

**Ultra-Low Phase Noise 100 MHz
Crystal Oscillator & Residual
Phase Noise Measurement
Systems**

Luke Richard Dummott

Doctor of Philosophy

University of York

School of Physics, Engineering and Technology

June 2024

Abstract

In this thesis, an ultra-low phase noise 100 MHz crystal oscillator is presented that demonstrates a phase noise of around -140 dBc/Hz at a 100 Hz offset. The architecture and each of its constituent components are described in detail. The oscillator performance was measured for three separate amplifier designs, including a differential amplifier in a quad parallel configuration. A method for measuring the flicker noise corner of the amplifiers using an LC resonator is also described and it is shown how the placement of an attenuator between the resonator and the input of the amplifier in an oscillator loop can be used as a simple way to reduce phase noise degradation in certain circumstances.

A set of digitally-controlled phase shifter and attenuator devices are presented that are designed to be used in the 100 MHz crystal oscillator and other applications. These are passive devices that utilise mechanical RF relays to switch in microstrip delay line elements in the case of the phase shifter, and resistive attenuator networks in the case of the attenuator, to minimise flicker noise. The complete design procedures for producing these devices are described in detail.

Finally, a low-cost residual phase noise measurement system is described that uses cross-correlation techniques to reduce the internal system noise floor. The system is designed primarily for characterising the residual phase noise performance of low power amplifiers operating at 100 MHz and exhibits a noise floor of -177 dBc/Hz when measuring a device operating at 0 dBm. A set of efficient digital signal processing algorithms is also presented, which includes the ability to vary the resolution bandwidth for different frequency bands in a single measurement.

Acknowledgements

Foremost, I would like to thank my PhD supervisors Prof. Jeremy Everard and Dr. Simon Bale for their support throughout and beyond this research. I realise that I am incredibly fortunate to have supervisors that are so willing to share their experience and expertise, and without their encouragement and technical assistance, this work simply would not have been possible.

Thanks also to my friends and colleagues St.John Gilbert and Stuart Kenny for the assistance they have provided in various aspects of the research. St.John especially has been a huge help, not only with specific technical aspects, but has put up with an almost constant stream of my (often inane) questions over the course of our PhDs. Although there are too many to name, I am also thankful to each of the technical, academic and administrative staff of the School of Physics, Engineering and Technology at the University of York who have been so helpful.

To my family, my sisters, nieces and especially my parents, Simon and Clara, for their love and support I will always be grateful.

Declarations

I declare that this thesis submitted for the degree of Doctor of Philosophy is an original work representing my own research and I am the sole author. All sources are represented appropriately in the text.

Outcomes from this research have been used to produce a pair of conference papers presented at the European Frequency & Time Frequency 2024. These papers are included in Appendices A and B.

Thesis Structure

This thesis is structured into three chapters, the first of which details the research carried out in producing an ultra-low phase noise 100 MHz crystal oscillator. The second chapter describes the design of low residual phase noise digitally-controlled phase shifter and attenuator devices, for use in 100 MHz oscillators and other applications. The final chapter outlines the design of a low-cost cross-correlation residual phase noise measurement system, designed primarily for measuring amplifiers used in the 100 MHz crystal oscillator, which also includes detailed descriptions of efficient digital signal processing techniques.

Contents

1	Ultra-Low Phase Noise 100 MHz Crystal Oscillator	19
1.1	Introduction	20
1.1.1	Existing Oscillators	21
1.2	Phase Noise Theory	22
1.2.1	Introduction to Phase Noise	22
1.2.2	Residual Phase Noise	23
1.2.3	Modelling Oscillator Phase Noise	25
1.3	Design	27
1.3.1	Block Diagram	27
1.3.2	Crystal Resonator	28
1.3.3	Spurious Resonance Rejection Filter	31
1.3.4	Electronic Phase Shifter	32
1.3.5	Digitally-Controlled Phase Shifter	36
1.3.6	Digitally-Controlled Attenuator	36
1.3.7	Amplifier	37
1.4	MAR-6 Amplifier	38
1.4.1	Design	39
1.4.2	Measurements	40
1.4.3	Crystal Oscillator Measurements	43
1.5	Differential Amplifier	45
1.5.1	Design	46
1.5.2	Model	48

1.5.3	Measurements	51
1.6	Flicker Noise Corner Measurement using an LC Resonator	55
1.6.1	Explanation	56
1.6.2	MAR-6 Measurement	59
1.6.3	Differential Amplifier Measurement	60
1.7	Improving Oscillator Phase Noise Using Attenuator Matching	63
1.7.1	Theory	63
1.7.2	Improved LC Oscillator Measurements	68
1.7.3	Crystal Oscillator Measurements	72
1.8	Residual Noise of Crystal Resonator	74
1.8.1	Comparison of Crystal Oscillator Measurements	75
1.8.2	Replacement Crystal	75
1.9	Quad Parallel Differential Amplifier	77
1.9.1	Theory of Operation	78
1.9.2	Design	81
1.9.3	Measurements	86
1.10	Conclusions and Further Work	91
1.10.1	Ultra-Low Phase Noise 100 MHz Crystal Oscillator Measure- ment	91
1.10.2	Further Work	92
2	Digitally-Controlled Phase Shifters and Attenuators	95
2.1	Introduction	96
2.2	Digitally-Controlled Phase Shifter	97
2.2.1	Theory of Operation	97
2.2.2	100 MHz 180° 6-Bit Phase Shifter	101
2.2.3	1.5 GHz 360° 6-Bit Phase Shifter	111
2.3	Digitally-Controlled Attenuator	117
2.3.1	Theory of Operation	117
2.3.2	100 MHz 6.3 dB 6-Bit Attenuator	120

2.3.3	100 MHz 31.5 dB 6-Bit Attenuator	125
2.4	Conclusions and Further Work	129
2.4.1	Further Work	130
3	Cross-Correlation Residual Phase Noise Measurement System	132
3.1	Introduction	134
3.2	Residual Phase Noise Measurement Theory	134
3.2.1	Single-Channel Residual Phase Noise Measurement	134
3.2.2	Cross-Correlation Residual Phase Noise Measurement	135
3.3	Single-Channel Residual Noise Measurement System	139
3.3.1	System Overview	139
3.3.2	Calibration Techniques	143
3.3.3	System Noise Floor	146
3.3.4	Limitations of Single-Channel Residual Noise Measurements	147
3.4	Cross-Correlation Residual Noise Measurement System	147
3.4.1	System Overview	147
3.4.2	Measurement Accuracy and Noise Floor	148
3.5	Digital Signal Processing Techniques for Residual Noise Measurements	156
3.5.1	Variable Window Size Frequency Banding Algorithm	158
3.5.2	Fast Fourier Transform vs. Discrete Fourier Transform	165
3.5.3	Optimising the FFT for Frequency Banded Data	168
3.5.4	Twiddle Factors	173
3.5.5	Simultaneous Fourier Transforms of Real Data	175
3.5.6	Cross-Correlation and Post-Processing	178
3.5.7	Complete System Flowchart	180
3.5.8	Results	183
3.6	Conclusions and Further Work	190
3.6.1	Further Work	190
4	Overall Conclusions	194

References	196
A EFTF 2024 Conference Paper: Ultra-Low Phase Noise 100 MHz Crystal Oscillator	200
B EFTF 2024 Conference Paper: A Low-Cost Residual Phase Noise Measurement System & Efficient Digital Signal Processing Techniques	205

List of Figures

1.1	Photo of the ultra-low phase noise 100 MHz crystal oscillator	19
1.2	Typical plot of the residual phase noise of an amplifier	24
1.3	Plot showing how the absolute phase noise of an oscillator relates to its open-loop residual phase noise and the bandwidth of the resonator	25
1.4	Block diagram of 100 MHz crystal oscillator	27
1.5	S21 frequency response of the KVG 100 MHz quartz crystal resonator measured using the SDR-Kits DG8SAQ VNWA	29
1.6	S21 broadband frequency response of the KVG 100 MHz quartz crystal resonator measured using the SDR-Kits DG8SAQ VNWA	29
1.7	Model of the KVG 100 MHz quartz crystal resonator measured using crystal analyser function of the SDR-Kits DG8SAQ VNWA	30
1.8	Circuit diagram of the 40 MHz bandwidth spurious resonance rejection filter	31
1.9	S21 frequency response of the spurious resonance rejection filter measured using the SDR-Kits DG8SAQ VNWA	32
1.10	100 MHz 25° electronic phase shifter	33
1.11	Circuit diagram of the electronic phase shifter	34
1.12	Plot of the absolute phase shift and insertion loss of the electronic phase shifter versus the DC bias voltage measured using the SDR-Kits DG8SAQ VNWA	35
1.13	Photo of MAR-6 amplifier	38
1.14	Circuit diagram of the MAR6 amplifier	39

1.15	Block diagram of the initial 100 Mhz oscillator setup that was used to measure the flicker noise of the MAR-6 amplifier	40
1.16	Gain compression of the MAR-6 amplifier measured using the Rohde & Schwarz FSWP	41
1.17	Residual phase noise of the MAR-6 amplifier measured using the cross-correlation system with 256 correlations at 0.125 Hz and 8,388,608 correlations at 100 kHz	43
1.18	Absolute phase noise measurement of the 100 MHz crystal oscillator with MAR-6 amplifier and 3 dB of attenuation after the amplifier made using the Rohde & Schwarz FSWP	44
1.19	Photo of differential amplifier circuit	45
1.20	Circuit diagram of the differential amplifier	46
1.21	PCB layout of the differential amplifier	48
1.22	Simplified model of a single end of the long-tailed pair differential amplifier	49
1.23	Gain compression of the differential amplifier with a bias voltage of -5 V measured using the Rohde & Schwarz FSWP	52
1.24	S-parameter measurements of the differential amplifier showing the input and output impedance at 100 MHz made using the Anritsu 37377C Vector Network Aanalyzer	53
1.25	Residual phase noise of the differential amplifier measured using the cross-correlation system with 256 correlations at 0.125 Hz and 8,388,608 correlations at 100 kHz	55
1.26	Plots showing how the absolute phase noise measurement can be used to estimate the residual noise of an amplifier	57
1.27	Schematic diagram of the LC resonator, including the equivalent series resistance of the inductor, that was used to replace the 100 MHz crystal for measuring the flicker noise corner of amplifiers	59

1.28	Absolute phase noise measurement of the 100 MHz LC oscillator with MAR-6 amplifier and 2.8 dB of attenuation after the amplifier made using the Rohde & Schwarz FSWP	60
1.29	Block diagram of the initial 100 MHz oscillator setup that was used to measure the flicker noise of the differential amplifier	61
1.30	Absolute phase noise measurement of the 100 MHz LC oscillator with differential amplifier and 0.2 dB of attenuation after the amplifier made using the Rohde & Schwarz FSWP	62
1.31	Simplified block diagram of the oscillator to demonstrate the effects of pre-amplifier and post-amplifier attenuation	64
1.32	Equivalent circuit of a resistive attenuator using pi-pad topology placed between source and load impedances	65
1.33	Plot of the the effective input impedance of a resistive attenuator versus the nominal attenuation when the output of the attenuator is terminated in either an open or short-circuit	66
1.34	Absolute phase noise measurement of the 100 MHz LC oscillator with MAR-6 amplifier and 2.5 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP	69
1.35	Absolute phase noise measurement of the 100 MHz LC oscillator with differential amplifier and 2.4 dB of post-amplifier attenuation made using the Rohde & Schwarz FSWP	70
1.36	Absolute phase noise measurement of the 100 MHz LC oscillator with differential amplifier and 1.8 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP	71
1.37	Absolute phase noise measurement of the 100 MHz crystal oscillator with MAR-6 amplifier and 5.4 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP	73

1.38	Absolute phase noise measurement of the 100 MHz crystal oscillator with differential amplifier and 0.8 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP	74
1.39	Absolute phase noise measurement of the 100 MHz KVG crystal oscillator with differential amplifier and 3.1 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP	76
1.40	Absolute phase noise measurement of the 100 MHz KVG crystal oscillator with MAR-6 amplifier and 5.4 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP	77
1.41	Photo of quad differential amplifier circuit	78
1.42	Block diagram of the quad parallel differential amplifier	79
1.43	Circuit diagram of the quad differential amplifier	83
1.44	PCB layout of the quad differential amplifier	85
1.45	Gain compression of the quad differential amplifier with a bias voltage of 7.8 V measured using the Rohde & Schwarz FSWP	86
1.46	S-parameter measurements of the quad differential amplifier showing the input and output impedance at 100 MHz made using the Anritsu 37377C Vector Network Analyzer	87
1.47	Residual phase noise measurement of the quad differential measured using the cross-correlation system with 256 correlations at 0.125 Hz and 8,388,608 correlations at 100 kHz	89
1.48	Absolute phase noise measurement of the LC oscillator with quad differential amplifier made using the Rohde & Schwarz FSWP to show the flicker noise corner	90
1.49	Absolute phase noise measurement (with smoothing function applied) of the 100 MHz crystal oscillator using the configuration described in Section 1.3 made using the Rohde & Schwarz FSWP	92
2.1	Digital phase shifter block diagram	97

2.2	Single relay stage of the digital phase shifter circuit corresponding to 1 bit	99
2.3	Photo of the 100 MHz 180° 6-bit phase Shifter	101
2.4	Generic schematic diagram for the digital phase shifter design	102
2.5	Microstrip layout of 100MHz digital phase shifter simulated in Momentum EM field solver	105
2.6	PCB layout for the 100 MHz 6-bit digital phase shifter	106
2.7	Plot of the phase shift of the 100 MHz digital phase shifter vs. the digital control code measured using the SDR-Kits DG8SAQ VNWA .	108
2.8	Plot of the insertion loss of the 100 MHz digital phase shifter vs. the digital control code measured using the SDR-Kits DG8SAQ VNWA .	109
2.9	Plot of the residual phase noise of the 100 MHz 180° digitally-controlled phase shifter when each of the relays is set to the '0' position (minimum phase shift) measured using the cross-correlation system	110
2.10	Photo of 1.5 GHz 180° 6-bit phase Shifter	111
2.11	PCB layout for the 1.5 GHz 6-bit digital phase shifter	113
2.12	Plot of the phase shift of the 1.5 GHz digital phase shifter vs. the digital control code measured using the Anritsu 37377C Vector Network Analyzer	114
2.13	Plot of the insertion loss of the 1.5 GHz digital phase shifter vs. the digital control code measured using the Anritsu 37377C Vector Network Analyzer	115
2.14	Plot of the residual phase noise of the 1.5 GHz 360° digitally-controlled phase shifter when each of the relays is set to the '0' position (minimum phase shift) measured using the HP 11848A system	116
2.15	Digital attenuator block diagram	117
2.16	Pi-pad resistive attenuation network	119
2.17	Single relay stage of digital attenuator circuit corresponding to 1 bit .	120
2.18	Photo of 100 MHz 6-bit digitally-controlled attenuator	121

2.19	Generic schematic diagram for the digital attenuator design	122
2.20	PCB layout for the 6.3 dB 6-Bit digital attenuator	123
2.21	Plot of the attenuation of the 6.3 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer	124
2.22	Plot of the phase shift of the 6.3 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer	125
2.23	Plot of the attenuation of the 31.5 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer	128
2.24	Plot of the phase shift of the 31.5 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer	129
3.1	Photo of the 100 MHz residual phase noise measurement system . . .	132
3.2	Block diagram of a single-channel residual phase noise measurement .	135
3.3	Block diagram of a cross-correlation residual phase noise measurement	136
3.4	100 MHz single-channel residual phase noise measurement system . .	139
3.5	Schematic diagram of the baseband amplifier used in the residual phase noise measurement system	141
3.6	Schematic diagram of the 5th-order Butterworth anti-aliasing filter used in the residual phase noise measurement system	142
3.7	Frequency response of the combined LT6018 amplifier and LC anti-aliasing filter measured on the SDR-Kits DG8SAQ Vector Network Analyzer	143
3.8	Plot of the calibration spur measured on the PicoScope when a -30 dBm calibration tone at various frequencies is injected into the signal path of the single-channel residual noise system	145

3.9	Internal noise floor of the single-channel residual phase noise measurement system	146
3.10	100 MHz cross-correlation residual phase noise measurement system .	148
3.11	System front-end setup used for the comparison of single channel and cross-correlation measurement techniques	149
3.12	Comparison of system front end noise floor (amplifier, filter, PicoScope) using single channel and cross-correlation measurement techniques	150
3.13	Comparison of system noise floor using single channel and cross-correlation measurement techniques	151
3.14	Comparison of the cross-correlation system noise floor for various DUT output power levels	152
3.15	Comparison of the residual noise of the MAR-6 amplifier when measured using the system described in this chapter versus the Rohde & Schwarz FSWP Phase Noise Analyzer	154
3.16	Demonstration showing a comparison of variable and fixed sampling window sizes for a theoretical 10 sec cross-correlation measurement at 200 kHz	161
3.17	Demonstration showing a comparison of different frequency banding styles for a theoretical 10 sec cross-correlation measurement at 200 kHz	162
3.18	Diagram showing a radix-2 Cooley-Tukey FFT for N=8	167
3.19	Diagram showing the calculation of partial FFTs for N=8	170
3.20	Complete flowchart of signal processing for a cross-correlation residual phase noise measurement	182
3.21	Plots showing the average processing time versus the size of the required output data, M, when using different methods to calculate a single Fourier transform with an input data size of 32,768 samples . .	184

3.22	Comparison of the average processing time to calculate the Fourier transforms of two sets of real data plus the cross correlation of the results for different input data sizes when using either a standard approach or the simultaneous method	185
3.23	Comparison of the processing time for a single data block of 16,777,216 captured samples versus the number of equal-ratio frequency bands that the measurement is split into when using either an FFT or PFFT algorithm	187
3.24	Photo of the low-cost Raspberry Pi Pico two-channel ADC capture device	193

List of Tables

2.1	100 MHz digital phase shifter design parameters	104
2.2	100 MHz digital phase shifter EM simulation results	105
2.3	1.5 GHz digital phase shifter design parameters	112
2.4	1.5 GHz digital phase shifter EM simulation results	113
2.5	Resistor values for 6.3 dB digital attenuator	123
2.6	Resistor values for 31.5 dB digital attenuator	127
3.1	A comparison of the the total measurement time for a variety of configurations versus the minimum theoretical time possible based on the desired minimum resolution bandwidth	189

Chapter 1

Ultra-Low Phase Noise 100 MHz Crystal Oscillator

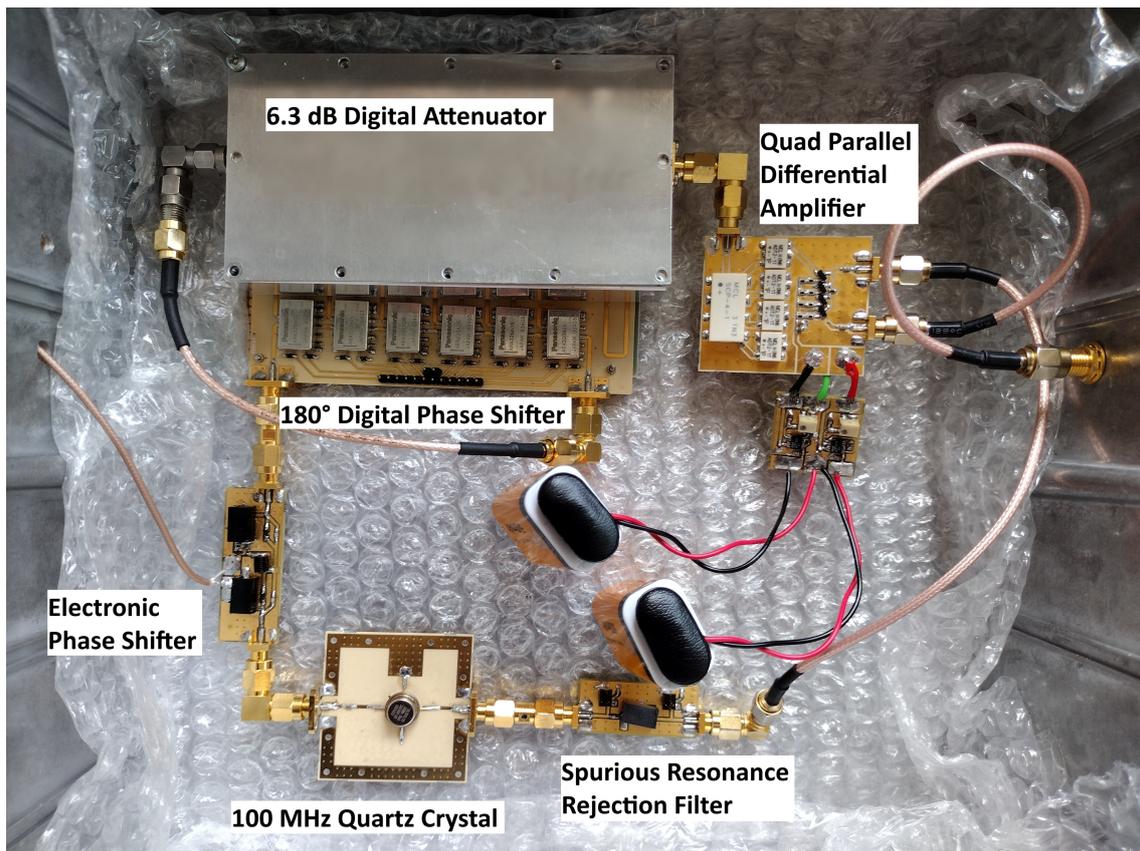


Figure 1.1: Photo of the ultra-low phase noise 100 MHz crystal oscillator

Abstract

In this chapter an ultra-low phase noise 100 MHz crystal oscillator that exhibits an absolute phase noise of around -140 dBc/Hz at a 100 Hz offset and a noise floor of around -168 dBc/Hz is presented. This matches the state-of-the-art close-to-carrier phase noise performance of existing oscillators, but here the construction of the oscillator and its constituent parts are described in detail.

Three separate low-noise amplifier designs that were considered for use in the oscillator are presented, one of which is a quad parallel differential amplifier design that is shown to improve the maximum available output power over a single differential amplifier without significantly affecting the gain or making the amplifier more sensitive to changes in the device parameters of the transistors.

A method for measuring the flicker noise of amplifiers where it cannot be measured directly due to the internal noise floor of the residual phase noise measurement system is described. This involves measuring the absolute phase noise of an oscillator with a relatively low Q LC resonator that incorporates the amplifier. It is shown that this method can be used to estimate the residual noise of various amplifiers.

Finally, it is shown that oscillator phase noise can be improved through the use of an attenuator in situations where the amplifier has an input impedance that is not matched to the rest of the system, or when the flicker noise of the amplifier is sensitive to the source impedance presented to it. In the examples given in this chapter, it is shown that the phase noise of various oscillators could be significantly improved by simply moving the attenuator included in the oscillator feedback loop from the output of the amplifier to the input to improve the matching between the input of the amplifier and the crystal resonator.

1.1 Introduction

Minimising phase noise and jitter is an important factor in designing oscillators used to provide a very stable reference frequency required for many modern systems. For

example, ultra-low phase noise oscillators operating at 100 MHz have applications in atomic clocks, RADAR and particle accelerator systems.

The aim of this research is to produce a 100 MHz crystal oscillator that exhibits less than -140 dBc/Hz single-sideband phase noise at an offset of 100 Hz from the carrier signal. It builds on previous done at York on crystal oscillators at 10 MHz [1].

The first section of this chapter describes the overall design of the 100 MHz crystal oscillator, including detailed descriptions of its constituent parts. The amplifier designs that were considered during the research are each outlined in their own section and include the commercially-available MAR-6 amplifier [12] in Section 1.4; a differential amplifier based on that designed for the 10 MHz oscillator [1] in Section 1.5; and a quad parallel configuration of the differential amplifier as used in the final oscillator design in Section 1.9. Additionally, a technique for measuring the flicker noise corners of amplifiers where they cannot easily be measured directly using a residual phase noise measurement system, such as that described in Chapter 3, is explained in Section 1.6; and a method for improving absolute phase noise by using attenuation in the oscillator feedback path is described in Section 1.7. Finally, the absolute phase noise measurements that were achieved from the final design of the oscillator are presented in Section 1.10 along with potential improvements and further work.

1.1.1 Existing Oscillators

An ultra-low phase noise 10 MHz crystal oscillator has been built previously at York [1] [2], which exhibits a phase noise performance of -148 dBc/Hz at a 10 Hz offset and a noise floor of around -160 dBc/Hz. The work in this chapter builds on the techniques used for this oscillator to achieve comparable performance from a crystal oscillator operating at 100 MHz.

Much analysis of the phase noise performance of crystal oscillators exists in the literature [2] [3] [4] and is applicable at a range of frequencies, including the 100 MHz that is the focus of this research.

There exists a number of commercially-available 100 MHz crystal oscillators that exhibit phase noise around the -140 dBc/Hz target of this research. One such example is the KVG O-40-ULPN-100M oven-controlled crystal oscillator [5] with a claimed maximum phase noise of -138 dBc/Hz at a 100 Hz and a noise floor of -185 dBc/Hz. This performance has also been confirmed by this group and this oscillator is used as the reference source for the cross-correlation residual phase noise measurement system described in Chapter 3. Descriptions of ultra-low phase noise 100 MHz oscillations with a similar performance exist in the literature [6], but they do not typically provide a complete and detailed design of the oscillator.

The research presented in this chapter differs in that it seeks to provide both a detailed description of the construction and the design procedures involved in producing 100 MHz crystal oscillators with a noise performance at, or ideally below, -140 dBc/Hz at a 100 Hz offset.

1.2 Phase Noise Theory

In this section, a background to the concepts of oscillator phase noise and how they relate to the design of an ultra-low phase noise 100 MHz crystal oscillator is provided. It is explained how the phase noise performance of an oscillator and each of its constituent parts are characterised. A model for calculating the theoretical phase noise of an oscillator is given, which is used to compare against the phase noise measurements made in this chapter.

1.2.1 Introduction to Phase Noise

Definition of Phase Noise

Phase noise is the random fluctuation in phase of a signal. In the frequency spectrum of the signal, this results in sidebands around the oscillating frequency and the phase noise is represented by the power spectral density of these sidebands. Phase noise is used as a way to describe the frequency stability of an oscillator. The lower

the power spectral density of the phase noise sidebands, the higher the short-term frequency stability of the oscillator.

Characterising Oscillator Phase Noise

The absolute phase noise of an oscillator is characterised by the power spectral density at an offset from the carrier frequency, referenced to the power of the carrier signal, and is measured in dBc/Hz. Typically, as phase noise is usually symmetrical about the carrier frequency (provided there is little or no AM noise), it is measured as the single-sided noise spectral density, L_f . The performance of oscillators are often compared using the single-sided noise spectral density at a particular offset frequency from the carrier.

1.2.2 Residual Phase Noise

Residual phase differs from the absolute phase noise of an oscillator in that it describes the additive phase noise of a two-port device, i.e. the phase noise that is added to a carrier signal due to being passed through the device, and is similarly characterised by the single-sided noise spectral density. Residual phase noise is a way to measure the phase noise contribution that a component placed in an oscillator loop, such as an amplifier, will have on the overall absolute phase noise of the oscillator.

Flicker Noise

Flicker noise is a property exhibited by nearly all electronic devices whereby there is a $1/f$ relationship in the noise spectral density moving away in frequency from the carrier. A typical residual phase noise plot of an electronic amplifier is shown in Figure 1.2, where f is the single-sided frequency offset from the carrier. The $1/f$ noise can be seen at frequencies close-to-carrier before the noise becomes dominated by flat white noise at higher offset frequencies. The transition point between the

$1/f$ and flat white noise is known as the flicker noise corner, f_c , and is one way of characterising the residual phase noise performance of a two-port device.

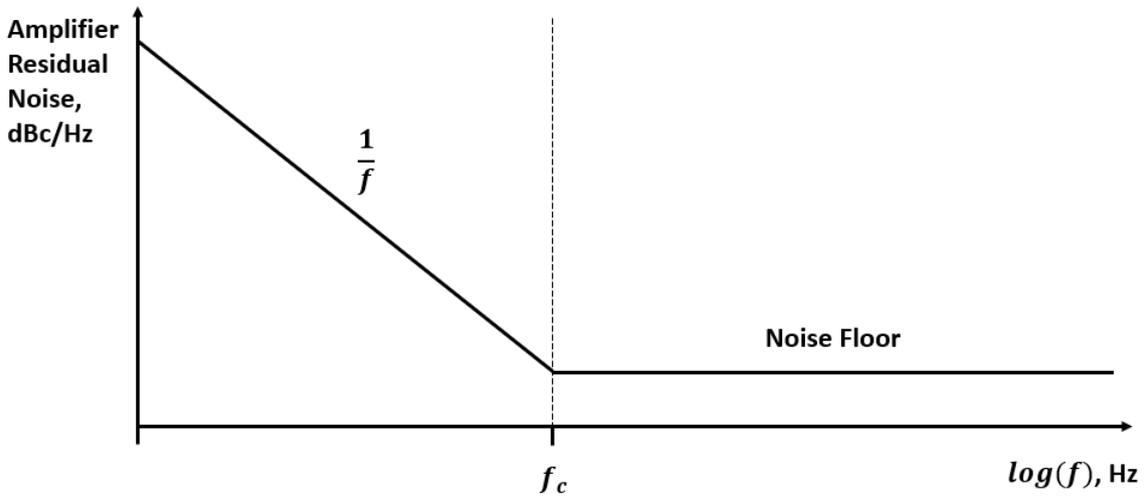


Figure 1.2: Typical plot of the residual phase noise of an amplifier

Noise Floor

The flat noise for far-out frequencies of an electronic device arises due to the thermal noise floor. The thermal noise floor in a given bandwidth, B is simply kTB , where k is the Boltzmann constant and T is temperature in Kelvin. The thermal noise floor in a 1 Hz bandwidth, i.e. the noise spectral density, in terms of dBm/Hz is then:

$$= 10 \log_{10}\left(\frac{kT}{1 \text{ mW}}\right) \approx -174 \text{ dBm/Hz} \quad (1.1)$$

The noise figure of the device describes the ratio of noise that is added to output signal of the device compared to the noise that results from the thermal noise at the input amplified by the gain of the device. The noise floor of the single-sided (i.e. 3 dB less than the total) noise spectral density of the device, L_{NF} can then be approximated from the following equation:

$$L_{NF} \approx -174 \text{ dBm/Hz} - P_{in} + NF - 3 \text{ dB} \quad (1.2)$$

Where P_{in} is the input power to the device and NF is the noise figure.

1.2.3 Modelling Oscillator Phase Noise

It is important to be able to predict the phase noise performance of an oscillator from its operating conditions and the characteristics of each of the components used in the oscillator. This makes the process of designing oscillators for excellent phase noise performance much easier and allows the measured noise performance to be compared with theory.

A Simple Model

A simple model for oscillator phase was first described by Leeson in [7], where it is shown that the flat white noise in the open-loop oscillator is transformed to $1/f^2$ noise inside the bandwidth of the resonator, f_{BW} , and similarly the flicker noise is transformed to $1/f^3$ noise at frequencies below the flicker noise corner. A typical single-sided phase noise plot of an oscillator can be seen in Figure 1.3.

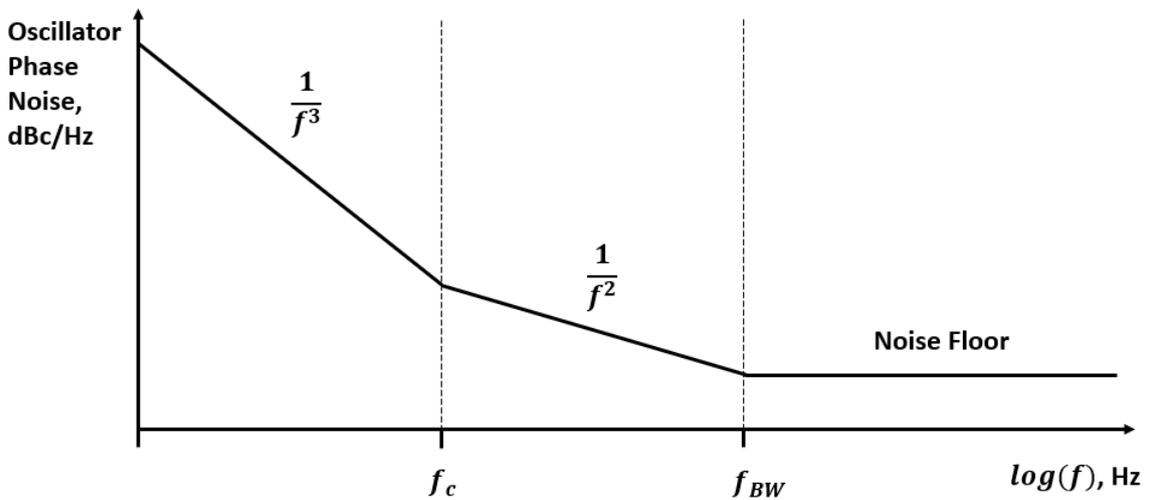


Figure 1.3: Plot showing how the absolute phase noise of an oscillator relates to its open-loop residual phase noise and the bandwidth of the resonator

An Equation for Calculating Oscillator Phase Noise

It is useful to have a more complex model that takes into account the measured characteristics of each of the components in the oscillator to produce a calculation of the absolute phase noise of the oscillator. Such a model is described in [1] and the complete phase noise equation is given as:

$$L(f) = 10 \log_{10} \left(\frac{F_2 k T}{C_0 2 P} + \left(1 + \frac{f_c}{f} \right) \left(\frac{F_1 k T}{2 P} \left(\frac{1}{[1 - \frac{Q_L}{Q_0}]^2} \right) + \frac{F_1 k T}{8(Q_0)^2 (\frac{Q_L}{Q_0})^2 + (1 - \frac{Q_L}{Q_0})^2 P} \left(\frac{f_0}{f} \right)^2 \right) \right) \quad (1.3)$$

Where P is the power available to the resonator,

Q_0 is the unloaded Q of the resonator,

Q_L is the loaded Q of the resonator,

F_1 is the noise figure of the sustaining amplifier,

F_2 is the noise figure of the output buffer amplifier,

and C_0 is the coupling coefficient relating the power into the resonator to the power into the buffer amplifier.

This model assumes that the amplifier is the dominant source of flicker noise.

Seeing as none of the oscillators that are presented in this chapter uses a buffer amplifier, the equation can be simplified as follows:

$$L(f) = 10 \log_{10} \left(\left(1 + \frac{f_c}{f} \right) \left(\frac{F_1 k T}{2 P} \left(\frac{1}{[1 - \frac{Q_L}{Q_0}]^2} \right) + \frac{F_1 k T}{8(Q_0)^2 (\frac{Q_L}{Q_0})^2 + (1 - \frac{Q_L}{Q_0})^2 P} \left(\frac{f_0}{f} \right)^2 \right) \right) \quad (1.4)$$

This equation will be used as the basis for the theoretical prediction of the phase noise of the oscillators described in this section.

1.3 Design

In this section, the overall design of the ultra-low phase noise 100 MHz crystal oscillator and its constituent parts are outlined in detail.

1.3.1 Block Diagram

A block diagram of the construction of the oscillator is shown in Figure 1.4. The oscillator has an output power of around 5 dBm and a control voltage can be used for narrow-band frequency tuning that allows the oscillator to be used as part of a phase-locked loop. The various elements that comprise the oscillator are described in detail in the following sections.

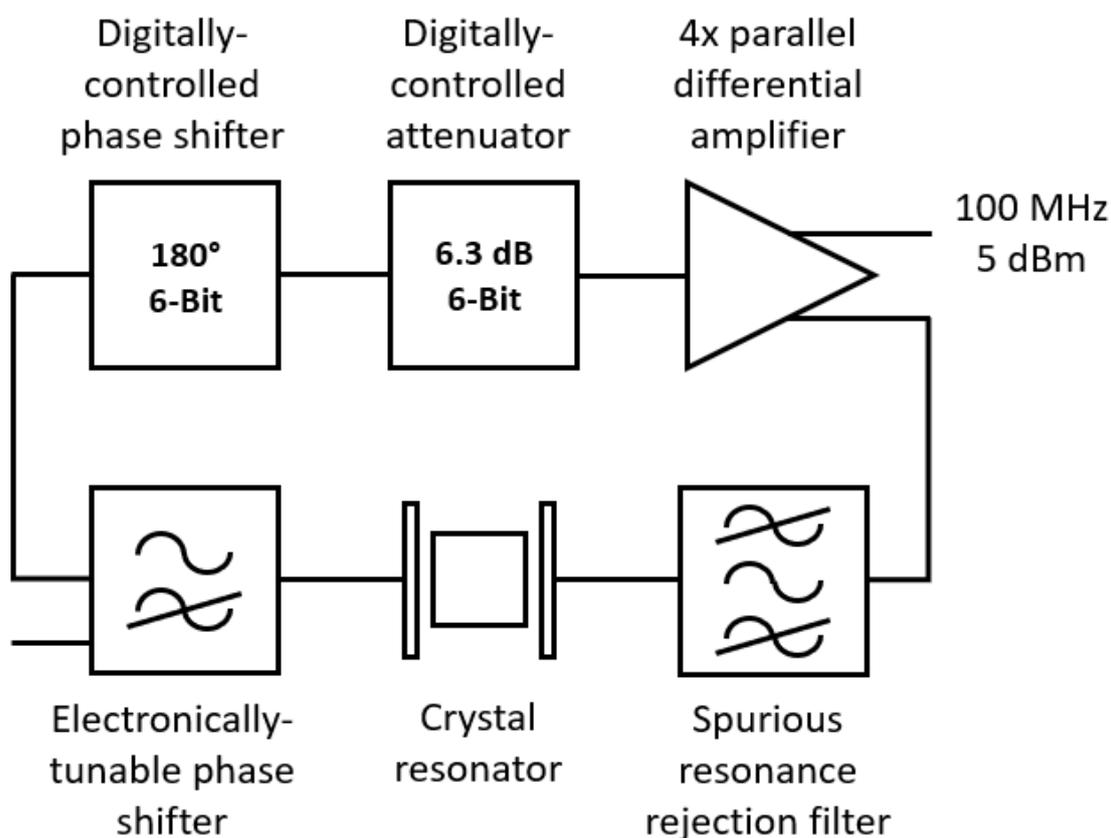


Figure 1.4: Block diagram of 100 MHz crystal oscillator

1.3.2 Crystal Resonator

The selection of the crystal resonator is important for achieving the lowest possible phase noise for a crystal oscillator. The crystal should have a high quality factor, low residual noise and, ideally, be capable of operating at a high drive level.

The KVG O5SC85105A 100 MHz quartz crystal resonator [8] was selected primarily as it has a specified maximum phase noise of -130 dBc/Hz at a 100 Hz offset when operating at the recommended drive level. Additionally, it has a relatively high recommended drive level of around 7 dBm. The effect that the crystal has on the oscillator phase noise measurements is explored in more detail in Section 1.8.

Frequency Response

The frequency response of the crystal was measured using the SDR-Kits DG8SAQ Vector Network Analyzer [9] and is shown in Figure 1.5. The crystal has an insertion loss of 4.86 dB. It can be seen that the parallel resonance of the crystal is very close to the series resonance and so the peak is slightly asymmetrical, through it is unclear if and how this could affect the phase noise that can be achieved with this crystal.

The broadband frequency response of the crystal in Figure 1.6 shows that there are additional spurious resonant frequencies present in the crystal both above and below the nominal frequency. These resonant frequencies could cause spurious oscillations when the crystal is placed in an oscillator loop.

The DG8SAQ also has a crystal analyser function that can be used derive an RLC model of a crystal and this is shown for the KVG crystal in Figure 1.7. The crystal has an unloaded Q of around 119,000 and an equivalent series resistance of around 81 Ω .

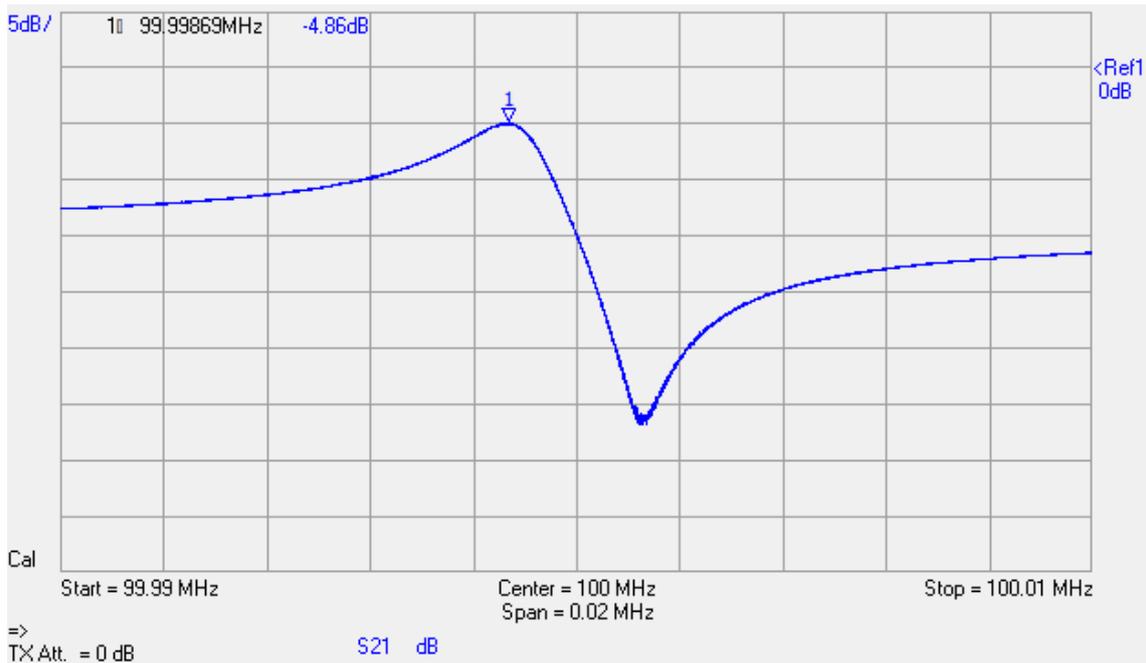


Figure 1.5: S21 frequency response of the KVG 100 MHz quartz crystal resonator measured using the SDR-Kits DG8SAQ VNWA

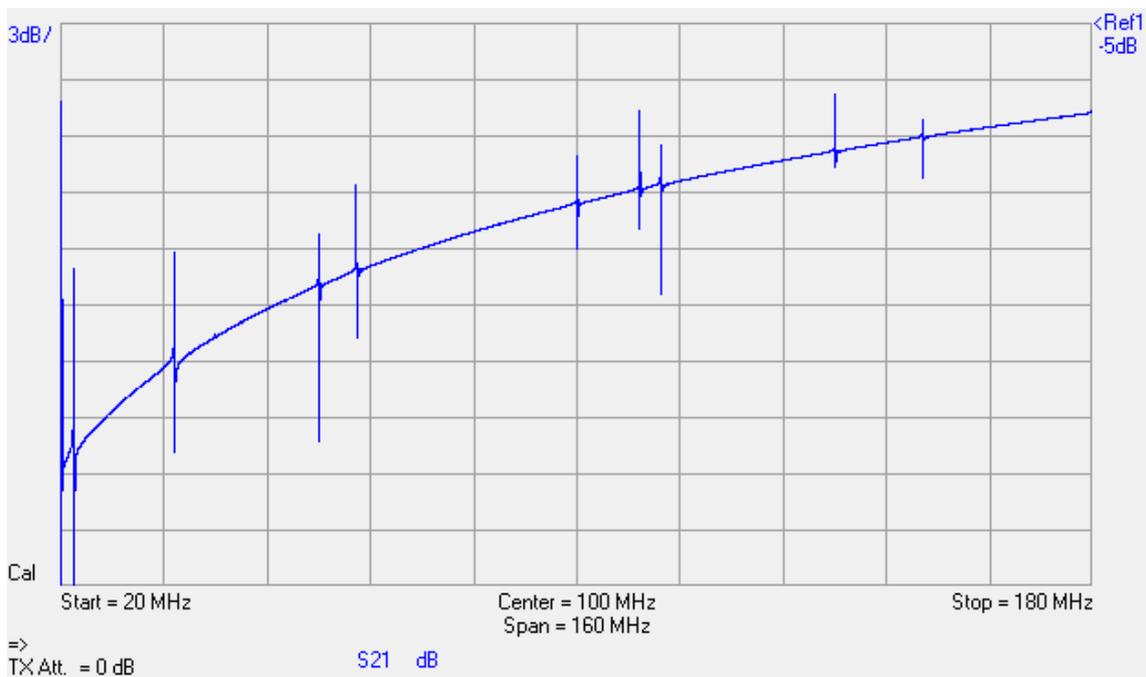


Figure 1.6: S21 broadband frequency response of the KVG 100 MHz quartz crystal resonator measured using the SDR-Kits DG8SAQ VNWA

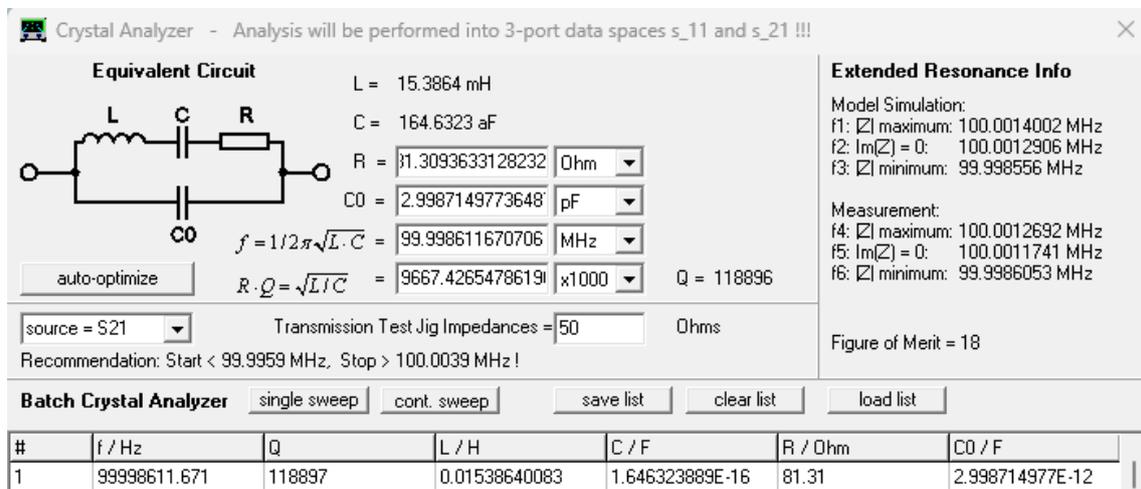


Figure 1.7: Model of the KVG 100 MHz quartz crystal resonator measured using crystal analyser function of the SDR-Kits DG8SAQ VNWA

1.3.3 Spurious Resonance Rejection Filter

The purpose of this filter is to prevent oscillations at unwanted frequencies that can occur due to spurious resonant frequencies present in the crystal. The bandwidth should be narrow enough to suppress these oscillations, but the insertion loss at the nominal frequency should be kept to a minimum to avoid excessive losses in the oscillator loop.

