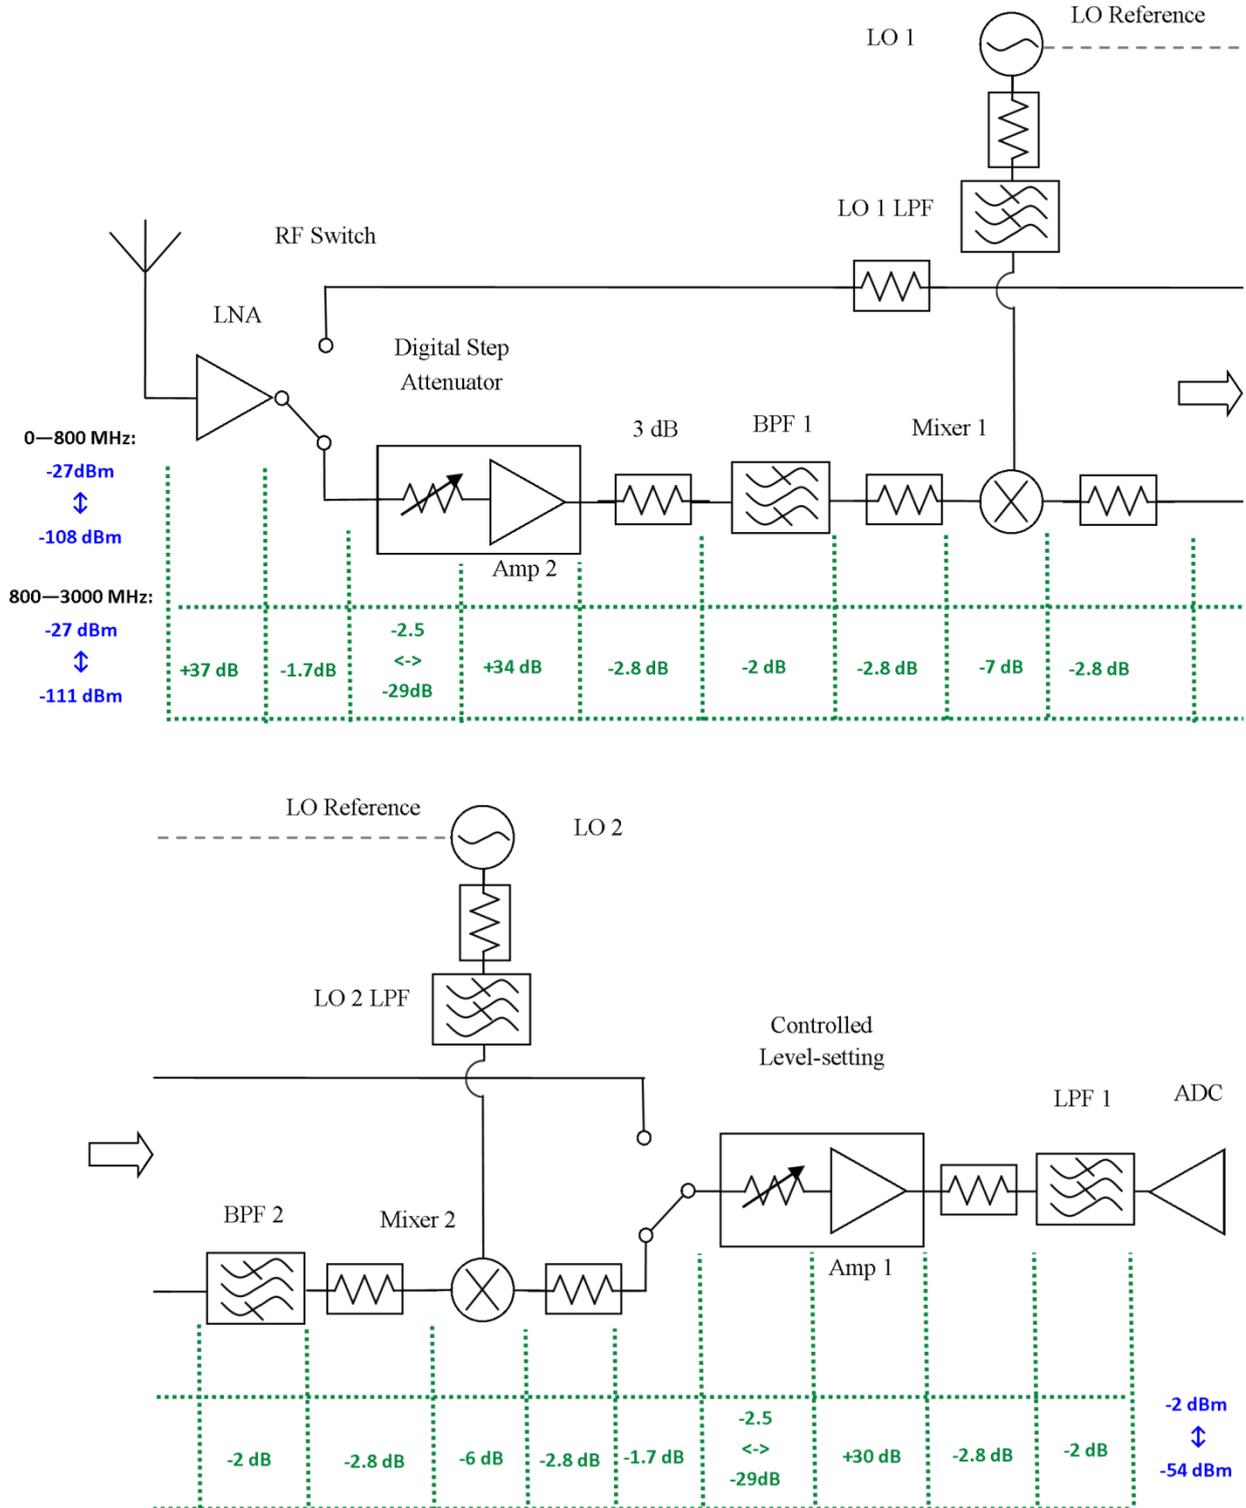


Figure A.3: Schematic of RATTY 2.2 with single-valued gain and dynamic range calculation. Note that signal-paths are split between the mixers, as indicated by the arrows. The figures enclosed in dashed lines are single-valued gain or loss approximations, whilst figures at the sides are estimates of the input and output power ranges. The limits on the right represent the minimum ADC range and the maximum input power is limited by the LNA IP1dB.

Appendix A.3: RATTY SU: Refined single-mixing FE topology

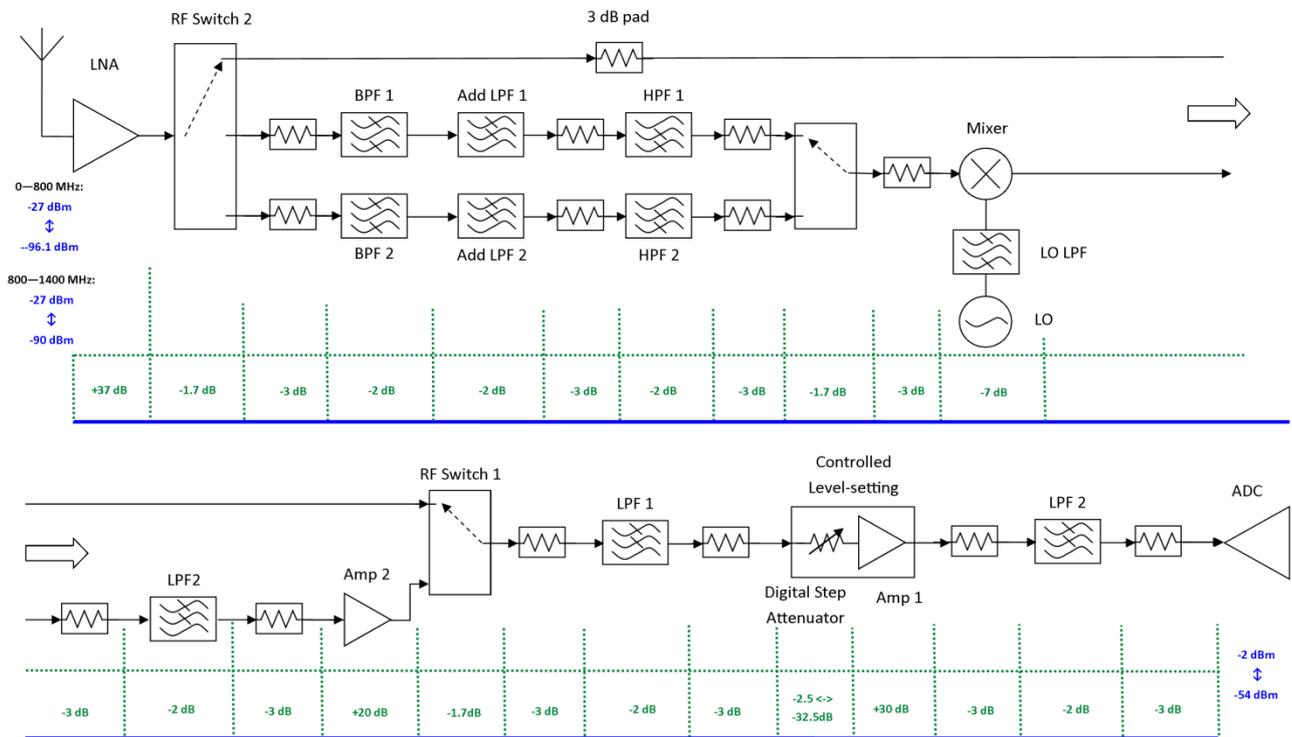


Figure A.4: Schematic of RATTY SU with single-valued gain and dynamic range calculation. Note that signal-paths are separated after the mixer, as indicated by the arrows. The figures enclosed in dashed lines are single-valued gain or loss approximations, whilst figures at the sides are estimates of the input and output power ranges. The limits on the right represent the minimum ADC range and the maximum input power is limited by the LNA IP1dB.

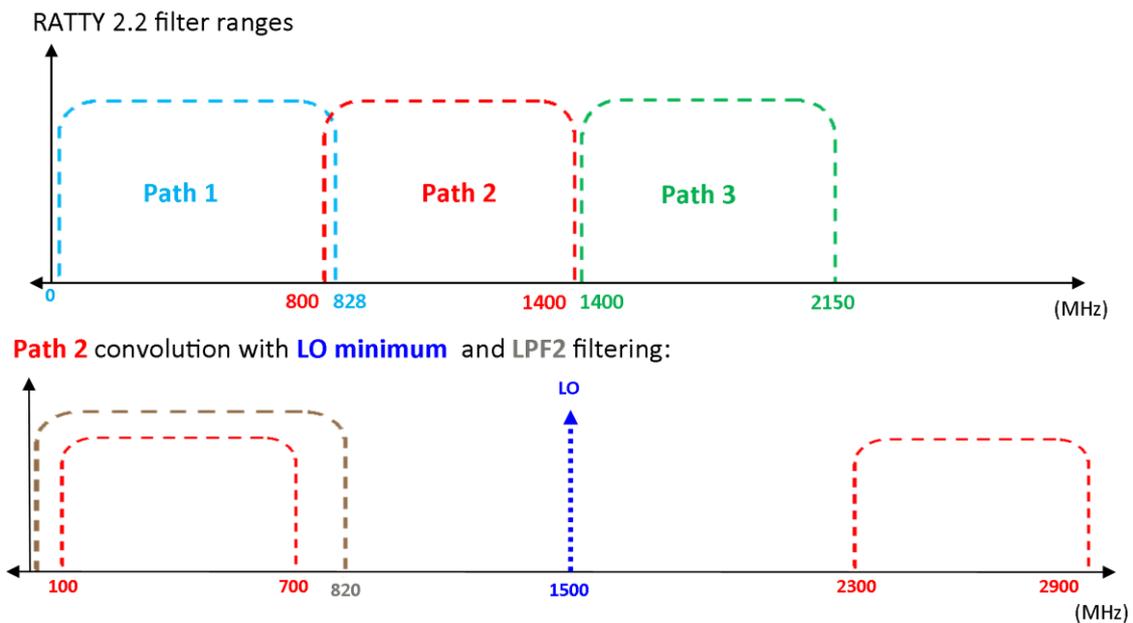


Figure A.5: Ideal FD representation of two stages of the FE signal processing in RATTY SU. The top figure shows the ideal filter ranges and the bottom figure represents the down-conversion in path 2.

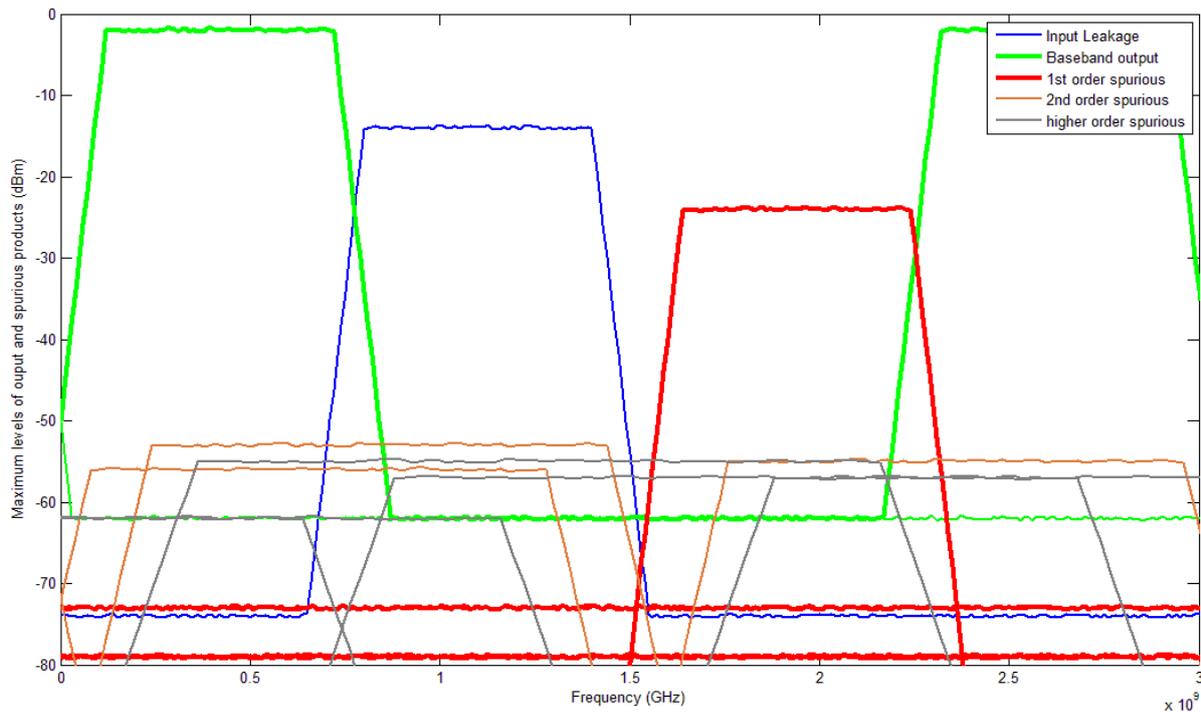


Figure A.6: Worst-case mixing-spurious analysis for RATTY SU path 2 (800 – 1400 MHz). This is for the maximum input level to the mixer. Note that there are three mixer spurious ranges that coincide with the output and are within 60 dBc of the maximum levels.