Circuit Diagram

The filter was designed as a 3rd-order LC band-pass filter centred at the desired operating frequency, 100 MHz, and with a bandwidth of 40 MHz. This bandwidth was chosen so as to suppress most of the spurious resonant frequencies highlighted in Figure 1.6 that occurred at frequencies above 120 MHz and below 80 MHz. The spurious resonances that are close to the nominal frequency are unlikely to cause spurious oscillations as the open loop phase shift of the oscillator at these frequencies will deviate from the 0° that is set for the nominal frequency of 100 MHz. The circuit diagram is shown in Figure 1.8.

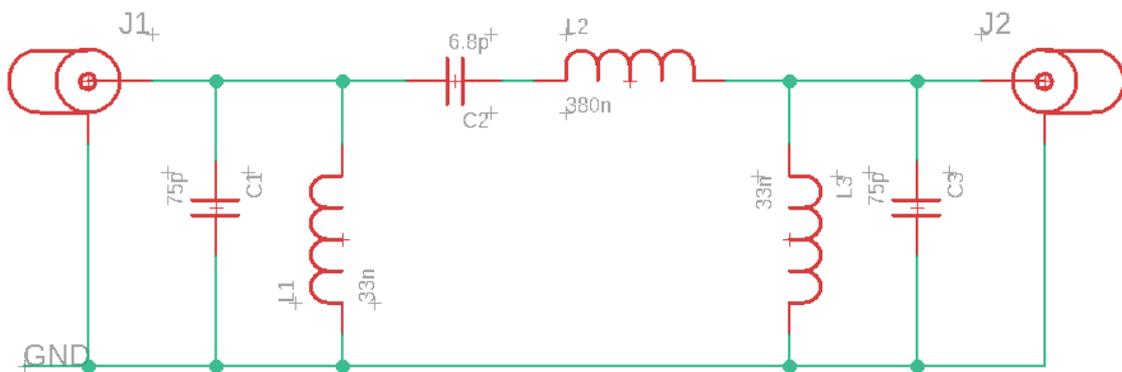


Figure 1.8: Circuit diagram of the 40 MHz bandwidth spurious resonance rejection filter

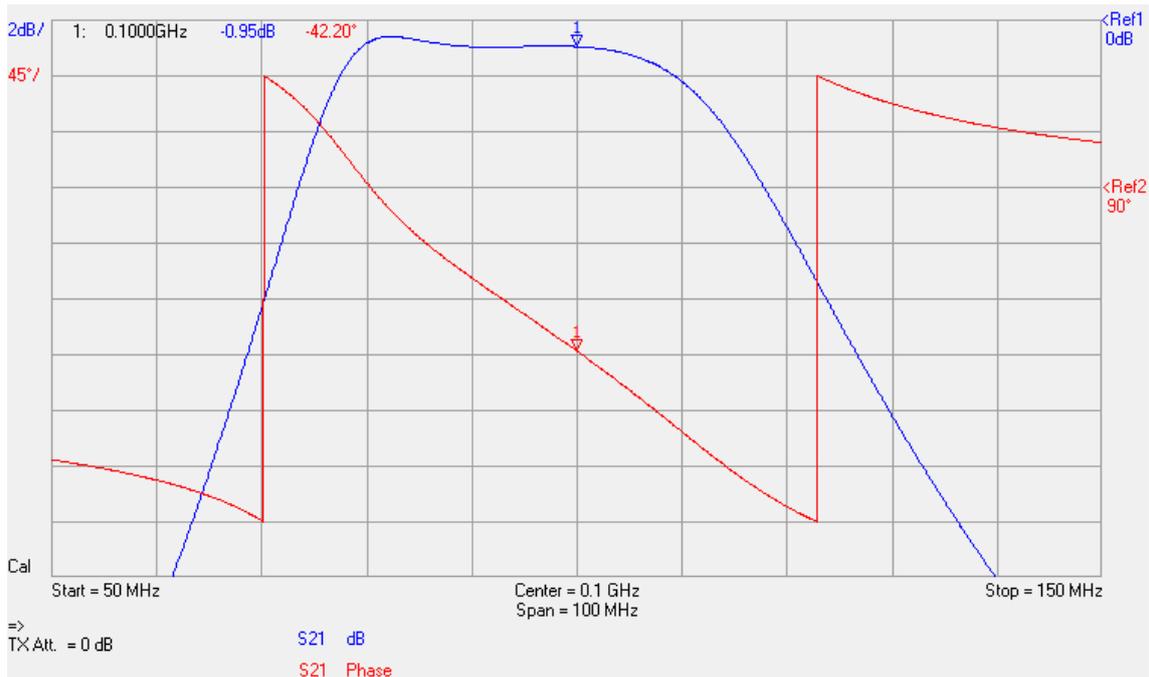


Figure 1.9: S21 frequency response of the spurious resonance rejection filter measured using the SDR-Kits DG8SAQ VNWA

Frequency Response

A measurement of the filter frequency response is shown in Figure 1.9. It can be seen that the filter should be sufficient to suppress oscillations outside the range of around 80 to 110 MHz. Any spurious frequencies inside this range should anyway be suppressed due to the setting of the oscillator loop phase. The insertion loss at the nominal frequency is 0.95 dB.

1.3.4 Electronic Phase Shifter

The electronically-tunable phase shifter allows for fine-tuning of the oscillator frequency by slightly adjusting the overall loop phase of the oscillator. As well as providing a method by which the oscillator can be incorporated into a system that requires frequency control, such as a phase-locked loop, it also allows the oscillator to be set as close as possible the resonant frequency of the crystal. It has been shown by Everard et al. that the noise performance of an oscillator degrades with

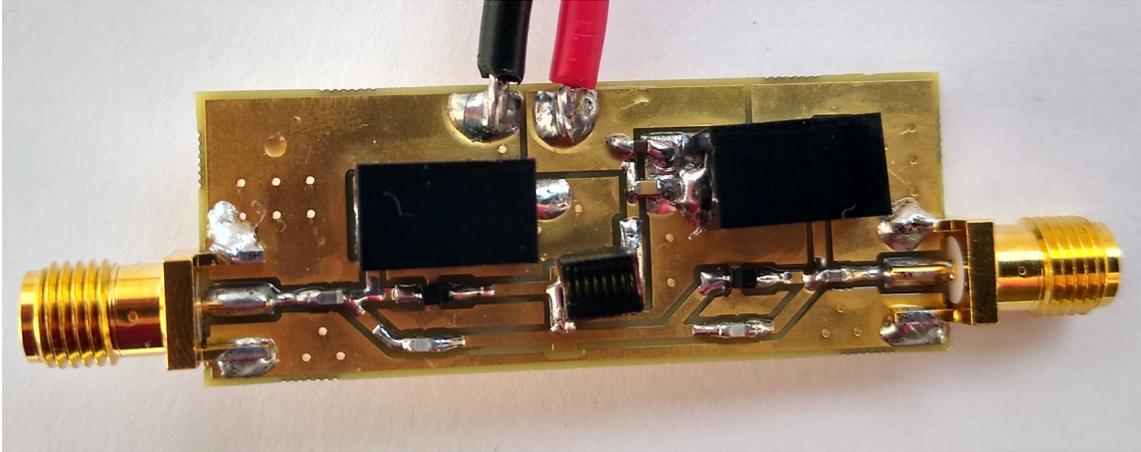


Figure 1.10: 100 MHz 25° electronic phase shifter

respect to the the open-loop phase error by a factor of $\cos^4 \theta$ [10] [11].

Circuit Diagram

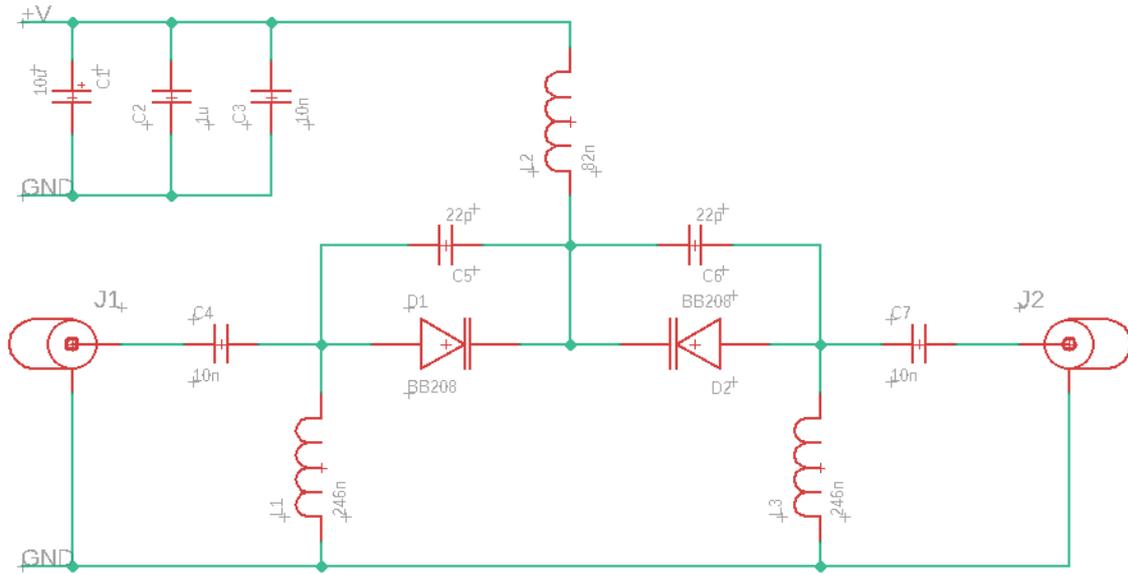


Figure 1.11: Circuit diagram of the electronic phase shifter

The circuit diagram of the electronic phase shifter is shown in Figure 1.11 and is based on the circuit described in [1]. It is a 5th order LC Butterworth high-pass filter where the capacitive elements of the circuits have been replaced with the parallel combination of a capacitor and a BB208 varactor diode that is biased through a DC voltage applied to the centre inductor. The bias voltage can be changed to adjust the capacitance of the varactors, and therefore, the cutoff frequency of the filter. A signal operating close to the cutoff frequency will then experience a change in phase according to the bias voltage. Of course, the signal will also experience a change in amplitude, which limits the usable phase range as the insertion loss of the device should be kept as low as possible. Similarly to the device described in [1], the cutoff frequency of the phase shifter was designed to be $(0.6 \cdot F_0)$, where F_0 is the oscillator frequency, as it was found to allow for a reasonable phase tuning range without a significant change in insertion loss.

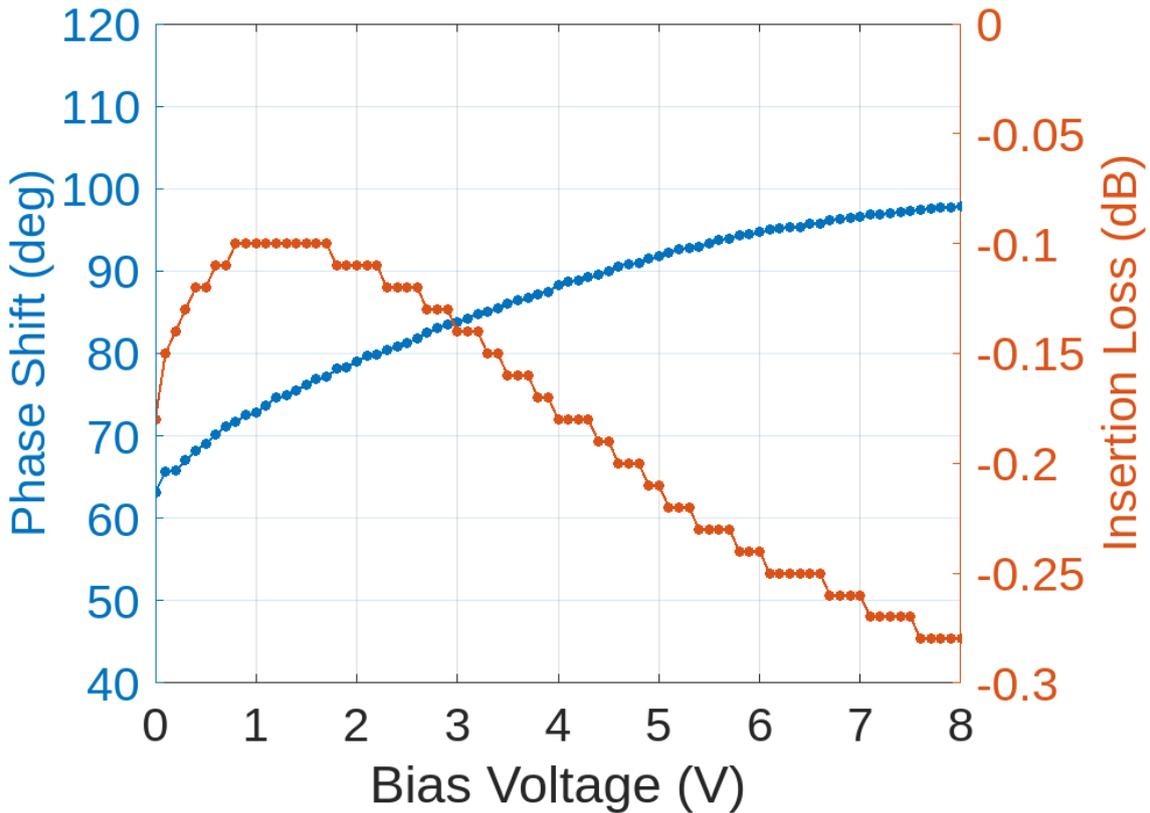


Figure 1.12: Plot of the absolute phase shift and insertion loss of the electronic phase shifter versus the DC bias voltage measured using the SDR-Kits DG8SAQ VNWA

Phase Shift and Insertion Loss

Shown in Figure 1.12 is the absolute phase shift and insertion loss of the phase shifter versus the applied bias voltage measured using the SDR-Kits DG8SAQ Vector Network Analyzer. The overall phase tuning range can be seen to be around 30° and across this range, the insertion loss does not change by more than 0.2 dB. The worst case insertion loss is around 0.3 dB.

Assuming the loaded Q of the crystal, Q_L , is in the region of $(0.5 \cdot Q_0) \approx 50,000$ the frequency tuning range, F_{tune} , at $F_0 = 100 \text{ MHz}$ can be calculated for a phase range of $\pm 30^\circ$ as follows:

$$F_{tune} = \left(2\pi \cdot \frac{\pm 15^\circ}{360^\circ}\right) \cdot \frac{F_0}{2Q_L} = \pm 262 \text{ Hz} \quad (1.5)$$

Therefore it should be possible using this design to adjust the frequency of the 100 MHz oscillator over a range of about 500 Hz without introducing significant losses into the loop.

1.3.5 Digitally-Controlled Phase Shifter

Compared to the electronic phase shifter, the digitally-controlled phase shifter operates over a much wider phase range and is used as a convenient way to set the open-loop phase shift of the oscillator to be around 0° without the need to add in different lengths of cable. It is used in conjunction with the electronic phase shifter as the limited resolution and response time of the digital phase shifter mean that it is not suitable as a method of oscillator frequency control.

Such a device is particularly necessary for the work described in this chapter, as it involves setting up the oscillators in a variety of different configurations using various amplifiers and resonators, each of which will affect the open-loop phase of the oscillator. The digital phase shifter can be used to make quick and repeatable adjustments to set loop phase to 0° .

The technical details and design procedure of the digitally-controlled phase shifters are described in Chapter 2. They were designed specifically to exhibit very low residual phase noise. The phase shifter used in the 100 MHz crystal oscillator is a 6-bit device with a 2.8125° resolution and 180° range, so that a total 360° can be achieved when used with a differential amplifier with complementary outputs. This device is described in more detail in Section 2.2.2.

1.3.6 Digitally-Controlled Attenuator

The oscillator also incorporates a digitally-controlled attenuator that operates similarly to the digitally-controlled phase shifter. The purpose of this attenuator is to allow for the loop power of the oscillator to be easily controlled so that the power

into the crystal resonator can be set to its optimum value even when using amplifiers with different gain compression characteristics. Having this control over the operating characteristics of the oscillator also allows measurements made using different oscillator configurations to be more easily compared.

The design of the digital attenuator is outlined in more detail along with the digital phase shifter in Chapter 2. The attenuator used in the oscillator is a 6-bit device with a 0.1 dB resolution and a range of 6.3 dB. It is described in more detail in Section 2.3.2.

The attenuator is placed in the oscillator loop at the input to the amplifier. It is shown in Section 1.7 how the placement of an attenuator at the the amplifier input can improve oscillator phase noise in situations where the input impedance of the amplifier is not necessarily $50\ \Omega$ or where the flicker noise of the amplifier is sensitive to its source impedance.

1.3.7 Amplifier

The amplifier used in the final version of the oscillator is a quad parallel differential amplifier, the design of which is described in detail in Section 1.9. This amplifier was designed with a configuration that allows it to have a high maximum output power, whilst not introducing too much excess gain, so that the crystal can be driven at its optimal input power without the amplifier operating in heavy saturation. It has a small-signal gain of around 7.5 dB and the 1 dB compression point is at around -3.5 dBm. It also has complementary outputs, one of which is used for the oscillator output, while the other is used in the feedback loop.

Other amplifier designs were considered for the oscillator and in the following sections of this chapter, the processes of selecting and designing the amplifiers; characterising their residual noise performance; incorporating them into the oscillator; and comparing the oscillator phase noise results to the theory are described.

1.4 MAR-6 Amplifier

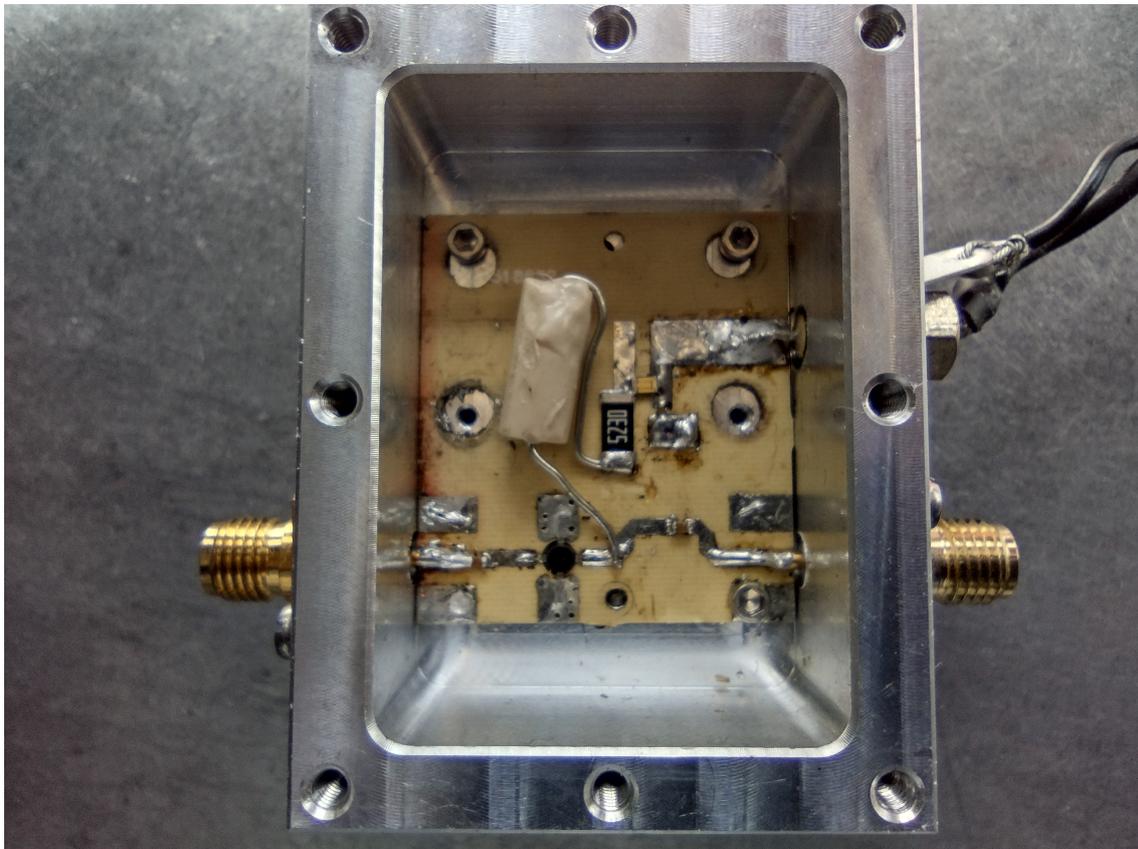


Figure 1.13: Photo of MAR-6 amplifier

The Mini Circuits MAR-6+ [12] is a commercially-available monolithic amplifier package that operates from DC to 2 GHz with a gain of around 22 dB. It was selected primarily for its stated low noise figure of 2.3 dB the fact that it is internally-matched to 50 Ω .

It is shown in this section that the high excess gain introduced by this amplifier in the oscillator loop limits the phase noise performance that can be achieved, however the residual and absolute phase noise measurements made using this amplifier provided good points of reference for the work on flicker noise measurement using an LC resonator, Section 1.6; improving phase noise using attenuator matching, Section 1.7; and the cross-correlation system described in Chapter 3.

1.4.1 Design

Circuit Diagram

The MAR-6 amplifier circuit was originally constructed by St. John Gilbert, another member of this research group, and the circuit diagram is shown in Figure 1.14.

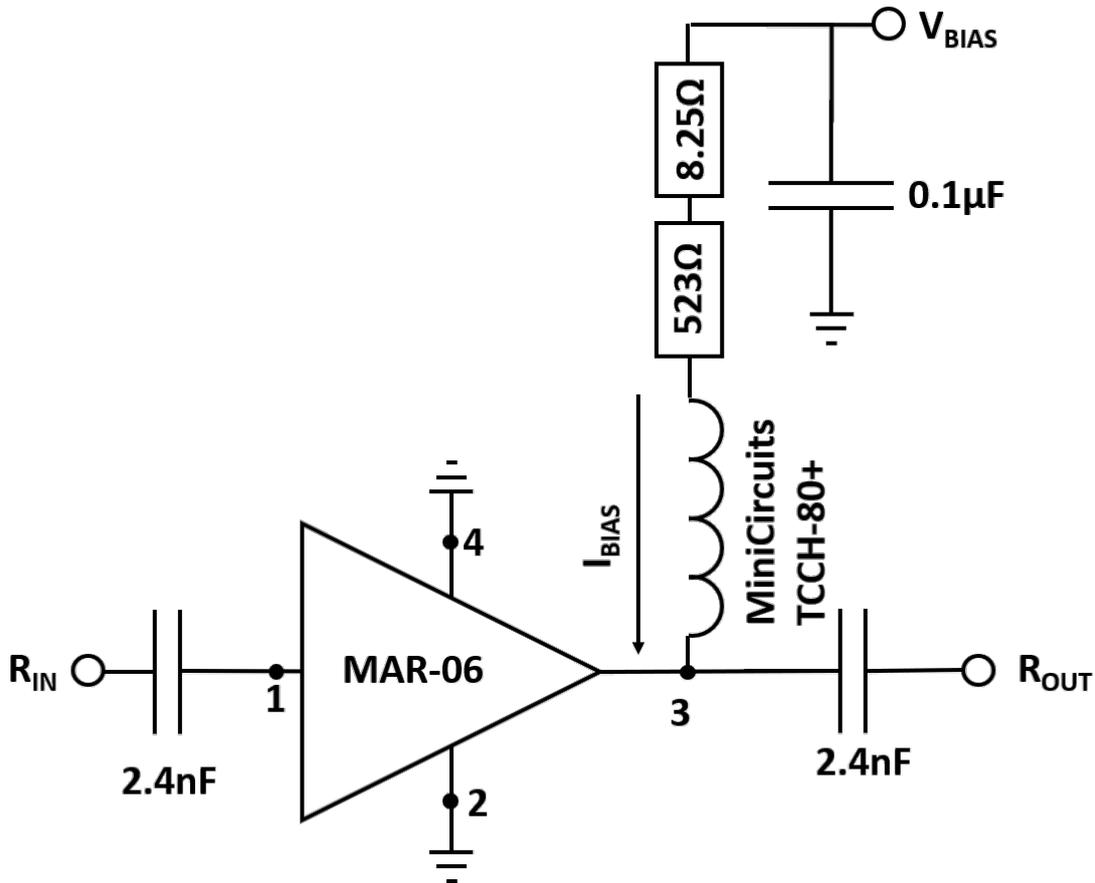


Figure 1.14: Circuit diagram of the MAR6 amplifier

The amplifier requires an external bias current on its output that is achieved through a series combination of resistors and an RF choke connected to a positive supply voltage, V_{bias} . The resistor values were selected according to the evaluation board circuit given in the datasheet to provide the optimum resistance for a 12 V supply voltage. 2.4 nF coupling capacitors on the input and output of the device are also included that have a negligible impedance at 100 MHz.

Oscillator Configuration

Unlike the differential amplifiers described in Section 1.5 and Section 1.9, the MAR-6 amplifier does not have complementary outputs and so there is no extra output that can be used as the oscillator output signal. This amplifier also has much more gain than is required to maintain oscillation, even more so than can be compensated for by the 6.3 dB digital attenuator. To address these issues, a configuration was devised whereby a 10 dB coupler was incorporated into the oscillator loop as shown in Figure 1.15.

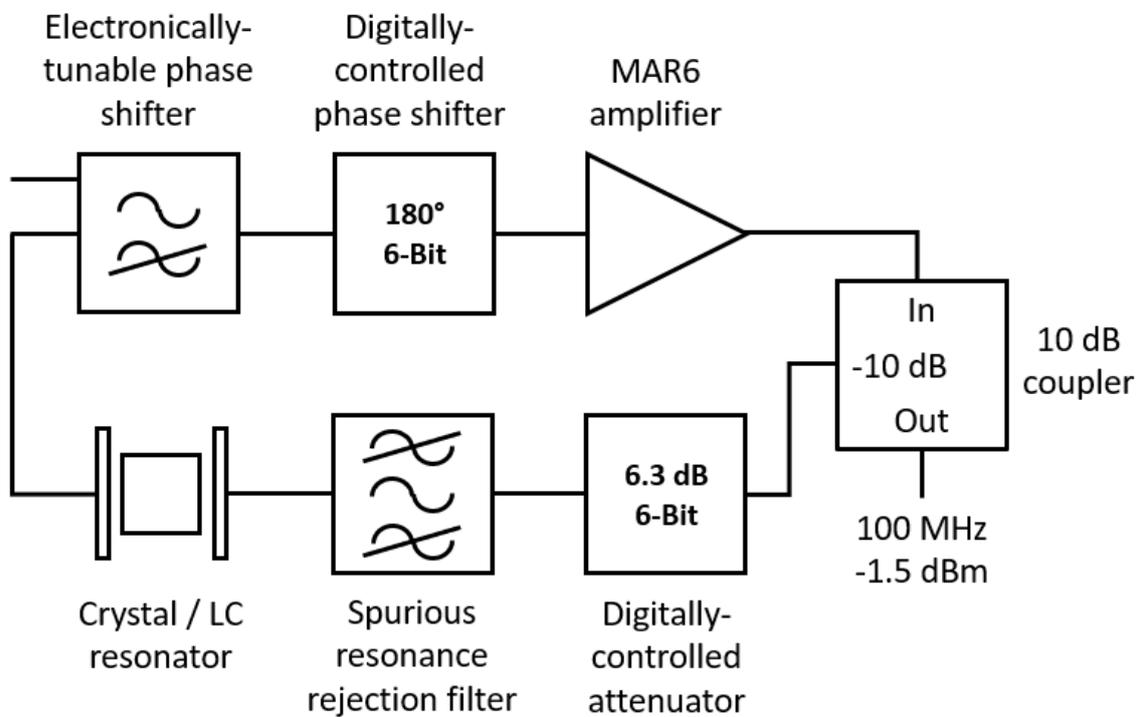


Figure 1.15: Block diagram of the initial 100 Mhz oscillator setup that was used to measure the flicker noise of the MAR-6 amplifier

1.4.2 Measurements

Gain Compression

It is important to understand the gain compression characteristics of the amplifier, as this will show the maximum power that can be driven in to the crystal resonator

and also the level of saturation the amplifier will be operating at for a given drive level. If the amplifier is operating under heavy saturation, this could significantly degrade its noise performance.

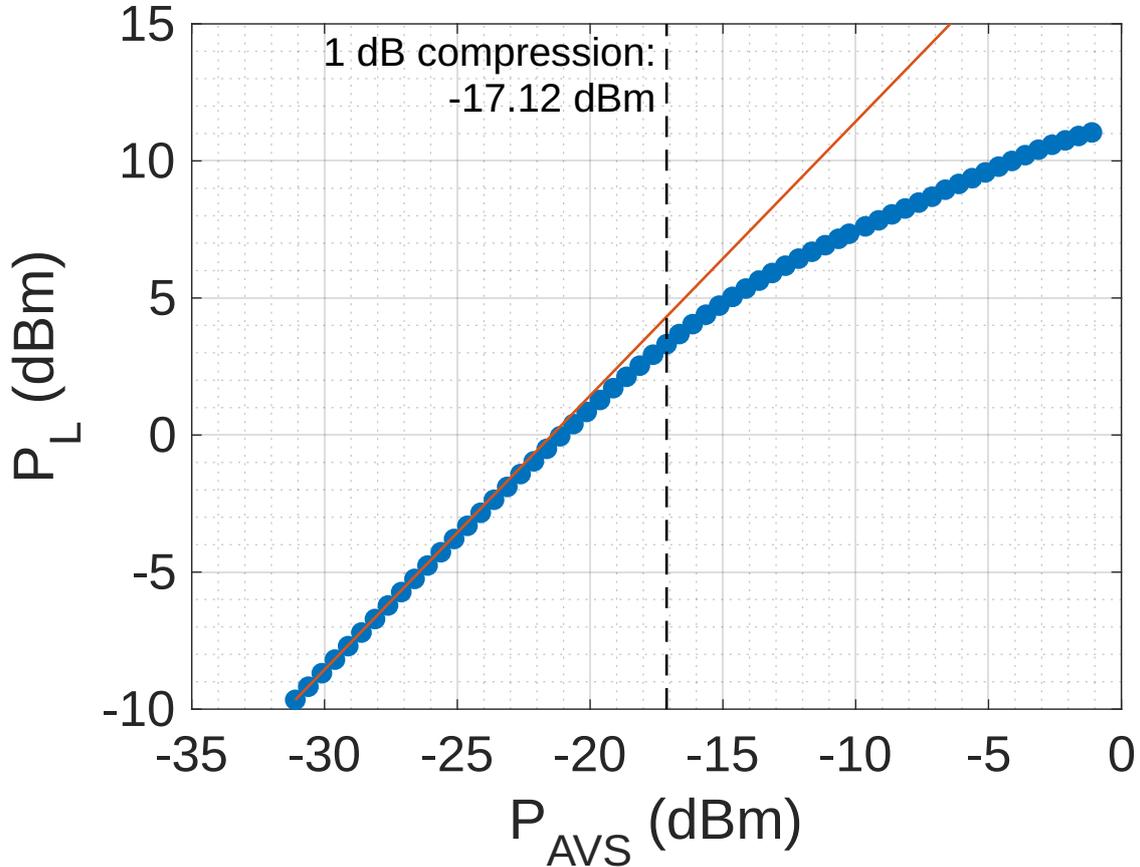


Figure 1.16: Gain compression of the MAR-6 amplifier measured using the Rohde & Schwarz FSWP

In Figure 1.16 the measured output power of the MAR-6 amplifier into a 50Ω load, P_L , is plotted against the source power applied to its input, P_{AVS} . Also shown on the plot is a straight line of constant gain that is extrapolated from the small signal gain of the amplifier, which allows the 1 dB compression point of the amplifier to be determined. The 1 dB compression point is the point at which the measured output power drops to 1 dB below the output power extrapolated from the small signal gain. Here it can be seen that the 1 dB compression point at 100 MHz occurs for a P_{AVS} of -17.12 dBm, which corresponds to an output power of around 3 dBm.

Noise Figure

The noise figure of the MAR-6 amplifier was measured using the Rohde & Schwarz FSWP Phase Noise Analyzer [13] and was found to be 1.9 dB. This is exceptionally low and even slightly lower than the typical value of 2.4 dB stated in the datasheet.

Residual Noise

The residual phase noise of the MAR-6 amplifier was measured using the cross-correlation system described in Chapter 3 and is shown in Figure 1.17. The output power of the amplifier was set so that it would be roughly the same as it would be when incorporated into the oscillator.

The flicker noise corner of this amplifier is exceptionally low at around 20 Hz, however it does not exhibit a particularly low noise floor at -151.5 dBc/Hz. This noise floor does tie in with the theory based on the operating conditions and the amplifier noise figure; it is simply the result of having such a high gain.

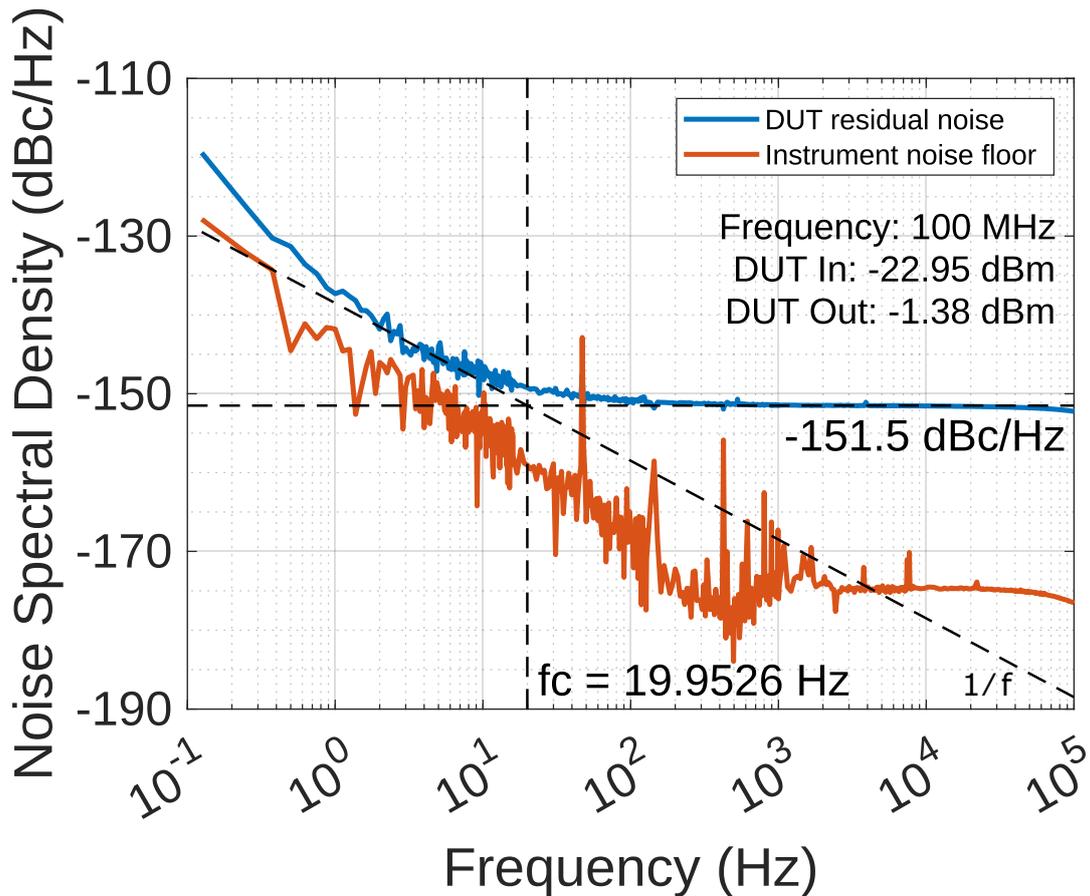


Figure 1.17: Residual phase noise of the MAR-6 amplifier measured using the cross-correlation system with 256 correlations at 0.125 Hz and 8,388,608 correlations at 100 kHz

1.4.3 Crystal Oscillator Measurements

Knowing the residual noise of the amplifier, it is possible to determine the theoretical absolute phase noise of the crystal oscillator incorporating this amplifier using the equation described in Section 1.2.3. This was compared to an actual measurement of the oscillator phase noise using the configuration described in Section 1.4.1 and the results are shown in Figure 1.18. At this point in the development of the oscillator, prior to the procurement of the KVG crystal resonator used in the final design, the crystal being used was a Nofech SC-cut 5th order quartz crystal with a Q of around 110,000.

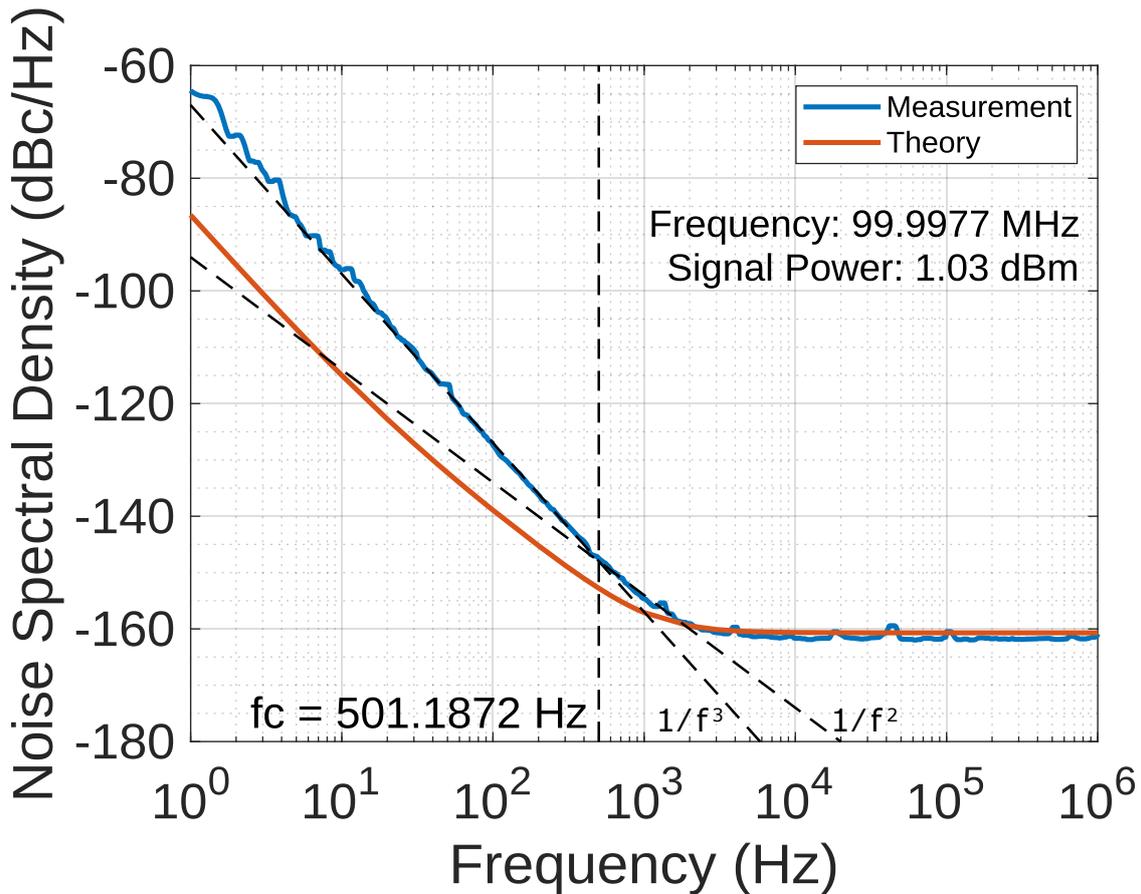


Figure 1.18: Absolute phase noise measurement of the 100 MHz crystal oscillator with MAR-6 amplifier and 3 dB of attenuation after the amplifier made using the Rohde & Schwarz FSWP

It can be seen that the oscillator exhibits an absolute phase noise of around -130 dBc/Hz at a 100 Hz offset, which is around 10 dB off from the target value. There is also a significant discrepancy between the theoretical and measured close-to-carrier noise; something that is discussed in greater detail in Sections 1.7 and 1.8.

Even if the measured phase noise was in agreement with the theory, it still would not quite reach the target phase noise. This is because the high excess gain of the MAR-6 amplifier requires a lot of extra attenuation in the oscillator loop, which effectively increases the noise figure of the amplifier, limiting the minimum phase noise that can be achieved. In the following section, an alternative amplifier circuit

is presented that is designed to have a similar maximum output power and similarly low flicker noise, but with greatly reduced gain.

1.5 Differential Amplifier

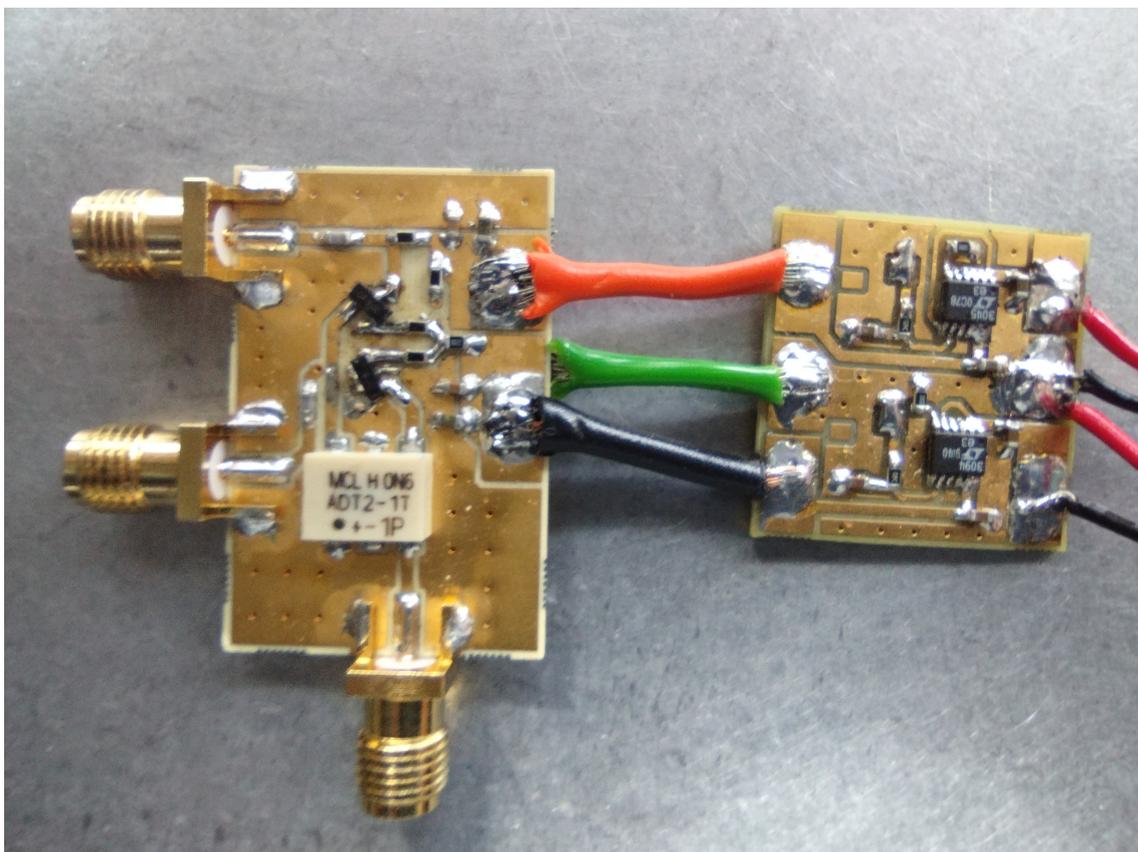


Figure 1.19: Photo of differential amplifier circuit

In this section, an alternative design is explored that is based around a long-tailed pair differential amplifier. This amplifier is designed to have a lower gain than the MAR-6 amplifier described in Section 1.4, whilst having a similar 1 dB compression point and a similarly low noise figure and flicker noise corner.

1.5.1 Design

Circuit Diagram

The differential amplifier circuit is based on a similar design that was used for the 10 MHz crystal oscillator designed at York [1] and is shown in Figure 1.20. It is based around a long-tailed pair differential amplifier where the differential inputs are driven by a Mini Circuits ADT2-1T tapped transformer with a 2-1 impedance transformation ratio. The bias current through the transistors is set by the negative bias voltage, V_{bias} , in conjunction with resistor R3. The amplifier has complementary outputs at the 50 Ω collector resistors, which ensure that the output impedance is roughly 50 Ω , along with 10 nF coupling capacitors, which have a negligible impedance at 100 MHz.

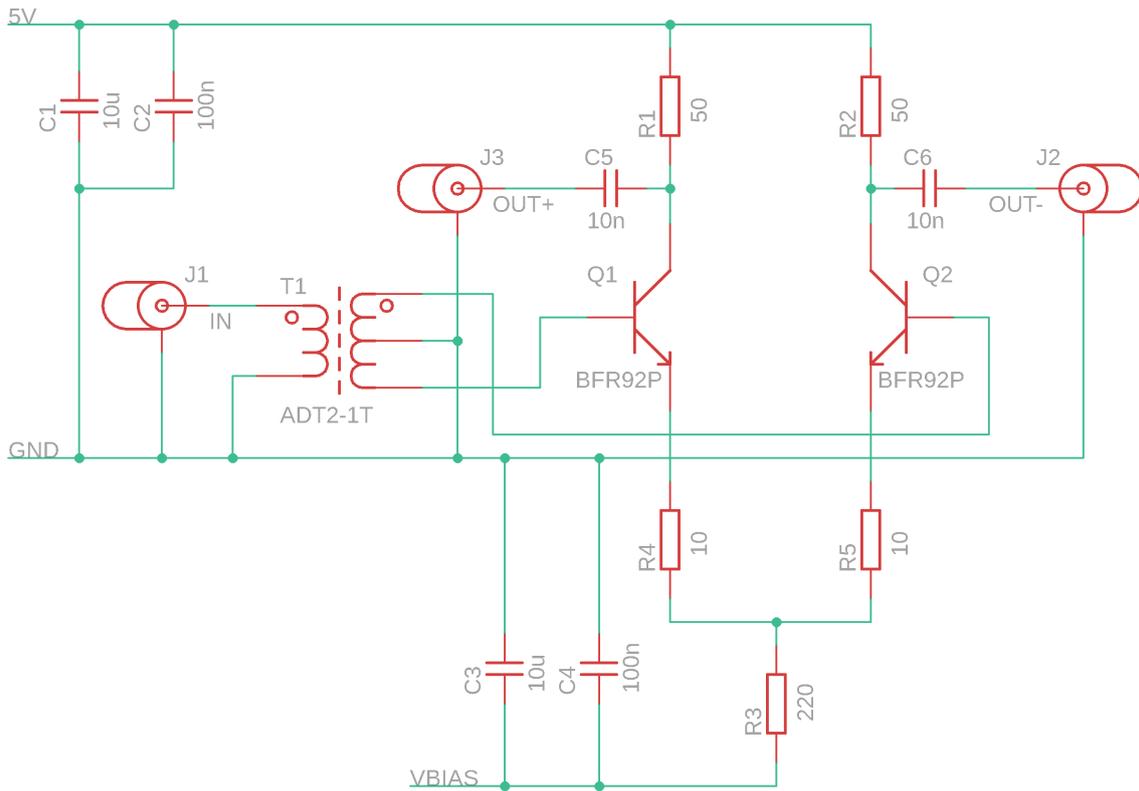


Figure 1.20: Circuit diagram of the differential amplifier

A separate board was built to provide low-noise regulated voltages for the 5

V supply as well as the negative V_{bias} . The board can be seen in Figure 1.19 and consists of an LT3045 positive regulator with a +5 V output and an LT3094 negative regulator with a -5 V output.