Appendix A.4: MATLAB script for mixing spurious analysis

This script was used for the initial spurious analyses using ideal filtering characteristics. It was also adapted to incorporate the simulated system responses to evaluate the spurious products of the mixer more accurately.

```
%This script analyses the spurious output of a single mixer
```

```
AnalBW = 12e9;           %BW used for analysis
pnts = (AnalBW*2/10e6)+1; %frequency resolution set to 10 MHz
FilRej = -60;           %BPF Rejection
InputPOW = 6;          %input level in dBm, used to adjust spu levels
```

```
%LO1 definition including first five harmonics, assuming higher harmonics are suppressed
```

```
M = 3;                 %Assess up to 3*RF
N = 5;                 %Assess up to 5*LO
LO1fund = 6350e6;
LO1 = zeros(2,N);
LO1(1,:) = (LO1fund:LO1fund:N*LO1fund);
LO1(2,1) = -6;
LO1(2,2) = -36;
LO1(2,3) = -56;
LO1(2,4) = -76;
LO1(2,5) = -96;
```

```
%Position of 0 Hz
pntzero = (pnts+1)/2;
```

```
%Setup of BPF1 (B1) ideal shape
%This was adapted to incorporate the simulated response of the FE
Blcnr1 = 800e6;        %Corner Frequencies
```

```

B1cnr2 = 3e9;
B1RO = 150e6;           %BPF1 Roll-off
B1rej = -60;           %BPF1 rejection ratio

BPF1 = ones(2,pnts) -1 +InputPOW;           %Adjust Filter passband amplitude to indicate
maximum input levels
%Setup filter from 0 Hz to max and then mirror for negative frequencies
BPF1(2,pntzero:pntzero+ (pnts+1)/2*(B1cnr1-B1RO)/AnalBW) = B1rej+ InputPOW;
BPF1(2,pntzero +(pnts+1)/2*((B1cnr2+B1RO)/AnalBW):pnts+1) = B1rej + InputPOW;
BPF1(2,pntzero +(pnts+1)/2*(B1cnr1-B1RO)/AnalBW:pntzero +(pnts+1)/2*(B1cnr1)/AnalBW) =
(0:-1*B1rej/((pnts+1)/2*(B1RO/AnalBW)):60)-60+InputPOW;           %filter roll-offs
BPF1(2,pntzero +(pnts+1)/2*(B1cnr2)/AnalBW:pntzero +(pnts+1)/2*(B1cnr2+B1RO)/AnalBW) =
(0:B1rej/((pnts+1)/2*(B1RO/AnalBW)):B1rej)+InputPOW;
BPF1(1,pntzero:pnts) = (0:10e6:AnalBW);

BPF1(2,:) = BPF1(2, :)+0.5*(rand(1,length(BPF1))-0.5);           %Add noise

for x = (1:(pnts -1)/2 )
    BPF1(2,pntzero - x) = BPF1(2,pntzero+x);
    BPF1(1,pntzero - x) = -1*BPF1(1,pntzero+x);
end

%Mixer 1 (M1) definition, mixer products in matrix form:
%MxN (M*RF +- N*LO)
%Spurious product lvls based on spurious table (-10 dbm input)
%Components <120 dBc ignored
M1lv1 = zeros(M,N);
M1RFIFiso = -20;           %Mixer RF to IF isolation, worst case
M1lv1(1,1) = -8;           %M1 conversion loss

%Factor in CL to dBc value
%Adjust spurs for input lvl: (n-1)*(abs(-10 - InputPOWER))dB; n = RF spur
%Adjustment based on Marki T3-08 Datasheet
M1lv1(2,1) = -70 + M1lv1(1,1) - (-10 - InputPOW);           %1*LO +- 2*RF
M1lv1(3,1) = -87 + M1lv1(1,1) - 2*(-10 - InputPOW);           %1*LO +- 2*RF

M1lv1(1,2) = -22 + M1lv1(1,1);           %2*LO +- 1*RF
M1lv1(2,2) = -67 + M1lv1(1,1) - (-10 - InputPOW);           %2*LO +- 2*RF
M1lv1(3,2) = -92 + M1lv1(1,1) - 2*(-10 - InputPOW);           %2*LO +- 3*RF

M1lv1(1,3) = -11 + M1lv1(1,1);           %3*LO +- 1*RF
M1lv1(2,3) = -69 + M1lv1(1,1) - (-10 - InputPOW);           %3*LO +- 2*RF
M1lv1(3,3) = -85 + M1lv1(1,1) - 2*(-10 - InputPOW);           %3*LO +- 3*RF

M1lv1(1,4) = -17 + M1lv1(1,1);           %4*LO +- 1*RF
M1lv1(2,4) = -65 + M1lv1(1,1) - (-10 - InputPOW);           %4*LO +- 2*RF
M1lv1(3,4) = -87 + M1lv1(1,1) - 2*(-10 - InputPOW);           %4*LO +- 3*RF

M1lv1(1,5) = -17 + M1lv1(1,1);           %5*LO +- 1*RF
M1lv1(2,5) = -67 + M1lv1(1,1) - (-10 - InputPOW);           %5*LO +- 2*RF
M1lv1(3,5) = -85 + M1lv1(1,1) - 2*(-10 - InputPOW);           %5*LO +- 3*RF

%Loop through matrix and setup all RF products: M*RF +- N*LO
M1Out = zeros(M,N,2,pnts);

for x = (1:M)
    for y = (1:N)
        M1Out(x,y,1,:) = x*BPF1(1,:)+LO1(1,y);           %Shift for convolution around LO
        %M1Out(1,1,1,:) = M1R1L1(1,:)+LO1(1,1);           %Apply spur table levels
        M1Out(x,y,2,:) = BPF1(2,:)+ M1lv1(x,y);
    end
end

%Plot Input spectra
h = plot(BPF1(1,:), BPF1(2, :)+ M1RFIFiso);           %Adjust input for mixer RF isolation,
worst case
set(h,'color',[0,0,1]);
set(h,'LineWidth', 2)
hold on

%Loop for plotting mixer products

```

```

labelFlag = 0;           %used to disable legend annotations

for x = (1:M)
    for y = (1:N)
        lenny = length(M1Out(x,y,1,:));
        h = plot(reshape(M1Out(x,y,1,:),1,lenny), reshape(M1Out(x,y,2,:),1,lenny));
        if x == 1
            if y == 1
                set(h,'color',[0,1,0]);
                set(h,'LineWidth',3);
                labelFlag = labelFlag + 1;
            else
                set(h,'color',[1,0,0]);
                set(h,'LineWidth',3);
                if labelFlag == 2
                    hAnnotation = get(h,'Annotation');           %remove from legend to
avoid double annotation
                    hLegendEntry = get(hAnnotation,'LegendInformation');
                    set(hLegendEntry,'IconDisplayStyle','off')
                else
                    labelFlag = labelFlag + 1;
                end
            end
        elseif x == 2
            set(h,'color',[1,0.5,0.2]);
            set(h,'LineWidth',2);
            if labelFlag == 3
                hAnnotation = get(h,'Annotation');           %remove from legend to
avoid double annotation
                hLegendEntry = get(hAnnotation,'LegendInformation');
                set(hLegendEntry,'IconDisplayStyle','off')
            else
                labelFlag = labelFlag + 1;
            end
        elseif x == 3
            set(h,'color',[0.5,0.5,0.5]);
            set(h,'LineWidth',2);
            if labelFlag == 4
                hAnnotation = get(h,'Annotation');           %remove from legend to
avoid double annotation
                hLegendEntry = get(hAnnotation,'LegendInformation');
                set(hLegendEntry,'IconDisplayStyle','off')
            else
                labelFlag = labelFlag + 1;
            end
        end
    else
        set(h,'color',[0.5,0.5,0.5]);
        set(h,'LineWidth',2);
    end
    h = plot(-1*reshape(M1Out(x,y,1,:),1,lenny), reshape(M1Out(x,y,2,:),1,lenny));
    if x == 1
        if y == 1
            set(h,'color',[0,1,0]);
            set(h,'LineWidth',2);
        else
            set(h,'color',[1,0,0]);
            set(h,'LineWidth',2);
        end
    elseif x == 2
        set(h,'color',[1,0.5,0.2]);
        set(h,'LineWidth',2);
    elseif x == 3
        set(h,'color',[0.5,0.5,0.5]);
        set(h,'LineWidth',2);
    end
    hAnnotation = get(h,'Annotation');           %remove from legend to avoid
double annotation
    hLegendEntry = get(hAnnotation,'LegendInformation');
    set(hLegendEntry,'IconDisplayStyle','off')
end
end
end

```

Appendix A.5: FE Component lists

Component name	Manufacturer (Component name / code)
Amp 1	WENTEQ (ABL0100-01-3010)
Amp 2	WENTEQ (ABL0100-10-1810)
Amp 4 (Figure 4.8)	Mini-Circuits (ZX60-3018G-S+)
Digital Step Attenuator	Mini-Circuits (ZX76-31-SP+)
HPF 1	Mini-Circuits (SHP-700+)
HPF 2	Mini-Circuits (VHF-1500+)
LNA	MITEQ (AFS3-00100300-12-10P-4)
LO	Valon Technology (5008 Dual Synthesizer)
LO LPF	Mini-Circuits (VLF-2850)
LPF 2	Mini-Circuits (VLFX-780)
Mixer	Marki Microwave (T3A3-08CKP)
RF Switch 1	Peregrine Semiconductor (PE42556 Flip Chip Eval. Pack.)
RF Switch 2	Mini_Circuits (ZSWA-4-30DR)
3 dB pad	Mini-Circuits (VAT-3+)

Table A.1: Component list for RATTY SU (Figure A.4). Note that the stripline filters developed at SU and MeerKAT components are not listed in the table. The stripline parts are described in section 3.3.

Component name	Manufacturer (Component name / code)
Amp 1	WENTEQ (ABL0300-01-3414)
Amp 2	WENTEQ (ABL0200-01-2718)
Digital Step Attenuator	Mini-Circuits (ZX76-31-SP+)
HPF 1	Mini-Circuits (SHP-700+)
HPF 2	Mini-Circuits (VHF-1500+)
LNA	MITEQ (AFS3-00100300-12-10P-4)
LPF 2	Mini-Circuits (VLFX-780)
RF Switch 2	Mini-Circuits (ZSWA-4-30DR)
3 dB pad	Mini-Circuits (VAT-3+)

Table A.2: Component list for RATTY SKA (Figure 3.12). Note that the stripline filters developed at SU and MeerKAT components are not listed in the table. The stripline parts are described in section 3.3.

Appendix B: RATTY SU mixing plan for improved spurious response

This section presents an alternative frequency plan for the RATTY SU design, which avoids the most significant spurious components and realises a 60 dB mixing SFDR at maximum input power.

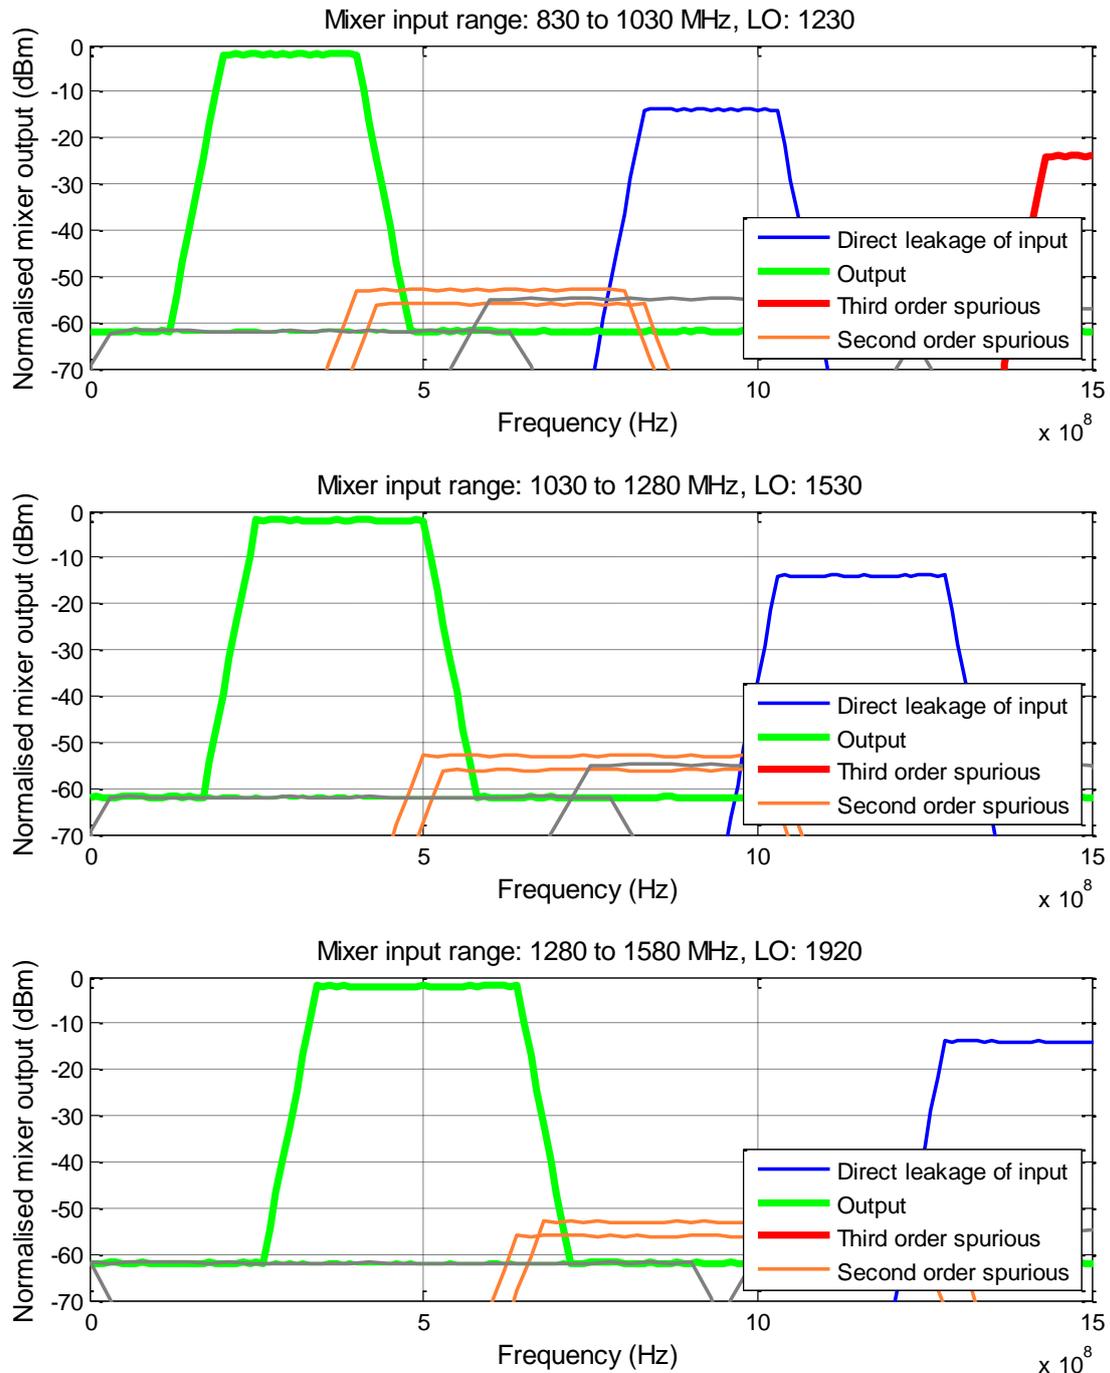


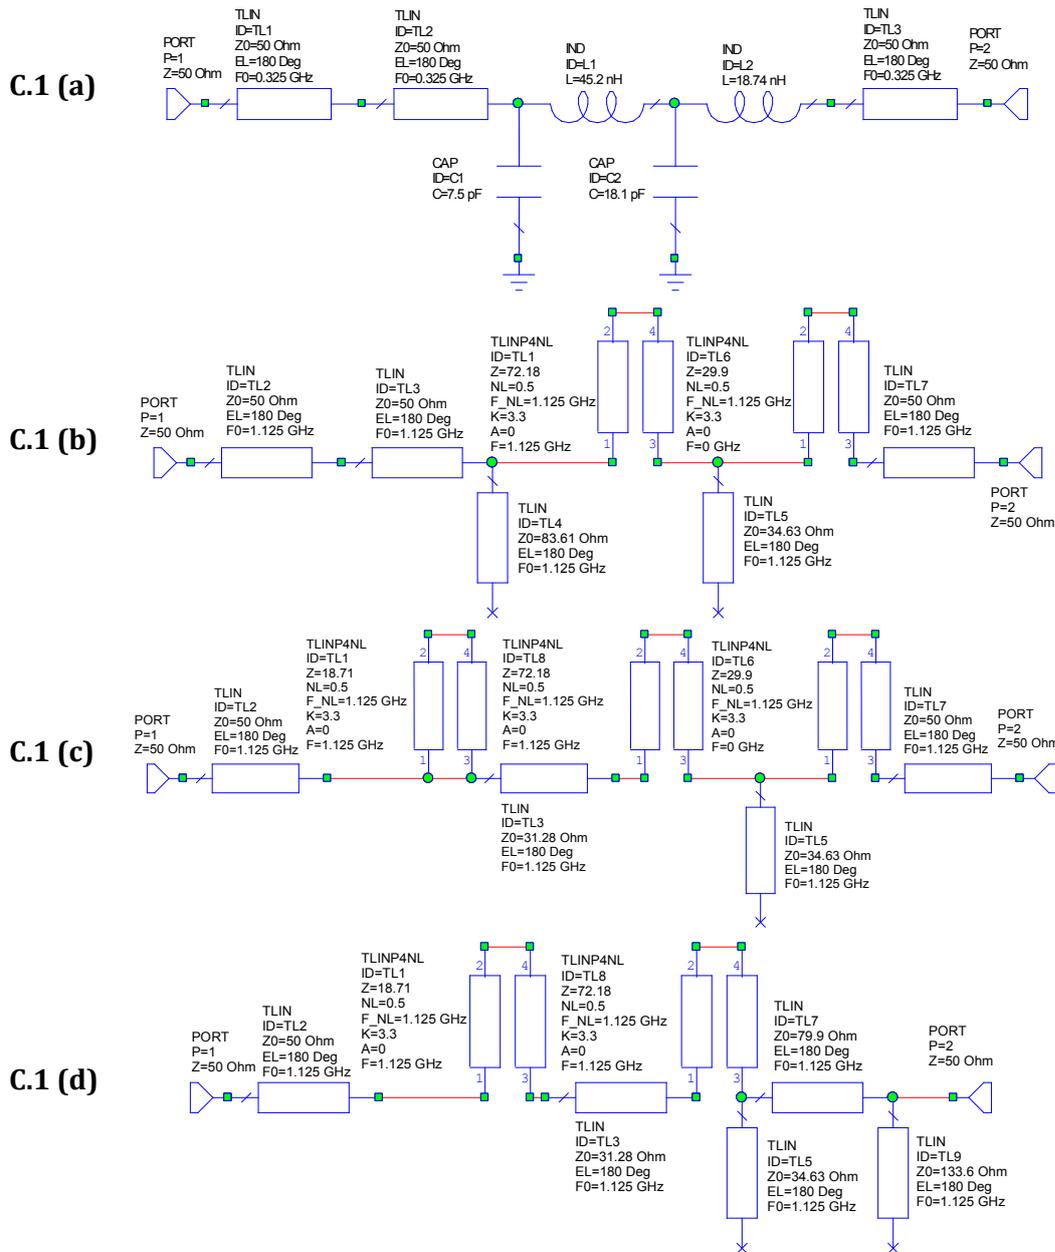
Figure B.1: Analyses of mixer spurious for improved SFDR. The RATTY SU path 2 bandwidth is covered in three ranges, which avoid the highest mixing products. Note that these bands do not include roll-off margins for filters, which will require further reductions in the bandwidths to maintain the SFDR.

Appendix C: Stripline filter designs

This appendix contains additional documentation of the stripline filter designs that were undertaken for the development of custom FE filters.

C.1 Transformation of a lumped-element to stripline filter.

This section documents the transformation of an initial, lumped-element filter into a stripline structure. Figures C.1 (a) – (f) are schematics of the stepwise transformation and below these is a MATLAB script used to calculate the impedance final values.



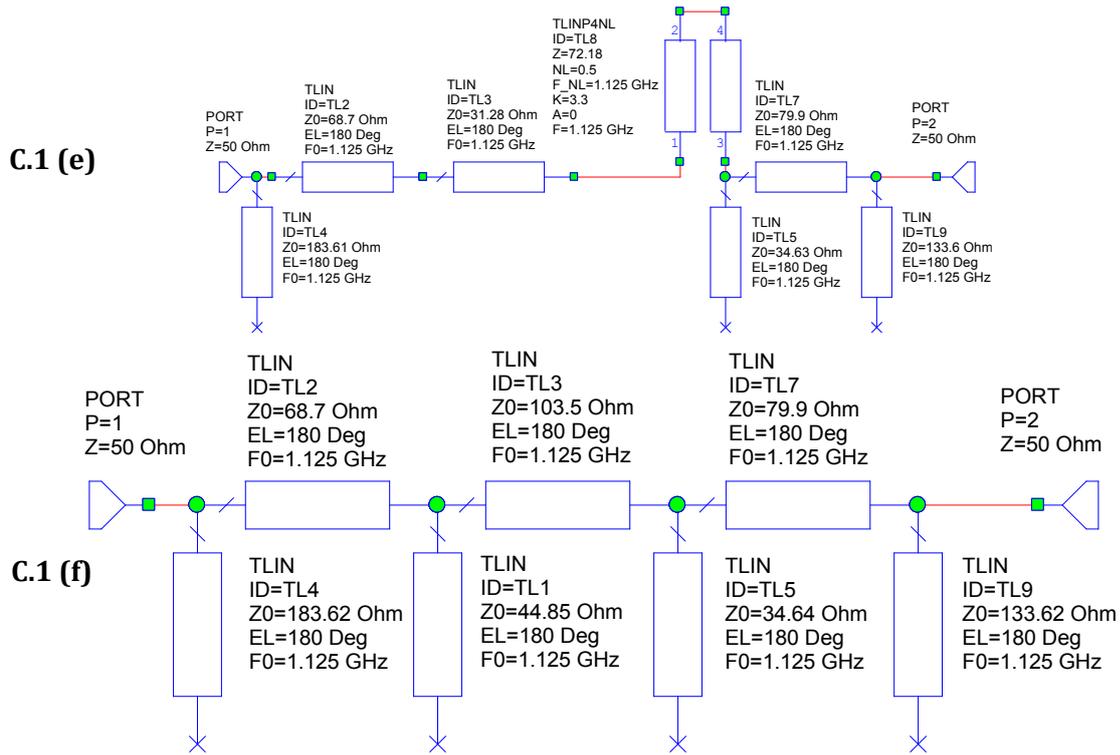


Figure C.1: Ideal stripline models (MWO) of a scaled, 4th order Butterworth filter. C.1 (b) is derived from (a) with the Richard's Transform. C.1 (b) to (f) show the application of the Kuroda identities described as transformations 1 to 4 in the script below, respectively.

MATLAB Code used for impedance calculations with the implementation of the first Kuroda identity.

```

Zo = 1;
omega1 = 1.28; %Scaling factor for bandwidth

ZC1 = 1/(0.7654 / omega1); %Butterworth element values
L1 = 1.8478/omega1; %Capacitance changed to impedance
ZC2 = 1/(1.8478/omega1);
L2 = 0.7654/omega1;

n1 = 1 + ZC1 / Zo; %Kuroda transforms
LT1 = Zo / n1; %Transformation 1
ZT1 = ZC1 / n1;

n2 = 1 + Zo / L2; %Transformation 2
ZT2 = n2 * L2;
ZCT2 = n2 * Zo;

n3 = 1 + Zo / LT1; %Transformation 3
ZT3 = n3 * LT1;
ZCT3 = n3 * Zo;

n4 = 1 + ZT1 / L1; %Transformation 4
ZT4 = n4 * L1;
ZCT4 = n4 * ZT1;

Z50C2 = 50*ZC2; %Split capacitance between sections and scale to 50 ohm
    
```

```

Z50CT2 = 50*ZCT2;
Z50CT3 = 50*ZCT3;
Z50CT4 = 50*ZCT4;

Z50ZT2 = 50*ZT2;    %Scale impedance of connecting strips
Z50ZT3 = 50*ZT3;
Z50ZT4 = 50*ZT4;

ZC1Avg = sqrt((Z50CT3*Z50CT2)); %Geometric average for symmetric filter
ZC2Avg = sqrt((Z50CT4*Z50C2));
Zl2Avg = sqrt((Z50ZT3*Z50ZT2));

```

C.2 Supplementary figures to RATTY SU filter designs.

This section contains additional figures for the path 2 filters in RATTY SU.

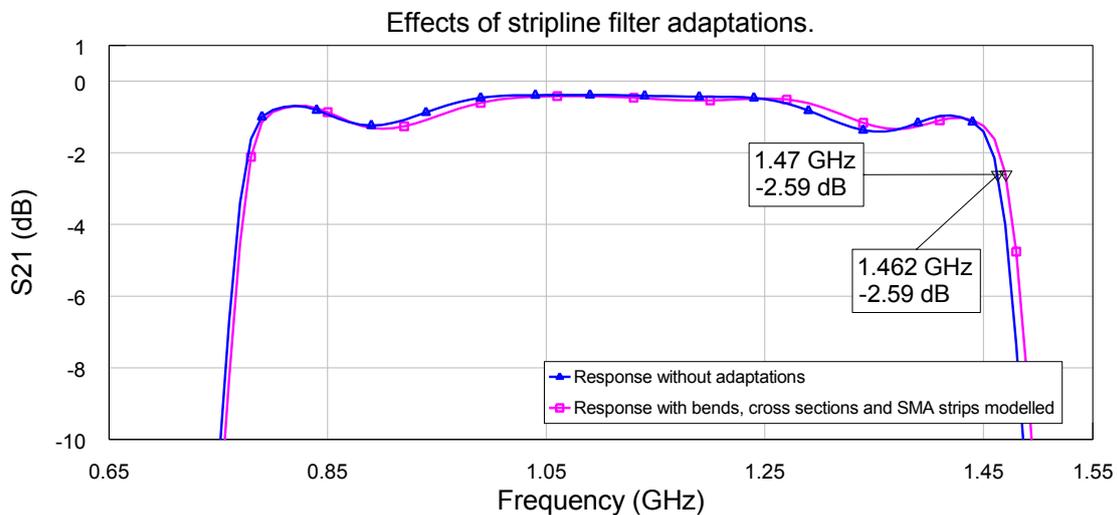


Figure C.2: Comparison of BPF 1 responses with and without bends, cross sections and input strips. Note that the models implement ideal transmission lines and that the responses are very similar.

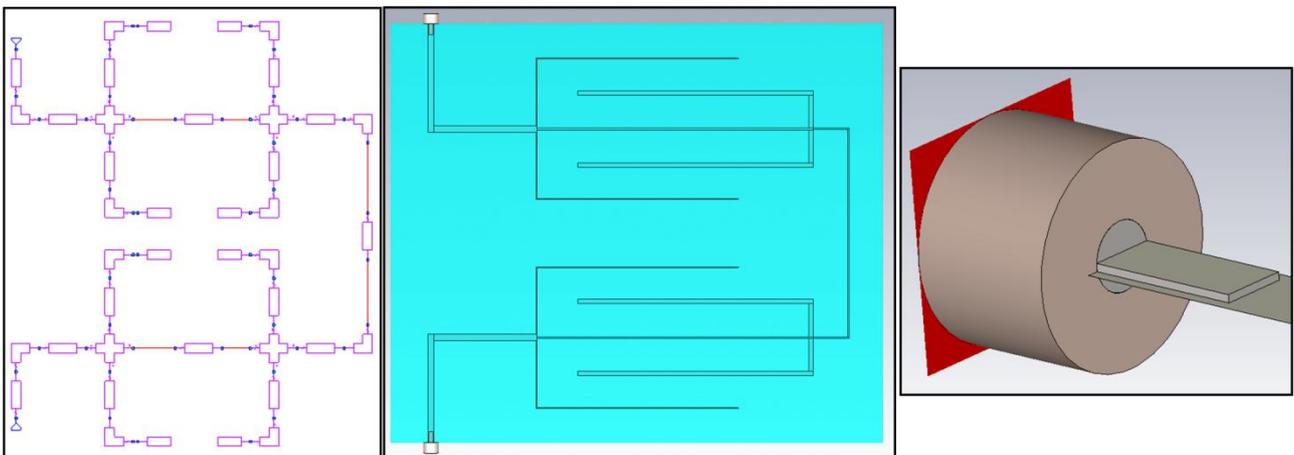


Figure C.3: MWO (left) and CST (middle) models of the final BPF 1 design and a flat pin SMA connector model (right). Note that the bends minimise the layout size, whilst retaining safe distances between lines and boundaries to avoid coupling between them.

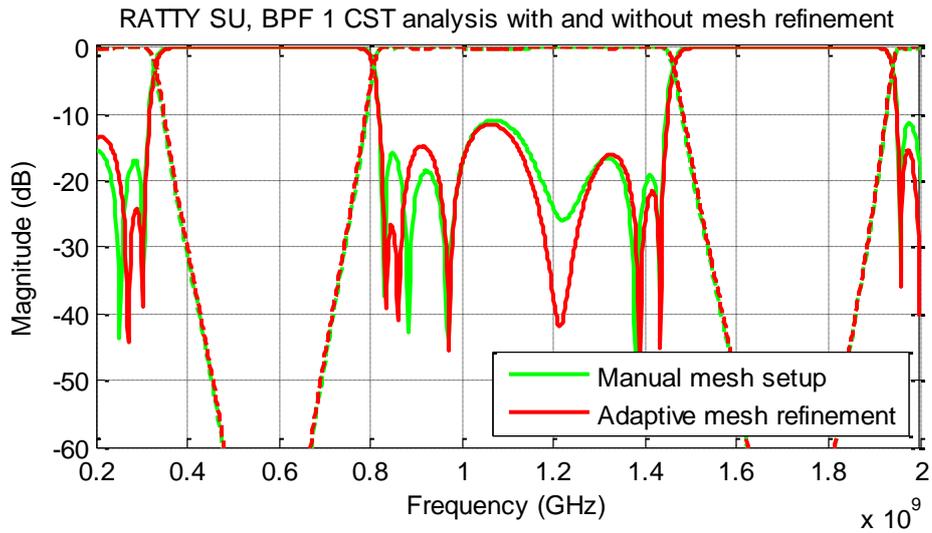


Figure C.4: Comparison of BPF 1 responses with and without CST adaptive mesh refinement. Note the S21 data is the broken lines and that the graphs agree, except for finer details.