The BFR92P bipolar transistors [14] were selected primarily due to their stated low noise figure of around 1.4 dB. Other low-noise RF transistors were considered, however it was found that transistors with high transition frequencies, f_T , would lead to amplifier instability in the GHz range, which was especially apparent in the quad parallel version of this amplifier described in Section 1.9. The BFR92P has a typical f_T of 5 GHz, which is high enough that a considerable amount of gain can be achieved at 100 MHz whilst preventing oscillations from occurring at much higher frequencies.

The amplifier also incorporates $10\ \Omega$ emitter feedback resistors R4 and R5. In his PhD thesis [15], and the paper [16], Theodoropoulos shows that the $1/f$ residual noise of silicon bipolar amplifiers can be significantly reduced through the inclusion of emitter feedback. This is because the $1/f$ noise arises in part due to the variation of the device's r_e parameter resulting from the base-emitter recombination flicker noise, the effect of which can be reduced if an external emitter resistance, R_E , that is significantly larger than r_e is added. For this amplifier, the value of r_e is around $2.5\ \Omega$.

PCB Layout

The PCB layout used for the differential amplifier is shown in Figure 1.21. It is constructed from 2-layer 0.5 mm I-Tera with gold plating and plated through holes. The bottom layer is used almost entirely as a ground plane.

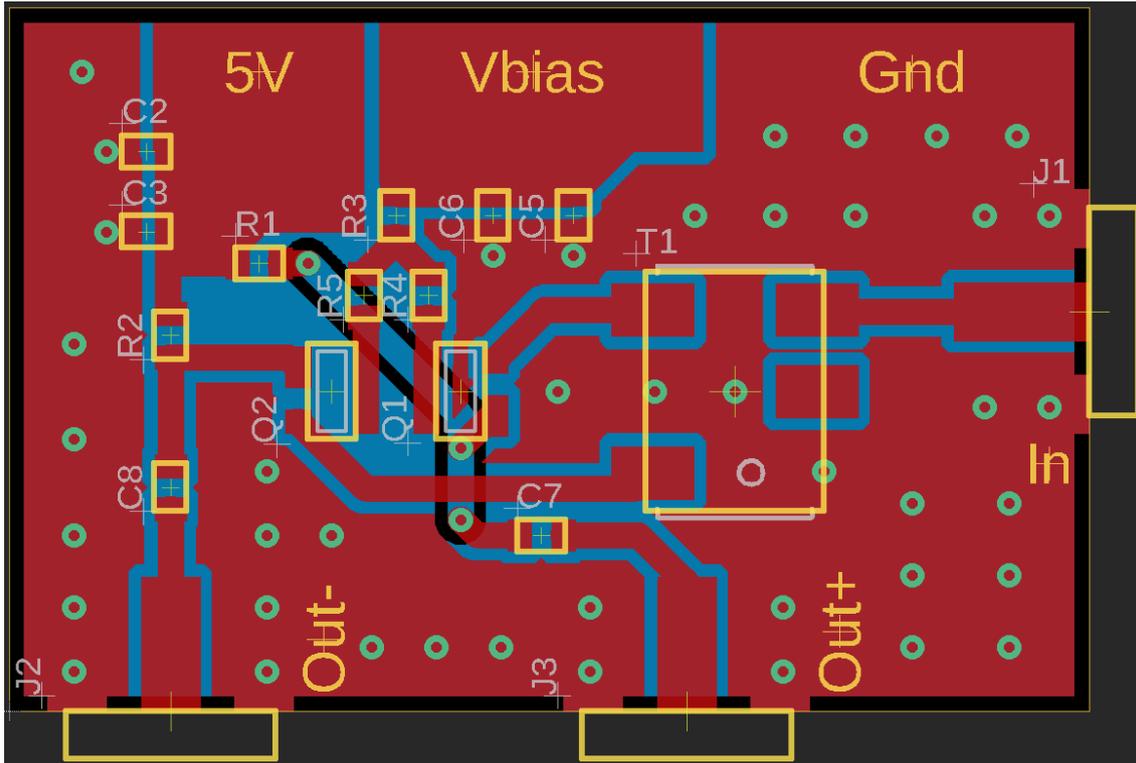


Figure 1.21: PCB layout of the differential amplifier

1.5.2 Model

A model to describe the operation of the differential amplifier can be derived by considering just a single end of the long-tailed pair, i.e. by considering each output as a common-emitter amplifier as shown in the left-hand diagram in Figure 1.22. Here, the emitter current, I_E , is equal to half the bias current set by V_{bias} and the resistor R_3 as it can be assumed to be split evenly between the two transistors provided they are well-matched. Since the tapped transformer that supplies the differential signals to the amplifier has an impedance ratio of 2, if the source impedance is 50Ω then the impedance presented to each of the transistors will also be 50Ω . An impedance transformation ratio of 2 corresponds to a voltage transformation ratio $\sqrt{2}$, therefore the input into each transistor can be replaced with a voltage source of $\frac{\sqrt{2}}{2} \cdot V_{\text{in}}$ with a 50Ω series resistance.

The right-hand side of Figure 1.22 shows the same circuit, but with an added

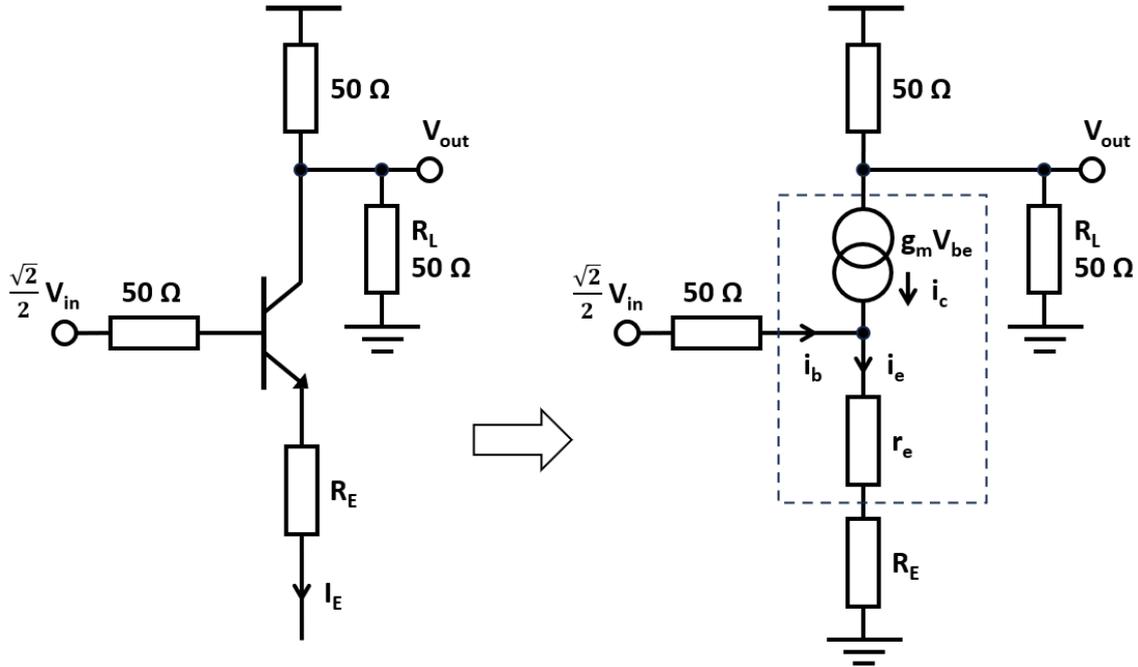


Figure 1.22: Simplified model of a single end of the long-tailed pair differential amplifier

ground signal to represent the virtual earth created by the differential amplifier configuration and with the transistor replaced by a T model for small signal analysis.

Gain

When the transistor is operating in the linear region, the emitter transresistance, r_e , relates the small-signal emitter current, i_e , to V_{be} . Assuming the thermal voltage, V_T , at room temperature is 25 mV, r_e can be approximated in terms of the quiescent emitter current, I_E , as:

$$r_e = \frac{V_T}{I_E} = \frac{25 \text{ mV}}{I_E} \quad (1.6)$$

V_{out} can be calculated quite simply from the current through the 50Ω collector resistor:

$$V_{out} = i_c \cdot 50 \Omega \quad (1.7)$$

The emitter and base currents can be determined from the transistors small-signal current gain parameter, h_{fe} , as follows:

$$\begin{aligned} i_b &= \frac{i_c}{h_{fe}} \\ i_e &= \frac{i_c \cdot (h_{fe} + 1)}{h_{fe}} \end{aligned} \quad (1.8)$$

V_{in} can then be calculated by summing the voltages across the emitter and base resistors:

$$V_{in} = \frac{\sqrt{2}}{2} \cdot \left(\left(\frac{i_c}{h_{fe}} \cdot 50 \Omega \right) + \left(\frac{i_c \cdot (1 + h_{fe})}{h_{fe}} \cdot (R_E + r_e) \right) \right) \quad (1.9)$$

The small-signal gain of the amplifier, A , is then:

$$A = \frac{V_{out}}{V_{in}} = \frac{\sqrt{2}}{2} \cdot \frac{50 \Omega}{\left(\frac{50 \Omega}{h_{fe}} \right) + \left(\frac{1+h_{fe}}{h_{fe}} \cdot (R_E + \frac{25 mV}{I_E}) \right)} \quad (1.10)$$

Which for $I_E = 10 mA$, $R_E = 10 \Omega$ and $h_{fe} = 100$ gives:

$$A = \frac{\sqrt{2}}{2} \cdot \frac{50 \Omega}{\left(\frac{50 \Omega}{100} \right) + \left(\frac{1+100}{100} \cdot (10 \Omega + \frac{25 mV}{10 mA}) \right)} = 8.6 dB \quad (1.11)$$

Input Impedance

For the amplifier to be matched to a 50Ω source impedance, it can be seen that the input impedance at the base of the transistor should also be 50Ω . Due to the 2:1 impedance transformation ratio of the tapped transformer, the impedance seen at the input to the amplifier will be equal to the impedance at the base of each of the transistors.

This impedance, Z_{in} can be calculated from the base voltage, i.e. the voltage across R_E and r_e , divided by the base current as follows:

$$\begin{aligned} Z_{in} &= \frac{i_e \cdot (R_E + r_e)}{i_b} \\ &= \frac{(h_{fe} + 1) \cdot i_b \cdot (R_E + r_e)}{i_b} \\ &= (h_{fe} + 1) \cdot (R_E + r_e) \end{aligned} \quad (1.12)$$

Which for $I_E = 10 \text{ mA}$, $R_E = 10 \Omega$ and $h_{fe} = 100$ gives:

$$Z_{in} = (100 + 1) \cdot (10\Omega + \frac{25 \text{ mV}}{10 \text{ mA}}) = 1262.5 \Omega \quad (1.13)$$

This suggests that the amplifier will be poorly matched at its input in a 50Ω system, however this model only considers real resistance and the impedance may be significantly affected by the internal capacitance of the transistor.

Maximum Available Output Power

The maximum available output power of the amplifier is limited by the the DC bias current through the emitter of each transistor, I_E . For a single transistor, the maximum output voltage occurs when no current is flowing through the emitter and the minimum occurs when all of the current ($I_E \cdot 2$) is flowing through it.

If we then consider the maximum power output signal to be a square wave set by the maximum output voltage swing into a 50Ω load, the amplitude of the fundamental frequency component can be calculated as:

$$V_{max} = \frac{4}{\pi} \cdot (I_E \cdot \frac{50 \Omega \cdot 50 \Omega}{50 \Omega + 50 \Omega}) \quad (1.14)$$

Which means the maximum available output power into a 50Ω load for $I_E = 10 \text{ mA}$ is:

$$P_{max} = \frac{V_{max}^2}{50 \Omega} = 1.56 \text{ dBm} \quad (1.15)$$

1.5.3 Measurements

Gain Compression

The gain compression of the differential amplifier with a bias voltage of -5 V was measured using a similar methodology to that for the MAR-6 amplifier in Section 1.4.2 and the results are shown in Figure 1.23.

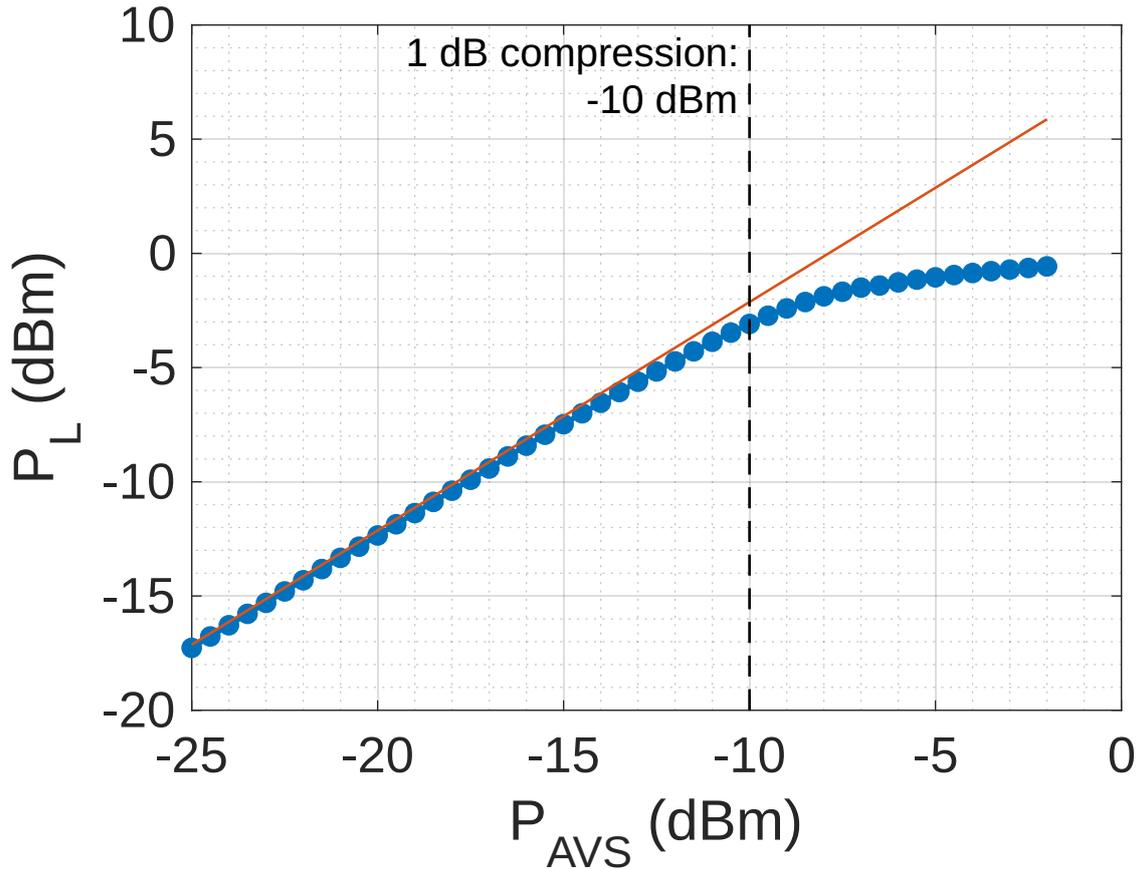


Figure 1.23: Gain compression of the differential amplifier with a bias voltage of -5 V measured using the Rohde & Schwarz FSWP

The small-signal gain was measured to be around 8 dB, which agrees with the value in Equation 1.11 that was calculated from the small-signal model. The maximum output power that was measured was around -0.5 dBm, which is approaching the maximum predicted by the model in Equation 1.15, but the input power was not increased any further as it can be seen that the amplifier is already saturating heavily at this point.

The 1 dB compression point occurs for a P_{AVS} of -10 dBm, which corresponds to an output power of around -3 dBm. This is slightly lower than that for the MAR-6 amplifier, which could limit minimum noise floor than can be achieved when placed in an oscillator, however this amplifier has a much lower excess gain and so does not require as much attenuation in the oscillator that would also degrade the noise

floor.

Input and Output Impedance

The input and output impedances of the differential amplifier were determined by measuring its S_{11} and S_{22} parameters at 100 MHz using the Anritsu 37377C Vector Network Analyzer and they are shown on the Smith chart in Figure 1.24. It can be seen that the previous assumption that the output impedance would be equal to the collector resistance was correct as the output impedance is almost exactly 50Ω .

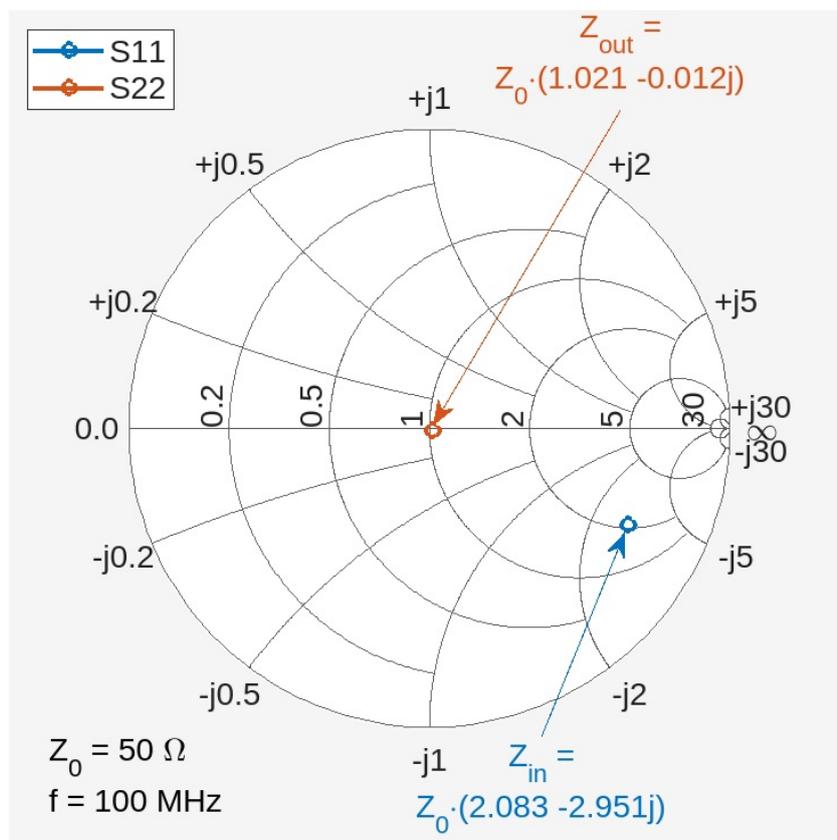


Figure 1.24: S-parameter measurements of the differential amplifier showing the input and output impedance at 100 MHz made using the Anritsu 37377C Vector Network Analyzer

Based on the simplified model outlined in Section 1.5.2 it was expected that the input impedance would be somewhat larger than 50Ω and this is confirmed by the

S_{11} measurement. The magnitude of the input impedance can be calculated to be $Z_0 \cdot \sqrt{2.083^2 + 2.951^2} = 180.6 \Omega$, which is less than that predicted by Equation 1.12, likely due to capacitive elements not being considered in the model. The value of S_{11} for optimum noise performance of a bipolar transistor, S_{opt} is not necessarily the same as that for optimum power transfer, and this is not provided in the datasheet for the BFR92P; so it is not clear whether this impedance mismatch will degrade the residual noise performance without further experimentation. This impedance mismatch, however could cause issues when the amplifier is placed in the oscillator loop, as is explored in Section 1.7.

Noise Figure

The noise figure of the differential amplifier was measured using the Rohde & Schwarz FSWP Phase Noise Analyzer and was found to be 3.1 dB.

Residual Noise

A residual noise measurement of the differential amplifier made using the cross-correlation system described in Chapter 3 is shown in Figure 1.25. The power at the amplifier output was set to be close to its 1 dB compression point, as this would be a reasonable point for oscillation to be able to be sustained without excessive saturation if the amplifier were placed in the oscillator loop.

The far out noise of the amplifier was measured to be -166 dBc/Hz, which is close to the theoretical noise floor from Section 1.2.2 based on its noise figure:

$$\begin{aligned} \text{NoiseFloor} &= -177 \text{ dB} - P_{AVS} + NF \\ &= -177 \text{ dB} + 8.9 \text{ dBm} + 3.1 \text{ dB} \\ &= -165 \text{ dBm} \end{aligned} \tag{1.16}$$

It is not possible to determine the close-to-carrier noise of the amplifier from this measurement, as the measured noise is not sufficiently above the internal noise floor of the measurement system. In the following section it will be shown how the flicker

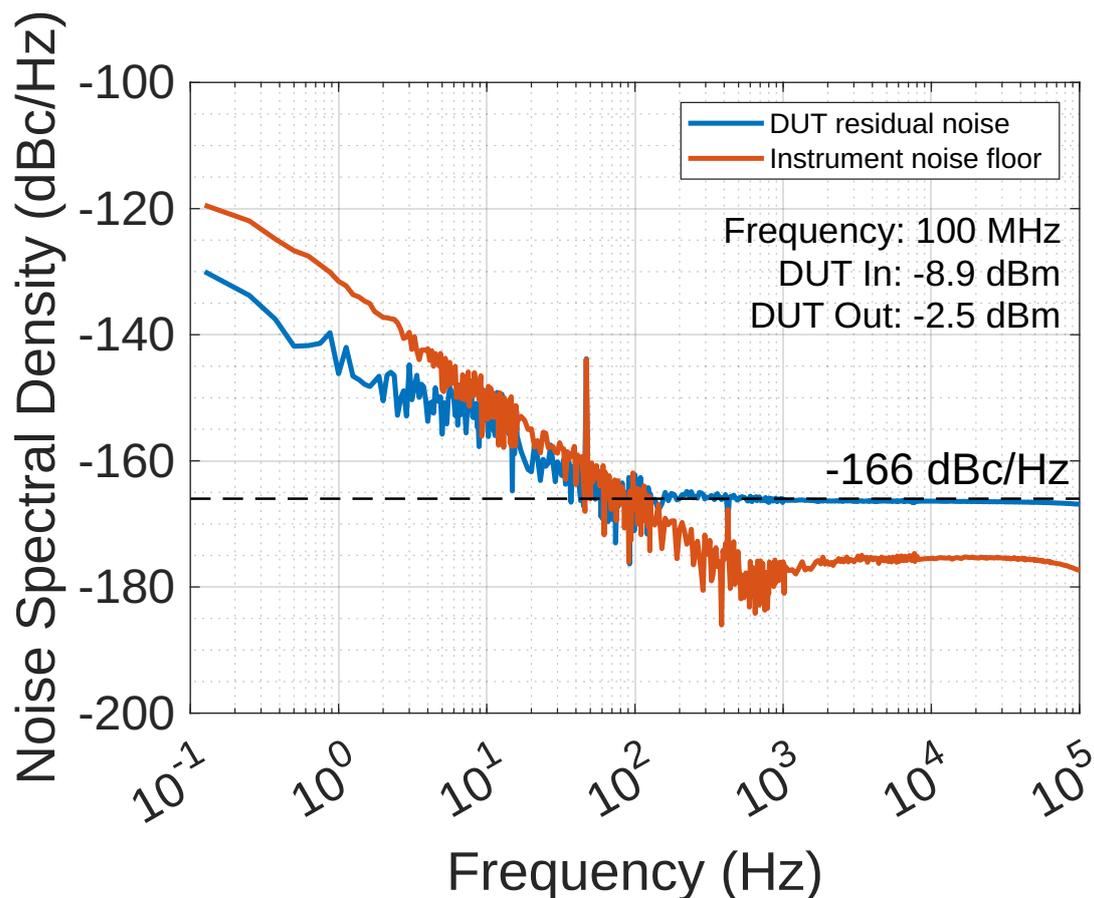


Figure 1.25: Residual phase noise of the differential amplifier measured using the cross-correlation system with 256 correlations at 0.125 Hz and 8,388,608 correlations at 100 kHz

noise corner can be determined from the absolute phase noise measurement of an LC oscillator incorporating the differential amplifier, and from this, the close-to-carrier residual noise can be estimated.

1.6 Flicker Noise Corner Measurement using an LC Resonator

In this section, a method for measuring the flicker noise corner of amplifiers whose close-to-carrier residual noise cannot be directly measured is presented, which in-

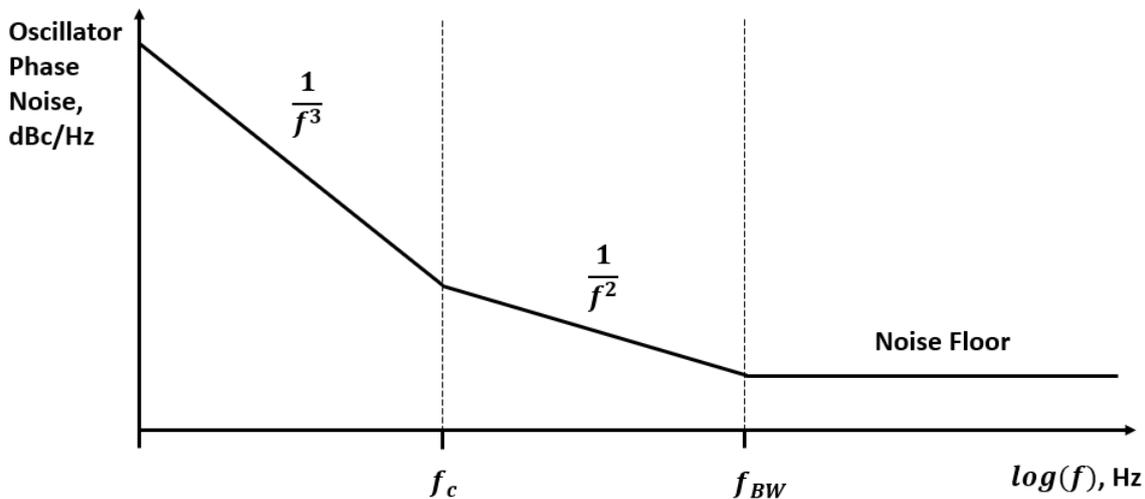
volves measuring the absolute phase noise of an LC oscillator that incorporates the amplifier. If the far-out noise floor of the amplifier is known, this can then be used to estimate the amplifier's residual phase noise.

1.6.1 Explanation

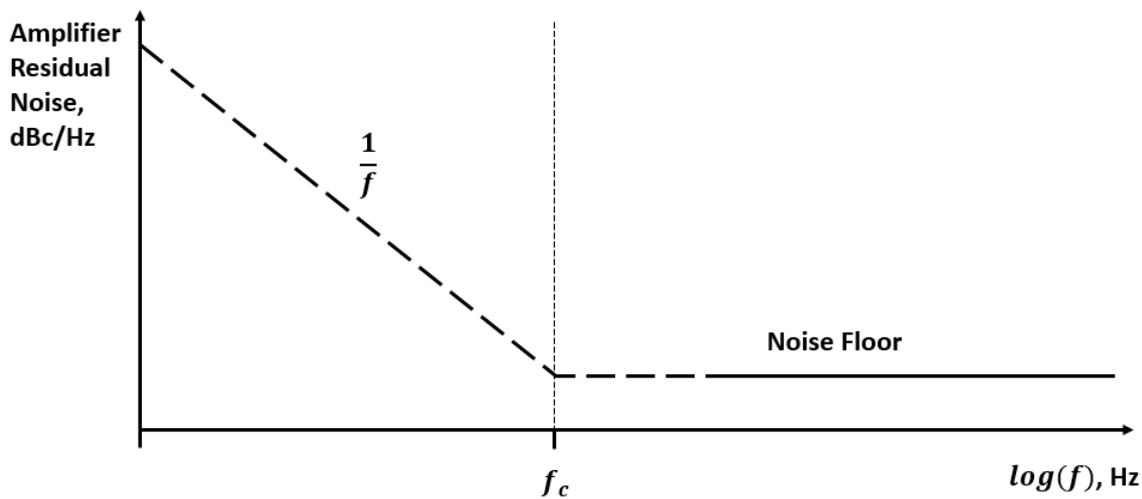
Oscillator Phase Noise Model

A model for oscillator phase noise, first described by Leeson [7], was presented in Section 1.2.3, where it can be seen that flicker noise corner of the amplifier corresponds to the $1/f^2$ to $1/f^3$ transition point in the absolute phase noise measurement.

It is shown in Figure 1.26 how an oscillator phase noise measurement can be used to estimate the close-to-carrier residual noise of an amplifier with a known noise floor. Figure 1.26 (a) shows how the flicker noise corner can be determined from the oscillator phase noise, and (b) shows how the close-to-carrier residual phase noise of the amplifier can then be extrapolated, which is represented by the dotted line.



(a) Absolute phase noise of an oscillator incorporating an amplifier with flicker noise corner, f_c



(b) Estimation of the close-to-carrier residual noise of the amplifier using f_c

Figure 1.26: Plots showing how the absolute phase noise measurement can be used to estimate the residual noise of an amplifier

LC Resonator

The model assumes that the amplifier is the dominant source of flicker noise in the oscillator loop, however it was suggested for the absolute phase noise measurement of the MAR-6 oscillator in Section 1.4.3 that the crystal resonator may be responsible for the higher than expected close-to-carrier phase noise.

By replacing the crystal with an LC resonator, any effect on the absolute phase noise due to the flicker noise of the crystal can be eliminated. Additionally, as the LC resonator would have a much lower Q, the transition point from flat white noise to $1/f$ noise in the oscillator bandwidth, f_{BW} , occurs at frequency that is likely to be much greater than the flicker noise corner of the amplifier, f_c . If these transition points were close together, it would be difficult to determine exactly where the $1/f^2$ to $1/f^3$ transition point, and hence the flicker noise corner, occurs.

The LC resonator was designed using the procedure outlined in Section 4.3 of Everard's book [10], which allows the resonator input and output to be matched to $50\ \Omega$. The schematic diagram of the resonator is shown in Figure 1.27 and includes the equivalent series resistance of the inductor, which is a CoilCraft Maxi Spring air core inductor with a Q of 114. The loaded Q of the resonator is set to be around 0.5 and the series resonant frequency is set to be as close as possible to nominal 100 MHz operating frequency of the crystal.

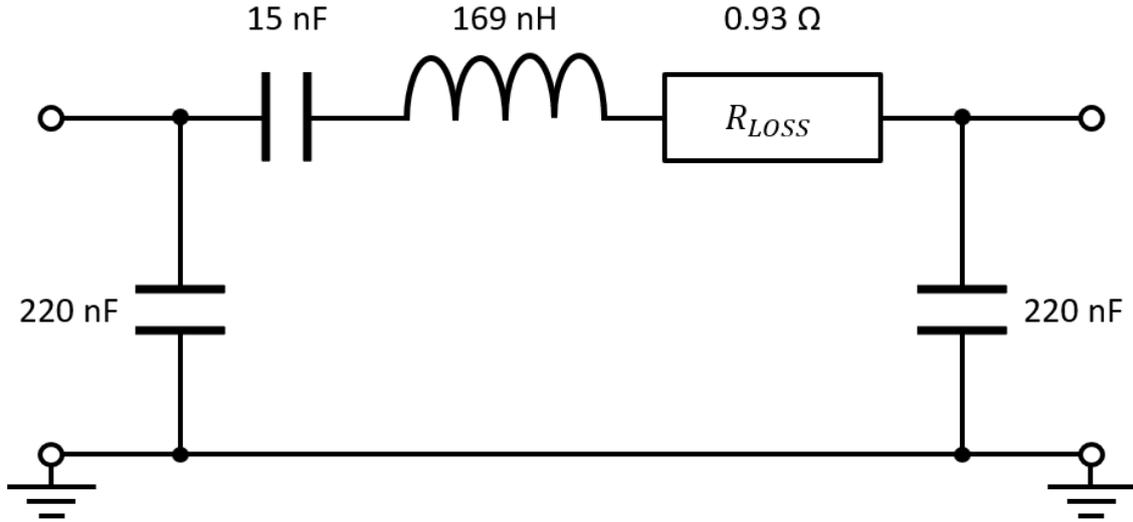


Figure 1.27: Schematic diagram of the LC resonator, including the equivalent series resistance of the inductor, that was used to replace the 100 MHz crystal for measuring the flicker noise corner of amplifiers

1.6.2 MAR-6 Measurement

To confirm the validity of this method, it was initially performed on an amplifier where the residual phase noise was already known, i.e. the MAR-6 as was shown in Section 1.4.2. This was done by simply replacing the crystal in the oscillator setup shown in Section 1.4.1 with the LC resonator, then adjusting the digital attenuator to give the same oscillator loop power as before.

The absolute phase noise of the MAR-6 LC oscillator is shown in Figure 1.28. The $1/f^2$ and $1/f^3$ have been superimposed onto the plot to show where they intersect, and hence where the flicker noise corner should be. The flicker noise corner measured using this method was around 12.6 Hz, which is somewhat close to the 20 Hz that was directly measured using the cross-correlation residual noise system.

The measurement is also compared with the theoretical phase noise that was calculated using the phase noise equation in Section 1.2.3. It can be seen that phase noise in the $1/f^2$ region is around 4 dB above the theory. Unfortunately, as the measurement frequency range is not large enough to see the noise floor of the oscillator, it is difficult to tell whether this discrepancy is due to the noise floor

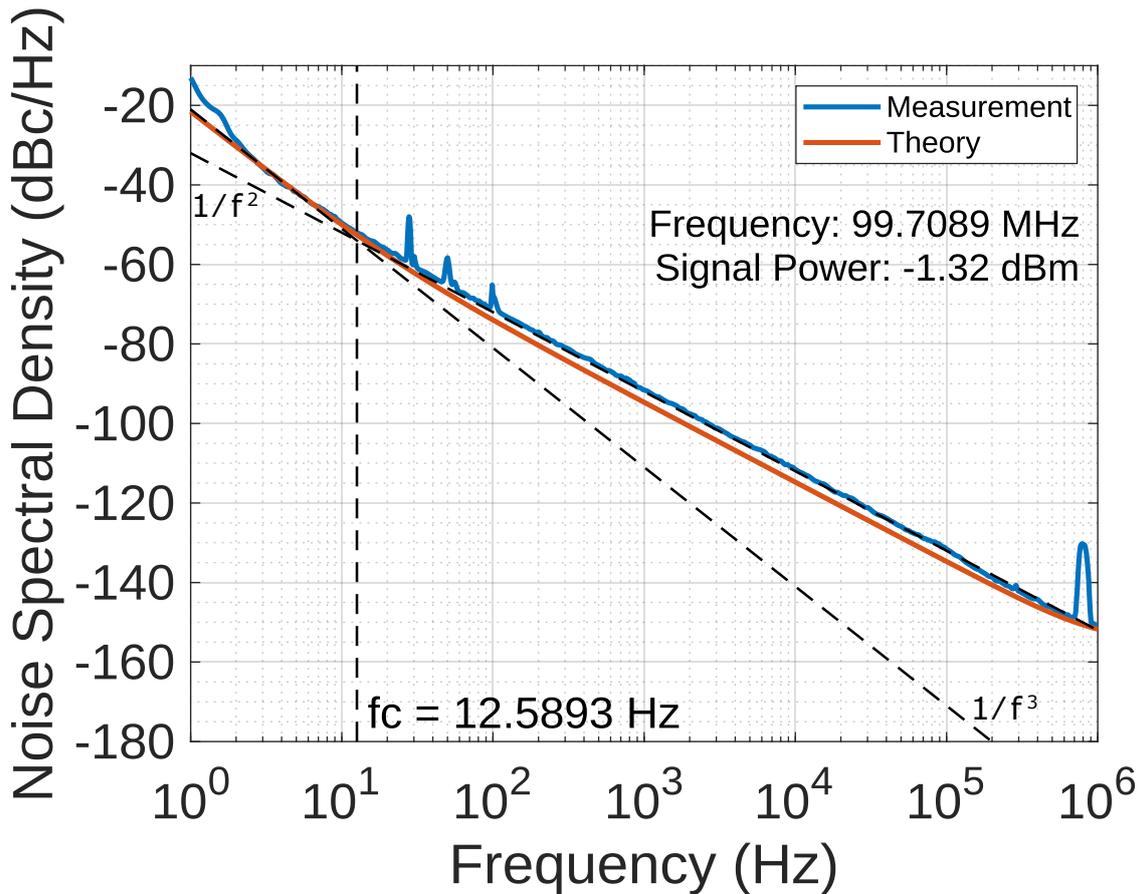


Figure 1.28: Absolute phase noise measurement of the 100 MHz LC oscillator with MAR-6 amplifier and 2.8 dB of attenuation after the amplifier made using the Rohde & Schwarz FSWP

being higher than expected or a deviation in the the flat white noise to $1/f$ transition frequency.

1.6.3 Differential Amplifier Measurement

Since a reasonable measurement was obtained for the flicker noise corner of the MAR-6 amplifier, it was decided to use this method to try to determine a rough value for the flicker noise corner of the differential amplifier, which could not be measured directly. The same methodology was used as for the MAR-6 amplifier and the oscillator was configured as per the block diagram in Figure 1.29.

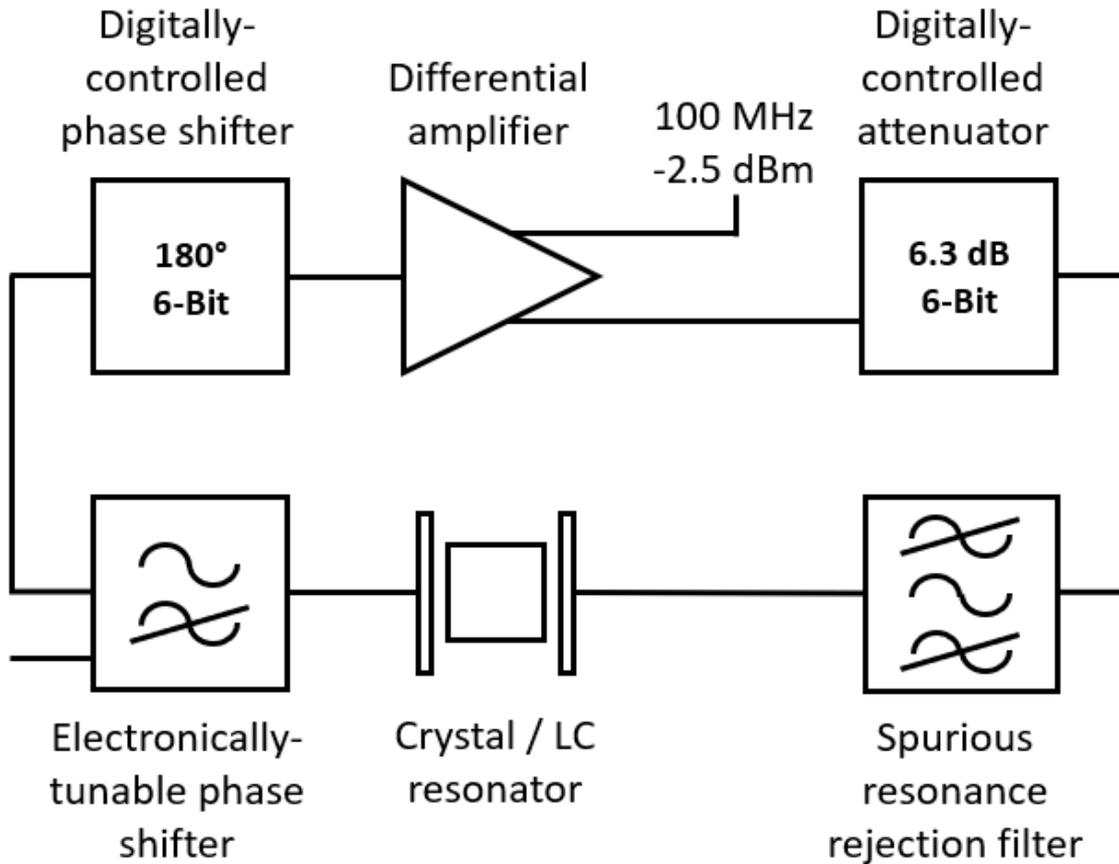


Figure 1.29: Block diagram of the initial 100 MHz oscillator setup that was used to measure the flicker noise of the differential amplifier

A plot of the measured absolute phase noise is shown in Figure 1.30. The theoretical phase noise is also shown in this plot, however since the actual flicker noise corner of the amplifier was not known, the $1/f^3$ part of the curve has been ignored.

It can be seen that there is a large discrepancy of around 10 dB between the measurement and theory in the $1/f^2$ part of the curve. Additionally, if the flicker noise corner of the amplifier were as high as 316 Hz as suggested by this measurement, then it can be seen from the residual noise measurement made in Section 1.5.3 that it would have been possible to measure this directly.

In the following section, it is suggested that the reason for this discrepancy is the impedance mismatch between the resonator and the input of the amplifier and it is shown how this can be addressed to bring these measurements closer to both

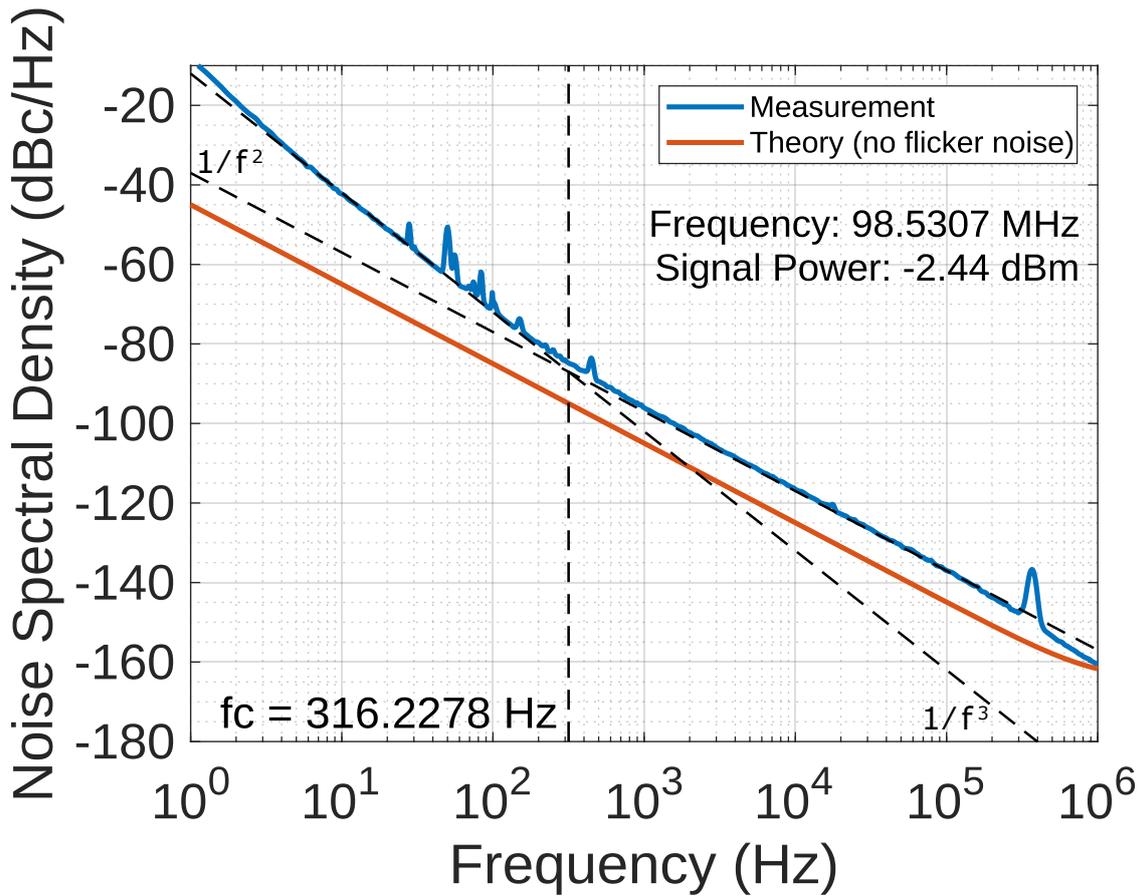


Figure 1.30: Absolute phase noise measurement of the 100 MHz LC oscillator with differential amplifier and 0.2 dB of attenuation after the amplifier made using the Rohde & Schwarz FSWP

the theory and the known characteristics of the amplifiers under test.

1.7 Improving Oscillator Phase Noise Using Attenuator Matching

In this section it is shown how a relatively simple change to the oscillator configuration, placing the attenuator between the resonator and the input to the amplifier, can result in significantly improved phase noise performance. It is also shown more generally how a resistive attenuator can be used to set the effective source and load impedances when placed between two devices to appear closer to the desired system impedance.