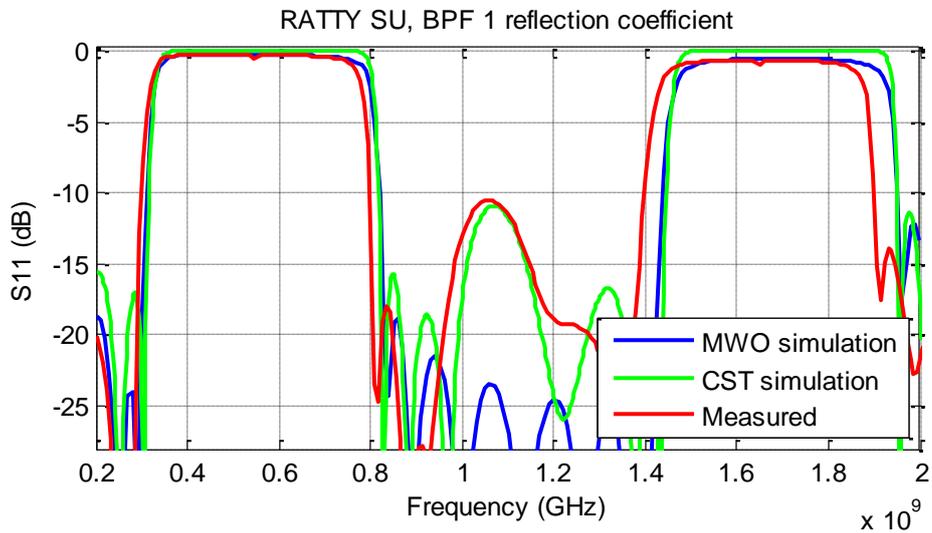


Figure C.5: Comparison of measured and simulated BPF 1 reflection coefficients. Note the downward shift in frequency for the measured filter. The figure shows reasonable agreement between the CST and measured response, with the exception of the shift and variations at lower values.

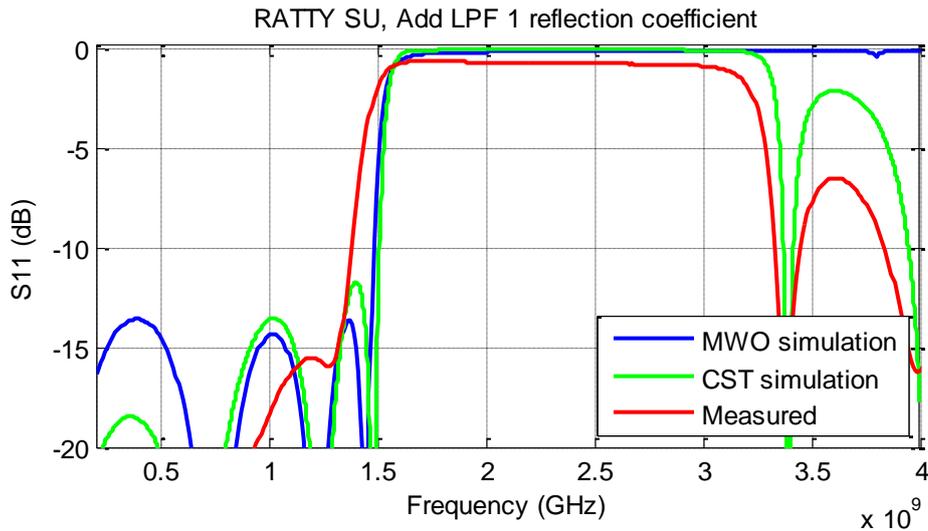


Figure C.6: Comparison of measured and simulated Add LPF 1 reflection coefficients. Note the downward shift in frequency for the measured filter. However, the figure shows reasonable agreement between the CST and measured response, with the exception of the shift and variations at lower values.

C.3. RATTY SU path 3 filters and measurements

This appendix documents the design of filters for a third signal path to the mixing XDM, which includes tests and analyses of the produced parts. The section concludes with a comparison of the measured and simulated gain function of the third signal path.

The filter chain in signal path 3 of the RATTY SU FE contains the combination of BPF 1, Add LPF 2 and HPF 2 (Figure 3.#). The BPF and LPF are stripline designs developed with the process described in section 3.3.1.1. Table C.1 lists the specifications of both parts that were used during optimisation. Note that these goals were set whilst the production of BPF 1 was underway and that they are intended to complement the ranges predicted by the optimised design of the path 2 filter.

Passband	1420 - 2070 MHz
Down-converted passband	125 - 775 MHz
60 dB out-of-band suppression	Below 1170 MHz and above 2320 MHz
Roll-off rate	60 dB / 250 MHz

Table C.1: Initial BPF 2 specifications.

There were several configurations assessed for implementing a 6th order structure on different substrates. The increase in the filter order makes the specifications less stringent within the production constraints. Tapered input lines are added, which improve the matching by making the change in impedance more gradual. The final design incorporates a substrate with $\epsilon_R = 10.9$, which lead to a significant reduction in the size of the configuration. Figure C.7 are comparative plots between the constructed and simulated forward transmission gains and reflection coefficients. The measured ranges of BPF 2 are listed in Table C.2 and the simulation models are shown in figure C.8.

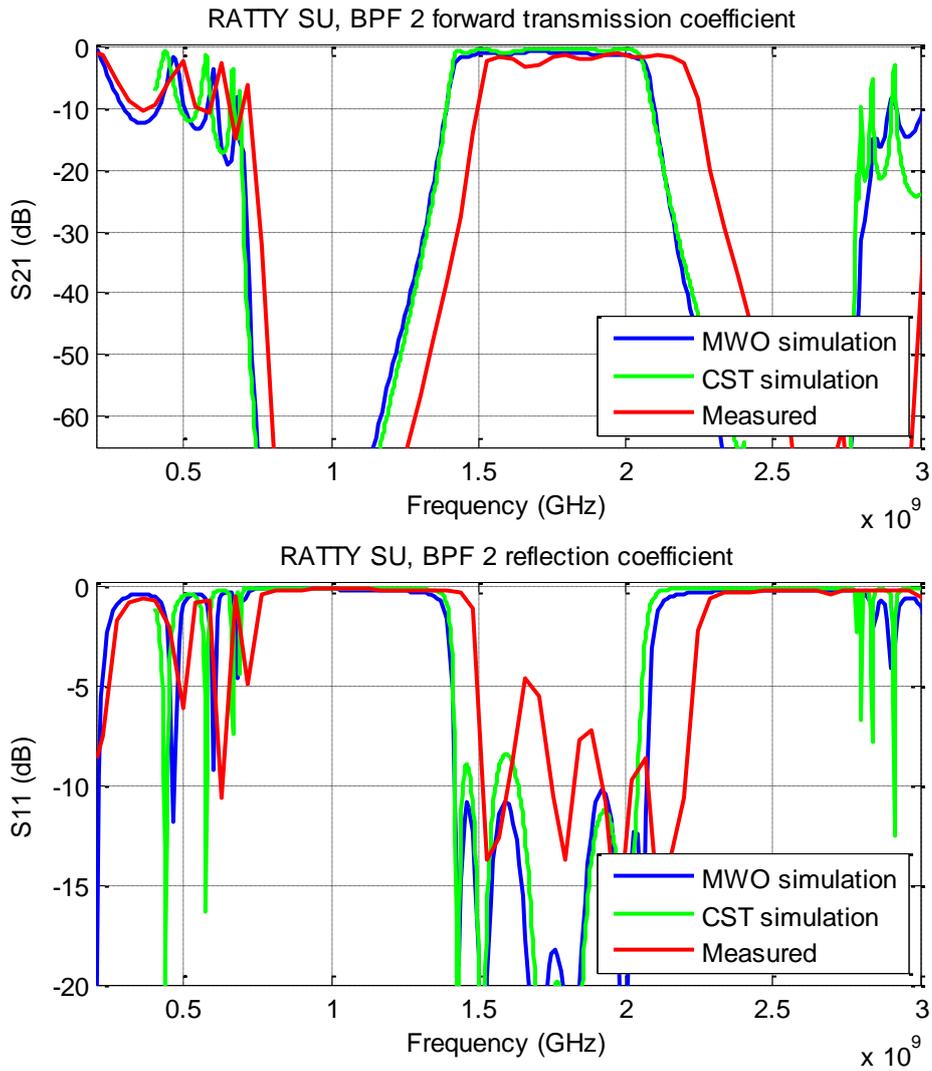


Figure C.7: Measured and simulated BPF 2 characteristics.

Passband	1520 - 2210 MHz
60 dB out-of-band suppression	Below 1280 MHz and above 2540 MHz

Table C.2: Measured BPF 2 ranges.

An upward shift occurred in the centre frequency and is attributed to the actual ϵ_R value being lower than that of the simulated substrate. This leads to the stub lengths being too short with wrong impedance values. Furthermore, the roll-off is non-symmetrical and the reflection coefficient is poor, which is apparent in the rippled passband. However, the filter is used to complete channel three and its refinement is a prospect for future work. In principle, the Add LPF 2 design is identical to BPF 2 and the results obtained for it are plotted in figure C.9. A simulated response of the filter chain in path 3 is shown in figure C.10 and is based on measured component characteristics, with the HPF 2 and buffering attenuation pads included.

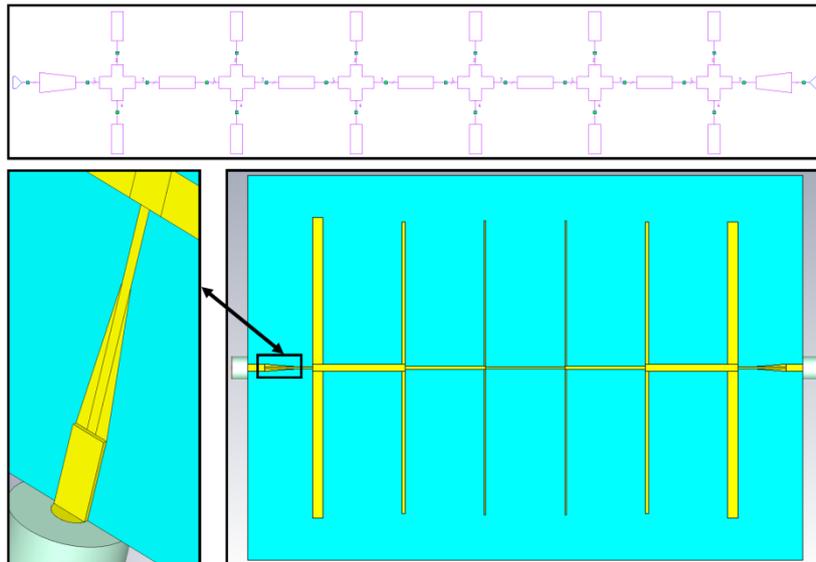
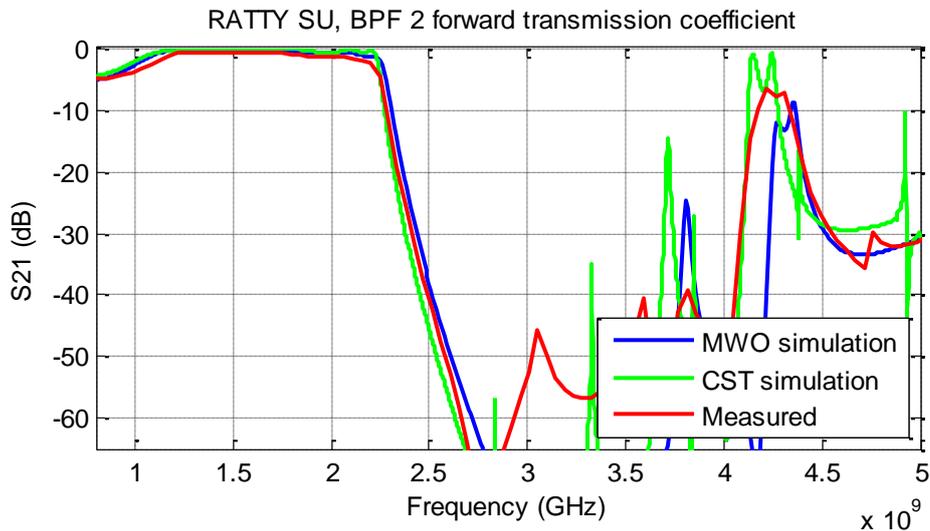


Figure C.8: The MWO (top) and CST models (bottom right) of BPF 2. Note the tapered input lines (Bottom left).

Although the filter chain does not comply with the specified rejection ratio for all frequencies, it will avoid spectral leakage in the mixing process with an appropriate selection of LO frequency (~ 2.3 GHz). However, the shift of the system response leaves a gap in the bandwidth coverage between 1.4 and 1.52 GHz. Furthermore, the lack of suppression above 3 GHz will leave this path vulnerable to out-of-band interference in these ranges and this requires an improved LPF. A high power signal within these unwanted filter responses can feature in the output bandwidth as a $|2*LO - 1*RF|$ spurious product. Nonetheless, the third signal path was assembled and tested in the mixing XDM and a comparison of the simulated and measured gain function of this path is shown in figure C.11. This figure shows good agreement between the simulated and SA-measured gain function, despite the latter being uncalibrated.



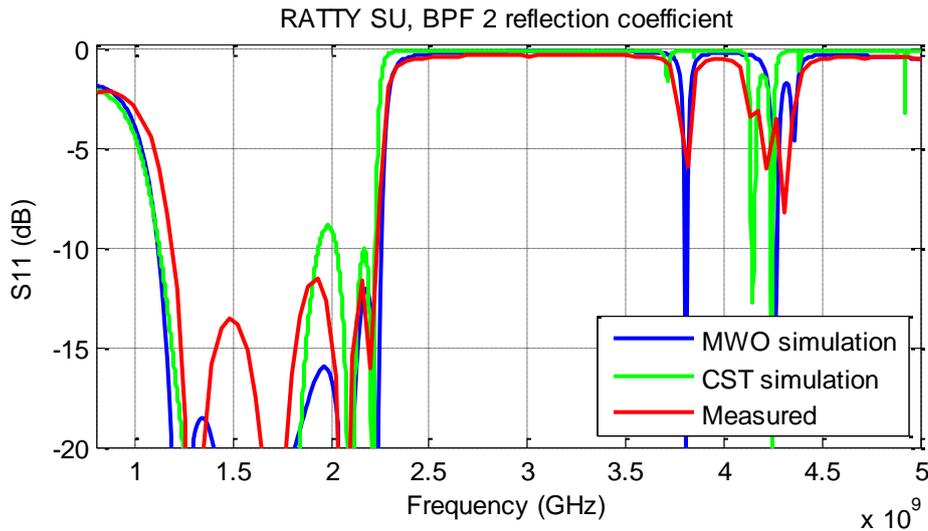


Figure C.9: Measured and simulated Add LPF 2 characteristics. Note that the effect of an inaccurate ϵ_R value results in variance between simulated and measured results, although the effects are less severe than in the BPF 2 case.

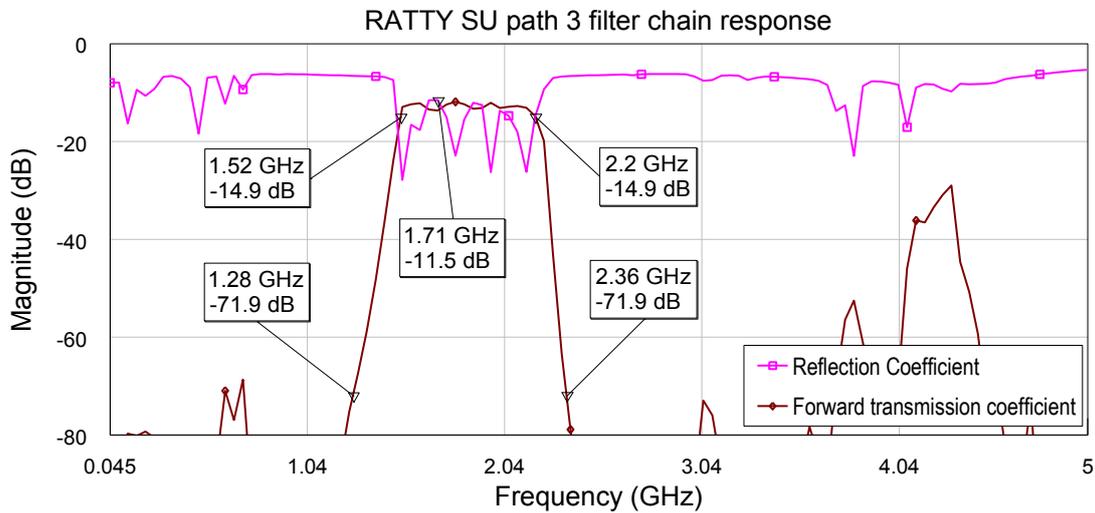


Figure C.10: Simulated response of the filter chain in signal path 3, RATTY SU. Note that the passband and out-of-band ranges are indicated.

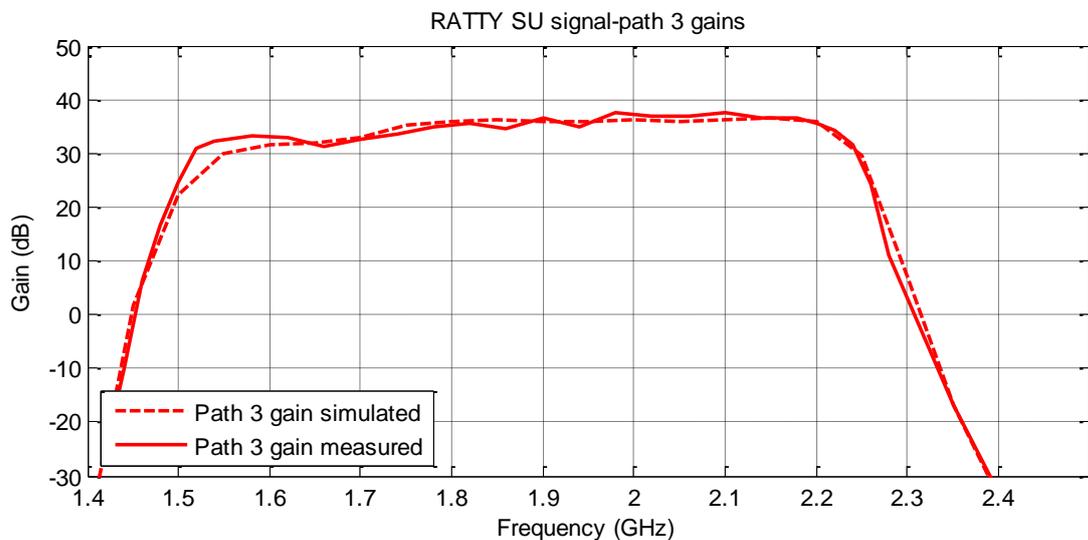


Figure C.11: Simulated response of the filter chain in signal path 3, RATTY SU. Note that the passband and out-of-band ranges are indicated.

C.4. Supplementary figures to RATTY SKA filter designs

This section contains additional figures for the RATTY SKA filter developments.

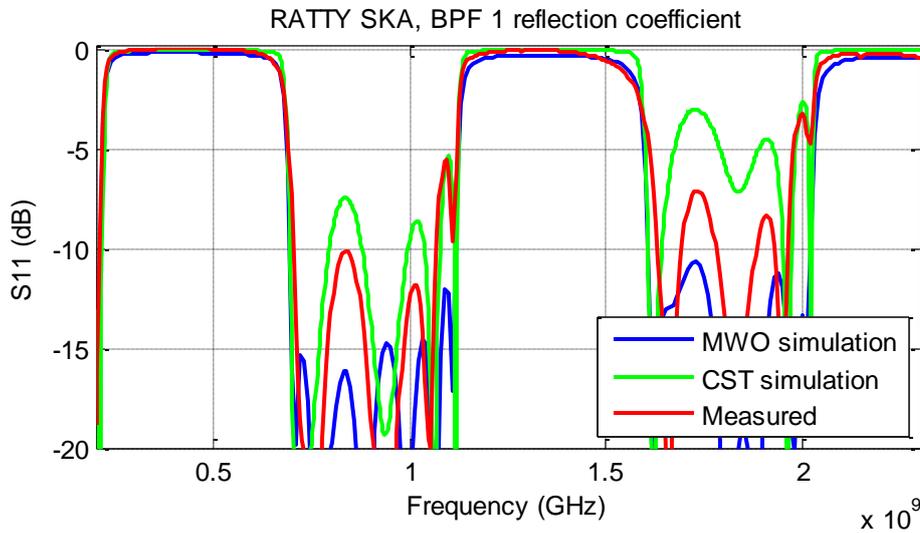


Figure C.12: Simulated and measured results of the BPF 1 reflection coefficient. Note that there is reasonable agreement between the CST simulation and the measured filter response.

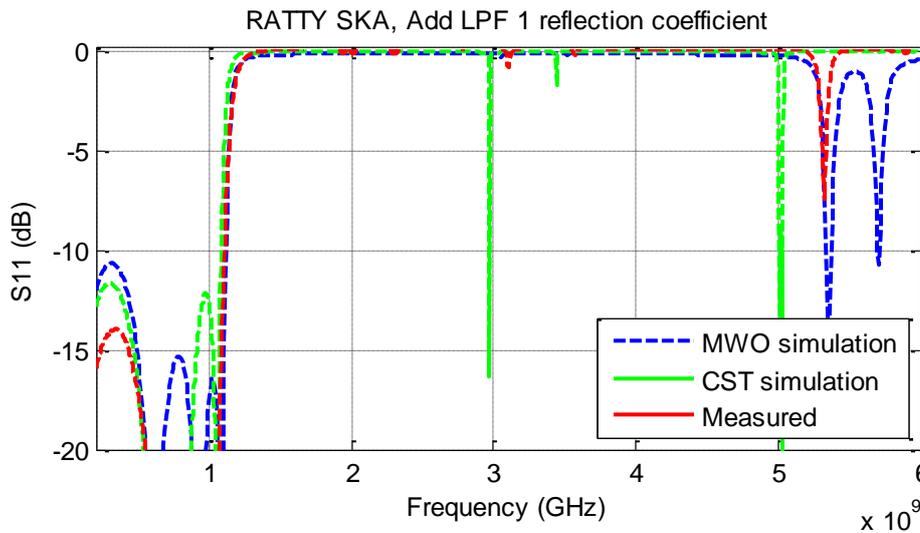


Figure C.12: Simulated and measured results of the Add LPF 1 reflection coefficient.

Appendix D: Component measurements and analyses

Appendix D.1: RATTY SU synthesiser

A Valon 5008 frequency synthesiser module is used as the LO source in the mixing FE. The device is programmed using a standard software package and a Universal Serial Bus (USB) connection. It is housed in a separate, shielded enclosure with dual-output SMA ports. Table D.1 contains several of the synthesiser specifications.

Frequency range	Power range	2 nd harmonic	3 rd harmonic	>3 rd harmonic	Port return loss
138 – 4400 MHz	-1 – 7 dBm	< 25 dBc	< 20 dBc	< 40 dBc	< -12 dB

Table D.1: Valon synthesiser specifications.

The LO filter was selected with the intention that a single output port could be used as an LO source, which enables two mixers to be operated from a single synthesiser. A single LPF is used to suppress the harmonics of multiple output frequencies. Also, two buffering 3 dB attenuation pads are included to improve the matching between the synthesiser and mixer, since the latter has poor return loss properties. The XDM mixing FE features only two mixing paths with 1.5 and 2.3 MHz LO frequencies, respectively. The most suitable and commercially available LO LPF was selected to suppress the harmonics of these output frequencies. This component only provides ~10 dB attenuation to the 2nd harmonic of the 1.5 GHz LO output, although > 25 dB suppression is applied to all products above 3.5 GHz.

After testing the synthesiser with the filtering, broadband leakage of the 10 MHz internal reference was discovered and this introduced the need for a LO HPF. A commercially available HPF was selected to suppress these emissions in the ADC-baseband range, where direct leakage of the signals would result in interference within the design SFDR. This choice was limited by the restriction of maximum 1.7 dB loss at 1.5 GHz, or the mixer would not have sufficient LO power. Figure D.1 shows the SA measured 1.5 GHz synthesiser output with and without filtering.

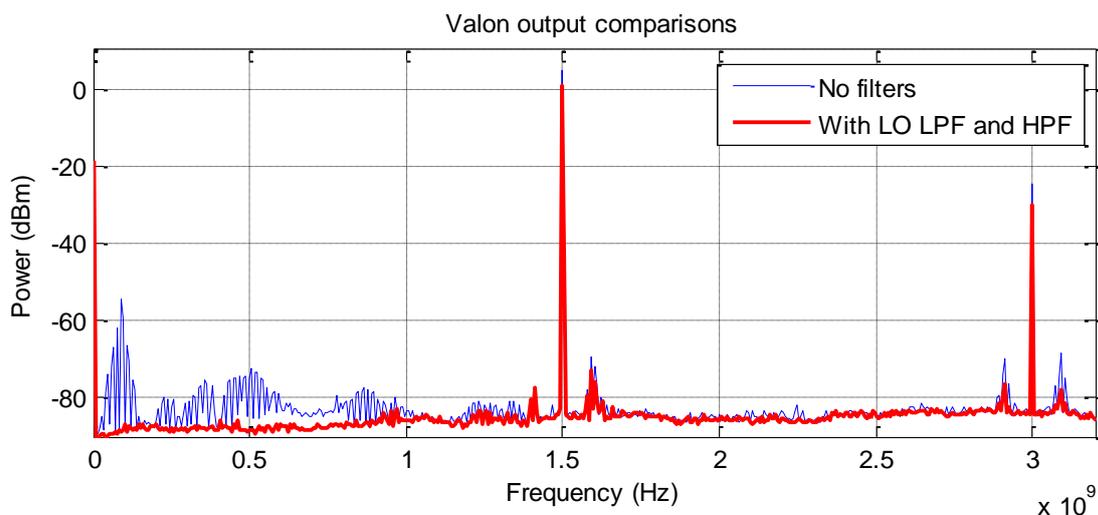


Figure D.1: SA measurements of 1.5 GHz synthesiser outputs with and without filtering.

The broadband leakage of the reference harmonics and the second harmonic of the fundamental output are apparent in the figure. As demonstrated by the second plot, these products are suppressed by the filters, although it is clear that the LPF is not effective in rejecting the 2nd harmonic. Note that there are also groups

of 10 MHz harmonics that occur around the fundamental and 2nd harmonic which cannot be reduced by this filtering. The components appear to be IMD products of the synthesising process and can also result in IMD products in the mixing process.

The filters improved the interference issues, although there is still detectable leakage of the reference source. This results from a lack of rejection in the 0.5 to 1 GHz range where the HPF suppression reduces from -48 to -16.8 dB. Ideally, notch filters are required that only pass the fundamental outputs and improving LO filtering must be considered in revised mixing FEs. This is essential for systems that follow a similar frequency-conditioning strategy and have greater sensitivity.

Appendix D.2: RATTY SU mixer characterisation

A set of two-port VNA measurements was performed to characterise the mixer ports, whilst a SA and SGs were used to do conversion loss, spurious and isolation tests. Two mixers were bought and the ports of both were characterised. The reflection coefficients showed some variation between the components and that the characteristics of the RF port was dependent on the LO frequency, since the both LO frequency were applied for these measurements. Nonetheless, the RF and IF ports were found to have reflection coefficients < -10 dB within the ranges of both RATTY SU paths. This suggests that some the buffering attenuation pads may be redundant and could be removed in revised systems. Figure D.2 is a plot of the RF port reflection coefficients for both mixers and LO frequencies.

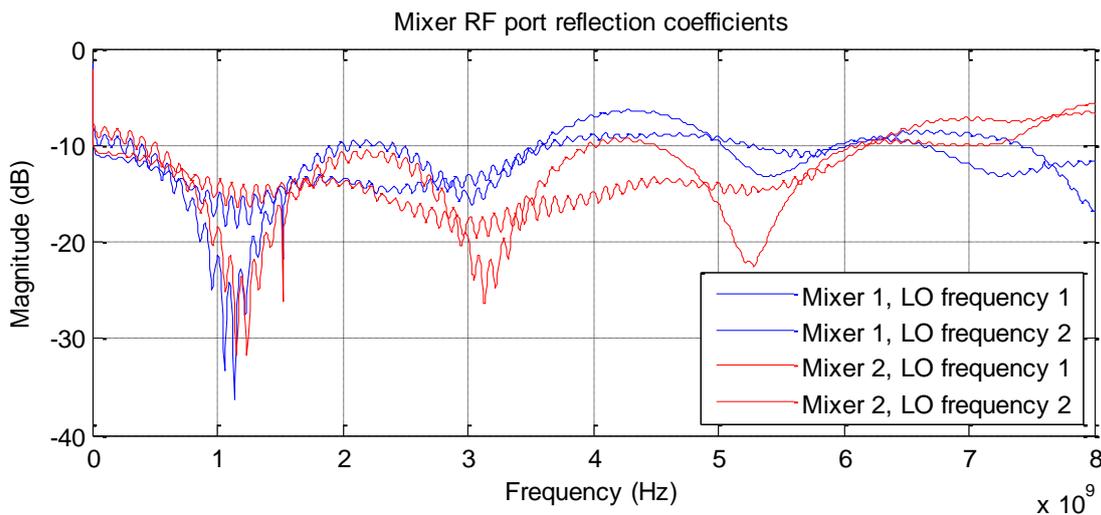


Figure D.2: Mixer RF-port, reflection-coefficient measurements at different LO frequencies.

As explained in section 3.4.1, the remaining mixer tests were performed in a configuration similar to the implementation context of RATTY SU. This included using the respective path filters at the RF input, buffering attenuators and the synthesiser with filtering. The SG output with input filtering was tested on the SA at several frequency points. All harmonics were found to be below 70 dBc. It is conceivable that some components may have been higher at the test frequencies used for the mixer assessment, although the setup was thought to be sufficient for the spurious tests not be influenced by input products within the design SFDR of 60 dBc. This assumption is further substantiated by the < -10 dBm LO leakage specified in the datasheet for the frequency ranges within which the highest LO harmonics would occur. Figure D.3 is a schematic depiction of the mixer-measurement layout.

Since the fundamental mixer products are used for the conversion loss calculation and these power levels are not affected by SG harmonic levels, the input filters were not used in these measurements. This simplifies the calibration of the mixer input and output powers, which then only takes into account the losses of the 3 dB attenuation pads and cables to the IF and RF ports. These measurements were conducted for both mixing signal paths and the results are displayed in figure D.4. Each frequency range was covered in 50 MHz steps using a SG input, although the SG was not calibrated and introduces some unknown errors.