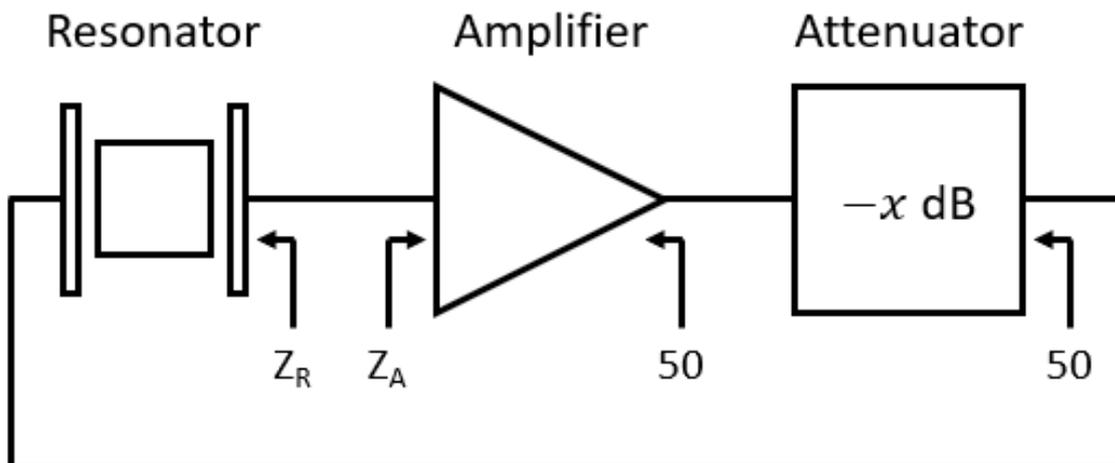
1.7.1 Theory

Source and Load Impedances

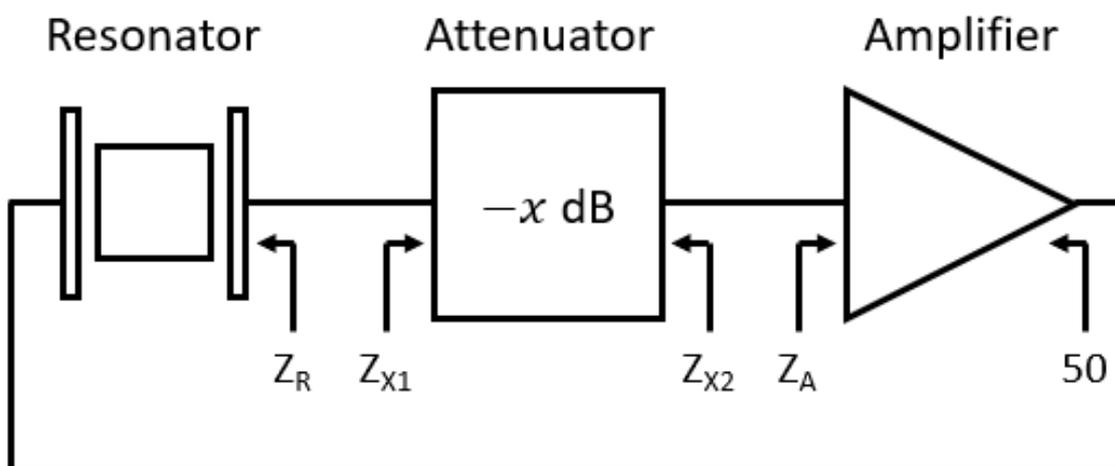
Figure 1.31 (a) shows a simplified block diagram of an oscillator, along with input and output impedances of the various elements, where an attenuator matched to $50\ \Omega$ is placed between the output of the amplifier and the resonator. This is similar to the configuration used for the oscillators in Sections 1.4, 1.5 and 1.6. The output impedance of the amplifier is assumed to be $50\ \Omega$, as is the case for the differential and MAR-6 amplifiers, and so the output impedance of the attenuator can also be assumed to be $50\ \Omega$. The source impedance presented to the amplifier is the output impedance of the resonator, Z_R , and the load impedance presented to the resonator is the input impedance of the amplifier, Z_A .

If Z_R deviates from $50\ \Omega$, it could degrade the phase noise performance of the amplifier if its flicker noise is sensitive to the source impedance presented to it. If Z_A deviates from $50\ \Omega$, it could affect the loaded Q of the resonator, degrading the phase noise.

Figure 1.31 (b) shows an alternative configuration where the attenuator is instead placed between the resonator and the input to the amplifier. Here, the source impedance presented to the amplifier is the output impedance of the attenuator, Z_{X2} and the load impedance presented to the resonator is the input impedance to



(a) Post-amplifier attenuation



(b) Pre-amplifier attenuation

Figure 1.31: Simplified block diagram of the oscillator to demonstrate the effects of pre-amplifier and post-amplifier attenuation

the attenuator, Z_{X1} . In the following sections, it will be shown that as the level of attenuation increases, the impedances Z_{X1} and Z_{X2} will become closer to 50Ω .

Resistive Attenuator Impedance Model

In Section 2.3, which presents a digitally-controlled attenuator, a resistive attenuator circuit that uses pi-pad topology is described. This circuit can be used as a model to show how such an attenuator transforms the source and load impedance presented to

a pair of devices when placed between them, as shown in Figure 1.32. Relating this to the oscillator block diagram in Figure 1.31, the resistors R_s and R_L correspond to the output impedance of the resonator, Z_R , and the input impedance of the differential amplifier, Z_A , respectively.

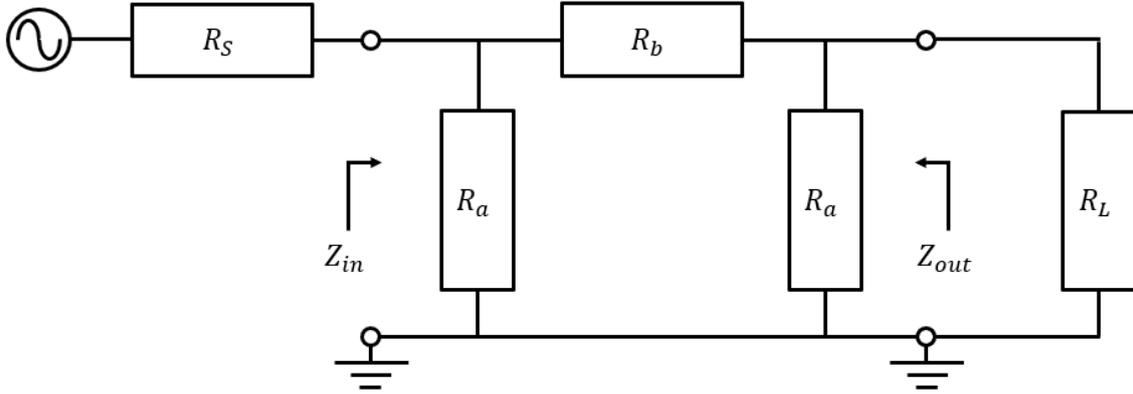


Figure 1.32: Equivalent circuit of a resistive attenuator using pi-pad topology placed between source and load impedances

As given in Section 2.3.1, the equations for calculating the pi-pad resistor values for a specified attenuation, $K = \frac{V_{out}}{V_{in}}$, are:

$$R_a = Z_0 \cdot \frac{1 + K}{1 - K} \quad (1.17)$$

$$R_b = Z_0 \cdot \frac{1 - K^2}{2 \cdot K} \quad (1.18)$$

The effective input impedance, Z_{in} , for a pi-pad attenuator terminated with a resistance R_L can then be determined from the resulting resistor network as follows:

$$\begin{aligned} Z_{in} &= R_a \parallel (R_b + (R_a \parallel R_L)) \\ &= \frac{R_a \cdot (R_b + (\frac{R_a \cdot R_L}{R_a + R_L}))}{R_a + (R_b + (\frac{R_a \cdot R_L}{R_a + R_L}))} \end{aligned} \quad (1.19)$$

And since the circuit is symmetrical, Z_{out} , the source impedance presented to the load resistor, is:

$$Z_{out} = \frac{R_a \cdot (R_b + (\frac{R_a \cdot R_S}{R_a + R_S}))}{R_a + (R_b + (\frac{R_a \cdot R_S}{R_a + R_S}))} \quad (1.20)$$

The potential for improving the source and load impedances using an attenuator can be seen by considering the worst case scenarios for achieving a 50Ω match at the attenuator input, i.e. when the attenuator is terminated in an open or short-circuit. Figure 1.33 shows the load impedance presented to a 50Ω source versus the nominal attenuation in both cases calculated using (1.19). It can be seen that Z_{in} converges towards 50Ω as the attenuation increases and that for an attenuation of just 10 dB, the effective input impedance is within around 20% of 50Ω .

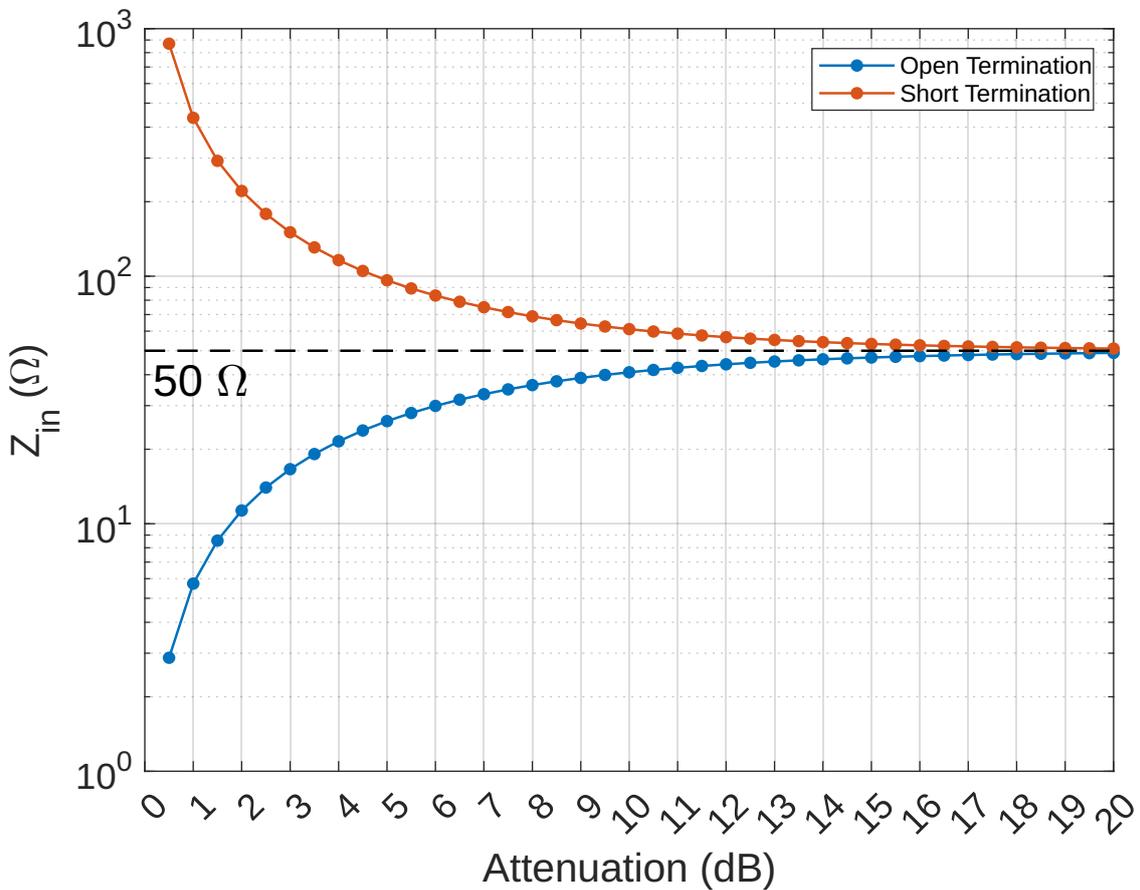


Figure 1.33: Plot of the the effective input impedance of a resistive attenuator versus the nominal attenuation when the output of the attenuator is terminated in either an open or short-circuit

Theoretical Improvement

Using the equivalent circuit model of an oscillator from Everard's book [10], if the loaded Q of the resonator is designed for $Q_L/Q_0 = 0.5$ in a 50Ω system ($R_{in} = R_{out} = 50 \Omega$):

$$\frac{Q_L}{Q_0} = \frac{R_{loss}}{R_{in} + R_{out} + R_{loss}} = 0.5 \quad (1.21)$$

$$R_{loss} = 100 \Omega$$

If the load impedance presented to the resonator is an impedance other than 50Ω , such as the input impedance of the amplifier, Z_A , the loaded Q is then:

$$Q_L = Q_0 \cdot \left(\frac{R_{loss}}{R_{in} + Z_A + R_{loss}} \right) \quad (1.22)$$

$$= 100 \cdot \left(\frac{100 \Omega}{50 \Omega + Z_A + 100 \Omega} \right)$$

If we compare this to the loaded Q when an attenuator is placed between the resonator and the amplifier, $Q_{L(att)}$:

$$Q_{L(att)} = 100 \cdot \left(\frac{100 \Omega}{50 \Omega + Z_{in} + 100 \Omega} \right) \quad (1.23)$$

The f_{BW} transition frequency in Figure 1.26, is set by the loaded Q of the resonator. And so for each case is:

$$f_{BW} = \frac{f_0}{2 \cdot Q_L} \quad (1.24)$$

$$f_{BWatt} = \frac{f_0}{2 \cdot Q_{L(att)}}$$

From which we can also see the relationship:

$$\frac{f_{BW(att)}}{f_{BW}} = \frac{Q_L}{Q_{L(att)}} \quad (1.25)$$

$$= \frac{Q_0 \cdot \left(\frac{R_{loss}}{R_{in} + Z_A + R_{loss}} \right)}{Q_0 \cdot \left(\frac{R_{loss}}{R_{in} + Z_{in} + R_{loss}} \right)}$$

$$= \frac{Z_{in} + 150 \Omega}{Z_A + 150 \Omega}$$

Since the the $1/f^2$ region of the phase noise model in Figure 1.26 (a) decreases at a rate of 20 dB per decade, we can therefore estimate the improvement in phase noise performance in the $1/f^2$ region that results from the change in the transition frequency, f_{BW} , when including the attenuator using the following equation:

$$\begin{aligned} &= 20 \cdot \log_{10}\left(\frac{f_{BW(att)}}{f_{BW}}\right) = 20 \cdot \log_{10}\left(\frac{Q_L}{Q_{L(att)}}\right) \\ &= 20 \cdot \log_{10}\left(\frac{Z_{in} + 150 \Omega}{Z_A + 150 \Omega}\right) \end{aligned} \quad (1.26)$$

1.7.2 Improved LC Oscillator Measurements

Here, the LC oscillator phase noise measurements when using the MAR-6 and differential amplifiers are compared for the different oscillator configurations; the original post-amplifier placement of the attenuator and the pre-amplifier placement, which is designed to improve the source and load impedances presented to the resonator and amplifier.

MAR-6 Amplifier

In Section 1.4.2 a measurement of the phase noise of the LC oscillator with the MAR-6 amplifier using the original post-amplifier attenuator configuration was performed. It was seen that the measured phase noise was close to the calculated theoretical phase noise with only a small discrepancy.

Figure 1.34 shows the phase noise using the pre-amplifier configuration, i.e. where the digitally-controlled attenuator has been placed between the resonator and the amplifier input. The amount of attenuation has also been adjusted so that oscillator loop power is similar to that for the post-amplifier measurement in Figure 1.28. The results show that the phase noise in the $1/f^2$ region has improved by around 3 dB and is very close to the calculated theoretical phase noise. The flicker noise corner derived from this measurement is also very close to the 20 Hz value that was measured using the cross-correlation residual noise system.

Although a particularly large improvement was not expected in the case of the

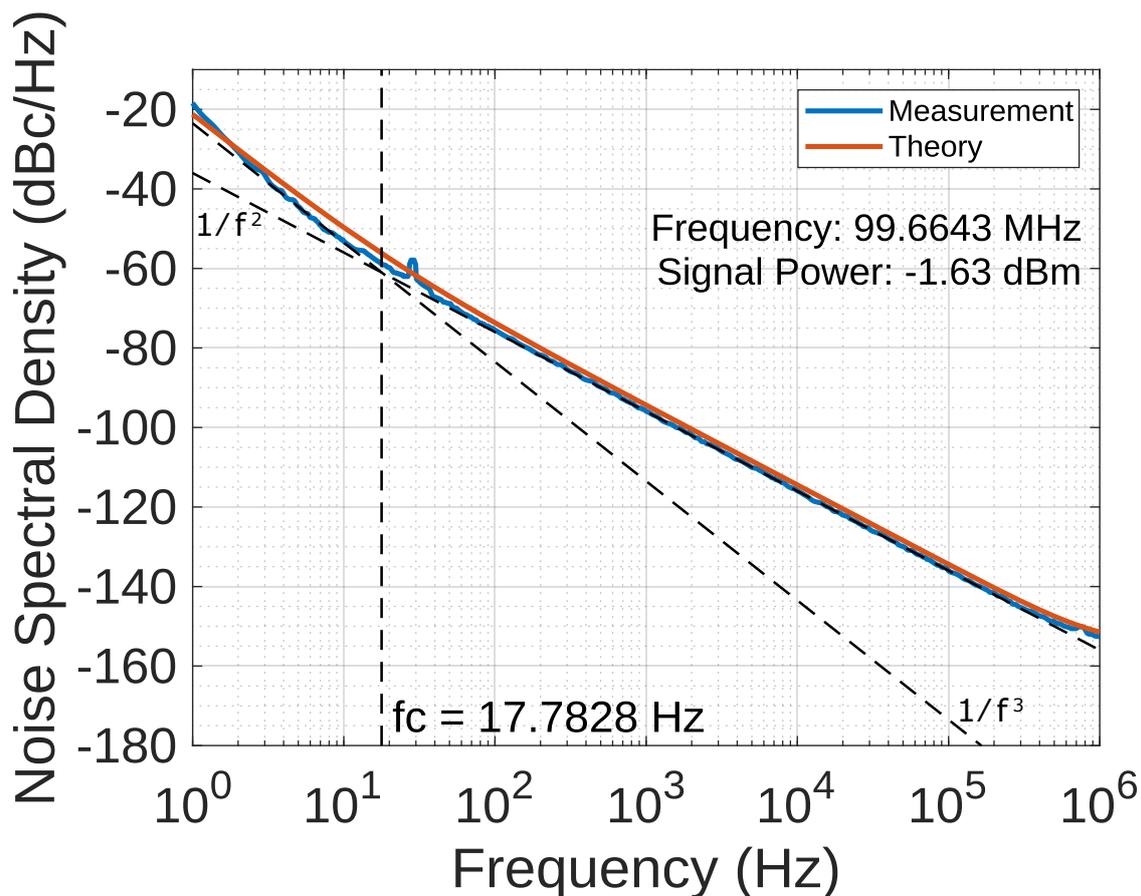


Figure 1.34: Absolute phase noise measurement of the 100 MHz LC oscillator with MAR-6 amplifier and 2.5 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP

MAR-6 amplifier, as the input to this amplifier is already internally matched to $50\ \Omega$, it appears that changing the LC oscillator to the pre-amplifier attenuator configuration has brought the measurement in line with the theory. This could be due to the MAR-6 amplifier being somewhat sensitive to the source impedance presented to it.

Differential Amplifier

Figure 1.35 shows the differential amplifier LC oscillator phase noise using the original post-amplifier configuration. Compared to the similar measurement that was

previously made in Section 1.5.3, some extra attenuation has been added, so the attenuator can appear closer to 50Ω . It can be seen that this measurement is already much improved over that in Figure 1.30, which suggests that the discrepancy in the original measurement was likely mostly due to the amplifier operating too far into saturation, degrading the phase noise. Even so, there is still a discrepancy in Figure 1.35 of around 3 dB from the theory in the $1/f^2$ region.

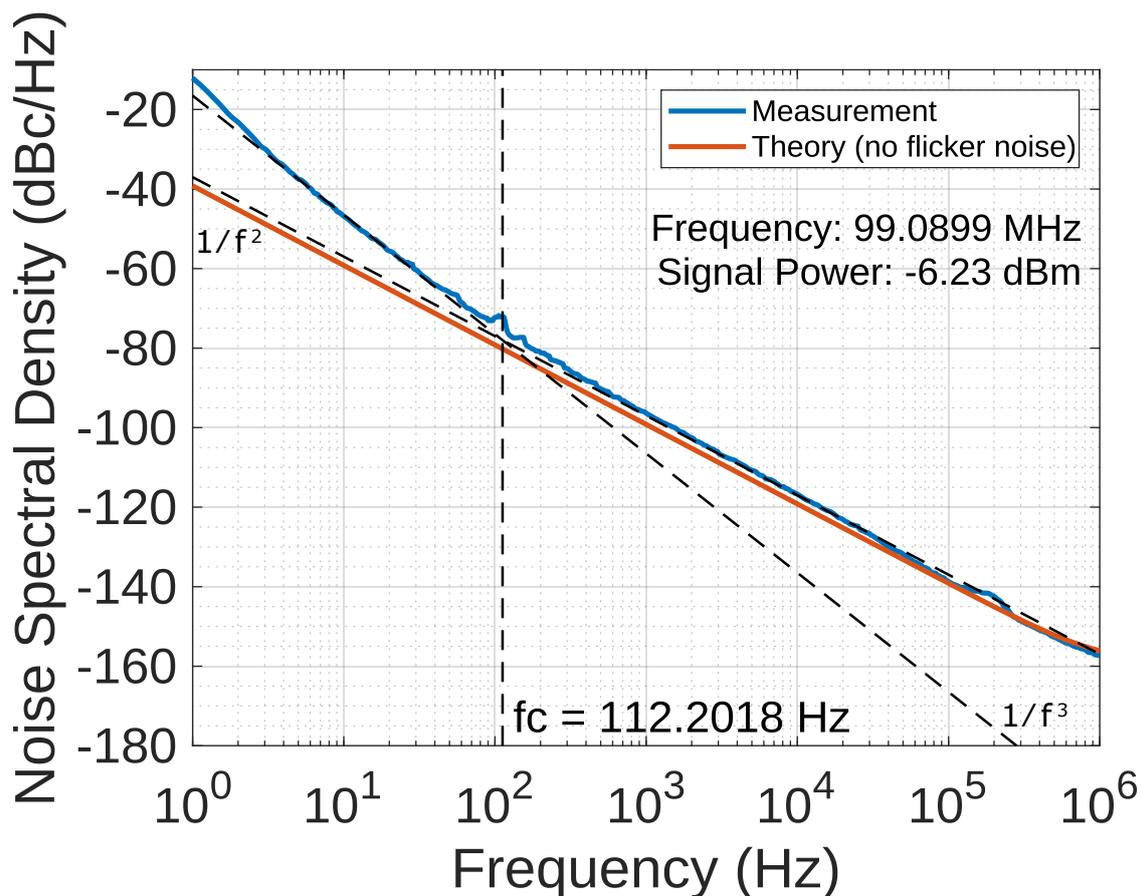


Figure 1.35: Absolute phase noise measurement of the 100 MHz LC oscillator with differential amplifier and 2.4 dB of post-amplifier attenuation made using the Rohde & Schwarz FSWP

Figure 1.36 shows the phase noise using the pre-amplifier configuration, as was done for the MAR-6 amplifier. It can be seen that the $1/f^2$ slope now matches closely to the theory and is improved over the phase noise for the post-amplifier

configuration by around 3 dB, even for a reasonably small 1.8 dB attenuation. The measured flicker noise corner is also significantly improved, and considering both that the rest of the measurement matches closely to the theory and that the flicker noise was successfully measured in this manner for the MAR-6 amplifier in Figure in 1.34, it gives confidence that 40 Hz is a good estimate for the flicker noise corner of the differential amplifier.

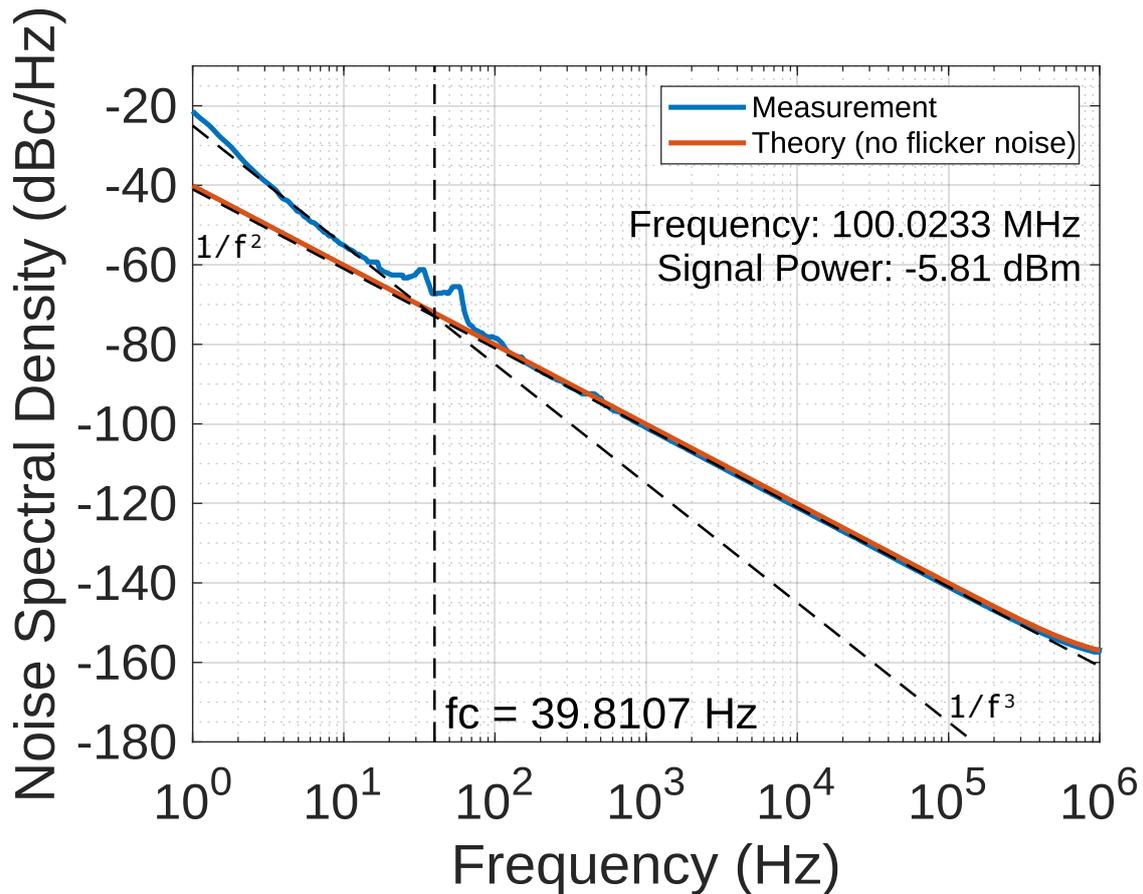


Figure 1.36: Absolute phase noise measurement of the 100 MHz LC oscillator with differential amplifier and 1.8 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP

1.7.3 Crystal Oscillator Measurements

Since an improvement was seen in the phase noise of the LC oscillators when switching to the pre-amplifier attenuator configuration, the same measurements were performed for the crystal oscillators to see if a similar improvement to the absolute phase noise could be achieved.

MAR-6 Amplifier

The phase noise of the oscillator with the MAR-6 amplifier was measured using the same operating conditions and pre-amplifier attenuator configuration as for Figure 1.34; the only difference being that the LC resonator was swapped out for the 100 MHz crystal. It can be seen that in this case there is a large discrepancy in the theoretical and measured close-to-carrier phase noise and there is no improvement over the original post-amplifier attenuator measurement shown in Section 1.4.3.

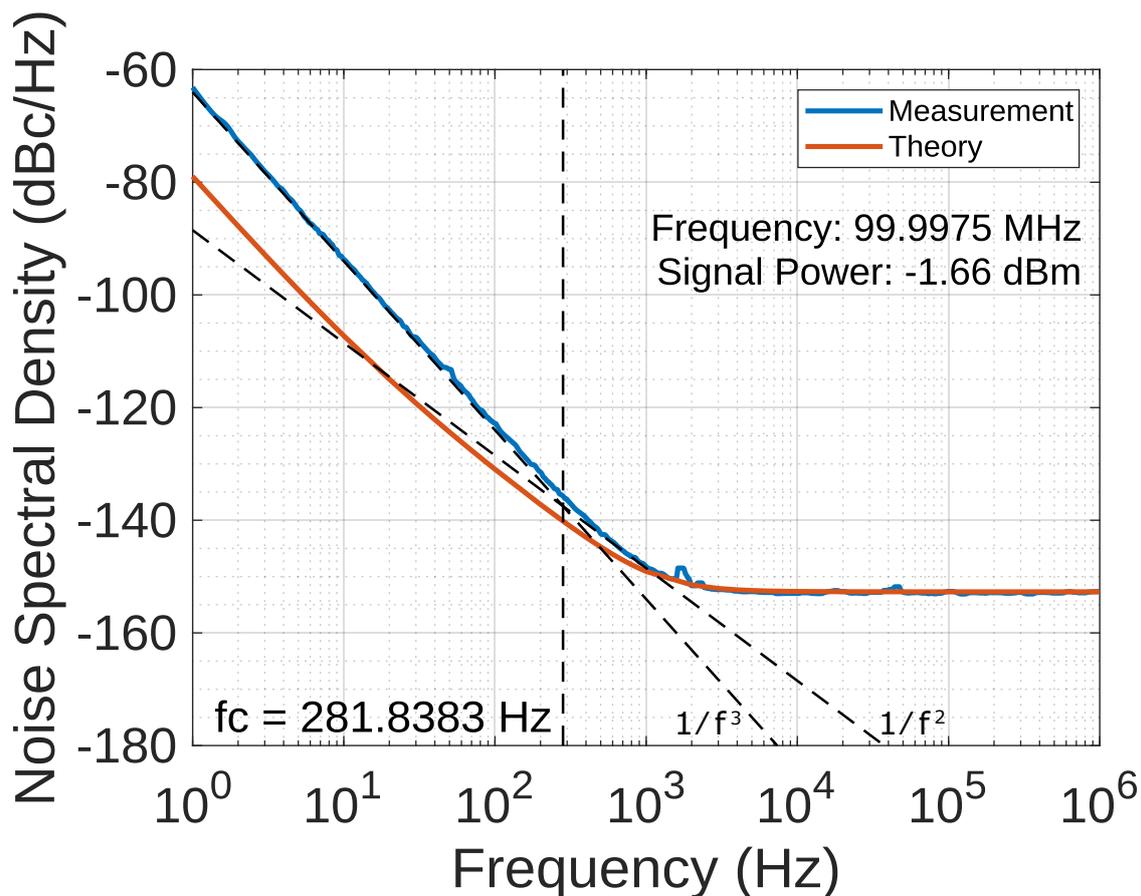


Figure 1.37: Absolute phase noise measurement of the 100 MHz crystal oscillator with MAR-6 amplifier and 5.4 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP

Differential Amplifier

A similar measurement was made for the differential amplifier crystal oscillator, which is shown in Figure 1.38. Unfortunately, any improvement that can be gained from the pre-amplifier attenuator configuration in this case will be likely reduced, as in this case the level of attenuation had to be relatively small to achieve the loop power level as that for LC oscillator measurement in Figure 1.36. It can be seen that, as for the MAR-6 oscillator, there is a large discrepancy in the close-to-carrier phase noise. In fact, the phase noise measured at offset frequencies below 100 Hz is nearly identical to that for the MAR-6 in Figure 1.37.

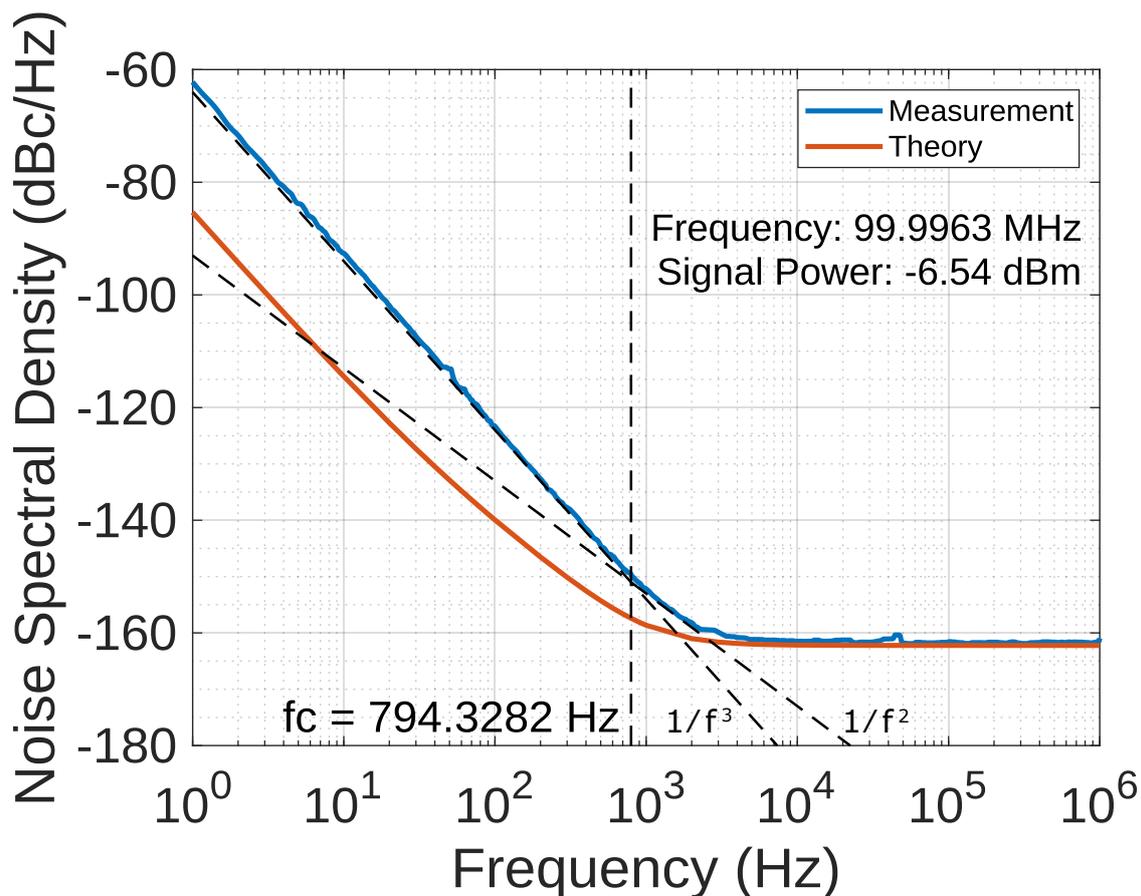


Figure 1.38: Absolute phase noise measurement of the 100 MHz crystal oscillator with differential amplifier and 0.8 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP

1.8 Residual Noise of Crystal Resonator

In this section, it is shown how it was determined that the discrepancies between the measured phase noise of the MAR-6 and differential amplifier crystal oscillators and the theoretical phase noise was likely caused by the residual phase noise of the original quartz crystal resonator that was used. The mechanism by which the crystal resonator can introduce phase noise into the oscillator has been shown in [17]. Measurements using a replacement KVG crystal resonator are shown to greatly improve the phase noise in both cases.

1.8.1 Comparison of Crystal Oscillator Measurements

Looking at the measurements made of the oscillators using the original resonator, a Nofech SC-cut 5th order quartz crystal with a Q of around 110,000, in Section 1.7.3, it can be seen that both measurements exhibited a large discrepancy from the theoretical phase noise close-to-carrier. This was despite the fact that it was shown in Section 1.7.2 that the pre-amplifier attenuator configuration could be used to bring the measurements in close agreement with the theory for the LC oscillator.

Additionally, it was noted that the close-to-carrier phase noise of both the MAR-6 and differential amplifier crystal oscillators was almost identical. This suggests that the source of the apparent increase in the flicker noise corner is due to something other than the amplifier.

It can be concluded from these observations that the oscillator phase noise is most likely being limited by the residual noise of the crystal resonator.

1.8.2 Replacement Crystal

To confirm suspicions that the crystal resonator was the cause of the limitation in phase noise performance, a replacement crystal was found in the form of the KVG O5SC85105A [8]. This crystal was chosen as it specifies a maximum phase noise of -130 dBc/Hz at a 100 Hz, which is lower than that measured with the original crystal. It also has a relatively high recommended drive level of 5 mW (7 dBm) that will allow for a lower noise floor to be achieved with an appropriate amplifier.

KVG Crystal Oscillator Measurements

The crystal oscillator measurements described in Section 1.7.3 were simply repeated after replacing the crystal resonator to see if this improved the phase noise.

The phase noise measurement of differential amplifier oscillator with the KVG crystal is shown in Figure 1.39. This appears to confirm that the crystal was the cause of the discrepancy in previous measurements, as with the KVG crystal, the

measurement is now very close to the theoretical phase noise with only a small discrepancy in the close-to-carrier phase noise.

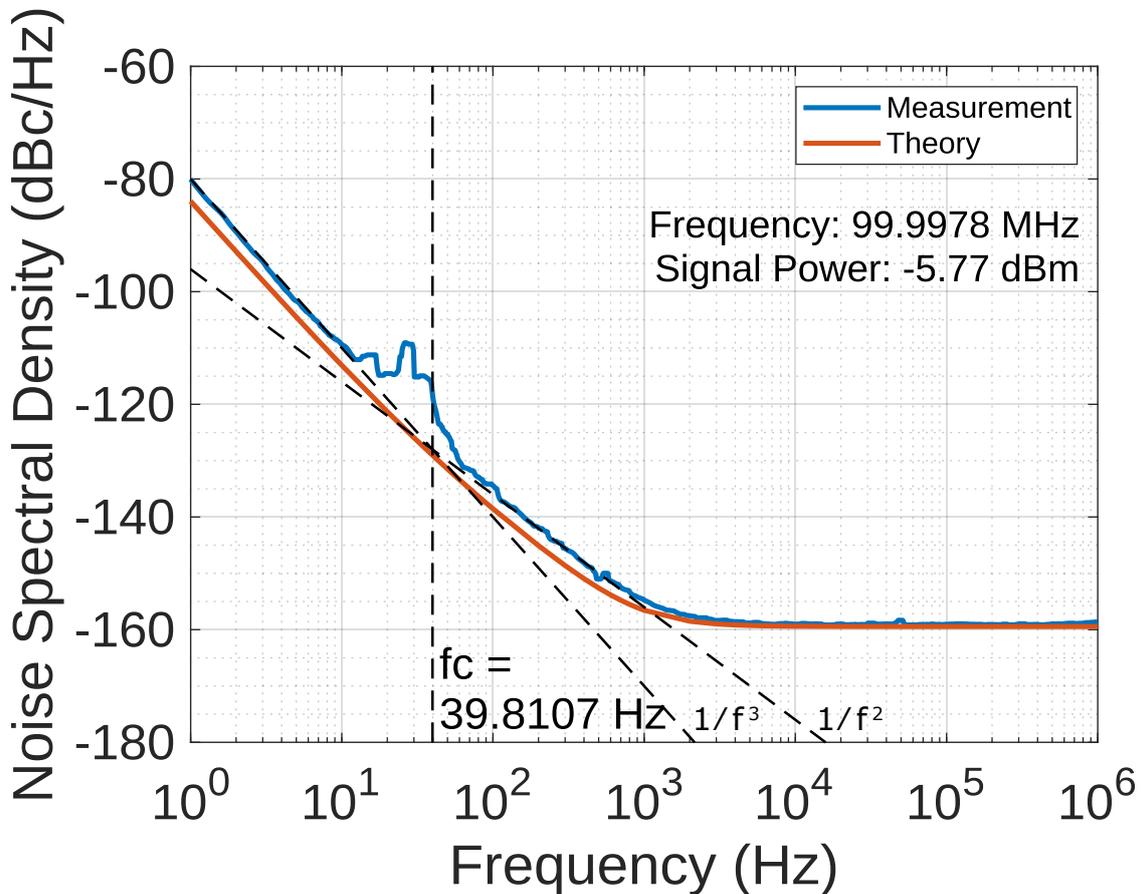


Figure 1.39: Absolute phase noise measurement of the 100 MHz KVG crystal oscillator with differential amplifier and 3.1 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP

The phase noise measurement of the MAR-6 oscillator with the KVG crystal in Figure 1.40 shows that it now matches up almost exactly with the theoretical phase noise.

Even though the crystal oscillator phase noise measurements now line up with the theory, it can be seen that the target of less than -140 dBc/Hz at a 100 Hz offset has still not been met. It is the case for both the amplifiers tested that the maximum oscillator loop power that can be achieved without degrading the flicker

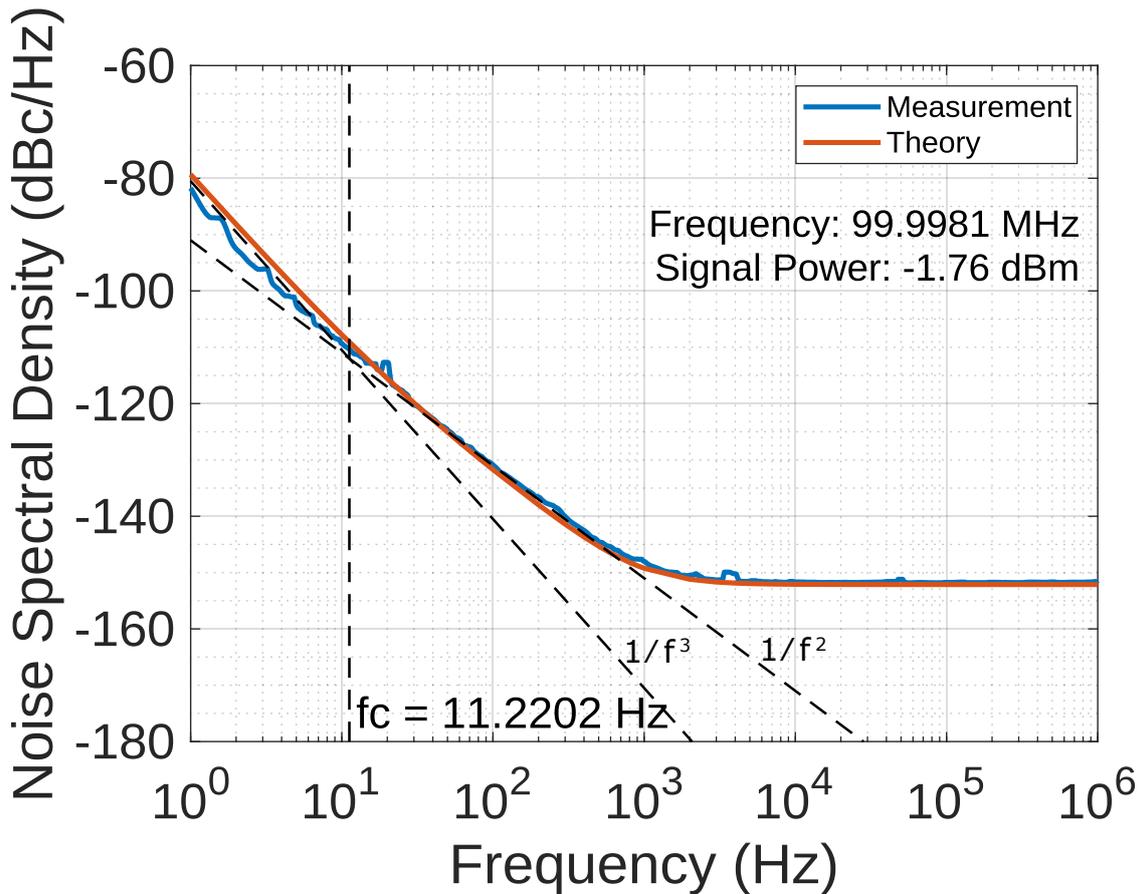


Figure 1.40: Absolute phase noise measurement of the 100 MHz KVG crystal oscillator with MAR-6 amplifier and 5.4 dB of pre-amplifier attenuation made using the Rohde & Schwarz FSWP

noise is considerably lower than the recommended drive level of the crystal, around 7 dBm. In the following section, an alternative amplifier design is explored that can be used to address this issue.

1.9 Quad Parallel Differential Amplifier

In this section, the final design of the amplifier used in the 100 MHz crystal oscillator is described, which uses a quad parallel configuration of the differential amplifier described in Section 1.5. It shown that this configuration can be used to increase the maximum available output power of the amplifier without significantly increasing

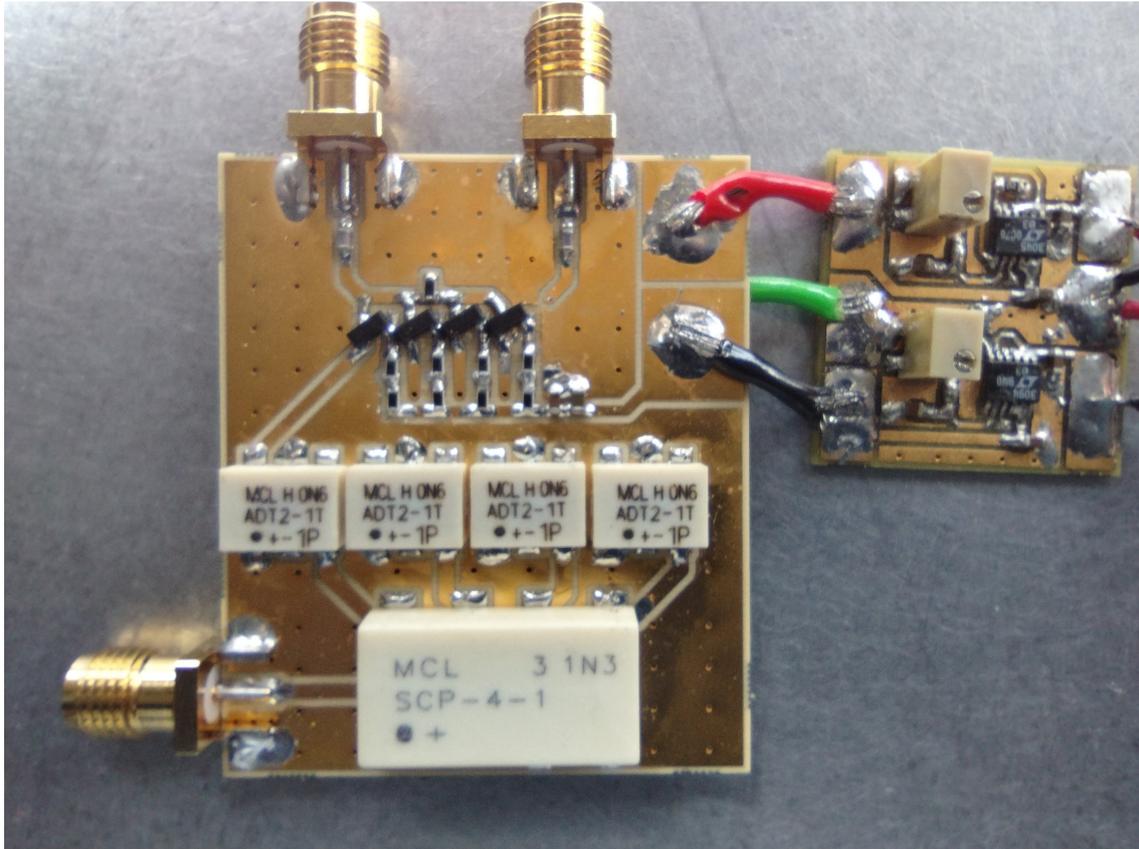


Figure 1.41: Photo of quad differential amplifier circuit

the gain or making the amplifier more sensitive to changes in the device parameters of the transistors. With a higher maximum output power, the oscillator loop power can be increased, and if the residual noise of the amplifier is not significantly degraded, the absolute phase noise of the crystal oscillator can be improved.

1.9.1 Theory of Operation

A block diagram of the operation of the quad parallel differential amplifier is shown in Figure 1.42. A 4-way power splitter is used to split the input signal between 4 identical differential amplifier circuits. These circuits use the same topology as for the differential amplifier in Figure 1.20, except that the collector resistors are removed, so that each output can now be modelled as a current source equal to the collector current in the transistor. Each differential amplifier provides complemen-

tary outputs as before, but only a single output is considered in this diagram. To retrieve a single output, the collectors of the corresponding transistors in each differential amplifier are connected in parallel, effectively summing the currents together to produce a voltage across a shared $50\ \Omega$ resistor.

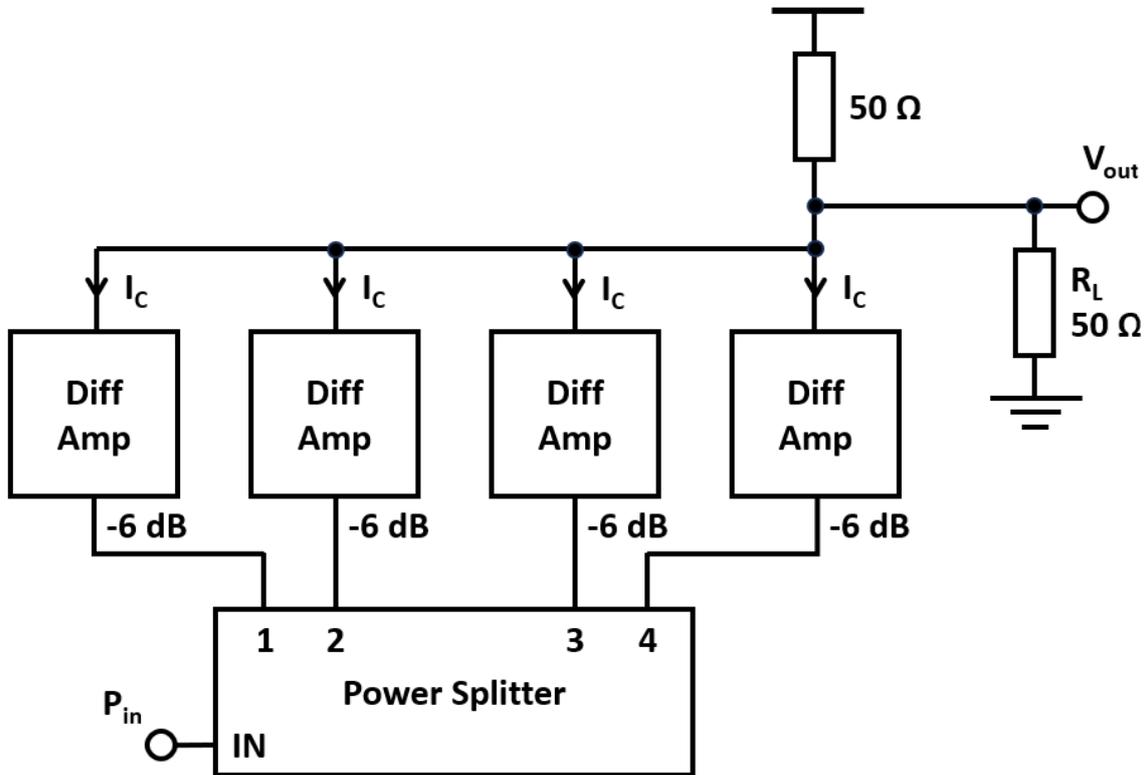


Figure 1.42: Block diagram of the quad parallel differential amplifier

Maximum Output Power

An equation for calculating the maximum output voltage swing of the differential amplifier in terms of the emitter bias current was shown in (1.15). By considering that for the quad differential amplifier the current into the load resistor has been effectively multiplied by 4, (1.15) can easily be modified to give the maximum output voltage swing of the quad differential amplifier in terms of the emitter bias current, I_E :

$$V_{Qmax} = \frac{4}{\pi} \cdot (4 \cdot I_E \cdot \frac{50 \Omega \cdot 50 \Omega}{50 \Omega + 50 \Omega}) \quad (1.27)$$

And hence the maximum available output power:

$$P_{Qmax} = \frac{V_{Qmax}^2}{50 \Omega} \quad (1.28)$$

Gain

For the same reasons as above, the ratio of the output voltage of the quad differential to the input voltage at one of the differential amplifier subcircuits (i.e. after the power splitter) is a factor of 4 times that for the single differential amplifier given in (1.10). Considering also that signal into each single amplifier has been reduced by a factor of 2 (6 dB) compared to the overall input signal, (1.10) can be modified to give:

$$\begin{aligned} A_Q &= \frac{V_{out}}{V_{in}} = \frac{4}{2} \cdot \frac{\sqrt{2}}{2} \cdot \frac{50 \Omega}{\left(\frac{50 \Omega}{h_{fe}}\right) + \left(\frac{1+h_{fe}}{h_{fe}} \cdot \left(R_E + \frac{25 mV}{I_E}\right)\right)} \\ &= \sqrt{2} \cdot \frac{50 \Omega}{\left(\frac{50 \Omega}{h_{fe}}\right) + \left(\frac{1+h_{fe}}{h_{fe}} \cdot \left(R_E + \frac{25 mV}{I_E}\right)\right)} \end{aligned} \quad (1.29)$$

Comparisons to Differential Amplifier

If the same bias conditions are used, the ratio of the maximum output power of the quad differential amplifier, P_{Qmax} , to the single differential amplifier, P_{Smax} , is simply:

$$\frac{P_{Qmax}}{P_{Smax}} = 20 \cdot \log_{10}\left(\frac{V_{Qmax}}{V_{Smax}}\right) = 20 \cdot \log_{10}(4) \approx 12 dB \quad (1.30)$$

And for the gain of the quad amplifier, A_Q , to the single amplifier, A_S :

$$\frac{A_Q}{A_S} = 20 \cdot \log_{10}(2) \approx 6 dB \quad (1.31)$$

It can be seen that although this configuration increases both the maximum output power and gain compared to the single differential amplifier, the maximum

output power is increased by a greater ratio. As the purpose of the quad amplifier is to increase the maximum available output power without affecting the gain or degrading the noise performance, it makes more sense to compare the configurations when they are operating at the same level of gain. This can be achieved reducing the bias current through the quad amplifier transistors and increasing the values of R_E .

If the quad amplifier is operating with $h_{fe} = 100$, $R_E = 20 \Omega$, $R_{bias} = 680 \Omega$ and $V_{bias} = 7.8 V$, the current through the emitter of each transistor is:

$$I_E = \frac{1}{2} \cdot \frac{V_{bias} - 0.6 V}{R_{bias}} \approx 5.3 mA \quad (1.32)$$

From (1.29) the gain of the amplifier is then:

$$A_Q = \sqrt{2} \cdot \frac{50 \Omega}{\left(\frac{50 \Omega}{h_{fe}}\right) + \left(\frac{1+h_{fe}}{h_{fe}} \cdot (R_E + \frac{25 mV}{I_E})\right)} \approx 8.87 dB \quad (1.33)$$

This is very similar to the calculated gain of the differential amplifier in (1.11).

Finally, the maximum available output power from (1.28) is then:

$$P_{Qmax} = \frac{\left(\frac{16}{\pi} \cdot (I_E \cdot \frac{50 \Omega \cdot 50 \Omega}{50 \Omega + 50 \Omega})\right)^2}{50 \Omega} = 8.09 dBm \quad (1.34)$$

This is around 6.5 dB higher than for the differential amplifier in (1.15), which shows that the quad parallel differential amplifier configuration can be used to increase the maximum available output power without changing the gain and whilst reducing the quiescent current through each of the transistors.

1.9.2 Design

Circuit Diagram

The complete circuit diagram of the quad differential amplifier is shown in Figure 1.43. The same transformers, BFR92P transistors and topology for the long-tailed pair circuits were used as for the differential amplifier described in Section 1.5, albeit

with the collectors of the transistors connected in parallel. The Mini Circuits SCP-4-1+ power splitter was selected for its low insertion loss, 0° outputs and small package size.

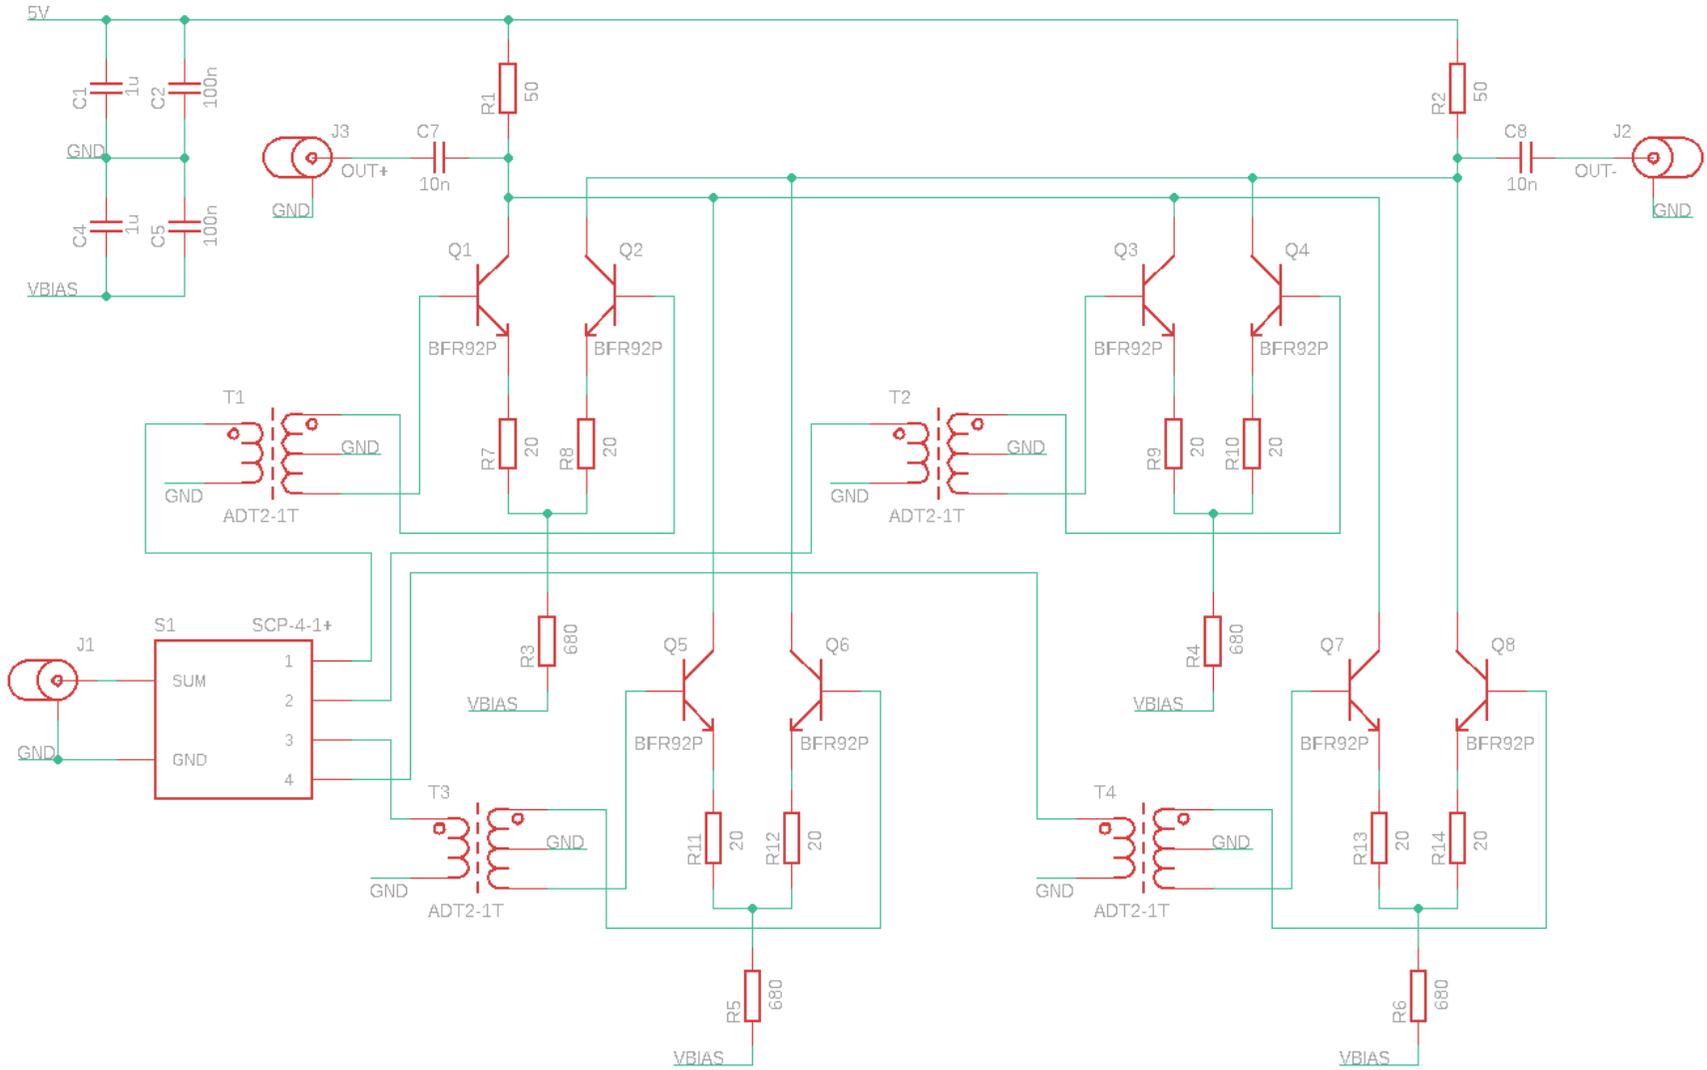


Figure 1.43: Circuit diagram of the quad differential amplifier

PCB Layout

The PCB layout used for the quad differential amplifier is shown in Figure 1.21. It is constructed from 2-layer 0.5 mm I-Tera with gold plating and plated through holes. The collectors of the transistors that are connected in parallel have been placed as physically close together as possible to avoid potential instability, and to achieve this, the devices have been placed on both sides of the board.

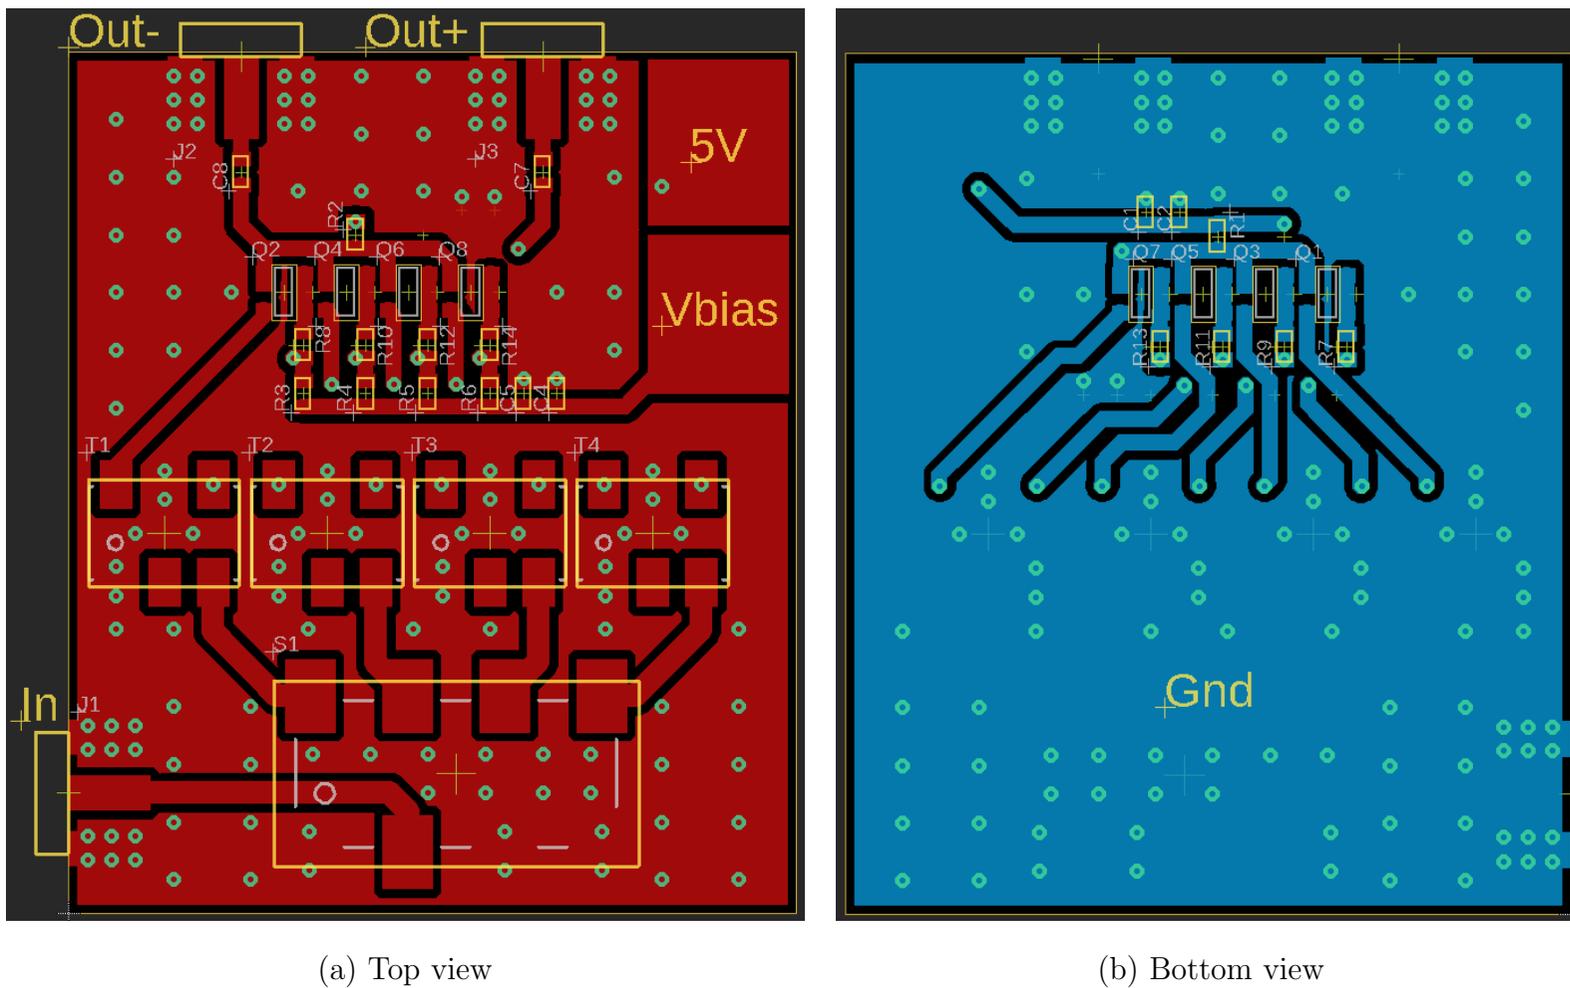


Figure 1.44: PCB layout of the quad differential amplifier

1.9.3 Measurements

Gain Compression

The gain compression of the differential amplifier with a bias voltage of -7.8 V was measured using a similar methodology to that for the other amplifiers described in this section and the results are shown in Figure 1.45.

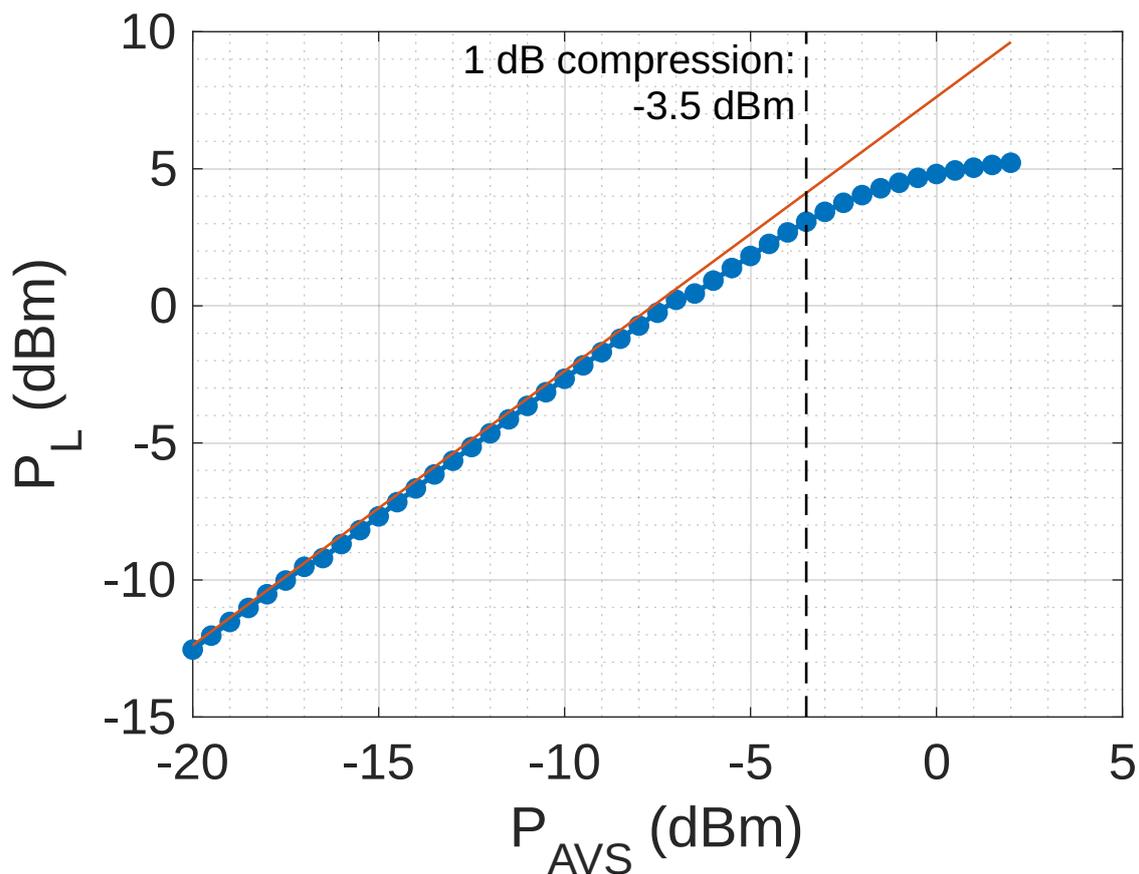


Figure 1.45: Gain compression of the quad differential amplifier with a bias voltage of 7.8 V measured using the Rohde & Schwarz FSWP

The small-signal gain was measured to be around 7.6 dB, which is around 1 dB less than the value in Equation 1.11, but this can be accounted for by the fact that the power splitter and transformer each have an insertion loss of around 0.5 dB at 100 MHz according to the datasheets.

As for the other amplifiers, the maximum available output power can't be easily

measured, but the 1 dB compression points can be compared. In this case it occurs for an input power of -3.5 dBm, which corresponds to an output power of around 3 dBm.

These results confirm that the quad parallel differential amplifier has a similar gain to the differential amplifier, but with a 1 dB compression point that is around 6.5 dB higher.

Input and Output Impedance

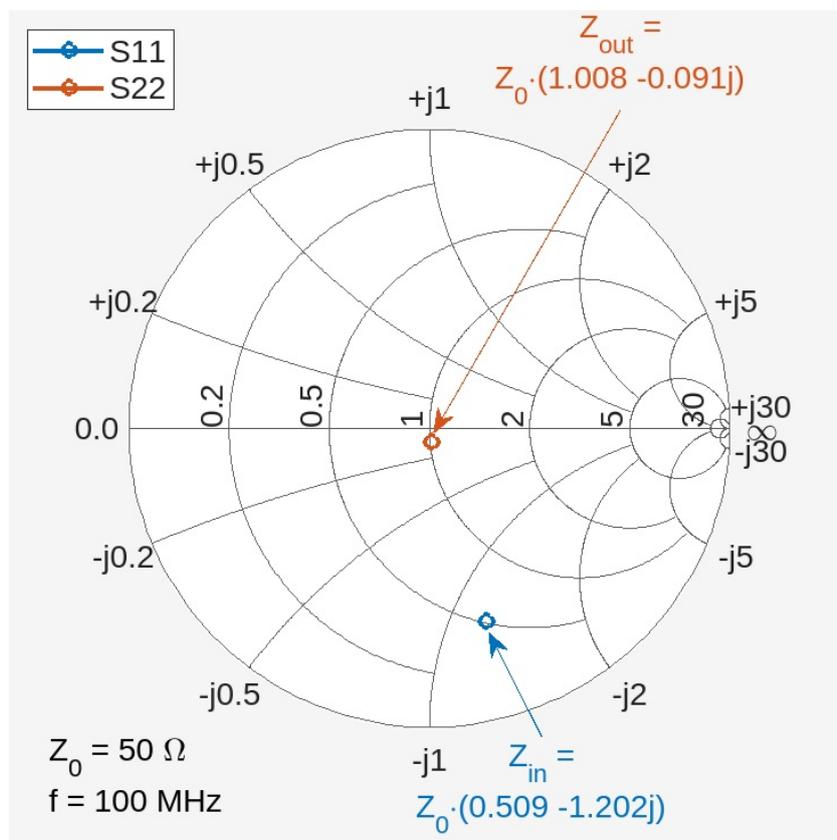


Figure 1.46: S-parameter measurements of the quad differential amplifier showing the input and output impedance at 100 MHz made using the Anritsu 37377C Vector Network Analyzer

As for the differential amplifier in Section 1.5.2, the input and output impedances of the quad differential amplifier were determined by measuring its S_{11} and S_{22}

parameters at 100 MHz and they are shown on the Smith chart in Figure 1.46.

It can be seen that the output is again well-matched to 50 Ω . The input impedance of the single differential amplifier was found to be poorly matched and this is also the case here; however, the magnitude of the input impedance can be calculated to be $Z_0 \cdot \sqrt{0.509^2 + 1.202^2} = 65.3 \Omega$, suggesting it could still be better matched than the single differential amplifier.

Noise Figure

The noise figure of the quad differential amplifier was measured using the Rohde & Schwarz FSWP Phase Noise Analyzer and was found to be 3.8 dB. The slight increase compared to the 3.1 dB noise figure of the single differential amplifier can be mostly attributed to the additional 0.5 dB insertion loss of the power splitter at the input of the amplifier.

Residual Noise

A residual noise measurement of the quad parallel differential amplifier made using the cross-correlation system described in Chapter 3 is shown in Figure 1.47. The power at the amplifier output was set to be close to the recommended maximum drive level of the 100 MHz KVG crystal (around 7 dBm), as this would be the desired operating condition when the amplifier is placed in the crystal oscillator loop.

As was the case for the residual noise measurement of the differential amplifier, the close-to-carrier noise of the quad amplifier cannot be measured as it is not sufficiently above the internal noise floor of the measurement system; however, it is possible to clearly see the far-out noise floor. This was measured to be -176.5 dBc/Hz, however, when compared to the theoretical noise floor from Section 1.2.2 based on its noise figure:

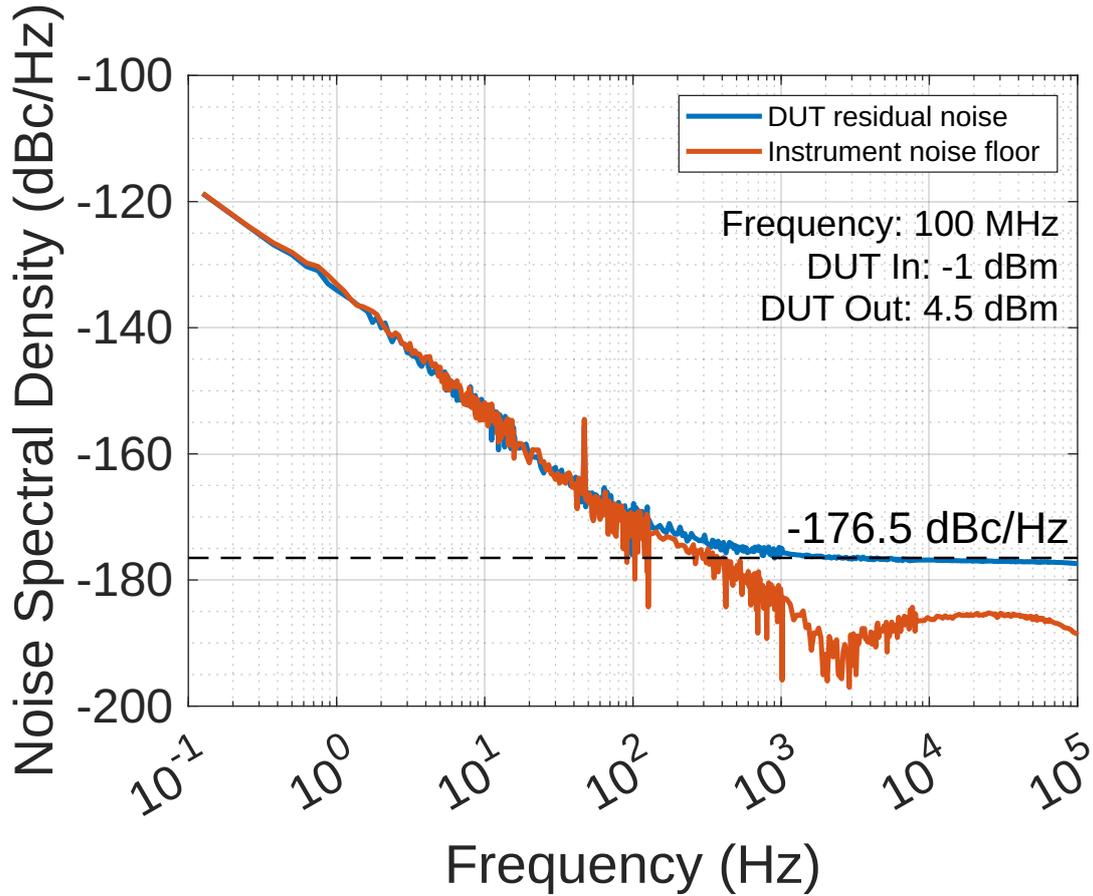


Figure 1.47: Residual phase noise measurement of the quad differential measured using the cross-correlation system with 256 correlations at 0.125 Hz and 8,388,608 correlations at 100 kHz

$$\begin{aligned}
 \text{NoiseFloor} &= -177 \text{ dB} - P_{AVS} + NF \\
 &= -177 \text{ dB} + 1 \text{ dBm} + 3.8 \text{ dB} \\
 &= -172.2 \text{ dBm}
 \end{aligned} \tag{1.35}$$

It can be seen that there is a discrepancy of around 4 dB, it is not clear how the measured noise floor can be less than the theoretical prediction. Since this low noise floor was not observed in the oscillator phase noise measurement using this amplifier in Section 1.10.1, it may be that there is an error in the residual noise measurement that requires further investigation.

It is possible that the parallel configuration of this amplifier is the reason for the unexpectedly low noise floor and that the theoretical noise floor calculation from Section 1.2.2 is not applicable in this case. The effect of this parallel configuration on the amplifier noise floor is something that should be investigated further.

Since the close-to-carrier noise could not be measured with the cross-correlation residual noise system, the flicker noise corner of the amplifier was measured using the LC oscillator method described in Section 1.6. Figure 1.48 shows the absolute phase noise of the LC oscillator with the quad differential amplifier operating at the same output power as that for the residual noise measurement in Figure 1.47.

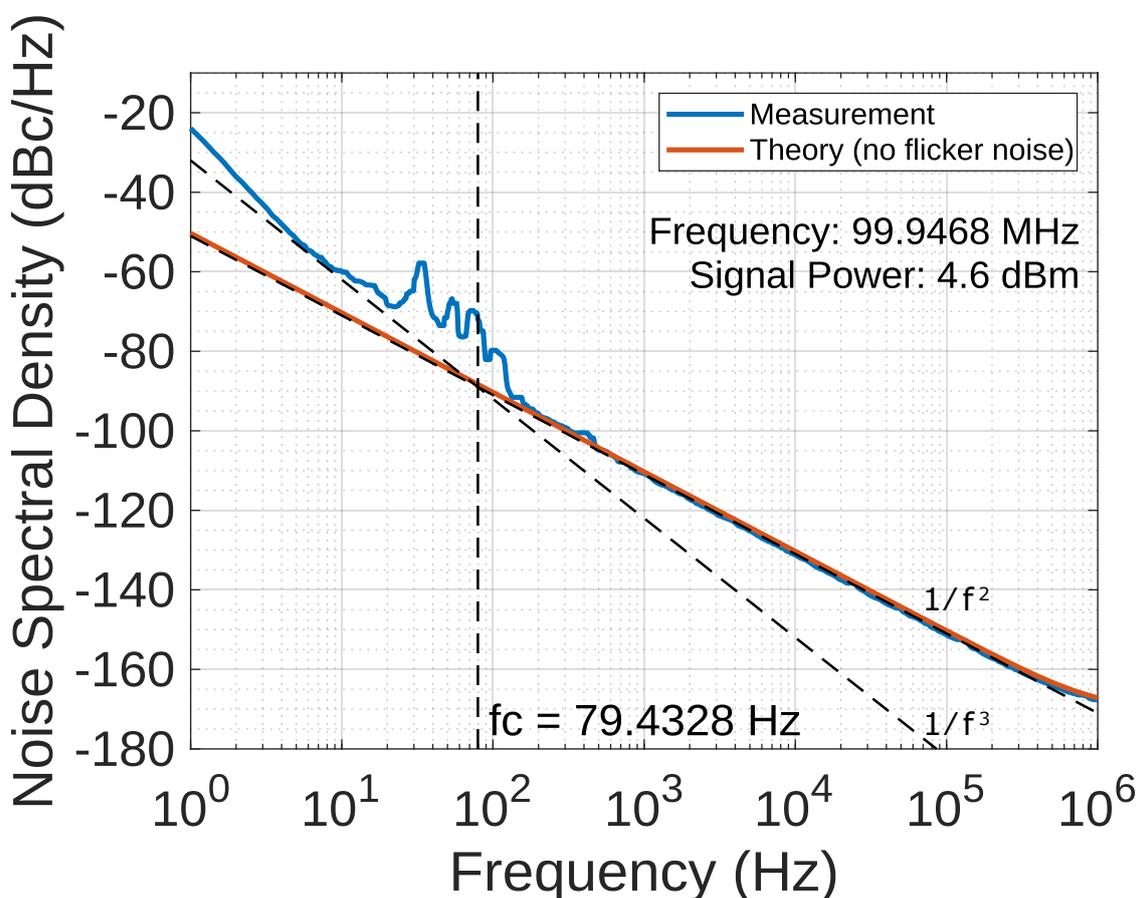


Figure 1.48: Absolute phase noise measurement of the LC oscillator with quad differential amplifier made using the Rohde & Schwarz FSWP to show the flicker noise corner

The flicker noise corner cannot be directly compared with that measured for the differential amplifier in 1.6.3, as the quad differential amplifier measurement has a much lower noise floor owing to its higher output power level. Even so, it can be seen the flicker is quite low at around 80 Hz.

1.10 Conclusions and Further Work

1.10.1 Ultra-Low Phase Noise 100 MHz Crystal Oscillator Measurement

To achieve the best possible phase noise performance of the crystal oscillator, the findings relating to low-noise amplifier designs, attenuator matching and residual phase noise of the crystal were incorporated into the final design, which is outlined in Section 1.3.

A low-flicker-noise quad parallel differential amplifier was presented in Section 1.9 that was shown to have the highest maximum available output power of those tested without excessive gain; however, the input was not well-matched to a 50 Ω system. In Section 1.7 it was shown that including attenuation in the oscillator loop at the input to the amplifier could be used to mitigate against this impedance mismatch, and so this was included in the design in the form of a digitally-controlled attenuator that doubled as a means to set the oscillator loop power. It was found in Section 1.8 that the KVG O5SC85105A 100 MHz quartz crystal resonator exhibited the lowest residual noise of those tested and had a particularly high recommended drive level for achieving a low oscillator noise floor.

Figure 1.49 shows the absolute phase noise of the 100 MHz oscillator when operating with a 4.64 dBm output power using the Rohde & Schwarz FSWP. A smoothing function was applied to the results and can be seen superimposed on top of the original measurement. This is to ignore spurs due to the mains frequency that occur at multiples of 50 Hz, which as one of the metrics used to characterise the phase noise performance is the absolute phase noise at a 100 Hz offset, can affect

the accuracy of the results.

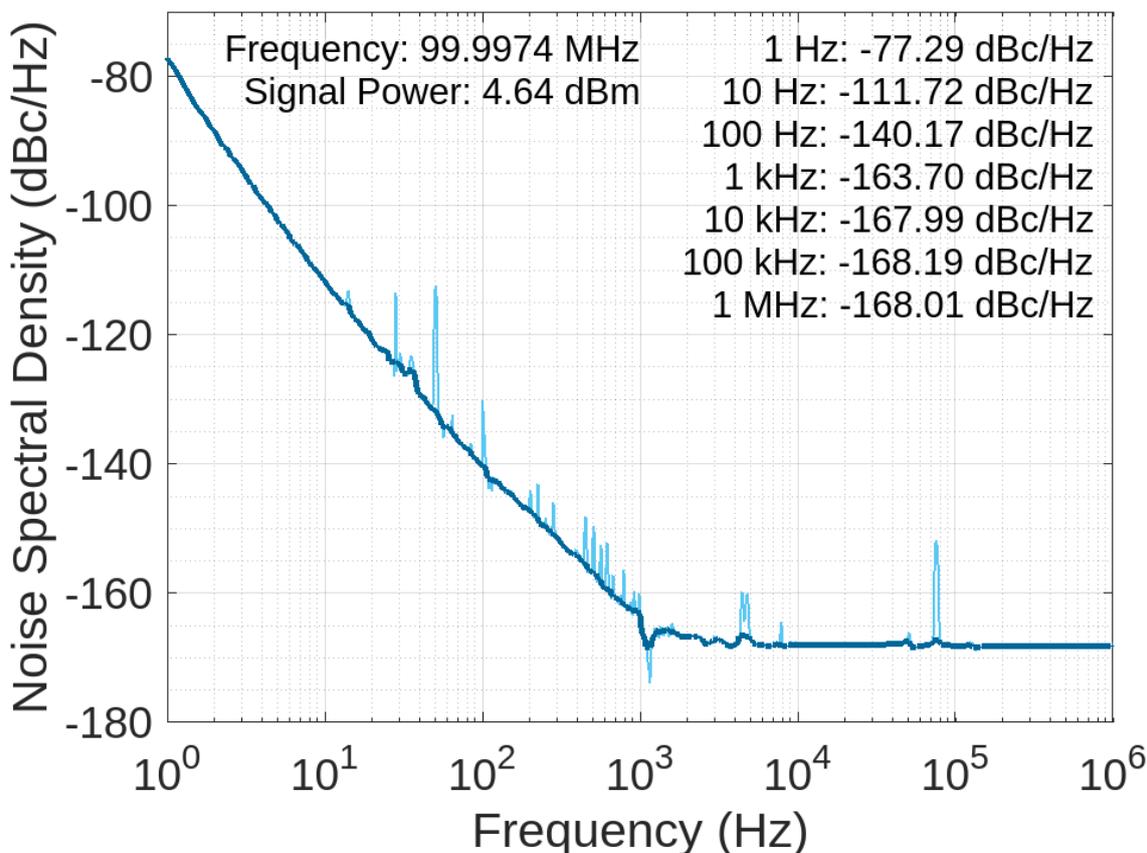


Figure 1.49: Absolute phase noise measurement (with smoothing function applied) of the 100 MHz crystal oscillator using the configuration described in Section 1.3 made using the Rohde & Schwarz FSWP

It can be seen that the oscillator exhibits -140.17 dBc/Hz at 100 Hz offset, which is just below the level that was the primary aim of this research. The far-out noise floor of the amplifier is around -168 dBc/Hz.

1.10.2 Further Work

Minor Improvements

It may be possible to further improve the absolute phase noise of the 100 MHz crystal by making some minor modifications to the quad parallel differential amplifier

described in Section 1.9. The minimum oscillator phase noise that was achieved using this amplifier was when it was operating with an output power of 5 dBm, 2 dB less than the recommended drive level of the crystal, and the amplifier was under around 3 dB of gain compression at this point. If the oscillator loop power can be increased by a further 2 dB, then the absolute phase noise should also improve by 2 dB. The maximum output power of the amplifier can be increased simply by increasing the bias current, but it is not clear if and how this will affect its residual noise.

Effects of Source Impedance on Amplifier Flicker Noise

It was found in Section 1.7 that the flicker noise of the differential amplifier could be improved by placing an attenuator at its input to make the source impedance presented to the amplifier appear closer to $50\ \Omega$, however this is not necessarily the optimal source impedance for minimum flicker noise. This could be investigated further by measuring the flicker noise of the amplifier at a variety of different source impedances, although the differential amplifier flicker noise could not be directly measured using available residual phase noise measurement systems, as shown in Section 1.5.3. Such an investigation could be possible using the LC oscillator flicker noise measurement technique described in Section 1.6 with an added impedance transformation network at the input to the amplifier.

Alternative Amplifier Designs

Of all the amplifiers tested, the Mini Circuits MAR-6 amplifier was the only one found to be well-matched to $50\ \Omega$ at its input. It also had a very low flicker noise corner of around 20 Hz. The drawback to this device was the high excess gain, which when combined with the relatively low 1 dB gain compression point, limited the absolute phase noise that could be achieved when placed in an oscillator.

Mini Circuits provide a range of amplifiers that use a similar technology, but with slightly different operating characteristics. Ideally, if used in an oscillator with

the KVG O5SC85105A crystal resonator, the amplifier would have a gain of around 8 dB and the 1 dB gain compression point would be around the recommended drive level of the crystal, 7 dBm. The MAR-2SM+ seems to be the most suitable device available, with a gain of around 12 dB and a 1 dB compression point of 7 dBm.

It is estimated that using the MAR-2SM+ and a suitable oscillator configuration it would be possible to improve the absolute phase noise of the oscillator by more 10 dB over the -130 dBc/Hz that was measured using the the MAR-6 amplifier.

Chapter 2

Digitally-Controlled Phase Shifters and Attenuators

Abstract

In this chapter, designs for low residual phase noise digitally-controlled phase shifter and attenuator devices are described and the detailed design procedure for operation at arbitrary frequencies is outlined in detail. These devices have been produced for various applications including the residual phase noise measurement system described in Chapter 3, where they are used to set signal power levels and phase shift without degrading the system noise floor.

The digitally-controlled phase shifter is a passive device that uses different lengths of microstrip transmission line switched with latching relays to alter its electrical length. A 6-bit digitally-controlled phase shifter with a range of 180° at 100MHz has been produced, which exhibits $<0.5\text{dB}$ insertion loss and a maximum phase error of around 0.7° . The same design procedure was also successfully applied at a higher frequency to produce a 6-bit 360° phase shifter at 1.5GHz with $<2.2\text{dB}$ insertion loss and a maximum phase error of around 1.75° .

The digitally-controlled attenuator operates in a similar way, but instead of transmission line delay elements, resistive attenuator networks are used so that the overall

attenuation of the device can be changed. Two 6-bit attenuators have been produced, one with a range of 6.3dB and a 0.1dB resolution and one with a range of 31.5dB and a 0.5 dB resolution. In both cases the maximum attenuation error was <0.1dB and the phase shift of the device did not change by more than 3° across its entire range of attenuation settings.

In all the residual noise measurements made of these devices the measured noise was not above the noise floor of the instrument used. This suggests these devices have exceptionally low residual noise, especially compared to similar silicon-based devices.

2.1 Introduction

To produce oscillators with exceptionally low phase noise it is helpful to have a flexible system where operating conditions, such as the loop power and open-loop phase, can be easily adjusted. However, it is important that any additional components used to produce such a system do not introduce significant phase noise, with the flicker noise of semiconductor devices being a particular concern. Low phase noise can be achieved through the use of passive devices, but these must be designed carefully so that they can be adjusted easily, and ideally automated, whilst keeping the unit to a reasonable size and cost.

The motivation for designing the digitally-controlled phase shifter and attenuator described in this chapter was primarily as a way to control the operating conditions of the ultra-low phase noise 100 MHz crystal oscillator described in Chapter 1. The digitally-controlled phase shifter is used to set the open-loop phase shift of the oscillator to around 0° , which may need to be adjusted whenever a component in the oscillator, such as the amplifier, is changed. The digitally-controlled attenuator is used to set the loop power of the oscillator, so that drive level into the crystal can be set to its maximum level before its phase noise begins to degrade. Variations of these devices are also used in a 1.5 GHz oscillator and the cross-correlation residual phase noise measurement system described in Chapter 3.

Devices that perform a similar function have been described in the literature that utilise micro-electro-mechanical systems (MEMS) integrated into coplanar waveguide [18], but here the focus is on producing a detailed design procedure that allows these devices to be built quickly, at minimal cost and for a wide variety of applications using off-the-shelf components.

2.2 Digitally-Controlled Phase Shifter

In this section, the design procedure for the digitally-controlled phase shifter is described in detail and example 6-bit devices operating at 100 MHz and 1.5 GHz are presented.

2.2.1 Theory of Operation

Concept

The digitally-controlled phase shifter operates by using a series of fixed phase shift elements that can be switched in or out of the signal path. The device consists of a number of 'stages', each comprised of a signal switch and phase shift element, where each stage represents one bit. The additional phase shift that can be introduced by each stage of the device is doubled in each subsequent stage, which makes it possible by setting the configuration of the switches to select a phase shift in steps the size of the first stage (the least significant bit). In other words, electrical length of the device is proportional to a binary number that represents the switch configuration of each stage.

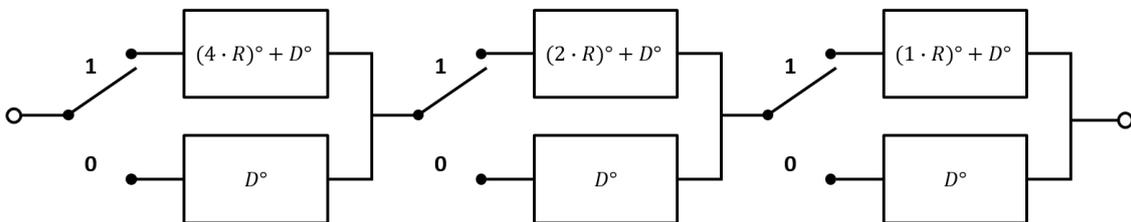


Figure 2.1: Digital phase shifter block diagram

A block diagram of a 3-bit phase shifter is shown in Figure 2.1. The smallest step size, the resolution of the device in degrees, is represented by R . There is also an additional phase shift, D that is present in both signal paths of each stage. Ideally the electrical length, and hence the phase shift, of each stage would be 0° when that stage is set to '0' and $(2^n \cdot R)^\circ$, where n is the bit number, when it is set to '1'. However, in any practical realisation of this device the pass-through signal paths (the '0' settings) will always have a non-zero electrical length, and so will also introduce a phase shift, D . It can be seen that the linearity of the device can be maintained by also adding D to each of the '1' signal paths as in the diagram. No matter the configuration of the switches there will always be a fixed phase shift of $(N \cdot D)^\circ$, where N is the number of stages, leaving the remaining phase shift as the $(2^n \cdot R)^\circ$ for each stage required for the correct stepping.