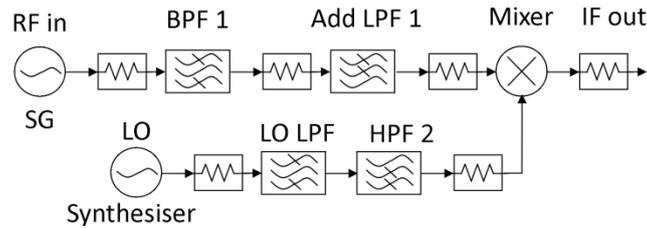


Figure D.3: Schematic showing the mixer-test configuration. Note that the HPF 1 component is not included as an input filter.

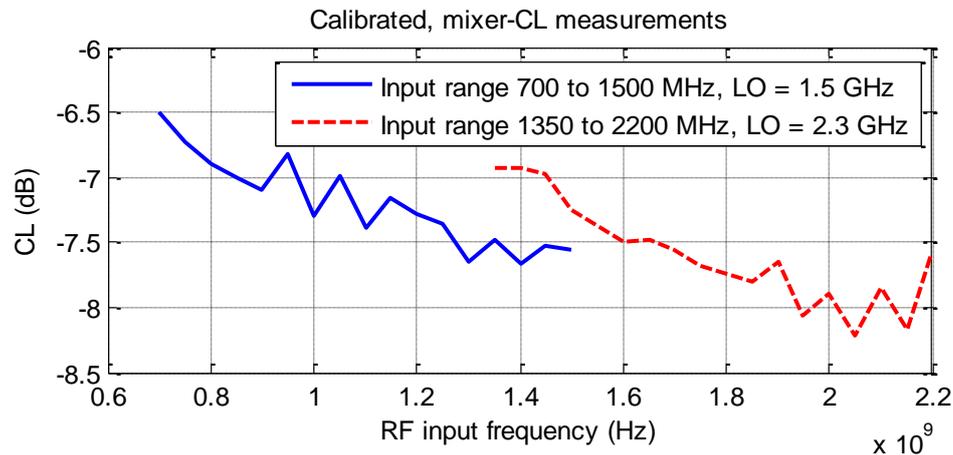


Figure D.4: Measured and calibrated conversion loss curves for mixing signal paths 1 and 2.

A maximum conversion loss of 7.6 and 8.2 dB was calculated for path 1 and 2, respectively, with ~ 1.2 dB variation across the bandwidths. A worst-case 8 dB figure was used in path 2 analyses, which is conservative when considering these results. These results are worse than the predicted levels, although the datasheet values are provided at a single IF, whilst the measurements are for a static LO frequency. This demonstrates the application dependency of mixer parameters. In particular, the figure shows that the conversion loss increase as the IF frequency becomes lower and this will reduce the flatness of the bandshape. The figure indicates that the conversion loss of the mixer increases as the IF becomes lower, which is evident in the variation measured across both bandwidths. This will affect the flatness of the bandshapes.

No calibration was required for the spurious or isolation calculations, since these parameters are specified relative to the RF output power. To determine the RF-to-IF isolation, the conversion loss is added to the output RF power levels and compared to the leaked signals. The minimum measured isolation between these ports was 25 dB, which is used in the spurious analysis to determine the maximum unwanted mixer products. The conversion loss was derived from the spurious measurements, which were also implemented a programmable SG. However, the bandwidth ranges could only be swept in sections, since the signals were spaced at 10 MHz intervals and various products would cause interference and obscure the low level spurious components. Figure D.5 shows two tone sweeps of path 2 and demonstrates the SFDR of the mixer.

Firstly, it is apparent from figure D.5 that a SFDR > 60 dBc was demonstrated for these path 2 ranges and no spurious components were observed above this level within the ADC baseband range. Furthermore, note that the square block highlights $1*LO - 2*RF$ spurious components of the RF output range on the left and that these products would be obscured if a single SG sweep was attempted. Lastly, the spurious elements on the far right and in the middle are $2*LO - 2*RF$ and $1*LO - 2*RF$ products of the output in the centre of the range.

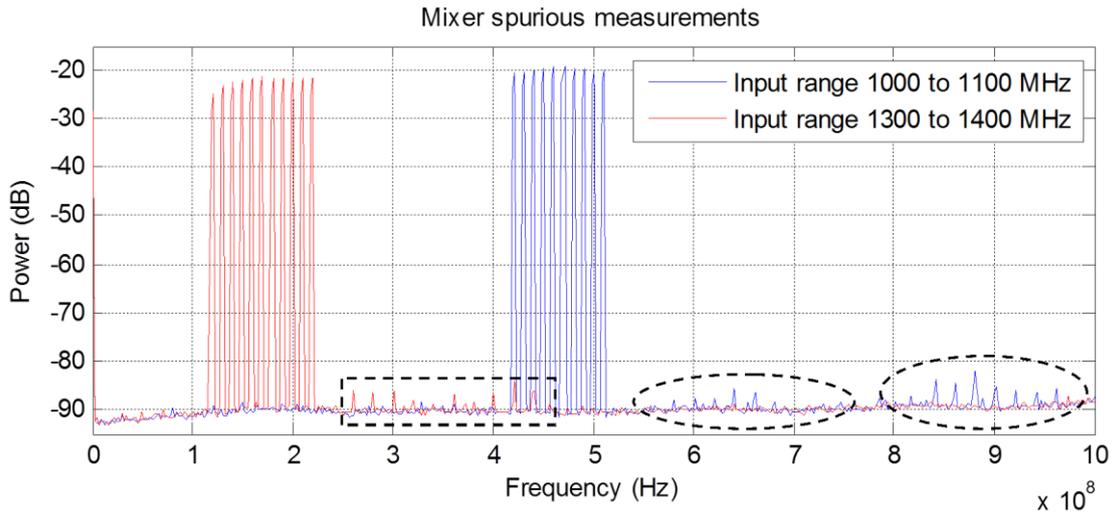


Figure D.5: Mixer spurious measurements for RATTY SU path 2 range. Note that the ellipses on the right highlight spurious from the 1 to 1.1 GHz inputs and the rectangle highlights spurious from the 1.3 to 1.4 GHz.

Appendix D.3: Supplementary documentation of FE analyses

This section contains figures and tables that are included as supplementary information derived during the FE analyses that are documented in section 3.4.

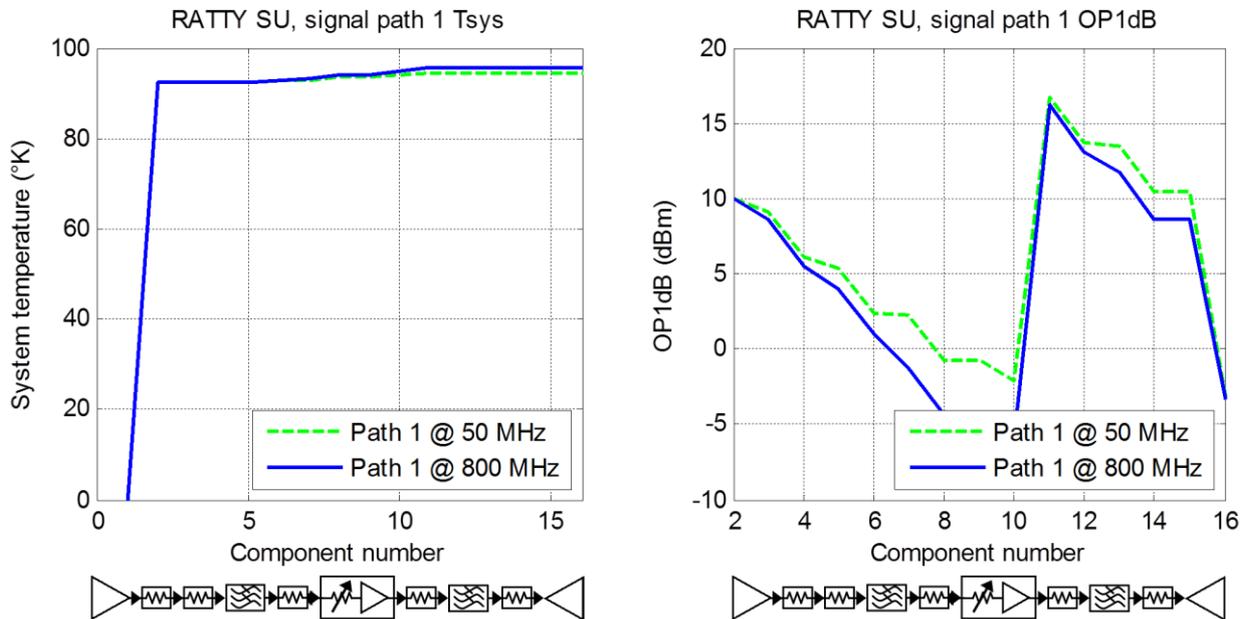


Figure D.6: Spread-sheet calculations for the T_{sys} (left) and OP1dB (right) parameters of the RATTY SU, signal path 1.

Note that the circuit diagrams indicating the components and position do not show switches. Also, the OP1dB curve indicates that compression point of the final amplifier (17 dBm) is reached before the LNA behaves non-linearly. This can be seen in the curve where the OP1dB does not increase by the full 30 dB gain of the last amplifier. Thus, an increase in attenuation is required to enable the measurement of high input powers. The final drop in OP1dB is due to the inclusion of the -2 dBm ADC limit.

RATTY SU, Path 2 (1100 MHz)				
Attenuator setting	Gain (dB)	NF (dB)	OP1dB (dBm)	Input noise power (dBm)
0 dB	38.6	2.5	11	-76.7
-10 dB	28.6	3.4	1.2	-75.9
-20 dB	18.6	7.8	-8.3	-71.5
-31 dB	7.6	17.44	-19.3	-61.8

Table D.2: Table of system parameters for RATTY SU path 2 at several step-attenuator settings. The decrease in the OP1dB in table D.2 coincides with the increase in attenuation, which indicates that the system is within its linear operating range. Also, note that the degradation of the NF will not lead to an increase in the noise floor level of measurements. This is due to the lowered gain.

RATTY SU, Path 1				
Parameters	AWR simulator		Spread-sheet calculations	
	Worst-case	Best	Worst-case	Best
NF	1.23	1.21	1.24	1.22
T_{Sys} (°K)	95.3	93.6	95.8	94
Gain (dB)	43.6	48	42.9	47.4
Input noise power (dBm)	-76.3		-77.9	-78
Maximum Gain variation (dB)	4.4		4.5	
Maximum OP1dB (dBm)	10.4		10.5	

Table D.3: Comparison of spread-sheet and simulator parameter calculations for RATTY SU path 1. Note that all the parameters in are good agreement.

Appendix E: RATTY script commands and help

The help output of the *r2_time.py* command is shown below as an example of this feature. Furthermore, table E.1 lists the commands for all the scripts. Note that the command shown below is in a –option format, which is used to display the help menu in this case. When nominal entries are required, the value is entered subsequent to the specifying the –option and all the option can be added for specific control settings.

r2_time.py – h:

-h, --help	Show this help message and exit
-v, --verbose	Enable debug mode.
-o, --plot_hist	Do not plot the histogram.
-u, --update	Do not update the plots (only plot a single capture).
-i, --plot_indep	Plot interleaved ADC independantly.
-t CAPTURE_LEN	Plot this many nano-seconds around the trigger point. Default:100
-a ANT, --ant=ANT	Choose an antenna calibration file. Note that this will auto-select y-axis units to be dBuV/m. Default:none
-n N_CHANS	Number of frequency channels to resolve in software. Default:512
-l TRIG_LEVEL	Ask the hardware to wait for a signal with at least this amplitude in mV before capturing. Default: 0 (disabled, just plot current input).
-f PLAY_FILE	Open an existing file for analysis.
-c CONFIG_FILE	Specify the configuration file to use.
-s, --csv_file	Output (convert) each timestamp to a separate CSV file.
-w, --wait_keypress	Wait for a user keypress before storing/plotting the update.
-p NUM_PLOT	De-activate the figure and record n datasets. Default allows plotting
-x ATTN1	Set level of Attenuator 1): pre-mixer
-z ATTN2	Set level of Attenuator 2): post-mixer
-A ATTN	Set levels of Attenuators 1 & 2 simultaneously. If no attenuator option selected, both attenuators set to 31dB
-F FREQ	For selecting freq channel: 1 = Ch1: 0-828 MHz; 2 = Ch2: 800-1400 MHz; Ch3: 1400-2100 MHz'

Command:	Use:
<i>r2_time.py</i>	TD operation.
<i>r2_spectrum.py</i>	FD operation.
<i>r2_init.py</i>	System check and FPGA programming.
<i>r2_init_serial.py</i>	FE control.

Table E.1: List of terminal commands for RATTY control.

Appendix F: Mixing XDM baseband SFDR

This section contains additional documentation of the RATTY SU SFDR assessment. There are several measurements that are similar to the tests performed and described in section 4.4.2 with spread-sheet comparisons. A plot of the worst-case SFDR performance in the baseband is also shown.

RF and IF attenuation, respectively (dB)	Fundamental (dBm)	2 nd IMD (dBm)	3 rd IMD (dBm)
0 15	-7.6	-39.1	-56.1
0 18	-10.6	-42.1	-59.2
3 15	-10.6	-45	-66.3
3 18	-13.6	-48.1	-69.5

Table F.1: Measured output levels used for calculating intercept points and SFDRs of the RATTY SU baseband path. Note that simulation OIP2 and OIP3 values are calculated from the simulated OP1dB with adjustments of 17.2 and 10, respectively. These values were calculated from the first measurement in the table. (Input power = -59 dBm)

RF and IF attenuation, respectively (dB)	Measurement OIP2 (dBm)	Measurement $SFDR_2$ (dB)	Simulated OIP2 (dBm)	Simulated $SFDR_2$ (dB)
0 15	23.9	31.5	23.9	31.5
0 18	20.9	31.5	21.9	32.5
3 15	23.8	34.4	23.9	34.5
3 18	20.9	34.5	21.8	35.4
RF and IF attenuation, respectively (dB)	Measurement OIP3 (dBm)	Measurement $SFDR_3$ (dB)	Simulated OIP3 (dBm)	Simulated $SFDR_3$ (dB)
0 15	16.7	48.5	16.7	48.7
0 18	13.7	48.6	14.7	50.6
3 15	17.3	55.7	16.7	54.7
3 18	14.4	55.9	14.6	56.5

Table F.2: Measurement and simulation based calculations of key intercept points and SFDRs for the RATTY SU baseband. There is good agreement between the measured and calculated parameters.

RF and IF attenuation, respectively (dB) [Input power (dBm)]	Measurement		Simulated	
	OIP2 (dBm)	$SFDR_2$ (dB)	OIP2 (dBm)	$SFDR_2$ (dB)
25 25 [-33]	6.7	21.3	6.7	29.3
31 31 [-33]	-3.4	23.2	0.8	29.4
0 20 [-54]	17.8	27.8	-20.3	30.3
RF and IF attenuation, respectively (dB) [Input power (dBm)]	Measurement		Simulated	
	OIP3 (dBm)	$SFDR_3$ (dB)	OIP3 (dBm)	$SFDR_3$ (dB)
25 25 [-33]	-0.7	27.8	1.6	32.5
31 31 [-33]	-12.6	28	-9.7	33.9
0 20 [-54]	12.2	44.4	13.1	46.3

Table F.3: Comparisons of measured and simulated OIPs and SFDRs for the mixing path. Note the varying attenuation settings and input powers. There is good agreement between measured and simulated parameters for these measurement configurations.