The total selectable phase range, Φ_{\max} , of the device in terms of the resolution and number of stages is:

$$\Phi_{\max} = R \cdot (2^N - 1) \quad (2.1)$$

We can then derive an expression for the total phase shift of the device in terms of x , which is the binary configuration of the switches represented by a decimal control code:

$$\Phi_{\text{total}} = (\Phi_{\max} \cdot \frac{x}{2^N - 1}) + (N \cdot D) \quad (2.2)$$

Microstrip Delay Line Method

There are various methods that can be used to realise the required phase shift elements for the device, whose suitability depends on the target frequency, phase range and resolution. At RF frequencies it is convenient to use delay lines made up of varying lengths of microstrip transmission lines. This is because they are low loss, can have a well-defined characteristic impedance and can be constructed with a very fine degree of precision. The disadvantage of this method is that for large phase

shifts at low operating frequencies the required length of microstrip line becomes very large, however such device specifications are not required for this research.

A generic single stage of the digital phase shifter corresponding to 1 bit is shown in Figure 2.2. A latching relay is used to switch between two microstrip transmission lines: one of length D , a fixed length between all stages that is ideally the shortest path between relay contacts; and one of length $D + L$, where L is calculated from the desired phase shift of the stage.

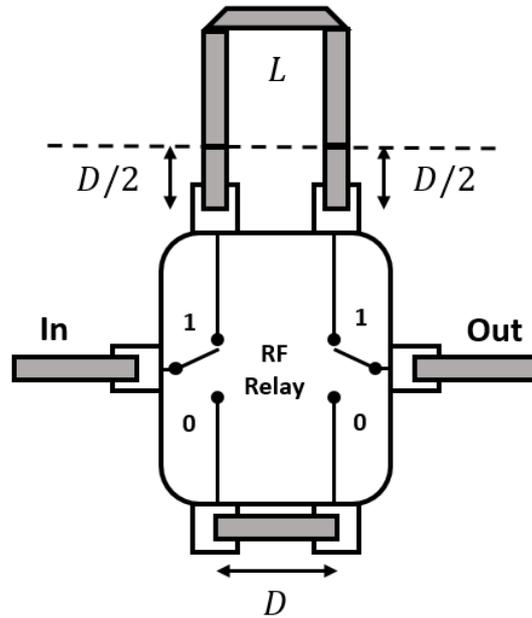


Figure 2.2: Single relay stage of the digital phase shifter circuit corresponding to 1 bit

To determine L the propagation velocity, v_P , of the signal must first be calculated. The propagation velocity in a vacuum is given by:

$$c = \frac{1}{\sqrt{\varepsilon_0 \mu_0}} \quad (2.3)$$

Where ε_0 and μ_0 are constants for the permittivity and permeability of free space respectively.

For a signal propagating along a microstrip transmission line the velocity is then:

$$v_P = \frac{1}{\sqrt{\varepsilon_0 \varepsilon_{\text{eff}} \mu_0 \mu_{\text{eff}}}} \quad (2.4)$$

Where ε_{eff} and μ_{eff} are the effective relative permittivity and permeability of the microstrip transmission line, i.e the values that would give the equivalent propagation velocity if the signal were propagating through a single dielectric material, and are given as a ratios of ε_0 and μ_0 .

The copper transmission line will not have a high magnetic permeability and so the approximation $\mu_{\text{eff}} = 1$ can be made. The equation then simplifies to:

$$v_P = \frac{c}{\sqrt{\varepsilon_{\text{eff}}}} \quad (2.5)$$

Using the propagation velocity, v_P , along with the operating frequency, f_0 , the length of microstrip line corresponding to a single wavelength can be calculated:

$$\lambda_{\text{TL}} = \frac{v_P}{f_0} = \frac{c}{f_0 \sqrt{\varepsilon_{\text{eff}}}} \quad (2.6)$$

The resolution of the phase shifter is set by the smallest microstrip length that can be switched in and out of the signal path, L_1 . Given the resolution, R , in degrees, this length can be calculated by taking a fraction of the wavelength:

$$L_1 = \frac{R}{360^\circ} \cdot \lambda_{\text{TL}} \quad (2.7)$$

This length is used in the stage of the device corresponding to the least significant bit, B_1 , and as shown earlier the length is doubled for each subsequent stage. Therefore we can derive a generic expression for the length, L in each stage corresponding to bit number, n :

$$L_n = \frac{2^{n-1} \cdot R}{360^\circ} \cdot \lambda_{\text{TL}} \quad (2.8)$$

The total length of microstrip track connected to the relay needed for the '1' signal paths as shown in Figure 2.2 is $L_n + D$.

The additional fixed minimum phase shift of the device due to the pass-through connections of length, D , was previously discussed and can be seen in Figure 2.2. It should be noted that there will of course also be further phase shift due to the connections between stages and the internal construction of the relays, but assuming the two states of the relay are symmetrical, this additional phase shift will remain constant regardless of the setting of the digital control code.

2.2.2 100 MHz 180° 6-Bit Phase Shifter

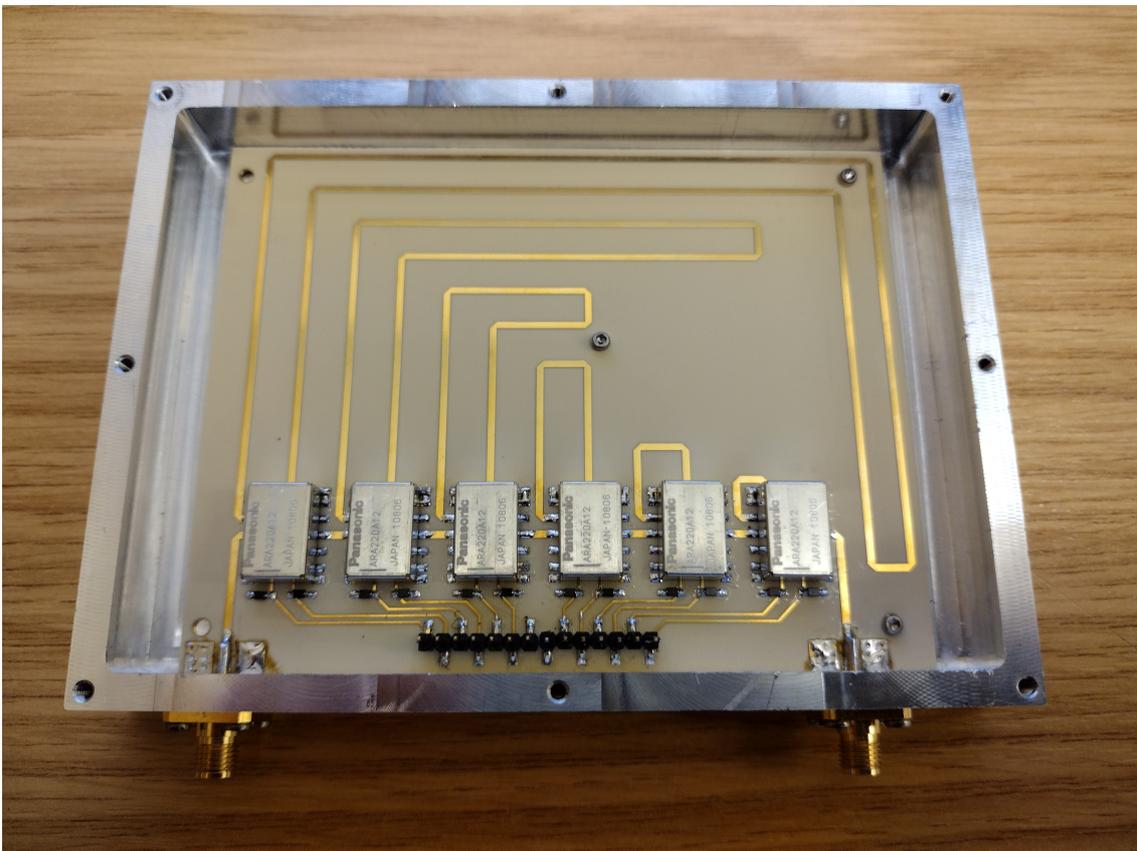


Figure 2.3: Photo of the 100 MHz 180° 6-bit phase Shifter

As shown in Figure 2.3, a 180° 6-bit digital phase shifter designed to operate at 100 MHz has been constructed using the design procedure outlined in Section 2.2.1. This phase shifter is intended to be used both as a convenient way to set the loop phase of the 100 MHz oscillator, described in Chapter 1, and to set the reference

and DUT signal paths of the residual phase noise measurement system, described in Chapter 3, to be in quadrature.

Control Circuitry

To switch between microstrip tracks as in Figure 2.2 Panasonic ARA220A12 12V DPDT dual-coil latching relays were used. These relays are capable of operating at frequencies up to 1 GHz with low insertion loss. The latching structure means that once the states of the relays have been set, the phase shifter can be disconnected from the power source and control logic that switches the relays whilst maintaining the current setting. This prevents additional noise from these sources being coupled into the signal path.

The schematic diagram for the phase shifter is shown in Figure 2.4. Each stage incorporates the dual-coil relay to switch between the two lengths of microstrip transmission line. There are two control signals per stage, B_{N+} and B_{N-} , that are used to charge the coils and set the respective bit for that stage. Applying +12V to the B_{N+} coil latches the relay so that the longer track length, $D + L_N$, is switched into the signal path; this state is represented as '1'. Applying +12 V to the B_{N-} coil latches the relay so that the short pass-through track, D , is switched into the signal path; this state is represented as '0'. The signals can be supplied from a suitable microcontroller with relay drivers. Flyback diodes are also added across the coils to eliminate voltage spikes.

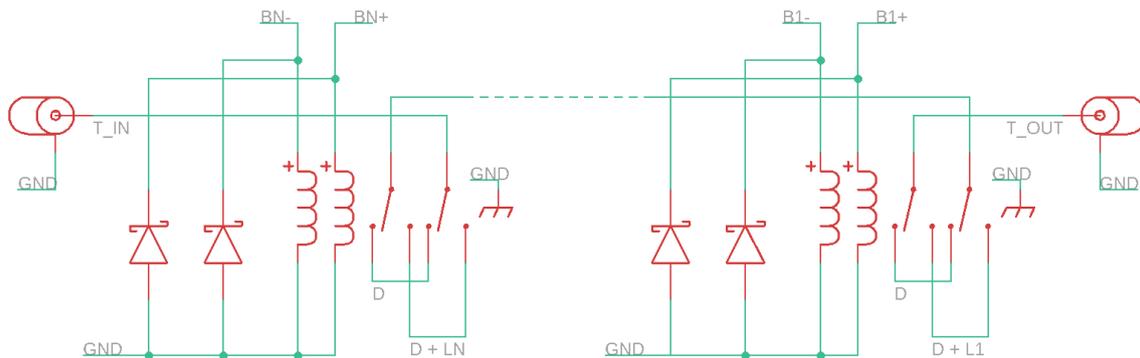


Figure 2.4: Generic schematic diagram for the digital phase shifter design

Microstrip Tracks

The PCB is constructed from 2-layer 0.5mm I-Tera with an 18 μm gold-plated copper layer. The LineCalc transmission line analysis and synthesis program from Advanced Design System (ADS) was used to calculate the width of microstrip tracks that would be required to give 50Ω characteristic impedance. This was found to be 1.14mm.

Hammerstad et al. [19] [20] have shown a set of equations that can be used for calculating the parameters of microstrip transmissions lines. Bahl and Trivedi [21] provide an approximate calculation for the effective permittivity, ϵ_{eff} , of microstrip transmission line when the width of the track is much greater than the thickness of the substrate (and ignoring the thickness of the track), which we can apply here:

$$\epsilon_{\text{eff}} = \frac{\epsilon_r + 1}{2} + \frac{\epsilon_r - 1}{2} \cdot \frac{1}{\sqrt{1 + \frac{12h}{w}}} \quad (2.9)$$

Where h is the thickness of the substrate (0.5 mm), w is the track width (1.14 mm) and ϵ_r is the relative permittivity of the substrate (0.345).

Plugging in the values gives us:

$$\epsilon_{\text{eff}} = 2.714485 \dots$$

Then from (2.6) we can then determine the length of microstrip line that corresponds to a single wavelength at an operating frequency of 100MHz:

$$\lambda_{\text{TL}} \approx 1.8196 \text{ m}$$

Finally, we can use this value along with (2.8) to calculate the required microstrip track lengths for each of the 6 stages of the phase shifter, the values of which are shown in Table 2.1.

In realising the actual design layout based on these calculations it is necessary that considerations be made that result in the microstrip tracks behaving differently from the theoretical predictions. Specifically, bends must be introduced into the tracks so that connections can be made between the relay contacts and the board

Bit	Calculated Track Length	Actual Track Length
1	14.215 mm	15.335 mm
2	28.431 mm	29.300 mm
3	56.863 mm	57.600 mm
4	113.725 mm	115.200 mm
5	227.450 mm	228.900 mm
6	454.900 mm	455.800 mm

Table 2.1: 100 MHz digital phase shifter design parameters

can be kept to a reasonable size, which may affect the electrical length; the tracks must similarly be kept fairly close together, which may introduce coupling effects between tracks; and the selected materials could produce significant attenuation when using such long lengths of transmission line.

For these reasons the proposed layout of the microstrip tracks, shown in Figure 2.5 was first simulated with the Momentum EM field solver to determine whether the electrical length of each track matched the theoretical calculations, and that the attenuation of and coupling between tracks was acceptably low. Only the tracks of length L_N adding the additional phase shift at each stage were simulated for the sake of simplicity, as it was assumed that the phase shift introduced by the relays would be the same regardless of which state it was in.

The results of these simulations are shown in Table 2.2. It can be seen that in all cases the simulated phase shift of each track was within 0.28° of the intended value, which equates to $<10\%$ of the resolution of the phase shifter. The insertion losses were considerably low; the total insertion loss for all tracks combined was <0.1 dB. The worst case coupling ratio between two adjacent tracks, the two longest tracks, was also simulated and found to be very low at -59.4 dB. These results suggest that the calculations used provide a very good approximation for determining the track lengths and that the effects of bends in the tracks and coupling between the tracks

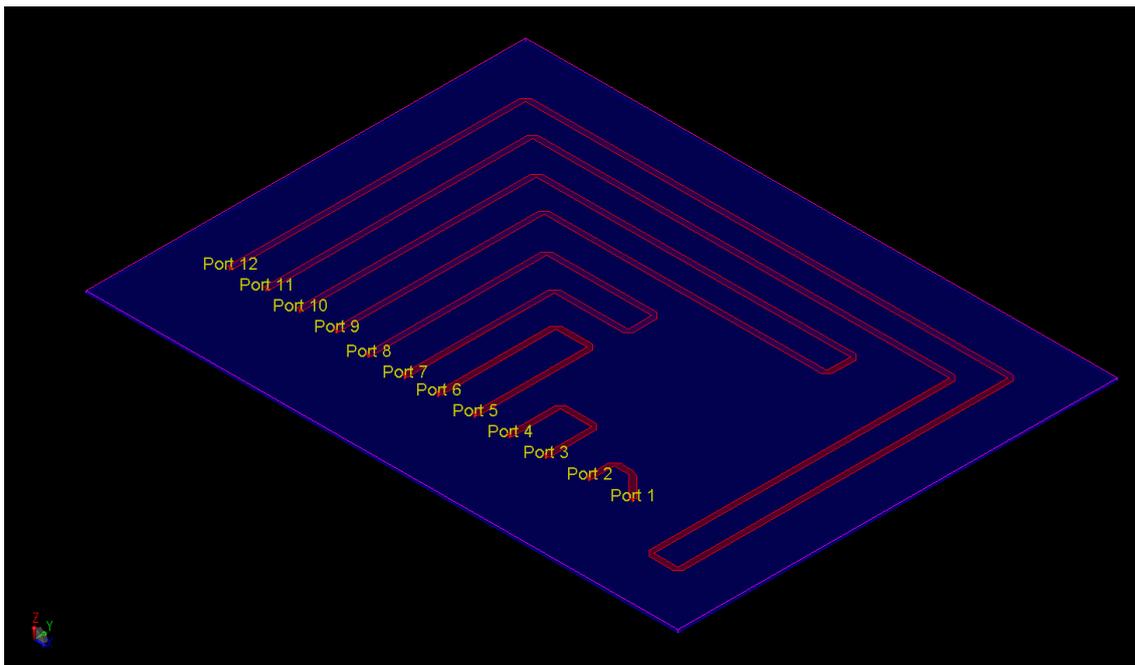


Figure 2.5: Microstrip layout of 100MHz digital phase shifter simulated in Momentum EM field solver

are minimal in this design scenario.

Bit	Nominal Phase Shift	Simulated Phase Shift	Simulated Insertion Loss
1	2.8125°	2.615°	0.0005463 dB
2	5.625°	5.442°	0.001183 dB
3	11.25°	11.108°	0.002661 dB
4	22.5°	22.217°	0.006315 dB
5	45°	44.835°	0.01574 dB
6	90°	89.871°	0.03223 dB

Table 2.2: 100 MHz digital phase shifter EM simulation results

It is however quite simple at this stage to adjust the length of each track such that the simulated electrical length is as close as possible to the intended values of the design, therefore maximising the accuracy of the device. Each of the microstrip

lengths of the simulation model was adjusted manually in small increments until the simulated phase shift was as close as possible to the nominal value. Table 2.1 shows the actual track lengths that were used in construction based on these adjustments.

Full Design

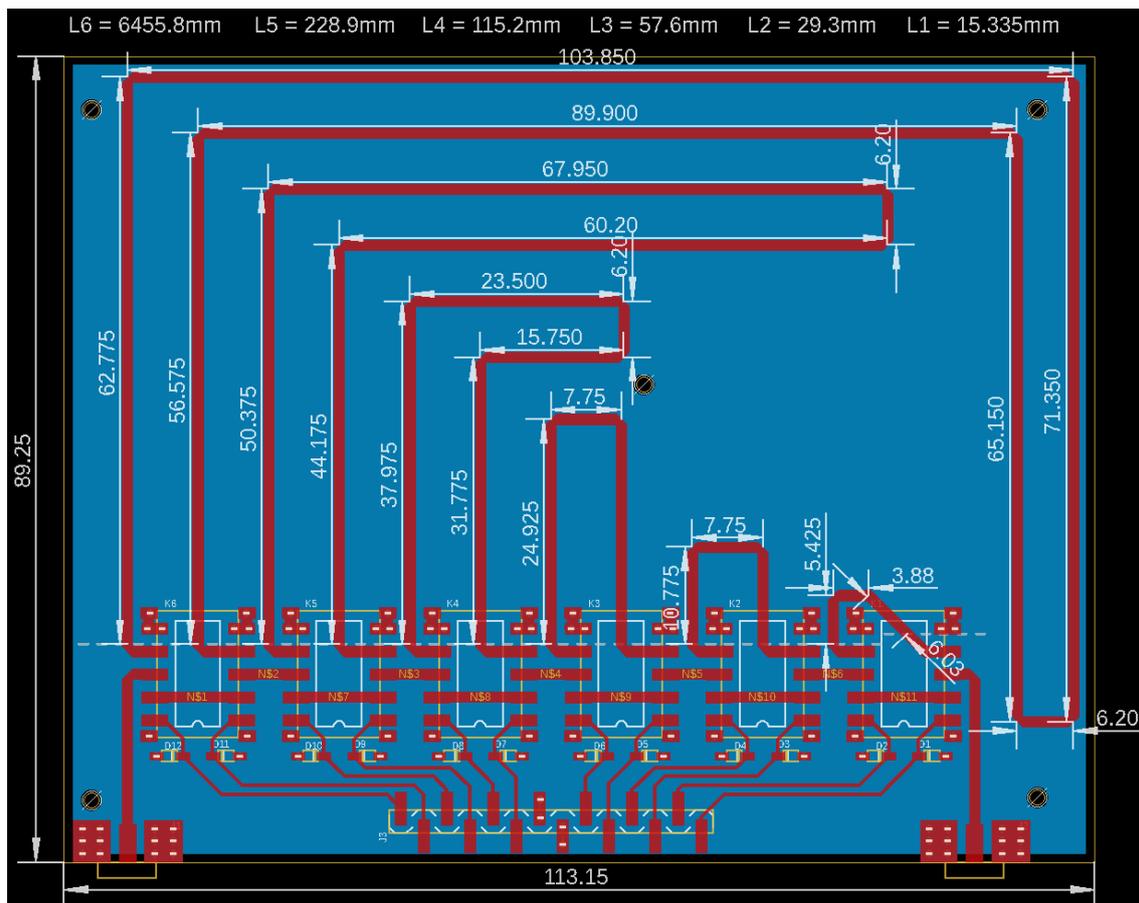


Figure 2.6: PCB layout for the 100 MHz 6-bit digital phase shifter

The full design of the PCB for the 100 MHz digital phase shifter is shown in Figure 2.6. The dotted lines are used to separate the microstrip tracks lengths representing the fixed lengths, D , and the lengths that are used to set the phase shift, L_N , which were determined in the previous section. This means that it would be possible to easily modify this design to produce a device with a different phase range and/or resolution by recalculating the values for L_N and adjusting the track

lengths above the dotted lines accordingly.

Phase/dB Measurements

To confirm the operation of the device, a network analyser was used to measure the absolute phase shift of the device for every possible combination of relay settings, which for a 6-bit device is $2^6 = 64$. A decimal number is used to represent the 6 'bits' of the device, where the most significant bit represents the the relay stage with the longest track length and hence largest phase shift. The nominal phase shift for each bit can be seen in Table 2.1. The device also has a minimum non-zero phase shift, i.e. the absolute phase shift when the relays are set to '0', which must be taken into account, but it is the linearity of the phase shift with respect to the control code setting that is the important function of the device.

These phase shift measurements are shown in Figure 2.7 and are compared with the expected results based on the nominal phase shift of each bit added to the measured minimum phase shift of the device, which in this case was -78.17° . It can be seen that the measured results are very close to the expected values across the whole range of the device. The worst case deviation from the nominal value was 0.6975° , which equates to 24.8% of the device resolution. In other words, we could say the maximum non-linearity error of the device is $<0.25*\text{LSB}$.

The same methodology was used to measure the total insertion loss of the device at each setting and is shown in Figure 2.8. It can be seen that insertion loss generally increases as more phase shift is introduced. This is to be expected as the longer track lengths in the signal path mean there will be more attenuation. The graph is not entirely linear, which could be due to the bends in the tracks causing additional attenuation or differences in the insertion losses of the relays. The worst case insertion loss was found to be 0.46 dB and the loss varies over a range of 0.22 dB depending on the control code setting. Ideally, the insertion loss would be as low as possible and remain constant over the full phase shift range of the device.

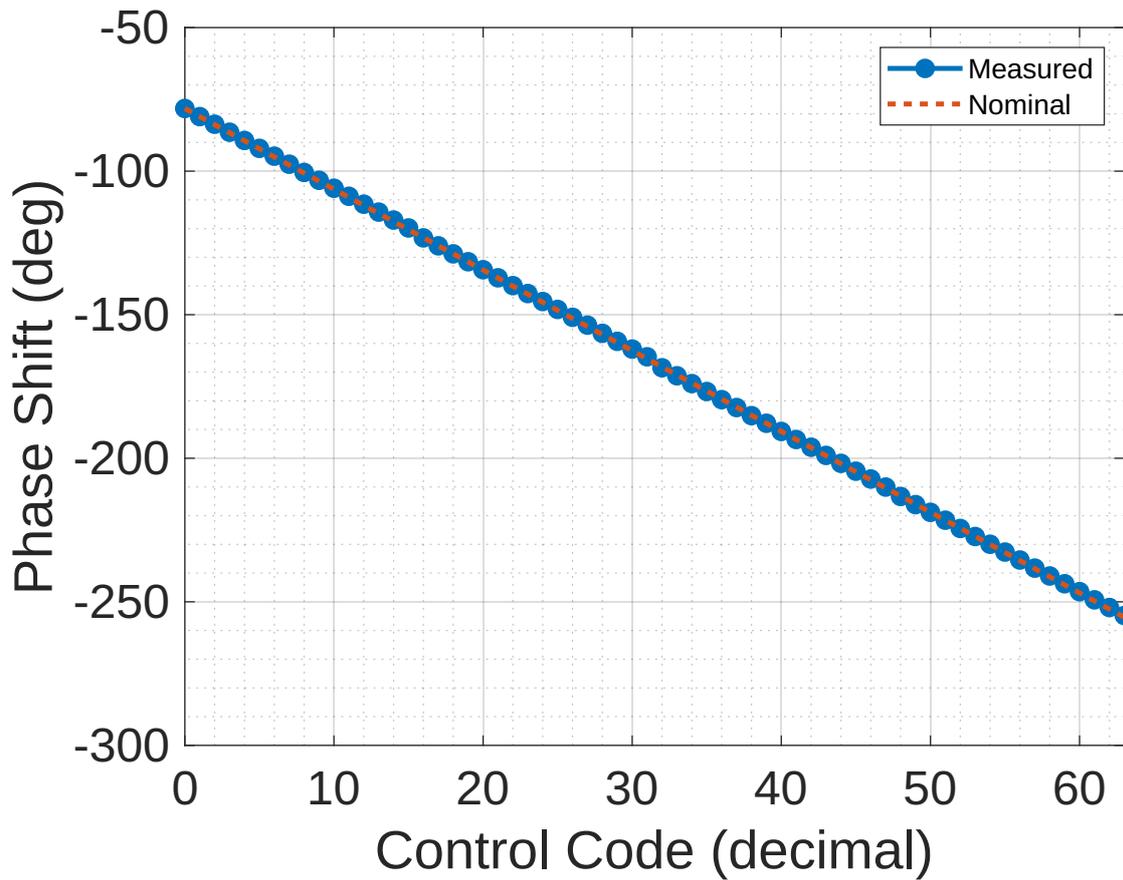


Figure 2.7: Plot of the phase shift of the 100 MHz digital phase shifter vs. the digital control code measured using the SDR-Kits DG8SAQ VNA

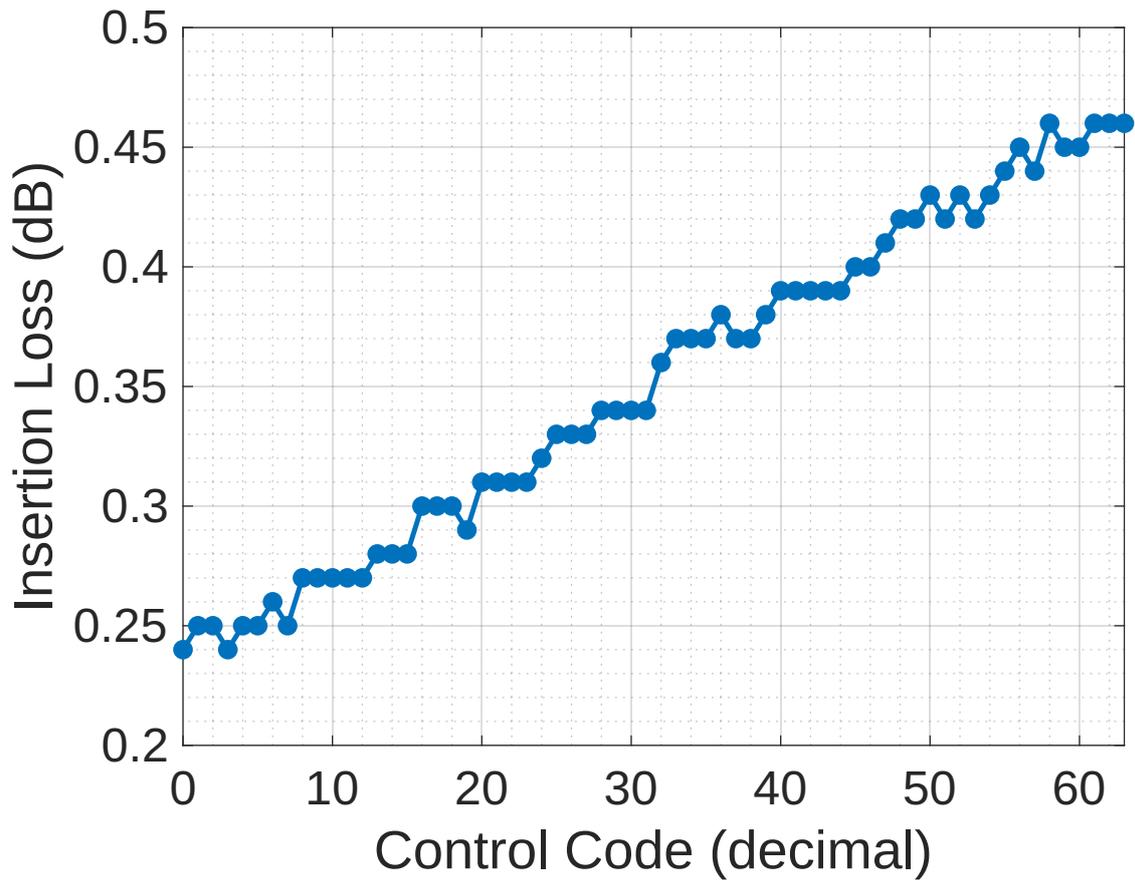


Figure 2.8: Plot of the insertion loss of the 100 MHz digital phase shifter vs. the digital control code measured using the SDR-Kits DG8SAQ VNA

Residual Phase Noise Measurement

The residual phase noise of the digitally-controlled phase shifter was measured using the cross-correlation system described in Chapter 3 and the results are shown in Figure 2.9. Each of the relays was set to the '0' position and the power into the phase shifter was set to 0 dBm, roughly what it would be when placed in the loop of the 100 MHz oscillator described in Chapter 1. The spurs that can be seen in the measurement at multiples of 50 Hz are due to mains interference.

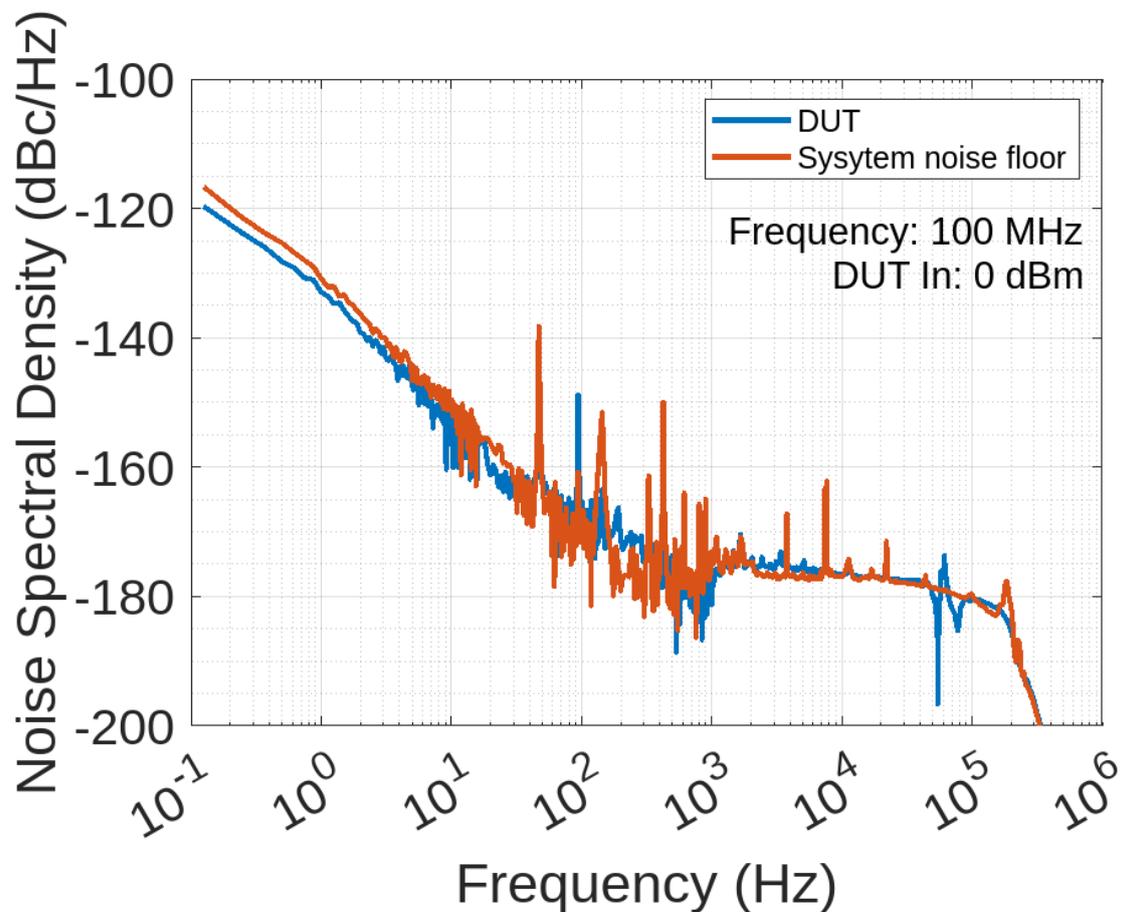


Figure 2.9: Plot of the residual phase noise of the 100 MHz 180° digitally-controlled phase shifter when each of the relays is set to the '0' position (minimum phase shift) measured using the cross-correlation system

The results show there is no significant difference between the measured residual noise of the device and the internal noise floor of the measurement system. This

suggests that the actual residual noise of the device cannot be accurately measured using this system, as it is not sufficiently above the system noise floor. However, this does confirm that the device has very low residual phase noise.

2.2.3 1.5 GHz 360° 6-Bit Phase Shifter

In this section, a 1.5 GHz phase shifter with a 360° range is described. The motivation for producing this device was so that it could be used to set the open-loop phase shift of a 1.5 GHz oscillator being built by this research group.

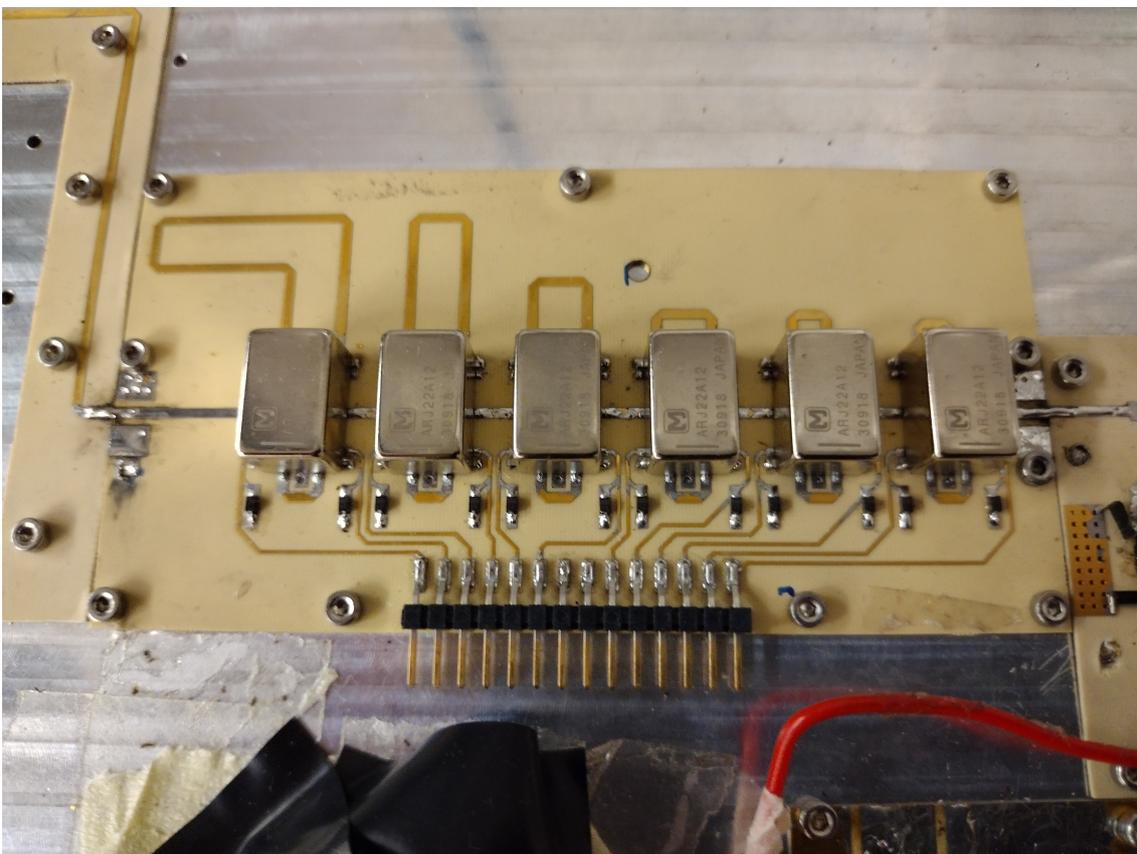


Figure 2.10: Photo of 1.5 GHz 180° 6-bit phase Shifter

Control Circuitry

The control circuitry used in this design is identical to that in the 100 MHz 180° 6-bit phase shifter (Figure 2.4), however it was necessary to change the relays so

that they would work at a higher operating frequency without excess attenuation. Panasonic ARJ22A12 12 V dual-coil latching relays were chosen for this design.

Microstrip Tracks

As this device is operating at a much higher frequency, it is to be expected that the required track lengths will be much shorter than those used for the 100 MHz phase shifter. The same PCB materials are used so we can easily recalculate the microstrip length that corresponds to a single wavelength at 1.5 GHz using (2.6):

$$\lambda_{TL} \simeq 0.1213m$$

We can then use (2.8) to determine the required lengths of L_N , which are shown in Table 2.3. Again the tracks were first confirmed in EM simulation and the lengths adjusted so that the phase shift of each track was as close as possible to the nominal values as described for the 100 MHz phase shifter in Section 2.2.2. The results of these simulations after the adjustment process are shown in Table 2.4.

Bit	Calculated Length	Track	Actual Track Length
1	2.0217 mm		1.9 mm
2	4.0433 mm		4.0 mm
3	8.0867 mm		7.9 mm
4	16.1733 mm		15.9 mm
5	32.3467 mm		32.0 mm
6	64.6933 mm		65.1 mm

Table 2.3: 1.5 GHz digital phase shifter design parameters

Bit	Nominal Phase Shift	Simulated Phase Shift	Simulated Insertion Loss
1	6°	5.879°	0.006716 dB
2	12°	12.068°	0.008683 dB
3	24°	24.083°	0.01528 dB
4	48°	47.928°	0.02734 dB
5	96°	95.931°	0.04212 dB
6	192°	168.060°	0.05259 dB

Table 2.4: 1.5 GHz digital phase shifter EM simulation results

Full Design

The full design of the PCB for the 1.5 GHz digital phase shifter is shown in Figure 2.11.

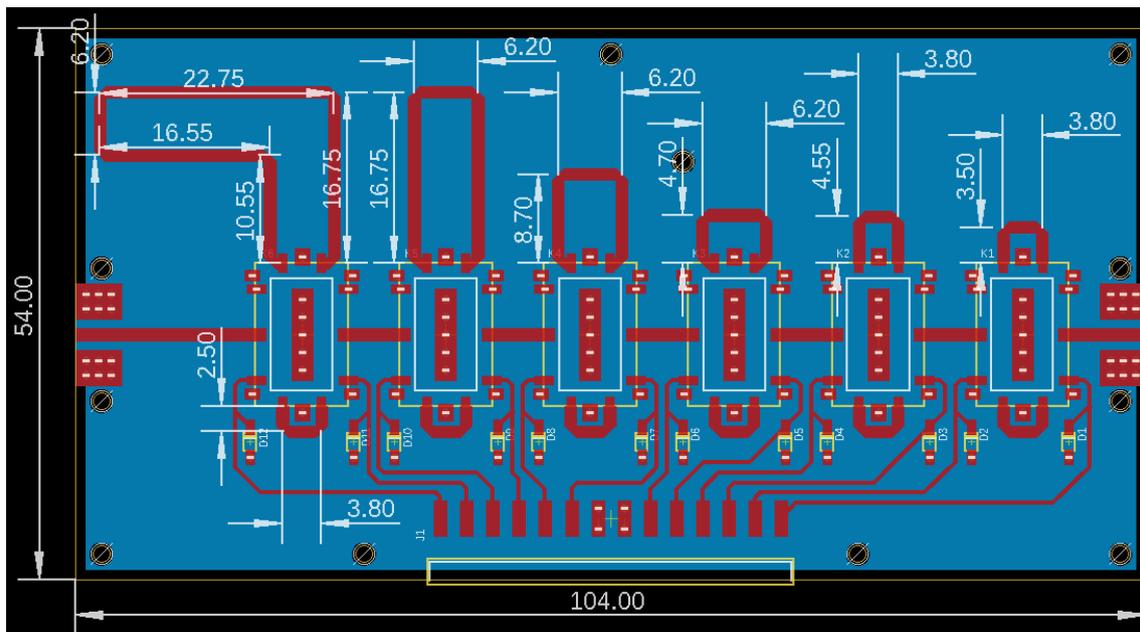


Figure 2.11: PCB layout for the 1.5 GHz 6-bit digital phase shifter

Phase/dB Measurements

The same phase shift measurement was performed as for the 100 MHz phase shifter and is shown in Figure 2.12. The actual performance of the phase shifter was again very close to the expected results based on its nominal phase shift values. The worst case deviation from the nominal value was -1.73° , which means the maximum non-linearity error of the device is $<0.3 \times \text{LSB}$; very similar in performance to the 100 MHz version.

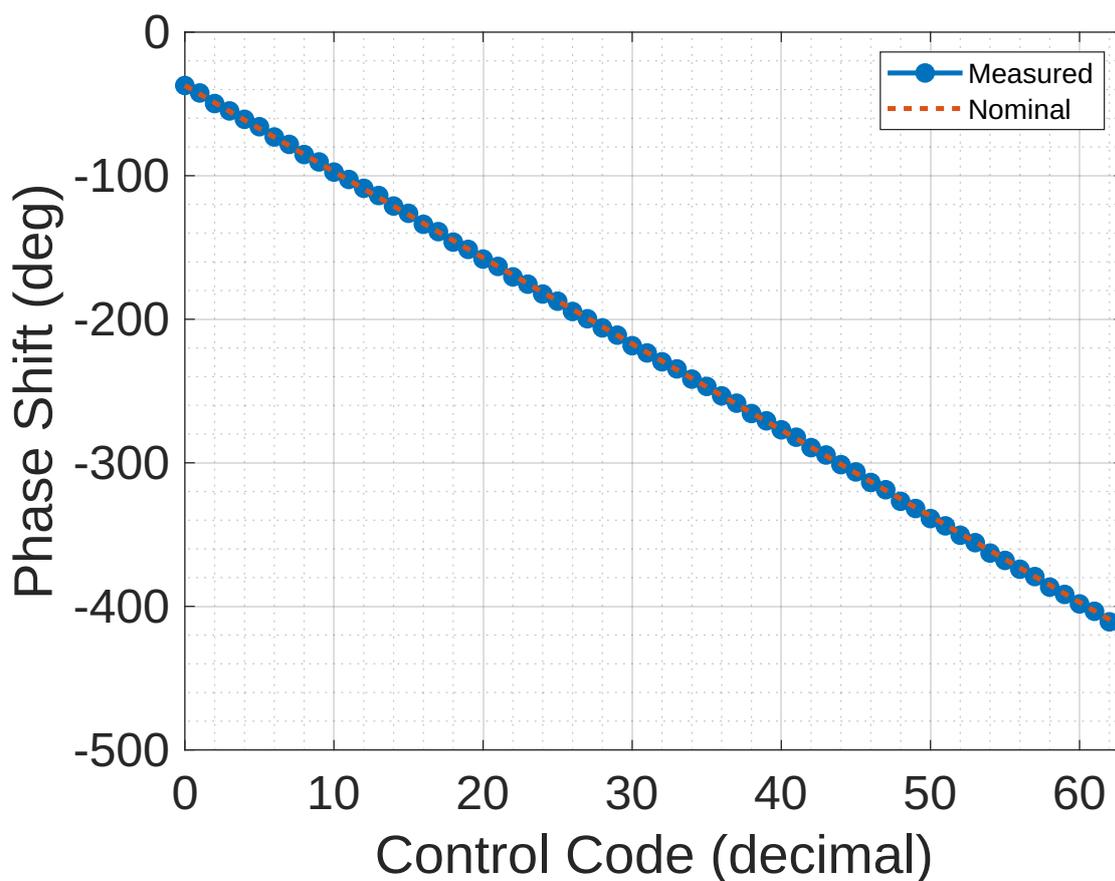


Figure 2.12: Plot of the phase shift of the 1.5 GHz digital phase shifter vs. the digital control code measured using the Anritsu 37377C Vector Network Analyzer

Figure 2.12 shows the insertion loss measurements at each setting of the device. The worst case insertion loss was found to be 2.13 dB. The loss also varies over a wider range of 0.46 dB and there is a significant step that can be seen when the

longest track length (bit 6) is switched in.

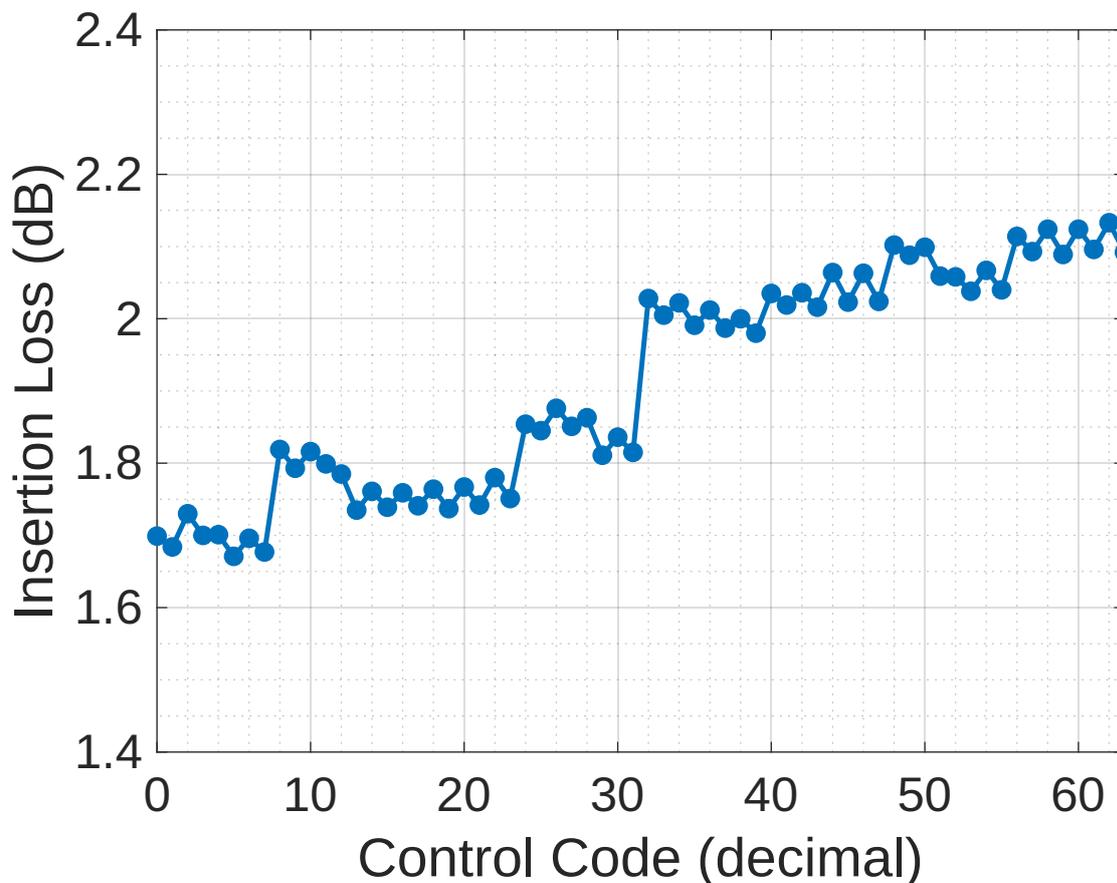


Figure 2.13: Plot of the insertion loss of the 1.5 GHz digital phase shifter vs. the digital control code measured using the Anritsu 37377C Vector Network Analyzer

Although this device can provide an accurate phase shift over a full 360° range at 1.5 GHz, it may not be suitable for applications where a very low insertion loss is required or that are very sensitive to changes in signal level.