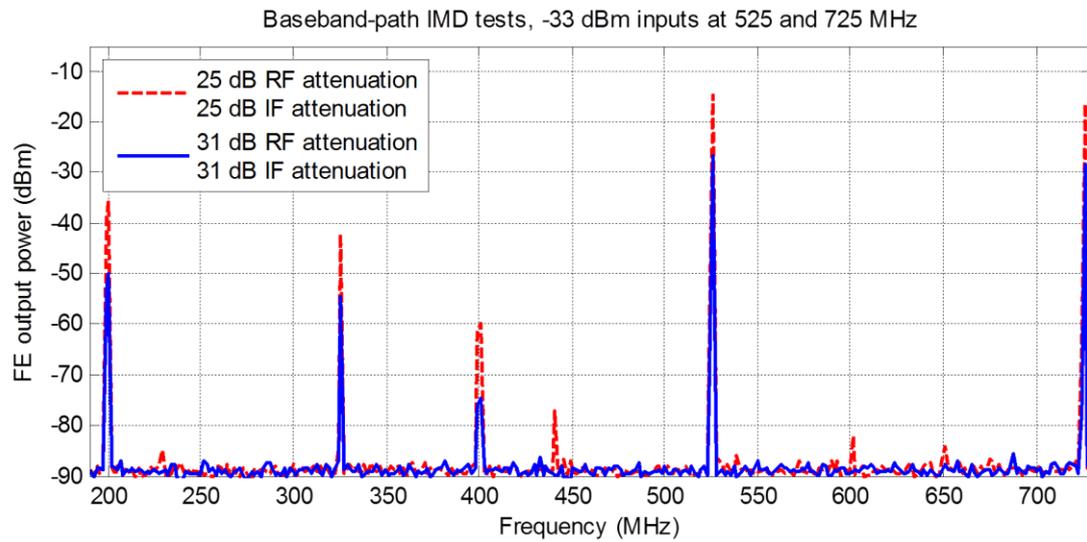


Figure F.1: Worst-case SFDR performance of the baseband path. Note that this is for near-maximum input power and that a SFDR of ~ 23 dBc is demonstrated.

Appendix G: RATTY design documents

This appendix contains two of the electronic references that were compiled during the development of the two experimental models. These require permission to view online and have been included here for ease of access.

G.1 RFI Measurement System - Requirements Specification for Ratty-2

Background

The design work for Ratty-2 has already progressed significantly, and therefore the Requirements Specification document does not start from a clean slate. It is nevertheless instructive to see what we're aiming for, to predict what we are going to achieve with the current design, and to influence the path going forward.

It is envisaged that the "core" of the RFI Measurement System can be used for a variety of purposes:

1. As a back-end to the Houwteq (ISSA) measurement equipment, to be used by *all* MeerKAT subsystems and Infrastructure to verify radiated RFI requirements.
2. To help with the design of the Digitiser and Receiver enclosures, by doing RFI measurements in a lab environment, before qualifying the design at Houwteq.
3. As a mobile measurement system to be used on-site.
4. As a permanent measurement system to be used on-site.
5. As a system to measure the stability of the Receiver. For this purpose, Ratty-2 will have to dump spectra at a rate of 1/20 ms (see R.005).

The objective of this document is to flesh out the requirements for Issue #1 above.

MeerKAT Radiated RFI Requirements

The radiated RFI requirements for both the Receiver and the Digitiser are shown in the graph below. The spectral line requirement (which is the more important one) is based on a RBW corresponding to a velocity resolution of 3km/s. The "3km/s RBW" is, conveniently, equal to $10-5f$, where f is the observing frequency. The discontinuity at 500 MHz is due to the fact that MeerKAT's frequency range will never go below 500 MHz, and the nearest other radio telescope (e.g. PAPER) is assumed to be 500 m away.

This graph limits the total amount of power that the Receiver and Digitiser are allowed to radiate within the specified RBW. In order to verify this requirement, a measurement system would typically measure the electric field strength at some distance r from the device. The graph below shows the threshold levels for the electric field strength, assuming a distance $r = 1\text{m}$. From the graph it can be seen that the SKA SA spectral line threshold levels (above 500 MHz) are about 80dB below the approximate noise floor of the Houwteq (ISSA) measurement equipment, assuming a RBW equal to the mil-std 461E measurement protocol (100kHz up to 1GHz, and 1MHz above 1GHz). They are also about 50dB below the

approximate noise floor of the FSH8 spectrum analyzer (above 500 MHz), assuming a RBW of 10kHz across the band.

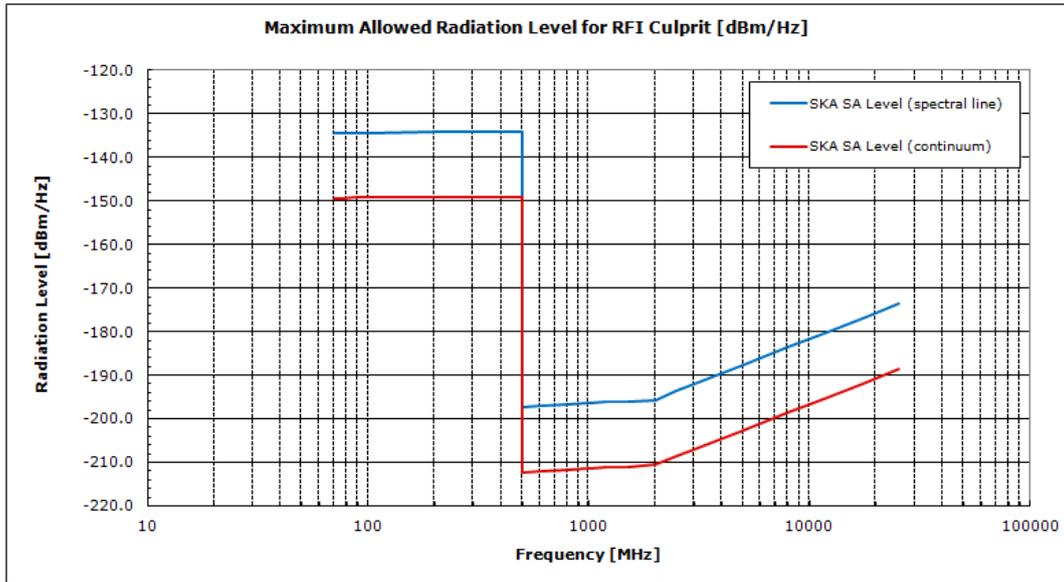


Figure G.1: Radiated RFI requirements for the Receiver and the Digitiser.

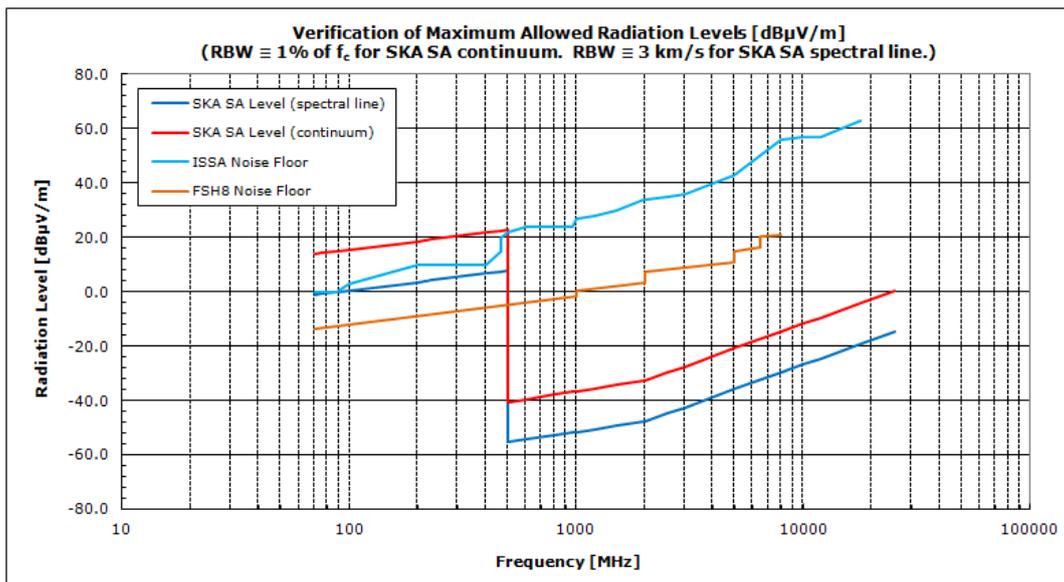


Figure G.2: Maximum electric field strength measured at a distance of 1m from the Receiver and Digitiser.

In summary, the RFI requirements are extremely demanding! This gives rise to the need for a measuring system that has exceptional sensitivity.

Advantages of a Time-Domain Measurement System

A time-domain measurement system, as opposed to a frequency-swept system, has the following advantages:

1. A large portion of the spectrum is seen at the same time. There is no need to sweep through the frequency range, which can take a very long time for narrow resolution bandwidths.
2. It is possible to integrate for a very long time (one hour, or even longer). This reduces the “noisiness” (or variance) of the measured result. It does NOT, however, lower the noise floor, which is given by $kT_{\text{sys}}B$, and which is independent of the integration time.
3. It is possible to sum neighbouring channels, so that the channel bandwidth, as a function of frequency, approximates the resolution bandwidth that corresponds to a velocity resolution of 3 km/s.
4. If the system has two input channels, which can be correlated, then the sensitivity of the measurement system could be increased significantly.
5. It is possible to accurately detect RFI spikes that have a low duty cycle.

Issue #2 warrants further clarification:

The SARAS levels are based on the premise that the detrimental amount of RFI is equal to 1/10th of the standard deviation of the correlator output samples, assuming an integration time of 2000 seconds. At first glance, it might seem possible that Ratty-2 could approximate this sensitivity by producing, say, 2000 magnitude spectra. Each magnitude spectrum would then consist of a 1-second integration period. This makes it possible to measure the standard deviation of the correlator output samples, since 2000 samples per channel are available for this calculation. This reasoning is correct, but one would only calculate the standard deviation of 1-second integration samples. In order to calculate the standard deviation of 2000-second integration samples, one would need to observe much, much longer than 2000 seconds, so that a sufficient population of 2000-second integration samples can be obtained.

Furthermore, if one assumes noise-like RFI with a bandwidth equal to or greater than a channel bandwidth, then the standard deviation of the RFI signal is also reduced when integrating the correlator output, making it difficult to detect the increase in standard deviation due to the RFI component.

In other words, it is not feasible to calculate the standard deviation of the measurement system's output, in order to measure RFI. Integration helps to reduce the noisiness of the measured samples, but the measurement system's ability to detect RFI depends on its noise floor, which is given by $kT_{\text{sys}}B$, and which is independent of the integration time. Putting it differently, if the RFI contribution to the calculated standard deviation is detectable, then the mean value of the RFI will also make it poke above the system's noise floor.

The measurement system's noise floor does not only depend on T_{sys} . It can also be reduced by correlating two independent input streams, whose noise contribution would decorrelate, but whose RFI contribution would correlate. In theory, the RFI-to-noise ratio would increase by n , where n is the number of independent samples being integrated, which would be equal to $n=2Bt_i$. (The factor of 2 is based on the minimum Nyquist sampling rate.) Here the integration period becomes important. The bandwidth is limited by the resolution bandwidth corresponding to a velocity resolution of 3 km/s.

Let's assume the following values:

- $B = 10 \text{ kHz}$ (This would be the required RBW at 1 GHz)
- $t_i = 10 \text{ s}$

The theoretical increase in sensitivity is equal to $12(3\text{dB} + 40\text{dB} + 10\text{dB}) = 26.5\text{dB}$. In other words, we would be able to detect RFI signals that are 26.5dB weaker than those that were detectable with the same system where no cross-correlation is deployed.

Detectable RFI Levels

Given all of the above, what levels of RFI can we expect to be able to measure with Ratty-2?

The current LNA envisioned for Ratty-2 has a noise factor of 1.2 dB. This converts to the following noise temperature:

$$\text{noise temperature} = 290 \times 1.2 = 348 \text{ K}$$

Therefore we can assume the following parameters for the Ratty-2 measurement system:

- $T_{\text{sys}} = 425 \text{ K}$ (300 K due to spillover from the anechoic chamber, 125 K from the Ratty-2 electronics)
- $G_r = 10 \text{ dB}$ (gain of receiving antenna)
- $r = 1 \text{ m}$ (distance from RFI Culprit to receiving antenna)
- $B = 10^{-5} \text{ f}$ (which corresponds to a velocity resolution of 3 km/s)
- $t_i = 10 \text{ s}$ (integration period)

We want to know the value of PRFI, which is the amount of RFI power, transmitted by the RFI Culprit within a bandwidth B, that can *just* be measured by the Ratty-2 measurement system.

The amount of power that is received by Ratty-2 is given by:

$$P_r = G_r \frac{c^2}{4\pi r^2} \text{PRFI}$$

The noise floor of Ratty-2 (within a channel bandwidth equal to B) is given by:

$$P_N = k T_{\text{sys}} B$$

Assuming that two independent RF channels are cross-correlated, the RFI-to-noise ratio will increase, and therefore the following power level can *just* be measured:

$$P_H = k T_{\text{sys}} B^2 t_i$$

We can now equate the received power, P_r , to the minimum measurable power level, P_H , in order to calculate the minimum measurable RFI power level PRFI:

$$\begin{aligned} P_r &= P_H \\ G_r \frac{c^2}{4\pi r^2} \text{PRFI} &= k T_{\text{sys}} B^2 t_i \\ \text{PRFI} &= k T_{\text{sys}} \frac{4\pi r^2}{c^2} \frac{B^2 t_i}{G_r} \end{aligned}$$

The graph below compares the levels, that are (should be) achievable with Ratty-2, with the Digitiser and Receiver requirements. Above 500 MHz, the Ratty-2 measurement system is still not sensitive enough by about 20 dB, if we assume two independent RF channels that are cross-correlated. If there is only a single RF channel, then the noise floor of the Ratty-2 system is about 50 dB above the required level (above 500 MHz). Furthermore, these numbers are fairly optimistic; it is quite possible that the Ratty-2 noise floor is worse than 425 K and that the gain of the receiving antenna is less than 10 dB.

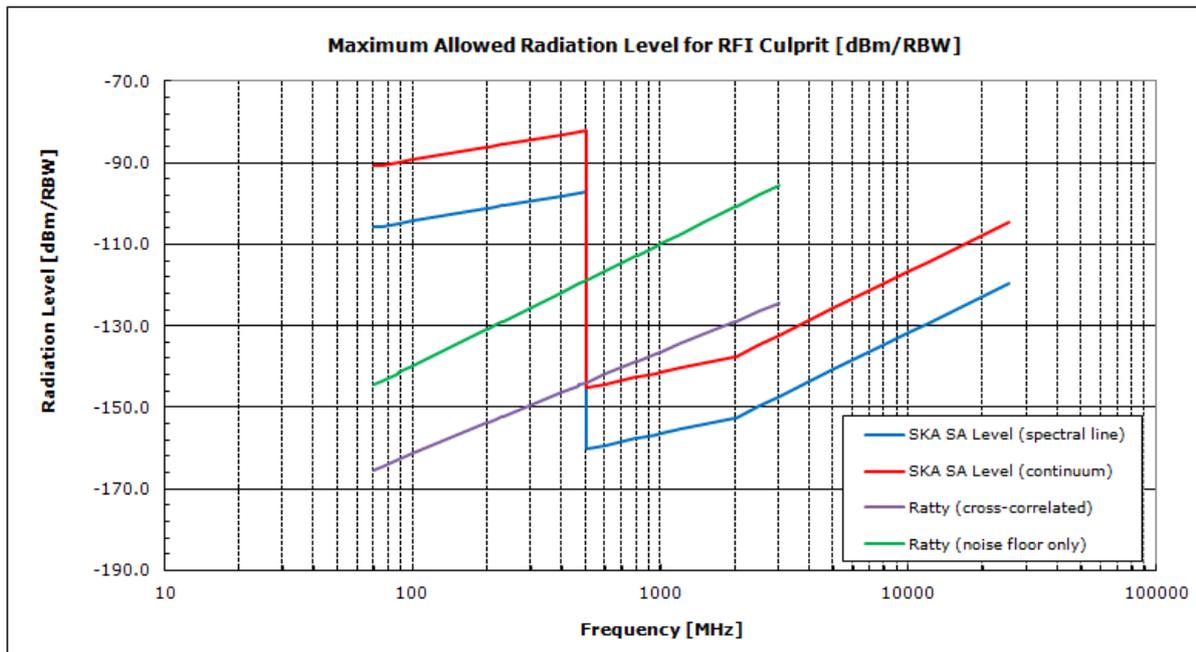


Figure 3: Theoretical sensitivity that Ratty-2 can achieve.