Residual Phase Noise Measurement

The residual noise of the 1.5 GHz phase shifter was measured as for the 100 MHz phase shifter in Section 2.2.2, although here the measurement was performed using system based on two HP 11848A units described in [22] [23], as this was the only available system capable of performing a measurement at 1.5 GHz. Again, it can

be seen that the measured residual noise of the device is not sufficiently above the internal noise floor of the measurement system. From this it can be similarly concluded that the 1.5 GHz phase shifter has very low residual noise.

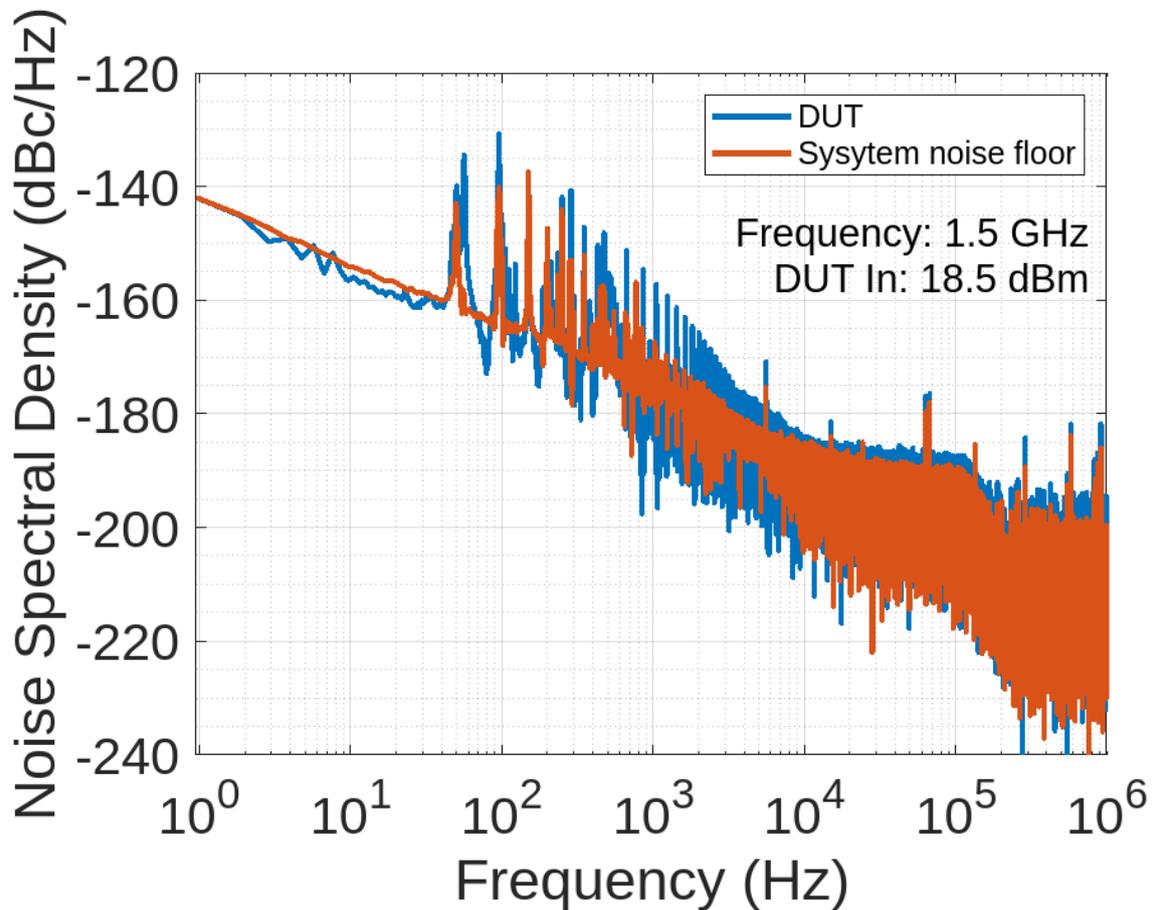


Figure 2.14: Plot of the residual phase noise of the 1.5 GHz 360° digitally-controlled phase shifter when each of the relays is set to the '0' position (minimum phase shift) measured using the HP 11848A system

2.3 Digitally-Controlled Attenuator

In this section, the design procedure for the digitally-controlled attenuator is described in detail and example 6-bit devices with 6.3 dB and 31.5 dB range are presented.

2.3.1 Theory of Operation

Concept

The concept for the digitally-controlled attenuator shown in Figure 2.15 is functionally very similar to that of the digitally-controlled phase shifter, though the delay line element used in each stage is replaced by an attenuation element. This allows a similar level of control over the total attenuation of the device as the previous designs had over the electrical length.

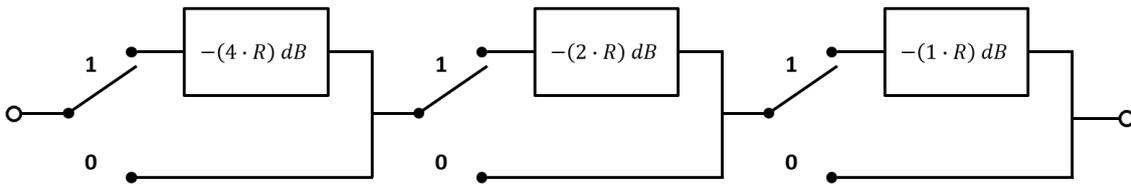


Figure 2.15: Digital attenuator block diagram

In one aspect the concept for the digitally-controlled attenuator is slightly simpler; there is no need to take into account a fixed attenuation in the '0' signal path of each stage in the way that fixed length, D , had to be considered for the phase shifter (Figure 2.1). This is because we can assume that a short pass-through connection between switch contacts or relay terminals would have negligible attenuation at the frequencies we are interested in.

The resolution of the device, R in dB is simply the attenuation at the first stage, and the attenuation in each subsequent stage is doubled. This gives us an equation for the total attenuation range, A_{\max} , of the device in terms of the resolution and number of stages, N :

$$A_{\max} = R \cdot (2^N - 1) \quad (2.10)$$

The settings of the switches in the device can be represented as an N-bit binary control code converted to decimal, x , and the attenuation, A_{total} , of the device can be set according to the equation:

$$A_{\text{total}} = (A_{\max} \cdot \frac{x}{2^N - 1}) + A_{\text{loss}} \quad (2.11)$$

Where A_{loss} is a fixed attenuation that arises due to the insertion loss of the components in the device.

Pi-Pad Method

There are many different methods that can be used for the attenuation elements, but as this design is intended to operate at frequencies <1 GHz it was decided to use resistor networks, which allow for a flexible and compact design. The pi-pad topology, Figure 2.16, provides a simple passive attenuator that can be easily incorporated into relay-switched microstrip layout that was used for the digital phase shifter.

In a matched system, the equations for calculating the pi-pad resistor values for a specified attenuation, $K = \frac{V_{\text{out}}}{V_{\text{in}}}$, are:

$$R_a = Z_0 \cdot \frac{1 + K}{1 - K} \quad (2.12)$$

$$R_b = Z_0 \cdot \frac{1 - K^2}{2 \cdot K} \quad (2.13)$$

For each stage of the device, n , the required attenuation, A in dB is:

$$A_n = R \cdot 2^n \quad (2.14)$$

Converting to a voltage transfer function gives us the value of K for each stage:

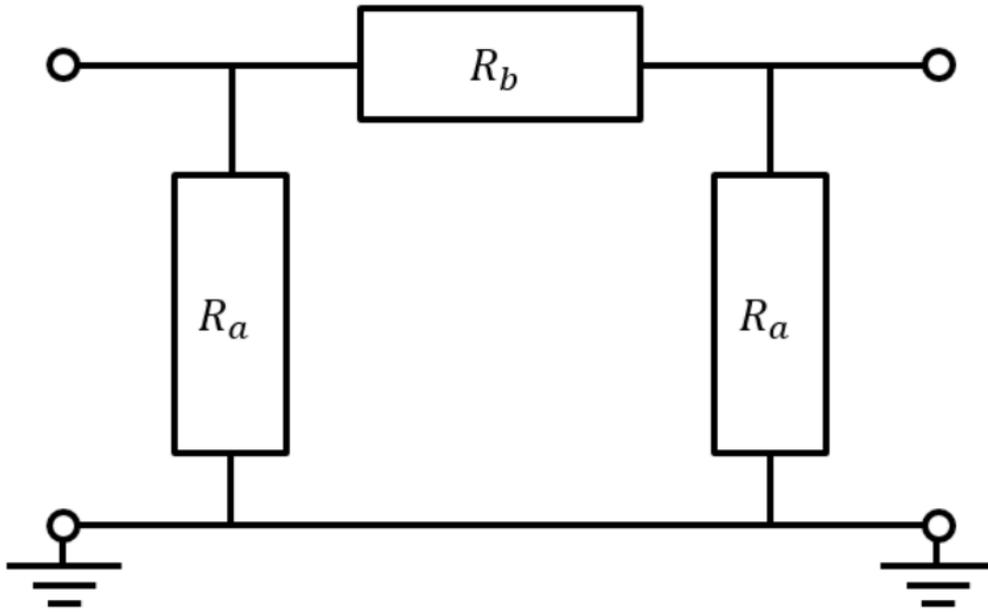


Figure 2.16: Pi-pad resistive attenuation network

$$K_n = 10^{\frac{A_n}{10}} = 10^{\frac{R \cdot 2^n}{10}} \quad (2.15)$$

Finally, we can provide a generic set of equations that can be used to calculate the required resistor values in the device for a given resolution:

$$R_{an} = Z_0 \cdot \frac{1 + (10^{\frac{R \cdot 2^n}{10}})}{1 - (10^{\frac{R \cdot 2^n}{10}})} \quad (2.16)$$

$$R_{bn} = Z_0 \cdot \frac{1 - (10^{\frac{R \cdot 2^n}{10}})^2}{2 \cdot (10^{\frac{R \cdot 2^n}{10}})} \quad (2.17)$$

The modified layout incorporating the pi-pad attenuator is shown in Figure 2.17. Here, the track lengths for the attenuator and pass-through signal paths are the same, so that the electric length of the device remains constant even as the attenuation setting is changed.

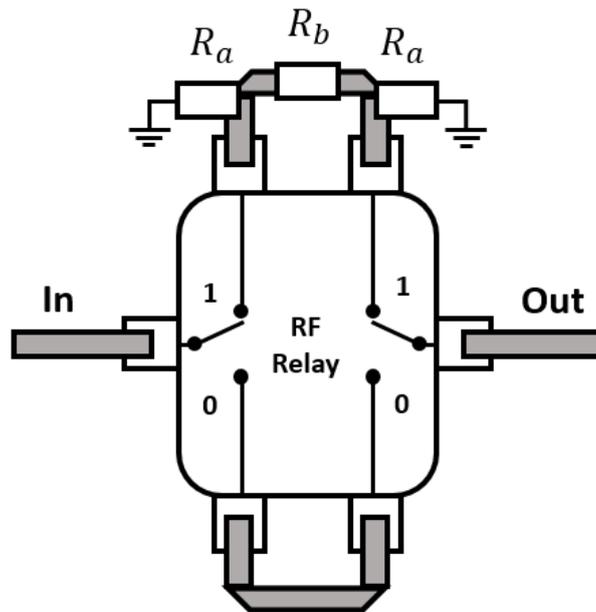


Figure 2.17: Single relay stage of digital attenuator circuit corresponding to 1 bit

2.3.2 100 MHz 6.3 dB 6-Bit Attenuator

In this section a 6.3 dB digitally-controlled attenuator with a resolution of 0.1 dB is presented. It was built primarily for controlling the loop power of the 100 MHz crystal oscillator described in Chapter 1 with a relative degree of precision, so that the oscillator can operate at the maximum possible power level without over-driving the crystal.

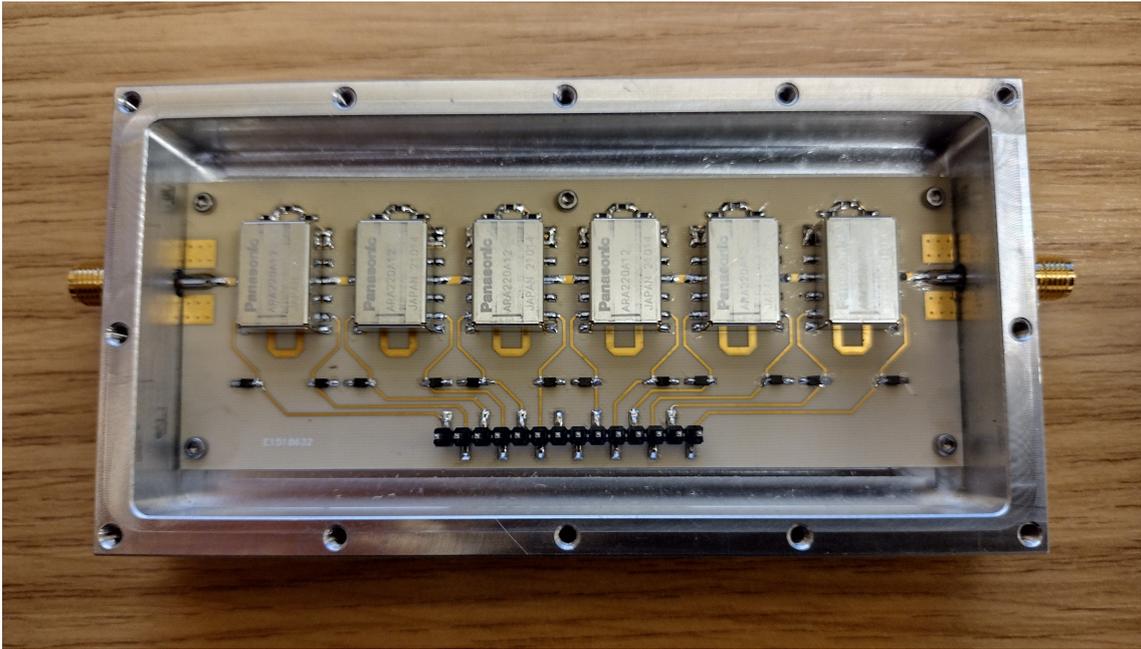


Figure 2.18: Photo of 100 MHz 6-bit digitally-controlled attenuator

Control Circuitry

The schematic diagram of the digital attenuator shown in Figure 2.19 is based on the one used for the 100 MHz digital phase shifter, the only difference being that the relays are used to switch in the pi-pad networks rather than different lengths of transmission line.

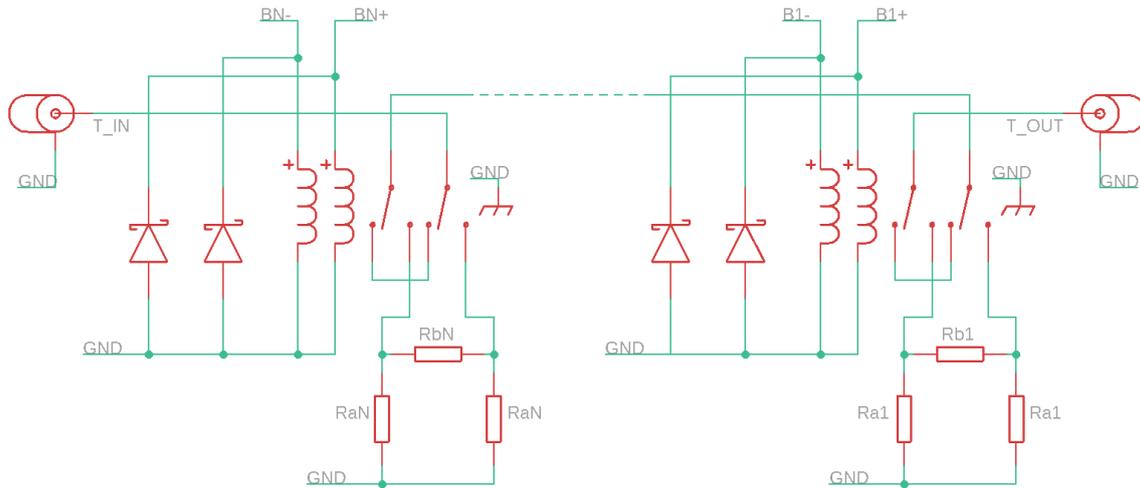


Figure 2.19: Generic schematic diagram for the digital attenuator design

Pi-Pad Networks

The required resistor values for the device were calculated using Equations (2.16) and (2.17) and are shown in Table 2.5. The closest available resistor values were selected for the actual circuit, which in all cases is $<5\%$ from the calculated value.

The PCB layout of the device is shown in Figure 2.20. It can be seen that the microstrip tracks for the attenuator and pass-through paths on either side of the relays are the same length to minimise changes in phase shift as the attenuation setting is changed. In addition, small surface mount resistors are used for the pi-pad networks to reduce the effects of parasitics on the electrical length.

Bit	Att.	Calculated		Actual	
		R_a	R_b	R_a	R_b
1	0.1 dB	8685.99 Ω	0.5757 Ω	8660 Ω	0.56 Ω
2	0.2 dB	4343.13 Ω	1.1514 Ω	4320 Ω	1.1 Ω
3	0.4 dB	2171.86 Ω	2.3034 Ω	2200 Ω	2.2 Ω
4	0.8 dB	1086.50 Ω	4.6117 Ω	1100 Ω	4.7 Ω
5	1.6 dB	544.402 Ω	9.2625 Ω	549 Ω	9.1 Ω
6	3.2 dB	274.497 Ω	18.8402 Ω	274 Ω	19.1 Ω

Table 2.5: Resistor values for 6.3 dB digital attenuator

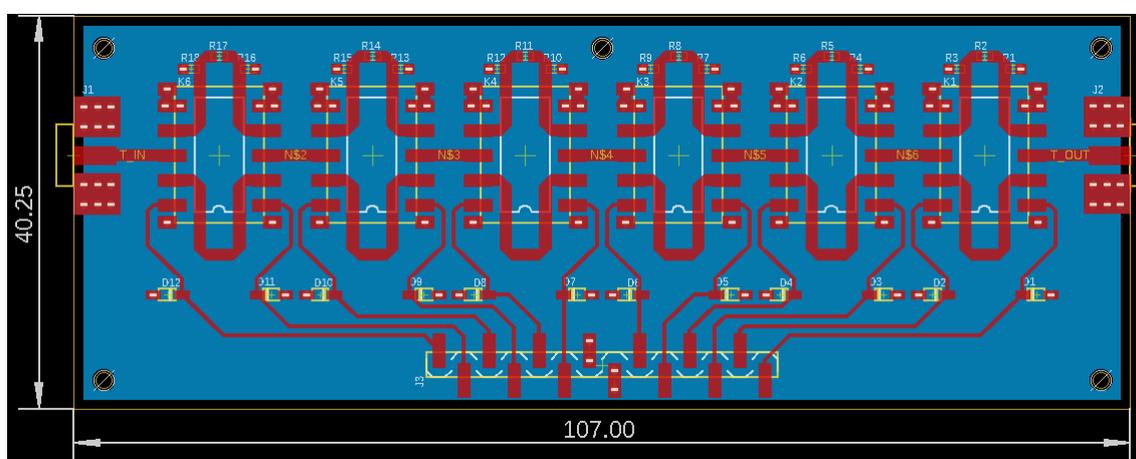


Figure 2.20: PCB layout for the 6.3 dB 6-Bit digital attenuator

Phase/dB Measurements

A similar methodology to that which was used for the digital phase shifters was also used for testing the operation of the digital attenuator. The device was switched through every possible control code setting and the phase shift and attenuation were measured using the Anritsu 37377C Vector Network Analyzer. Figure 2.21 shows the attenuation results and compares them to the expected results based the nominal range of the device (0 to 6.3 dB) plus a fixed non-zero attenuation, i.e the insertion loss of the device when the relays are set to '0', which was measured to be 0.367 dB. The measured results are very close to the expected results across the

full range of the device; the maximum deviation was 0.054 dB or 54% of the device resolution. The maximum non-linearity error of the device is then $<0.55 \cdot \text{LSB}$.

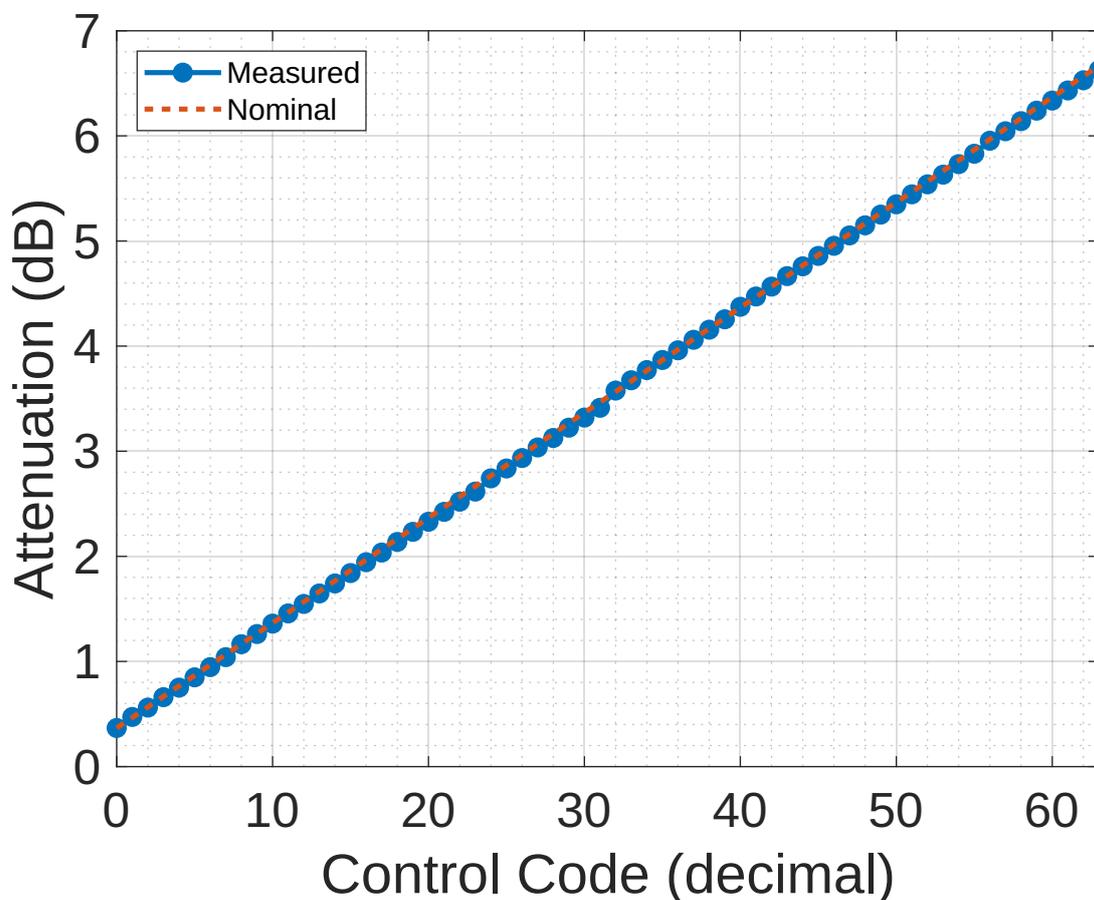


Figure 2.21: Plot of the attenuation of the 6.3 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer

The phase results are shown in Figure 2.22. Ideally, the phase should remain constant even as the amount of attenuation is changed. The results show that the phase shift of the device varies within a range of around 2° at 100 MHz, which although not much, might still have an effect in applications that are very sensitive to signal phase. It can be quite easily seen from the plot that the phase is mostly affected by the number of relays that are in the '1' position at any given time. For example, the least amount of phase shift occurs when the control code is 0 (pass-

through connections only) and the greatest amount of phase shift occurs when the control code is 63 (all attenuators switched in). This is backed up by the fact that there is a large jump in the phase shift that occurs when the code changes from 31 \rightarrow 32, as the number of attenuators switched in is suddenly changing from 5 to just 1.

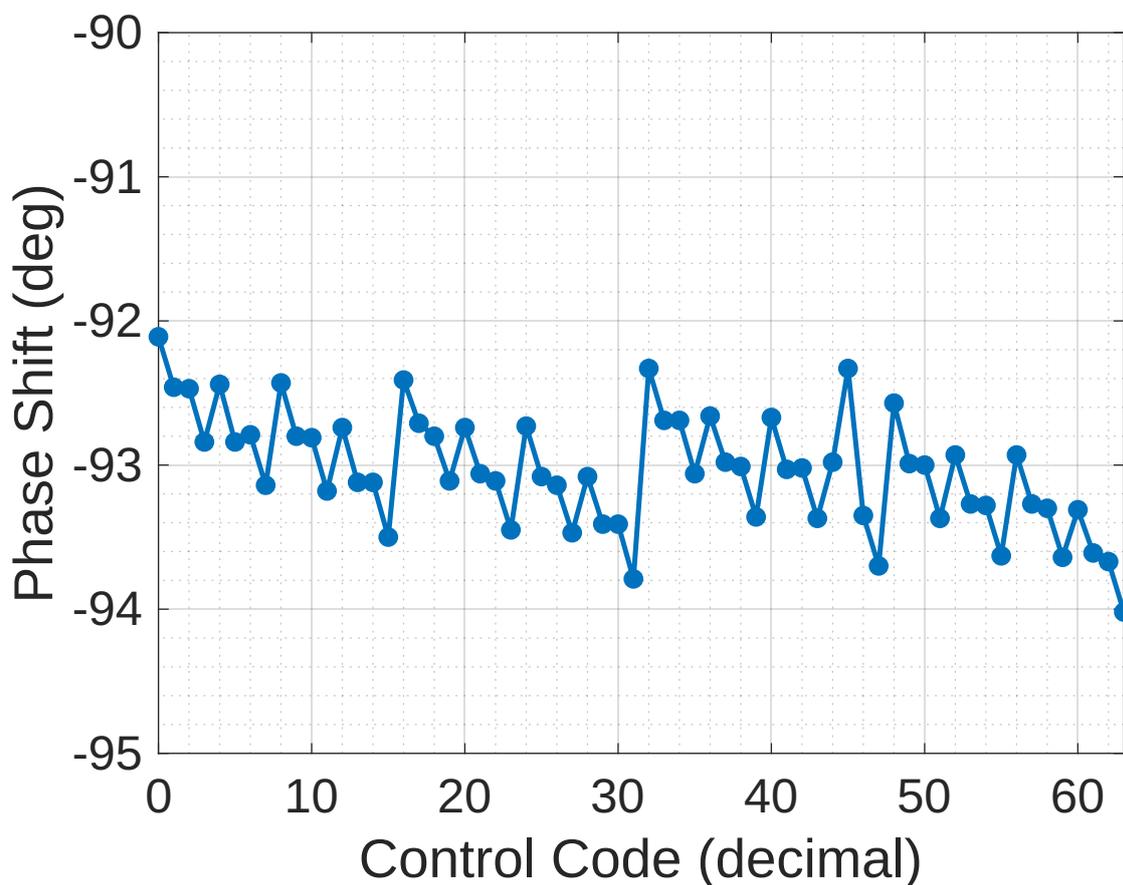


Figure 2.22: Plot of the phase shift of the 6.3 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer

2.3.3 100 MHz 31.5 dB 6-Bit Attenuator

The digital attenuator was designed to be flexible; the attenuation range of the device along with its resolution depends only on the resistor values used in the pi-

pad networks, therefore a variety of different attenuators for specific applications can be built using the same PCB layout.

An example of a relatively high resolution device with a narrow range was shown to be achievable with the 6.3 dB attenuator described in the previous section. Here, a device with a much wider range (31.5 dB) is described and is used to demonstrate how the target range and resolution affects the performance of the device.

The motivation for constructing this 31.5 dB digital attenuator is so that it can be used as part of the cross-correlation residual noise measurement system described in Chapter 3. Specifically, it will be used for setting the signal power level into the device under test, and due to the large variation in device operating conditions that might be required, an attenuator with as wide a range as possible while still maintaining a reasonable resolution would be ideal. It was decided that the device should have a maximum resolution of 0.5 dB, which for a 6-bit attenuator results in a total attenuation range of 31.5 dB.

Design

In redesigning the digital attenuator for the increased range of 31.5 dB, the only necessary change was to recalculate the resistor values of the pi-pad networks using equations (2.16) and (2.17). These calculated values are shown in Table 2.6 along with the actual resistor values that were used in the design.

Bit	Att.	Calculated		Actual	
		R_a	R_b	R_a	R_b
1	0.5 dB	1737.66 Ω	2.87983 Ω	1740 Ω	3 Ω
2	1 dB	869.548 Ω	5.7692 Ω	866 Ω	5.6 Ω
3	2 dB	436.212 Ω	11.6149 Ω	432 Ω	11.5 Ω
4	4 dB	220.971 Ω	23.8484 Ω	221 Ω	23.7 Ω
5	8 dB	116.143 Ω	52.8445 Ω	115 Ω	52.3 Ω
6	16 dB	68.8339 Ω	153.777 Ω	68.1 Ω	154 Ω

Table 2.6: Resistor values for 31.5 dB digital attenuator

Phase/dB Measurements

The results of the attenuation of the device versus the control code are shown in Figure 2.23. As was the case with the 6.3 dB attenuator, the results for this higher range device are very close to the expected values. The largest deviation from the nominal value that was measured was 0.073, or 14.6% of the LSB. Although the maximum non-linearity error of this device when expressed as a percentage of the LSB is considerably less than that for the 6.3 dB attenuator (54%), the deviation as an absolute value in dB is actually fairly similar for both the 6.3 dB and the 31.5 dB devices; 0.054 dB and 0.073 dB respectively. The important conclusion that can be made from this is that redesigning the digital attenuator for a much larger range of attenuation does not seem to degrade the performance in an appreciable way.

The plot of the device phase shift in Figure 2.24, however does show some unusual behaviour at higher levels of attenuation. It could be seen for the 6.3 dB attenuator that change in phase shift corresponded closely with the number of attenuators that had been switched into the signal path, and although we see this same relationship in the plot for the 31.5 dB attenuator for low levels of attenuation, above the decimal control code '30' (15 dB) this starts to become much less predictable. This could be a limitation of the network analyser that was used to measure the device; the

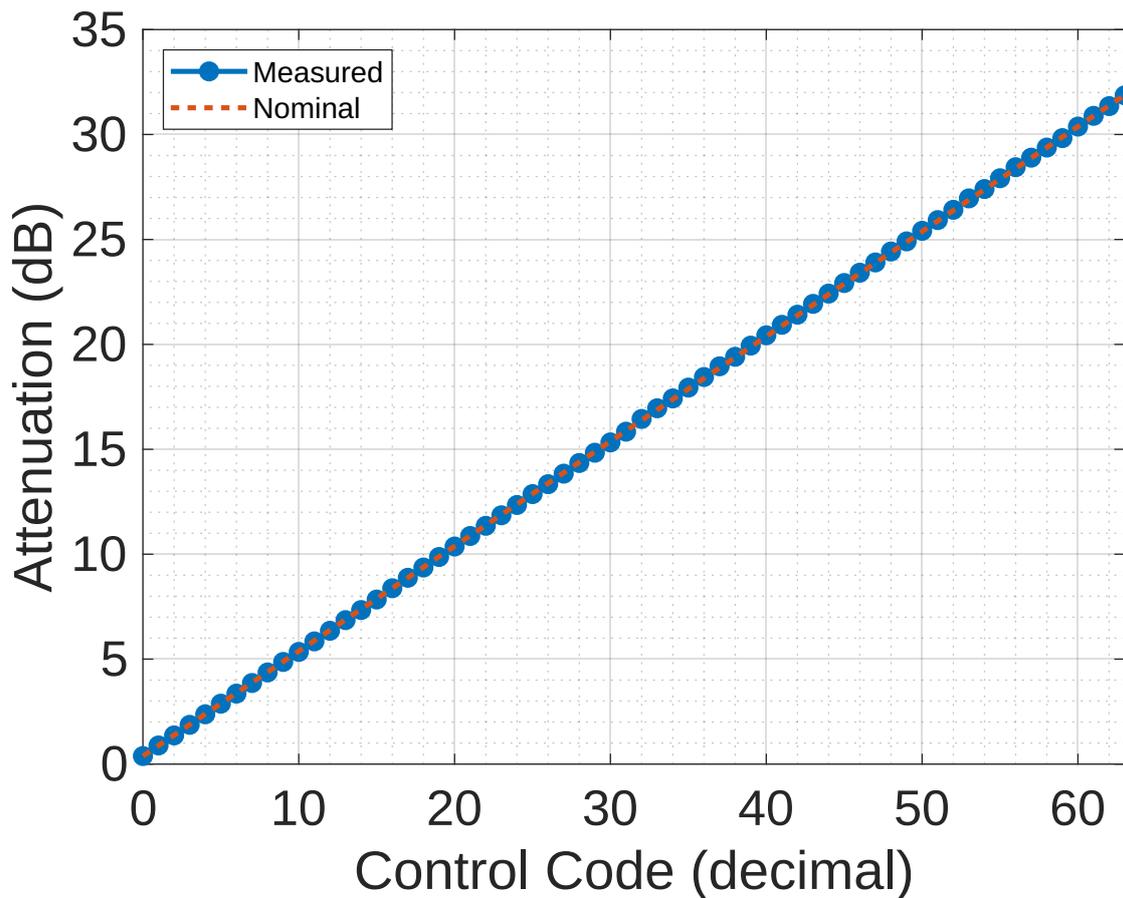


Figure 2.23: Plot of the attenuation of the 31.5 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer

accuracy of phase measurements may have been affected by the decreased signal levels. It could also be a limitation of this digital attenuator design that higher levels of attenuation introduce more and less predictable phase. Even so, the phase shift varies within a range of around 3° at 100 MHz.

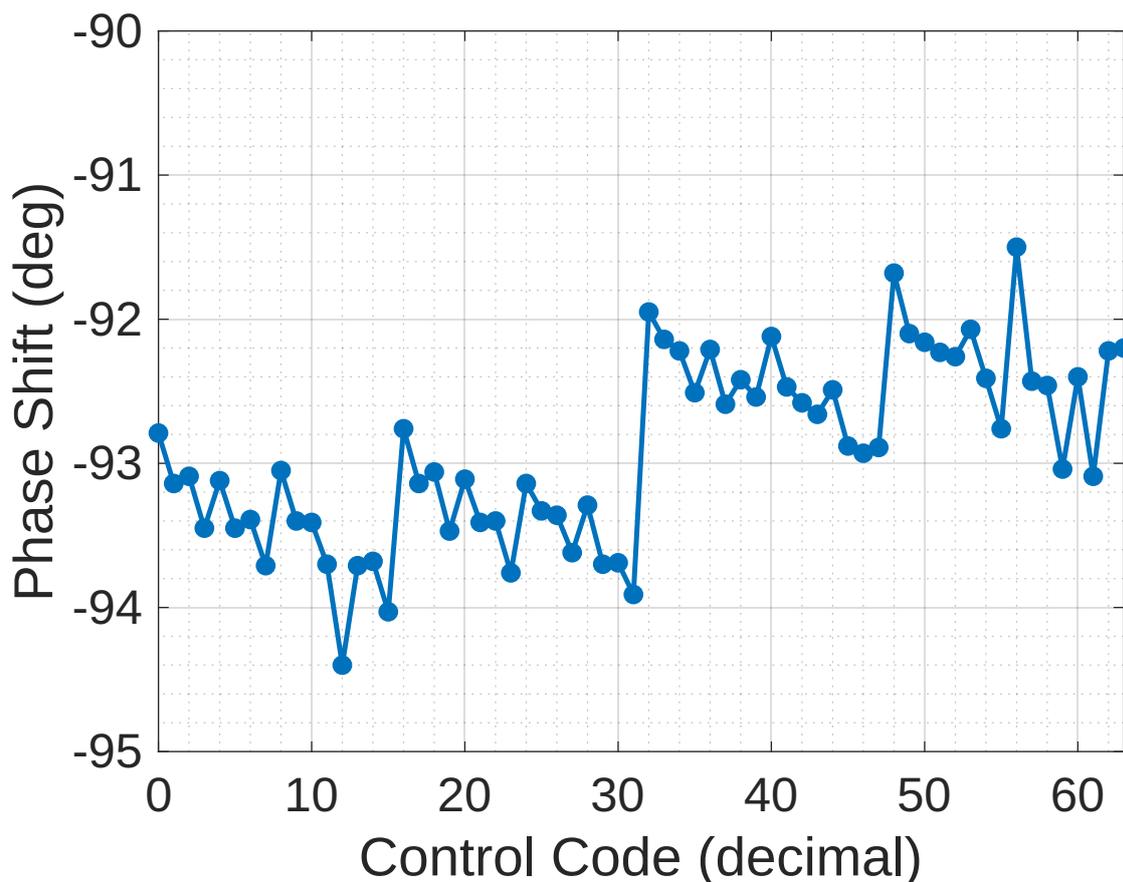


Figure 2.24: Plot of the phase shift of the 31.5 dB digital attenuator vs. the digital control code measured at 100 MHz using the Anritsu 37377C Vector Network Analyzer

2.4 Conclusions and Further Work

A set of digitally-controlled phase shifters and attenuators has been produced that allow for accurate and linear control of the phase shift /attenuation. These include 6-bit phase shifters operating at 100 MHz and 1.5 GHz with a range of 180° and 360° respectively and 6-bit attenuators with a total range of 6.3 dB and 31.5 dB. In all cases, the maximum non-linearity error of the devices was no worse than $0.55 \times \text{LSB}$. The residual phase noise of the devices is low enough such that it could not be measured with any measurement systems that were available. Generic design procedures for designing these devices have been described in detail and examples

with various phase shift / attenuation ranges at various operating frequencies have been presented.

2.4.1 Further Work

Minor Improvements

One of the issues highlighted for the digitally-controlled phase shifters was that the calculated lengths of microstrip tracks required for the desired phase shift often deviated from the length that was found from EM simulation. It was suggested that this was mainly due to the bends in the tracks that were not taken into consideration. To streamline the design procedure, the effect of the track bends could be characterised and incorporated into the length calculation, making it possible to produce phase shifters with different parameters without needing to perform EM simulations.

For the digitally-controlled attenuator it was found that the phase shift varied somewhat for different levels of attenuation and this was especially noticeable for higher levels of attenuation. For low levels of attenuation the layout of the board could be adjusted so that phase shift between the 'on' and 'off' signal paths is as low as possible using the measured results as a reference. More investigation is needed into why the results are less predictable for higher levels of attenuation.

Fixed Phase Shifter

Although these device are very useful in the testing and optimisation phases of a design, as they can be easily changed, consideration should also be made for when further adjustments are no longer necessary. At this point it makes sense to remove unnecessary components, to reduce insertion loss, weight and cost. This is easy enough in the case of the digitally-controlled attenuator, as the device can simply be replaced with a single, fixed pi pad attenuator network. For the phase shifter, it would be possible to create a separate PCB layout including only the required phase shift for a given system, but this could be expensive and time-consuming, especially if a minor adjustment needs to be made later. Work has been ongoing to produce

a fixed version of the phase shifter that simply replaces the relays with selectable solder bridges, so that it easily could replace a digitally-controlled version with a fixed version that has the benefit of getting rid of unwanted losses from the relays.

Chapter 3

Cross-Correlation Residual Phase Noise Measurement System

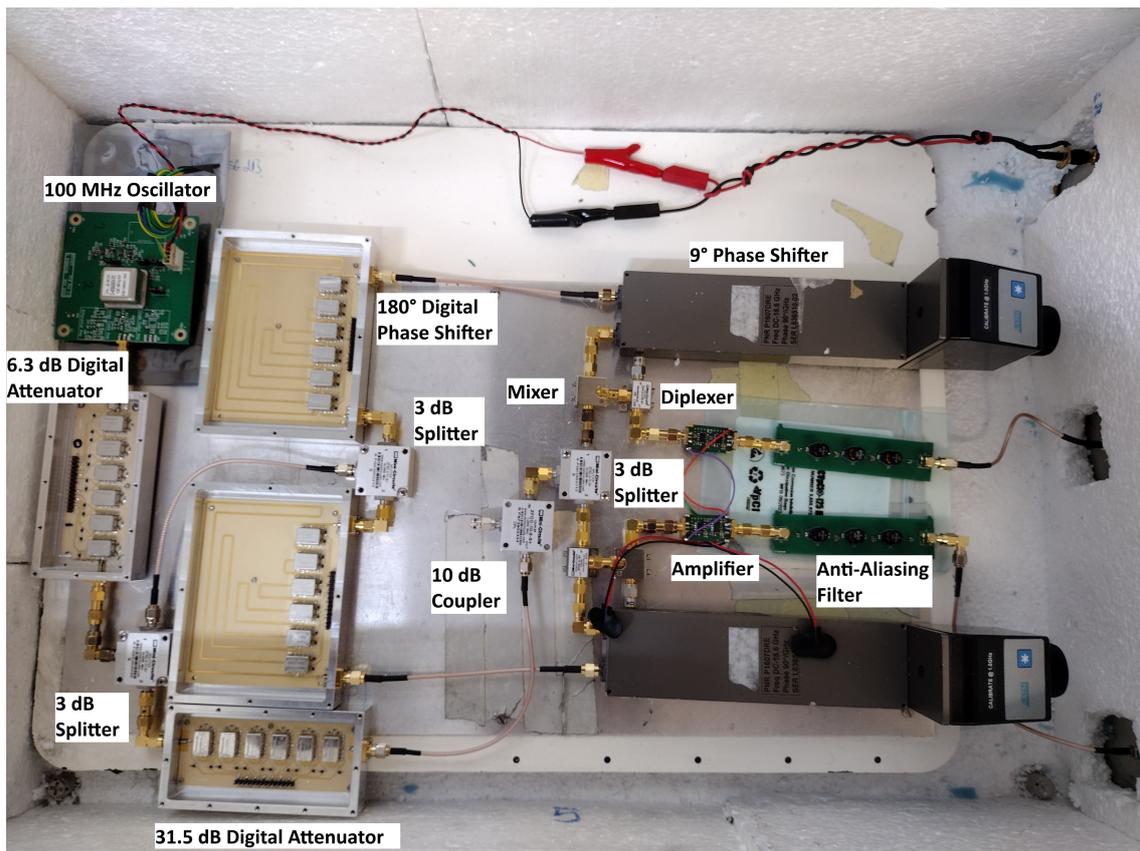


Figure 3.1: Photo of the 100 MHz residual phase noise measurement system

Abstract

In this chapter a residual phase noise measurement system utilising cross-correlation techniques is described. The system is designed to be a low-cost open-source alternative to other commercially-available systems that is built from readily-available components and can be easily modified to perform measurements at a wide range of operating frequencies and power levels.

A 100 MHz low-power measurement system has been built that exhibits an internal noise floor of -177 dBc/Hz when measuring a device operating at a 0 dBm output power. The cross-correlation configuration used has been shown to suppress the internal noise floor of the system by up to 20 dB compared to a similar system in a single-channel configuration. This system has been used to successfully measure the residual noise of the amplifiers described in Chapter 1, which could not be measured with the commercially available Rohde & Schwarz FSWP Phase Noise Analyzer [13] due to its internal noise floor.

A set of efficient digital signal processing techniques for residual phase noise measurements are also described in detail and their performance is measured and compared for a variety of measurement settings. A C++ software library that is intended to be released open-source has been written which enables the raw capture data to be processed efficiently such that residual noise data can be calculated in real-time without significantly increasing the measurement time even on inexpensive hardware. The software includes an algorithm to vary the resolution bandwidth for different frequency bands, so that a large number of cross-correlations can be performed to suppress far-out noise, whilst still maintaining a small resolution bandwidth close-to-carrier. It is demonstrated that this algorithm can be implemented such that the processing time is not significantly increased compared to processing the data without frequency banding. This software has been used to perform a residual phase noise measurement with sub-Hz close-to-carrier resolution bandwidth and over 3,000,000 cross-correlations at a 100 kHz offset in around 13 minutes.

3.1 Introduction

The motivation for this work is the need to measure the residual phase noise of very low-noise amplifiers for use in the 100 MHz crystal described in Chapter 1. It is based on previous work done at York to build a cross-correlation residual noise measurement system [22] [23], but here the system is designed to be low-cost and operate at 100 MHz with a DUT output power of around 0 dBm. The digital signal processing aspects of the system have also been greatly expanded upon to reduce measurement times.

This chapter is structured into three sections. The first section provides an introduction to residual phase noise measurement techniques and explains how cross-correlation can be used to suppress the internal noise floor of the system. The second section outlines the hardware designs for the single-channel and cross-correlation measurement systems and compares their measured internal noise floors. The third section is focussed on the digital signal processing techniques that have been utilised and includes an in-depth explanation and derivation of each of the techniques as well as a review of their relative performance in different measurement scenarios.

3.2 Residual Phase Noise Measurement Theory

In this section a simplified method for measuring residual noise is described and it is shown how the cross-correlation method can be used to suppress the internal noise floor of the system.

3.2.1 Single-Channel Residual Phase Noise Measurement

The idea behind performing residual noise measurements using this method is to measure the difference in phase between a reference signal and the same signal when it has been passed through a device under test (DUT). Spectral analysis of the resulting phase delta reveals the amount of phase noise that is added to the signal by the DUT, i.e. its residual noise.

An implementation of this method is shown in Figure 3.2. Here, a power splitter is used to split the signal between the DUT and reference paths which are then fed into the input of a frequency mixer. An additional 90° phase delay is introduced into the reference path so that the signals at the LO and RF ports of the mixer are in quadrature. In this configuration the mixer acts as a phase detector where any deviation in phase introduced by the DUT is converted to a baseband signal at the output of the mixer. As well as the phase difference, the mixer also produces a signal that is the the sum of the frequencies at the inputs, which must be removed using a suitable low-pass filter.