In summary, the radiated MeerKAT RFI requirements for the Digitiser and Receiver (and for the Antenna Positioner), cannot be verified directly with Ratty-2.

It is, however, conceivable, that the requirements are verified by measuring a device's radiation with the "lid open", and then subtracting the enclosure's shielding efficiency (measured separately) from these results.

Requirements for Ratty-2

The requirements that follow are very top-level, and have not been analysed in detail. The "system" refers to the Ratty-2 system, excluding the receiving antenna, i.e. excluding any RF front-end components.

Performance Requirements

R.001: The system shall utilise LNAs with a noise figure equal to 1.2 dB or less. The goal for the system's noise floor is 100 K or less.

R.002: The system shall have two independent RF channels, which can be cross-correlated to maximise the RFI-to-noise ratio.

R.003: The system shall be able to conduct measurements in the frequency range from 70 MHz to 3 GHz.

Functional Requirements

R.004: The system shall produce magnitude spectra of integrated data, with a specified integration time and channel bandwidth.

R.005: The integration time shall be configurable between a minimum value of 100 ms and a maximum value of 3600 s.

R.006: The channel bandwidth shall be equal to, or less than, the resolution bandwidth corresponding to a velocity resolution of 3 km/s. As a goal, the channel bandwidth shall be a function of the observing frequency, and shall be equal to the resolution bandwidth corresponding to a velocity resolution of 3 km/s.

R.007: The system shall be controlled and monitored using a laptop, which can be removed during the measurement period, so that no laptop-generated RFI enters the system.

R.008: The system shall store all measured data on a harddrive for further analysis.

Enclosure Requirements

R.009: The system shall be housed in 19" rack-mountable enclosures.

RFI Requirements

R.010: The system shall not add any self-generated RFI to the data stream, that can be measured using the system itself with the input terminated with a noise source.

Interface Requirements

R.011: The system shall have N-type connectors at the RF input, so that it can easily interface to the Houwteq RF front-end.

G.2 SKA_SA_RATTY2_design_doc

RF line-up requirements

1. Frequency range: 50-1670 MHz consisting of 3 bands
 - a. 0-828MHz
 - b. 750-1200 MHz (TBC)
 - c. 900-1670 MHz
2. Input power range: -24 to -118dBm
3. Output power range:
 - . Limited by the ADC
 - a. Max input power to the ADC = 0dBm
 - b. Min input power to the ADC = -60dBm
4. Gain
 - . Provide sufficient adjustable gain that will ensure that the input power to the ADC = -1dBm (goal) -7dBm (required) for the entire input power range i.e. 17dB @ -24dBm and

111dB @ -118dBm. This results in a variable gain requirement of 94dB i.e. 3x standard 6 bit mini-circuit attenuators! Maybe we should consider adding an external attenuator that can be added/removed by the operator. Another option would be to switch in a pre-amp for high sensitivity measurements (phase 2?)

a. Optimal input power to the ADC is chosen for best SFDR based on results obtained in the following

doc: <https://docs.google.com/a/ska.ac.za/document/d/1xECWuT1xU9akLqAhn2Z8jwv6kT7KKGlXXYUcWApCkJ4/edit>

5. Filters

. -60dB rejection @ fs & fs/2 to ensure that the SFDR of the ADC dominates the spurious response of the system.

a. add LPF on Tellumat filter

6. Switches

. >3 ports

a. Look at Hittite switches

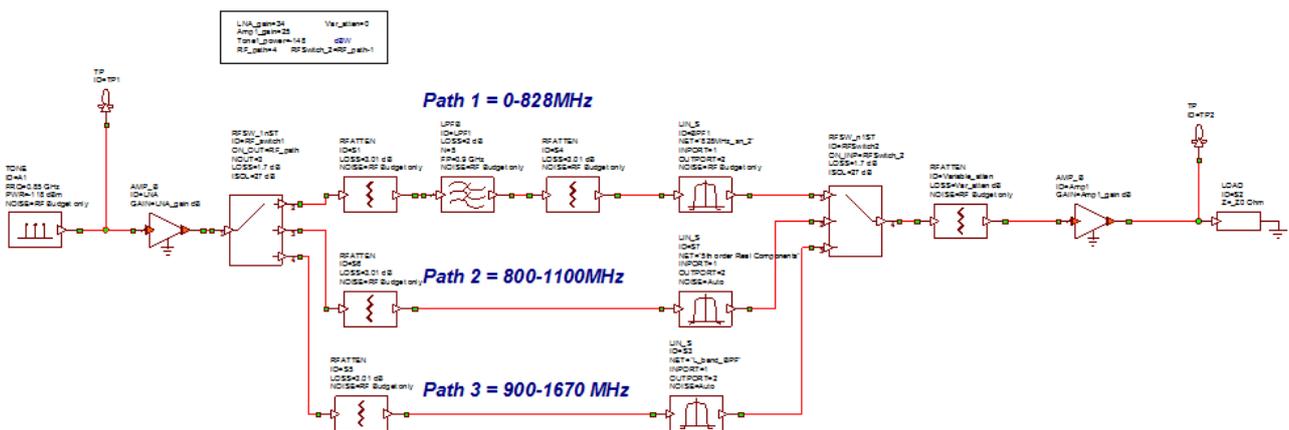
b.

DSP requirements

1. Phase 1 # of channels = 2[^]
2. Channelisation = $f_c \cdot 1e-5 = 1\text{kHz}$ for $f_c=100\text{MHz}$ and 20kHz for $f_c = 2\text{GHz}$.
3. # Channels = $900\text{MHz}/1\text{kHz} = 900\text{k}$ (approx 2[^]20) for nyquist zone 1 and 45k (approx 2[^]16) for nyquist zone 2

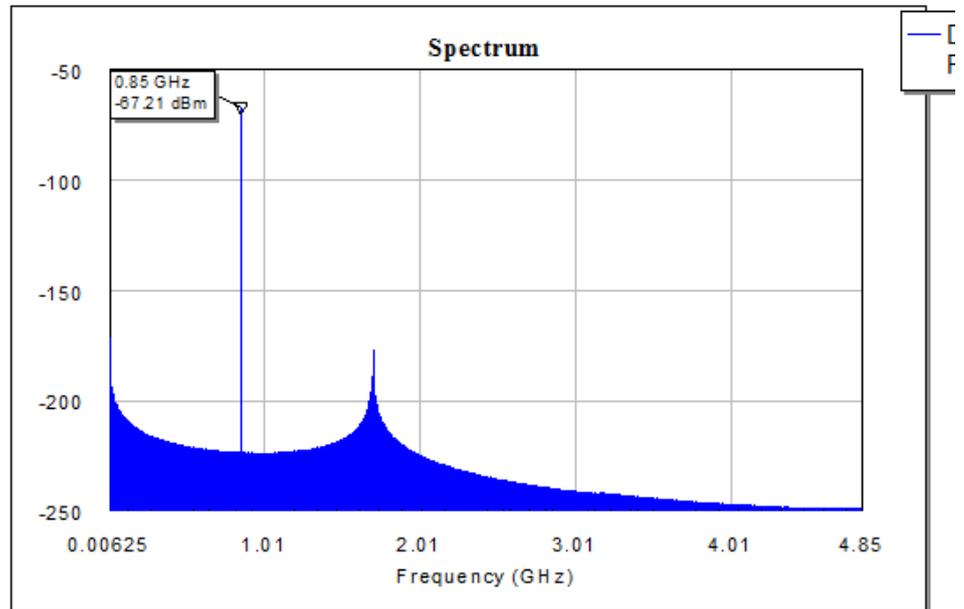
RF line-up

The proposed line-up for 3 channels is given below. The 3rd channel filter design has not yet been adopted for the band indicated above. The total gain needs to be increased to meet the requirements given above.



Preliminary results

1. Tsys @ 850MHz = 136K

Spectrum**Parts**

<u>Common</u>	<u>US</u>	<u>SKA</u>	<u>Ordered</u>	<u>Additional parts</u>
LNA	1	1	4	
4 way switch	1	2	2	4
3dB attenuator	16	7	35	
Variable atten	1	2	4	2
Tellumat filter	1	1	3	1
2 way switch	2	0	4	
920MHz filter	2	1	4	2
1GHz amp	1	0	2	2

2GHz amp 25dB	0	2	2	
1 GHz amp 2	1	0	2	

Parts List

Green -- arrived.

Part No.	Description	Specification	TTL Ctr	Supply (V)	Due
1	LNA: MITEQ (AFS3-00100300-12-10P-4)	100 - 3000 MHz Gain: 37 dB Gain Flatness: ± 1 dB P1dB: 10 dBm, min NF: 1.2 dB, max VSWR (in/out): 2:1, max	-	15 V (125 mA)	28/11/2012
2	RF Switch 1: Peregrine Semiconductor PE42556 Flip Chip SPDT UltraCMOS RF Switch (Evaluation Package)	0.009 - 13500 MHz Isolation: 26 dB Return Loss: 13 dB Insertion Loss: 1.7 dB P1dB: 30dBm, 1MHz-13GHz	1 (4V max)	3.3 V (21.5 uA) (3 - 3.6 V)	1/12/2012
3	3 dB Pads: Mini-Circuits (VAT-3+)	0 - 6000 MHz	!	!	-
4	Mixer: Marki Microwave (T3A3-08CKP) Triple Balanced Mixer	LO/RF: 10 - 8000 MHz IF: 1 - 5000 MHz Conversion Loss: 10.5 dB, max Leakage: LO-RF: -8.5 dBm, max LO-IF: -13 dBm, max Isolation: RF/IF: 30dB Input P1dB: 17 dBm	!	5 V (145 mA)	!
5	LO: Valon Technology (5008 Dual Synthesizer)	137.5 - 4400 MHz P(out): 0/3/6/8 dBm typ	(USB)	5 V (340 mA)	-
6	LO LPF: Mini-Circuits (VLF-2850)	3 dB Point: 3300 MHz	!	!	-

7	LPF2: Mini-Circuits (VLFX-780)	3 dB Point: 920 MHz	█	█	-
8	AMP1: WENTEQ (ABL0100-01-3010w)	0.01 – 1 GHz Gain: 30 dB, min Gain Flatness: 1 dB P1dB: 16 dBm, min NF: 1.2 dB, max	█	12 V (110 mA)	-
9	Step Attenuator: Mini-Circuits (ZX76-31-SP+) 5 bit, 31 dB, Serial Controlled, Single Supply	0 - 2400 MHz Insertion Loss: 2.8 dB, max P(in): 24 dBm, max NO DC!!	3 (5 V max)	3 V (1.5 mA) (3.3 V max)	-
10	AMP2: WENTEQ (ABL0100-10-1810)	0.1 – 1 GHz Gain: 16 dB, min Gain Flatness: 1 dB P1dB: 10 dBm, min NF: 1 dB, max	█	12 V (60 mA)	-
11	AMP3: WENTEQ (ABL0200-01-2718)	0.1 – 2 GHz Gain: 25 dB, min Gain Flatness: 1 dB P1dB: 20 dBm, min NF: 2 dB, max	█	12 V (160 mA)	-
12	LPF1: Tellumat filter (SKA)	3 dB Point: 828 MHz	█	█	-
13	BPF1: Prof. P Meyer & ARB	3 dB Points: 828 - 1420 MHz	█	█	-
14	BPF2: Prof. P Meyer & ARB	3 dB Points: 1420 - 2170 MHz	-	-	-
15	ADC: SKA	Input range: 0 – (-60) dBm : 1.8 Gsa/s	-	-	?
16	RF Switch 2: Mini-Circuits (ZSWA-4-30DR)	0 - 3000 MHz Isolation: 26 dB, min Insertion Loss: 2.3 dB, max P1dB: 23dBm, 10MHz-3GHz	4 (3.5 V min High)	+5V (4 mA max) -5V (20 mA max)	25/10/2012
17	HPF 1: Mini-Circuits (SHP-700+)	700 – 3000 MHz	█	█	█
18	HPF 2: Mini-Circuits	1220 – 4600 MHz	█	█	█

	(VHF-1200+)				
	SKA amp: WENTEQ (ABL0300-01-3414)	0.01 – 3GHz Gain: 34 dB, min Gain Flatness: 0.75 dB P1dB: 11dBm, min NF: 1.4 dB, max	█	█	█

Part Ordering List

Part No.	Part Name	Manufacturer	Who?	No.	Quote (VAT incl.)	Total
1	Front End LNA	Miteq (AFS3-00100300-12-10P-4)	US	4	R15766	R 63064
3	3dB Pads	Mini-Circuits (VAT-3+)	US	35 + 15	R142	R 4970 + R2130
7	LPF2	Mini-Circuits (VLFX-780) 920 MHz	US	4 + 2	R 570	R 2280 + R1140
8	AMP1	WENTEQ (ABL0100-01-3010) 0.01 – 1 GHz	US	2	R 1499	R 2998
9	Step Atten	Mini-Circuits (ZX76-31-SP+)	US	4 + 4	R1054	R 4216 + R4216
11	AMP3	WENTEQ (ABL0200-01-2718) 0.1 – 2 GHz	US	2	R 1841	R 3682
16	RF Switch 2	Mini-Circuits (ZSWA-4-30DR)	US	2 + 4	R1620	R3240 + R6480
-	SKA amp	WENTEQ (ABL0300-01-3414)	US	2	R2200	R4400

Potential future upgrades:

Extended, frequency coverage to ~2GHz:

- 50MHz-828MHz (1.8Gsps)
- 750-1050MHz (1.3Gsps)
- 900-1670 MHz (1.712Gsps)
- 1.5-2.1GHz (1.45Gsps)

Or, extended, overlapping frequency coverage to beyond 3GHz (requires a 7-way switch, or switch chain + programmable VALON):

- 50MHz-828MHz (1.8Gsps)
- 750-1050MHz (1.3Gsps)
- 900-1670 MHz (1.712Gsps)
- 1.55-2.1GHz (1.45Gsps)
- 1.9-2.6GHz (1.8Gsps)
- 2.5GHz - 2.9GHz (1.2Gsps)
- 2.8GHz-3.5GHz (1.8Gsps)

High speed readout

- via SPEAD protocol
- separate 1GbE port.
- Requires small 3-port 1GbE switch: first for control and monitoring (PPC), second for data traffic (FPGA), third for laptop
- Optional spectrum readout down to less than 1ms (as opposed to 1s dumps now) (recording laptop/computer limited)
- Time-domain readout at of 500ms data in 40s (as opposed to 10 minutes now)
- Leverage JP's EMSS work.