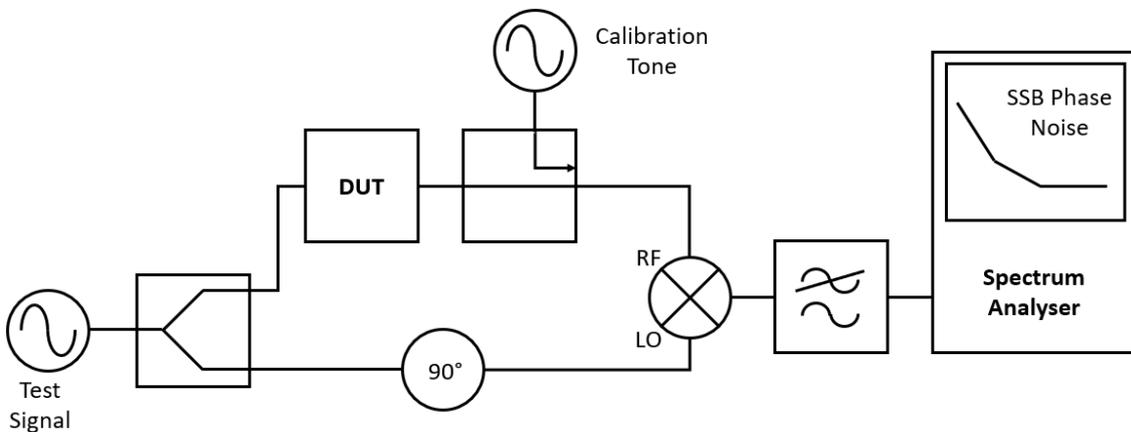


Figure 3.2: Block diagram of a single-channel residual phase noise measurement

A spectrum analyser is then used to see the noise spectral density of the DUT phase noise at a frequency offset from the carrier signal, however the system must be calibrated in order to express the residual noise as single-sideband residual phase noise referenced to the power of the carrier signal in dBc/Hz. The calibration procedure is shown in Section 3.3.2.

3.2.2 Cross-Correlation Residual Phase Noise Measurement

A limitation of the single-channel residual noise measurement system is that any phase noise introduced due to the constituent parts of the system, such as the mixer, filter or the internal noise floor of the spectrum analyser, will also appear in

the measured residual noise and could potentially mask the actual residual noise of the DUT. Furthermore, an actual implementation of this system would likely include additional signal conditioning elements, such as an amplifier to increase the level of the output signal from the mixer, which will also contribute to the system noise floor. Cross-correlation techniques can be used to suppress the internal noise of the system [24].

Simplified System Overview

The cross-correlation configuration shown in Figure 3.3 seeks to suppress the noise contribution of these elements by splitting the measurement over two separate channels, each with its own mixer, amplifier etc. with the idea being that the only things shared between the channels are the DUT and reference signals. By taking the cross-correlation of the resulting noise spectral densities for each channel, the uncorrelated noise in each channel due to the other elements in the system is suppressed while the residual noise due to the DUT present in both channels is left.

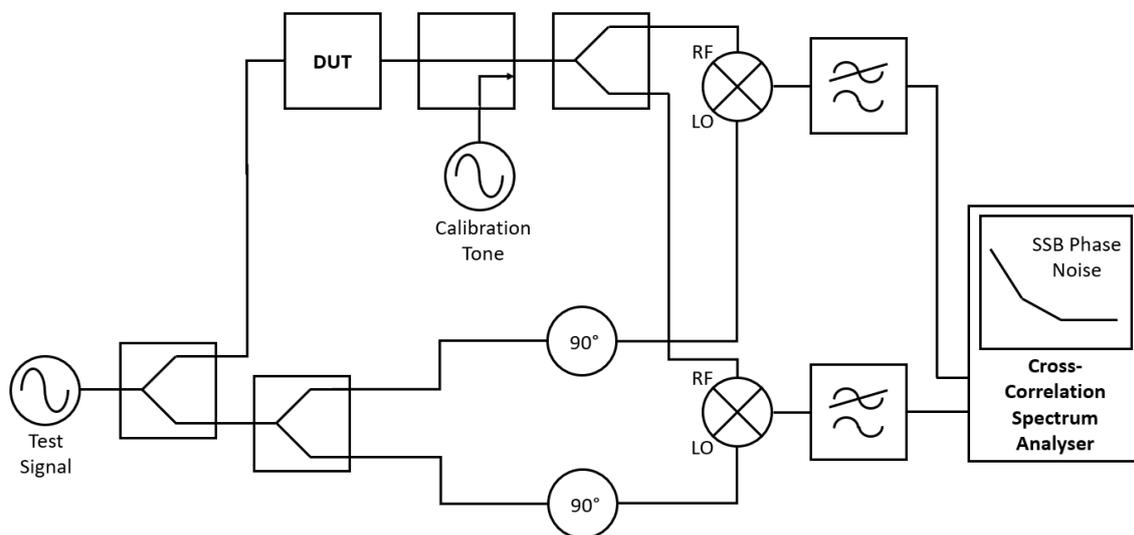


Figure 3.3: Block diagram of a cross-correlation residual phase noise measurement

Derivation

In the model described by Rubiola & Vernotte [25] the discrete time-domain signals captured at the inputs to the cross-correlation spectrum analyser can each be considered to consist of a correlated component that is shared between the two channels, representing the noise of the DUT; and an uncorrelated component that is statistically-independent from the noise present in the other channel, representing the system noise. We can write this for the two channels, x and y , at discrete time intervals, n , as follows:

$$\begin{aligned} x_{n=0 \rightarrow (N-1)} &= a_n + c_n \\ y_{n=0 \rightarrow (N-1)} &= b_n + c_n \end{aligned} \quad (3.1)$$

Where a is the uncorrelated noise in channel x , b is the uncorrelated noise in channel y and c is the correlated noise shared between the channels.

The discrete Fourier transform can then be applied to give the frequency spectrum of each channel for a set of discrete frequency points, k :

$$\begin{aligned} X_{k=0 \rightarrow (N-1)} &= A_k + C_k \\ Y_{k=0 \rightarrow (N-1)} &= B_k + C_k \end{aligned} \quad (3.2)$$

The cross-spectrum of the two channels can then be calculated by performing the cross product of the frequency spectrum of one channel with the complex conjugate of the other and taking the average over M measurements. This is converted to power spectral density by multiplying by the number of samples in sampling window, N , and dividing by the sampling frequency, F_s :

$$\begin{aligned} S_{XY} &= \frac{N}{F_s} \cdot \frac{1}{M} \sum_{m=1}^M [X_m Y_m^*] \\ &= \frac{N}{F_s} \cdot \frac{1}{M} \sum_{m=1}^M [(A_m + C_m) \times (B_m + C_m)^*] \\ &= \frac{N}{F_s} \cdot \frac{1}{M} \sum_{m=1}^M [(A_m B_m^*) + (A_m C_m^*) + (C_m B_m^*) + (C_m C_m^*)] \end{aligned} \quad (3.3)$$

As more measurements are averaged the results will tend towards an expected value, which can be represented as:

$$\mathbb{E}\{S_{XY}\} = \frac{N}{F_s}(\mathbb{E}\{AB^*\} + \mathbb{E}\{AC^*\} + \mathbb{E}\{CB^*\} + \mathbb{E}\{CC^*\}) \quad (3.4)$$

And since the expected value of the product of two statistically-independent signals is 0:

$$\mathbb{E}\{AB\} = \mathbb{E}\{AC^*\} = \mathbb{E}\{CB^*\} = 0 \quad (3.5)$$

Therefore the expectation of the cross-correlation of channels X and Y corresponds to the power spectral density of the DUT.

$$\mathbb{E}\{S_{XY}\} = \frac{N}{F_s} \cdot CC^* \quad (3.6)$$

Noise Floor Suppression

Although the expected values for the terms in (3.5) is 0, a real system would be limited by the number of measurements, and hence the number of averages, that can be performed in a given time frame. These terms represent the system noise floor.

If the expectation in (3.4) is replaced with the average over m measurements, we have:

$$\begin{aligned} \langle S_{XY} \rangle_m &= \frac{N}{F_s} (\langle AB^* \rangle_m + \langle AC^* \rangle_m + \langle CB^* \rangle_m + \langle CC^* \rangle_m) \\ &= \frac{N}{F_s} \cdot CC^* + O\left(\frac{1}{\sqrt{m}}\right) \end{aligned} \quad (3.7)$$

Where $O\left(\frac{1}{\sqrt{m}}\right)$ denotes a result in the order of $\frac{1}{\sqrt{m}}$.

Therefore, the internal noise floor of the system can be suppressed proportionally to the square root of the number of measurements at each frequency point. Expressed in terms of dB this is then:

$$\begin{aligned}
 & 10 \cdot \log_{10}\left(\frac{1}{\sqrt{m}}\right) \\
 & = -5 \cdot \log_{10}(m)
 \end{aligned} \tag{3.8}$$

In other words, the cross-correlation method can be used to suppress the noise floor of the system along with the number of measurements performed at a rate of 5 dB per decade.

3.3 Single-Channel Residual Noise Measurement System

In this section a single-channel residual noise measurement system is described in detail and the internal noise floor of the system is measured.

3.3.1 System Overview

A block diagram of the single-channel residual phase noise measurement system is shown in Figure 3.4.

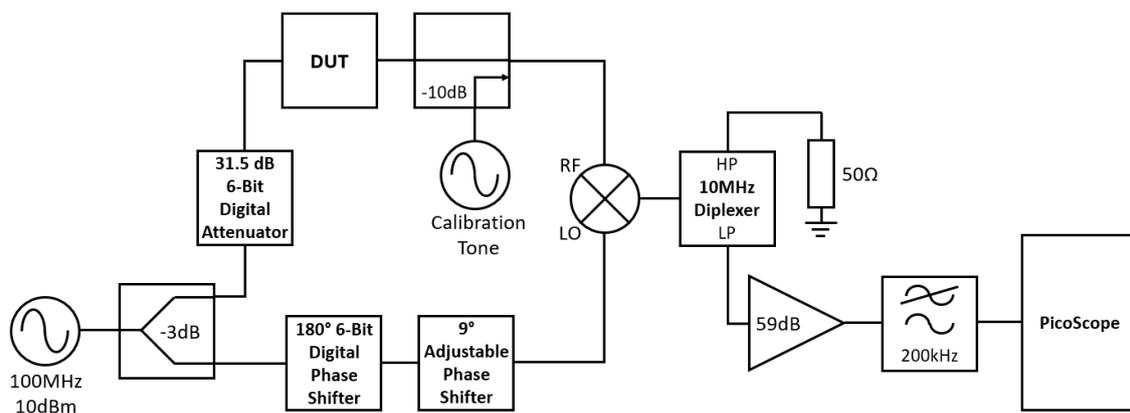


Figure 3.4: 100 MHz single-channel residual phase noise measurement system

Mixers

To optimise the system for the desired DUT output level of 0 dBm, a low-power mixer, the Mini-Circuits ZX05-1L-S+ [26], has been selected. This mixer is designed to operate with an LO power of 3 dBm and at frequencies between 2 to 500 MHz. Using such a low power mixer allows the RF and LO ports to be as close to saturation as possible to achieve the best performance even when measuring low-power device and also removes the need for additional amplifiers at the inputs to the mixer.

Reference Oscillator

A 100 MHz reference oscillator is required to provide the carrier signal for the reference and DUT arms of the system and should have an output power of around 10 dBm so that the power is around 3 dBm at the LO port of the mixer. A KVG O-40-ULPN-100M oven-controlled crystal oscillator [5] has been selected as it exhibits very low phase noise: around -138 dBc/Hz at a 100 Hz offset with a far-out noise floor of around -185 dBc/Hz. It is important that the reference oscillator has as low noise as possible, so that it does not limit the internal noise floor of the system.

Digital-Controlled Phase Shifters and Attenuator

In order to easily set the phase shift of the reference signal path to be roughly in quadrature with the DUT path, two 6-bit digitally-controlled phase shifters are used, which provide a total range of 180° and a resolution of 2.8125° . The phase shift can then be fine-tuned using a 9° mechanical phase shifter. The digital phase shifters were designed specifically to exhibit low residual noise and are described in detail in Section 2.2.

A 6-bit digitally-controlled attenuator, as described in Section 2.3, is also used to adjust the power level of the carrier signal at the input to the DUT. This allows the DUT to be measured at the desired operating conditions. The attenuator has a resolution of 0.5 dB and a total range of 31.5 dB to accommodate high-gain amplifiers.

Signal Conditioning

In this system the unwanted RF portion of the mixer outputs is removed by using a Mini-Circuits ZDPLX-2150-S+ diplexer with the high frequency output terminated with 50Ω and low frequency output connected to the rest of the baseband circuitry. This has the benefit of reducing potential reflections of the RF signal that could interfere with measurements.

The amplitude of the baseband output from the mixers will typically be far too low to be directly measured using a spectrum analyser, and so here a 59 dB amplifier is used to buffer the signal such that it can be accurately measured. The amplifier circuit is shown in Figure 3.5 and is based around two LT6018 operational amplifiers in an inverting configuration each with a DC voltage gain of 30, giving a total gain of 59 dB. The input is matched to the 50Ω output of the mixer via a 50Ω resistor connected to the virtual earth of the amplifier. The reason for using two amplifiers in series is that the bandwidth of the amplifier would not be large enough if a single amplifier with the same gain is used.

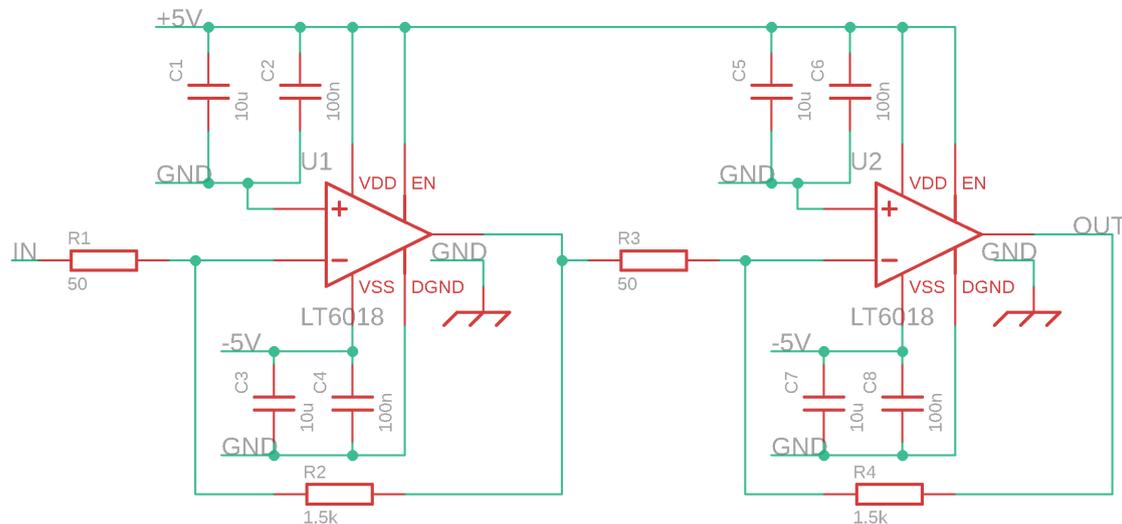


Figure 3.5: Schematic diagram of the baseband amplifier used in the residual phase noise measurement system

It is also important to filter out any frequencies above the Nyquist frequency (half

the sampling rate) prior to being sampled by the PicoScope, as these frequencies will be aliased into the measured noise spectral density. For this reason, an anti-aliasing filter is used at the input to the PicoScope. The circuit for the anti-aliasing filter is a 5th-order Butterworth low-pass LC filter with a cutoff frequency of around 250 kHz, as shown in Figure 3.6, and was originally designed by Elliot Lamb for his work on ultra-low noise noise floor phase noise measurement instrumentation [27].

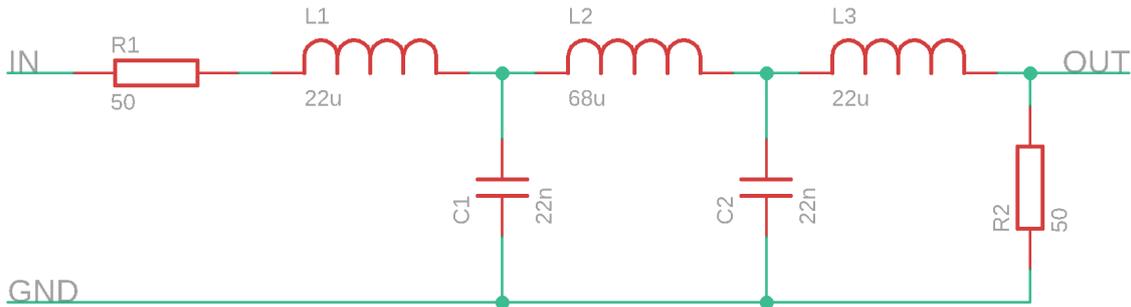


Figure 3.6: Schematic diagram of the 5th-order Butterworth anti-aliasing filter used in the residual phase noise measurement system

Since any frequencies present in the captured signal that are above the Nyquist frequency will be folded back around the Nyquist frequency and added to the measured noise, the filter roll-off should be sufficiently steep so that any frequencies that are folded back into the measurement have been greatly attenuated while also making sure the frequency response inside the measurement bandwidth is as flat as possible. As each of the amplifier stages also has a single-pole roll-off at around 500 kHz due to the limited gain-bandwidth product of the LT6018, the amplifier also contributes to the anti-aliasing. The response of the combined amplifier and anti-aliasing filter measured using the SDR-Kits DG8SAQ Vector Network Analyzer [9] is shown in Figure 3.7.

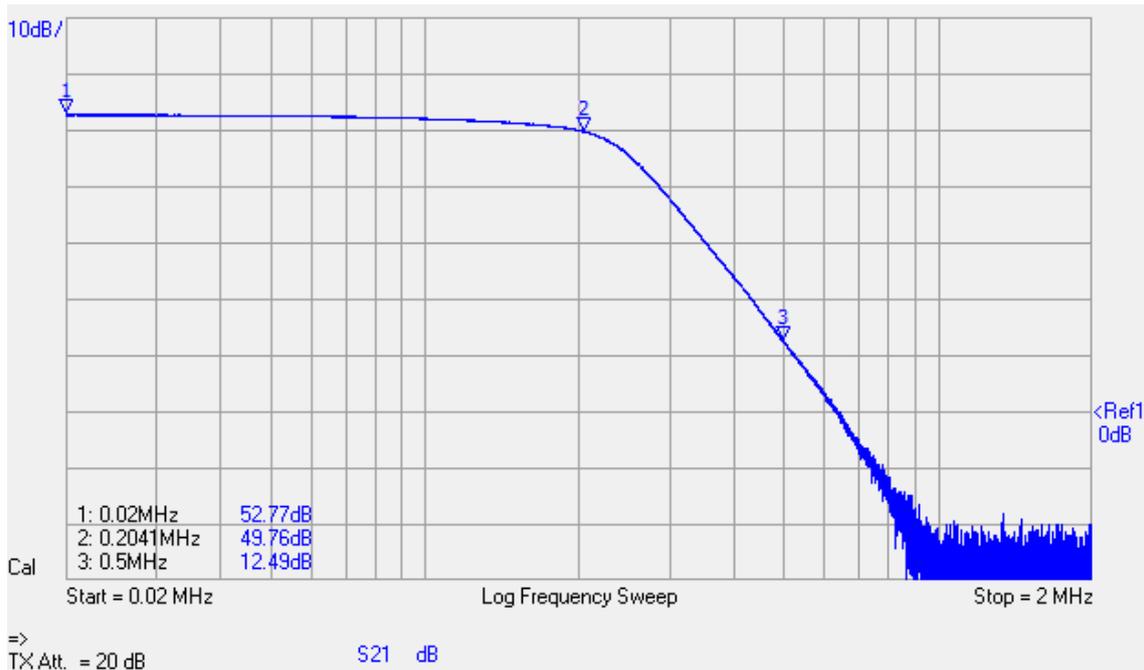


Figure 3.7: Frequency response of the combined LT6018 amplifier and LC anti-aliasing filter measured on the SDR-Kits DG8SAQ Vector Network Analyzer

Digital Signal Processing

Once a block of samples has been captured using the PicoScope and transferred to a connected PC, the samples are processed to calculate the measured noise spectral density in dBc/Hz. More measurements can be taken and averaged to produce a cleaner plot. The digital signal processing involved is described in detail in Section 3.5.

3.3.2 Calibration Techniques

The calibration process involves injecting a tone into the system at the output of the DUT of a known offset frequency and power level in relation to the carrier signal. By observing the amplitude of the peak due to the calibration tone on the spectrum analyser (not the spectral density as for noise measurements) at the system output, a reference can be obtained to convert the absolute noise spectral density measured on the spectrum analyser in dBV/Hz to dBc/Hz. Ideally, the calibration process would

be performed across the full range of offset frequencies that are to be measured.

The following signals are measured in order to calibrate the system:

Calib_{tone} - Power of the injected calibration measured at the RF port of the mixer in dBm.

Calib_{RF} - Power of the RF carrier signal measured at the RF port of the mixer in dBm.

Calib_{out} - Power of the baseband calibration spur measured on the spectrum analyser in dBV.

S_φ - Double-sideband noise spectral density measured on the spectrum analyser in dBV/Hz.

The single-sideband power spectral density can then be calculated by converting from the double-sideband power spectral density (-6 dB) and applying the calibration values as follows:

$$L(f) = S_{\phi} + (\text{Calib}_{\text{tone}} - \text{Calib}_{\text{RF}}) - \text{Calib}_{\text{out}} - 6 \text{ dB} \quad (3.9)$$

To verify that the calibration is valid across the desired offset frequency range of the system, a calibration tone of known offset frequency from the carrier was injected into the system at the RF port of the mixer and the amplitude of the resulting baseband spur was measured on the PicoScope. A flat-top window was applied to the captured data samples, so that the measured peak amplitude would be accurate even in situations where the spur falls between frequency bins. This was repeated at various offset frequencies and the results are shown in Figure 3.8 as a series of superimposed spectra.

The frequency setting of the signal generator for each measurement is shown on the graph, but due to the slight mismatch between the operating frequency of the signal generator and the reference oscillator used in the measurement system, the measured spurs are slightly offset from their nominal frequency, which is especially apparent at lower frequencies. This has no effect on residual phase noise measurements, however.

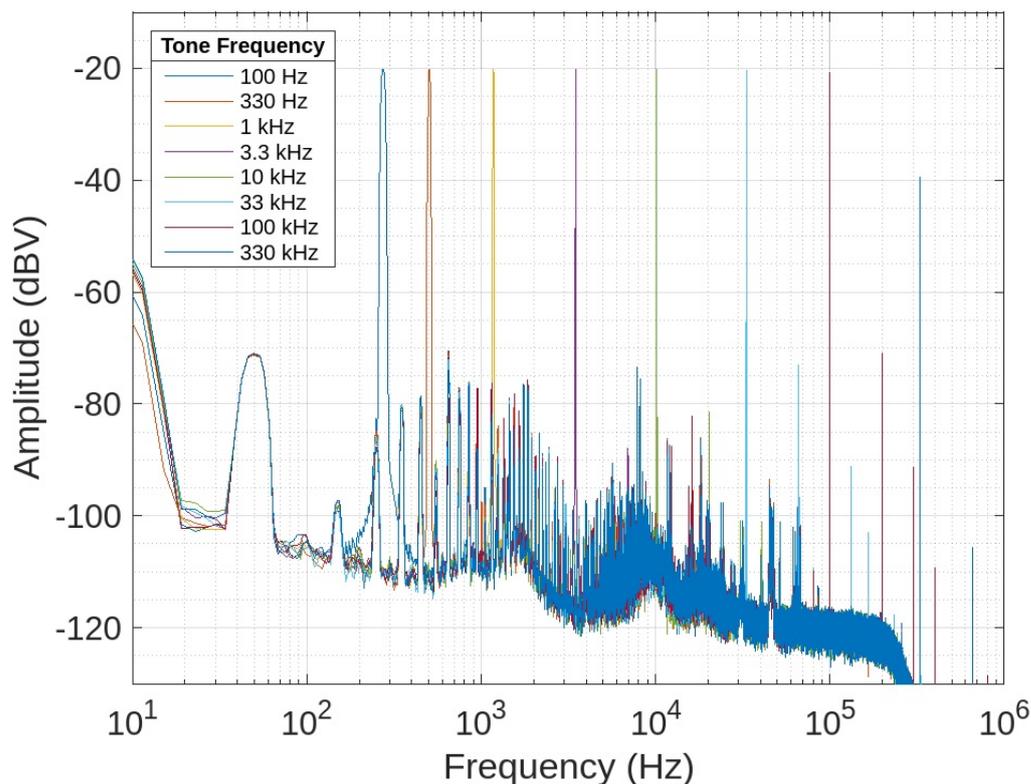


Figure 3.8: Plot of the calibration spur measured on the PicoScope when a -30 dBm calibration tone at various frequencies is injected into the signal path of the single-channel residual noise system

The peaks show very little deviation in amplitude within the desired measurement range (< 100 kHz) with the maximum being around 0.6 dB, which means that performing the calibration procedure at a single frequency should be sufficient to produce a reasonably accurate residual noise measurement. The peak amplitude drops significantly when the offset frequency is greater than 100 kHz due to the combined effect of the anti-aliasing filters and the limited bandwidth of the amplifiers, so valid measurements can not be made at these frequencies.

3.3.3 System Noise Floor

The internal noise floor of the single channel system was measured by placing a straight-through connection in place of the DUT and setting the measurement to run for around 30 minutes with a minimum resolution bandwidth of 0.125 Hz. The results are shown in Figure 3.9. It can be seen that the single-channel system has a minimum noise floor in the valid offset frequency range (up to around 100 kHz) of around -166 dBc/Hz.

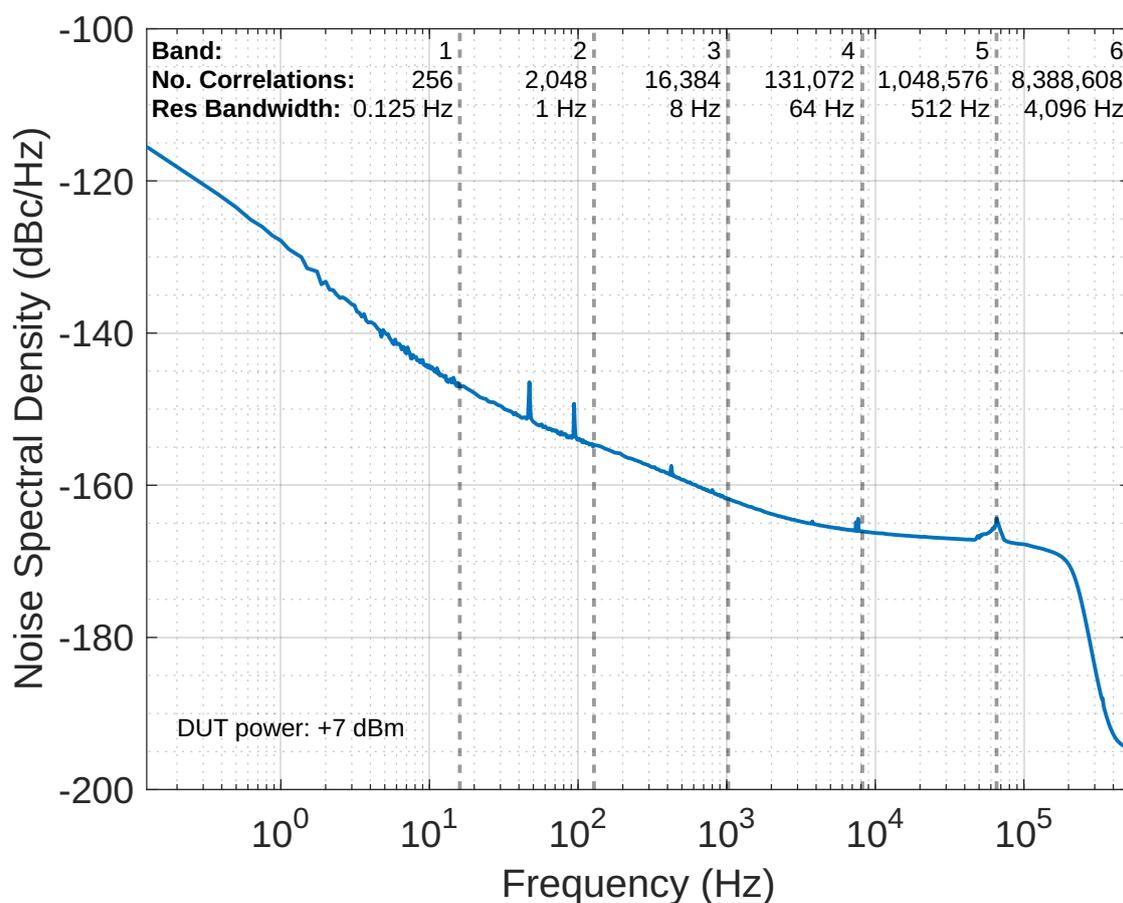


Figure 3.9: Internal noise floor of the single-channel residual phase noise measurement system

3.3.4 Limitations of Single-Channel Residual Noise Measurements

The main limitation of the single-channel residual noise system is the internal noise floor. Without any way to suppress the noise from the various elements that make up the system, it is very difficult to measure the residual noise of particularly low-noise devices or of devices operating at low power levels.

3.4 Cross-Correlation Residual Noise Measurement System

In this section the residual noise measurement system is set up in a cross-correlation configuration. The internal noise floor of the system is measured for various DUT power levels and compared to the noise floor of the single-channel system. The accuracy of the system is also demonstrated by comparing a measurement of an amplifier with both its theoretical residual noise and with an existing residual noise measurement system; the Rohde & Schwarz FSWP.

3.4.1 System Overview

The block diagram of the cross-correlation residual noise system is shown in Figure 3.10. Here, the DUT and reference arms of the signal path have each been split so that there are now two baseband channels that can be measured simultaneously with the PicoScope. Using the technique described in Section 3.2.2, the noise of the mixers, filters and amplifiers can be suppressed.

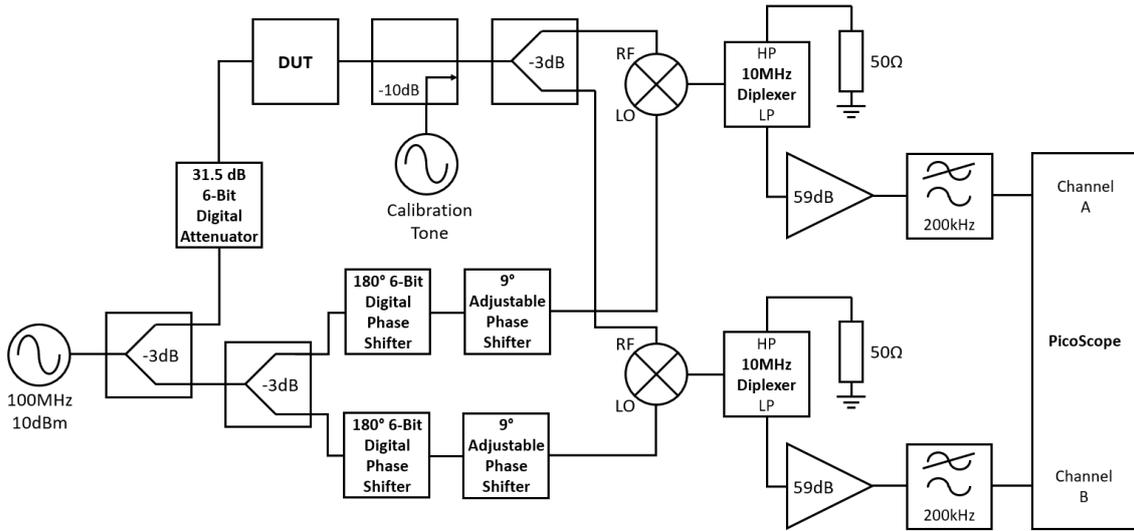


Figure 3.10: 100 MHz cross-correlation residual phase noise measurement system

3.4.2 Measurement Accuracy and Noise Floor

Measurements with this system were made using a similar methodology to that described for the single channel system in Section 3.3. The cross-correlation digital signal processing techniques used are described in Section 3.5.

Single Channel vs. Cross-Correlation Noise Floor Comparison

The simplest way to verify the ability of the system to suppress the uncorrelated noise between the two channels is to look only at the baseband front-end of the system, i.e. the amplifier and anti-aliasing filter for each channel going into the PicoScope as shown in Figure 3.11. This way any correlated noise between channels can be kept to a minimum, so that suppression of the uncorrelated noise can be more easily seen. The battery-powered amplifiers are effectively acting as independent noise sources. For this reason, the system front-end can have a major impact on the minimum achievable noise floor of residual noise measurements and so it is important that the system is able to suppress this noise effectively.

A comparison of the measured noise spectral density for single channel and cross-correlation configurations are shown in Figure 3.12. The single channel measurement

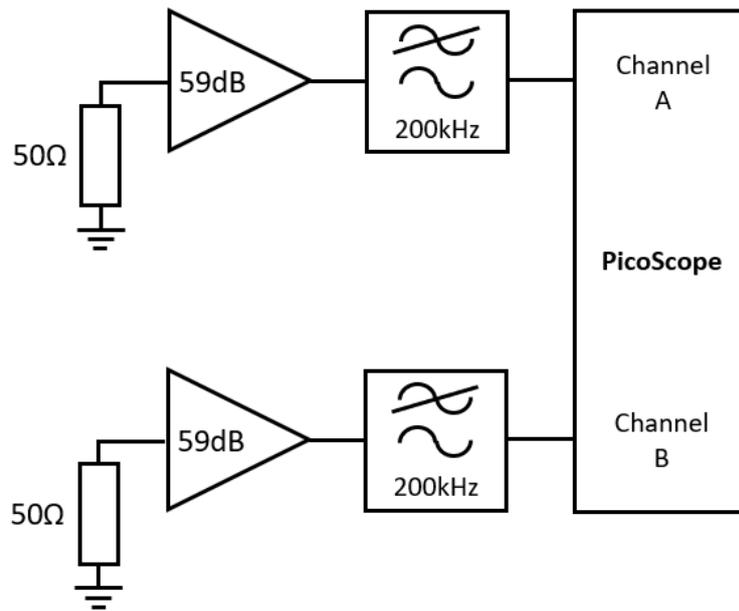


Figure 3.11: System front-end setup used for the comparison of single channel and cross-correlation measurement techniques

uses the same configuration, but only the samples from Channel A are used to calculate the residual noise. As there is no carrier signal in this configuration, the results are given in terms of absolute noise spectral density, dBV/Hz. The frequency banding techniques described in section 3.5.1 are used here so that more correlations are performed at higher offset frequencies, or more averages in the case of the single channel measurement.

It can be seen that using cross-correlation techniques allows suppression of the noise floor across the full range of the measurement with a maximum suppression of 30 dB. The suppression increases along with the number of correlations in each band, however the cross-correlated noise floor begins to increase at offsets above 10 KHz. This could be due to internal correlated noise between channels in the PicoScope.

Since the ratio of increased correlations with each subsequent band is 8, we would expect the cross-correlated noise floor to decrease each time by a factor of $10 \cdot \log_{10}(\sqrt{8}) \approx 4.516...dB$ and the suppression in the first band to be $10 \cdot \log_{10}(\sqrt{256}) \approx$

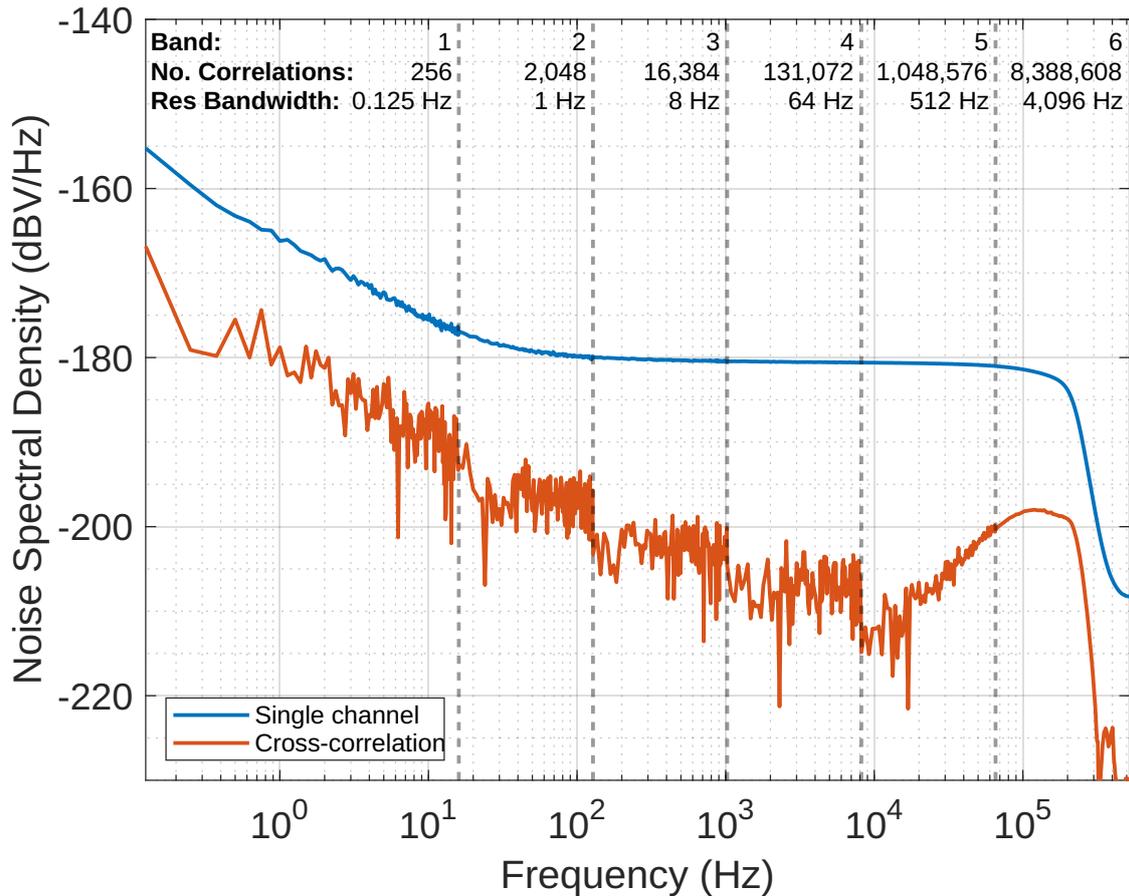


Figure 3.12: Comparison of system front end noise floor (amplifier, filter, PicoScope) using single channel and cross-correlation measurement techniques

12.041...dB. Although it is hard to determine an exact value for the suppression in each band, it can be seen these predictions hold roughly true, at least up to 10 KHz.

The same measurement was performed, but this time using the full system setup as shown in Figure 3.4 and Figure 3.10. To measure the noise floor, a straight through connection was used as the device under test and the power level at this point was +7 dBm. The system was calibrated so that the results can be presented in terms of noise spectral density in relation to the power of the carrier signal at the input to the DUT, dBc/Hz. The results are shown in Figure 3.13.

As for the previous measurement of the baseband part of the system, the cross-correlation measurement of the full system also exhibits noise floor suppression com-

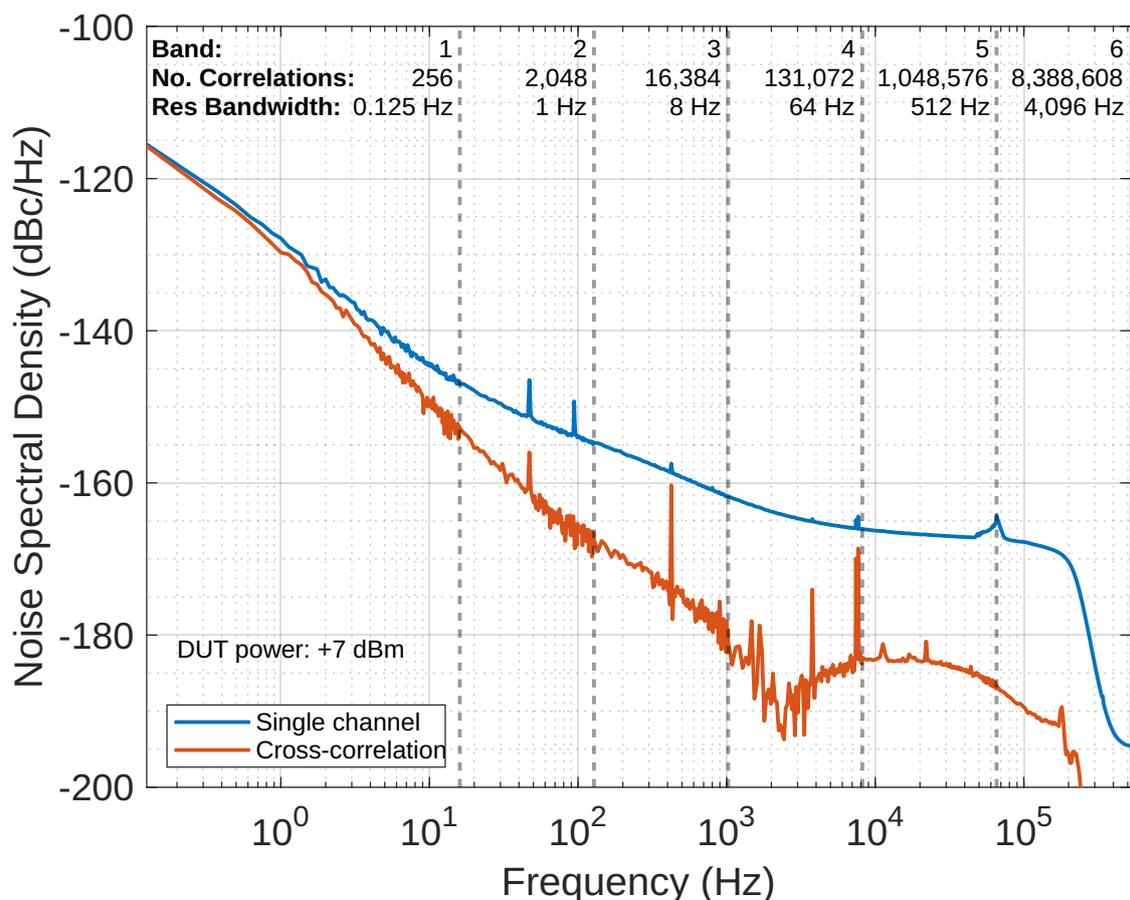


Figure 3.13: Comparison of system noise floor using single channel and cross-correlation measurement techniques

pared to the single channel measurement, however by not as much, especially at close-to-carrier frequencies. There is still significant suppression up to 20 dB for higher offset frequencies, but this reduces to <5 dB for frequencies below 10 Hz.

It should also be noted that there is a dip for the cross-correlation measurement at around 2-3 kHz, which may be the result of cross-correlation collapse [28]. Therefore the measured noise spectral density for this range may actually be below the actual noise floor of the system.

DUT Power Level Noise Floor Comparison

Further measurements were made of the cross-correlation system to see how the power level at the output of the device under test (and hence the power level at the RF ports of the mixers) affect the internal noise floor of the system. Again, the DUT in this case is a straight-through connection and the power at this point is set using the 6-bit digital attenuator.

The measurement is shown in Figure 3.14. For all of the power levels tested there was very little deviation in the relative noise spectral density at offset frequencies below 1 kHz. The far out noise floor scales with the DUT power level such that increasing the power by 7 dB also decreases the noise floor by 7 dB.

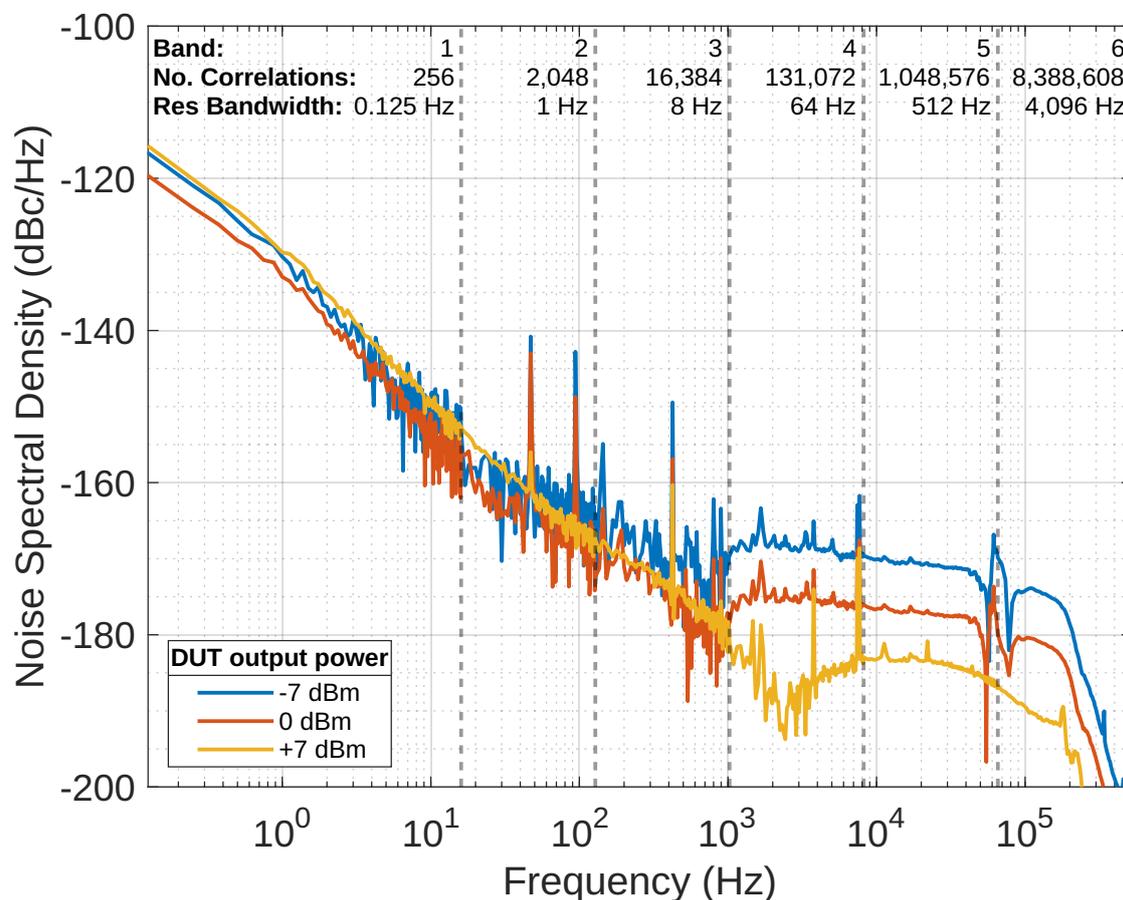


Figure 3.14: Comparison of the cross-correlation system noise floor for various DUT output power levels

Amplifier Measurements

To confirm the ability of the system to accurately measure residual noise, a measurement should be performed using a DUT with a known noise floor. If an amplifier operating at an input power of P_{AVS} in dBm has a measured noise figure, NF, then the noise floor of the amplifier in dBc/Hz can be calculated using (3.10) from Section 1.2.2 as follows:

$$NoiseFloor = -177 \text{ dB} - P_{AVS} + NF \quad (3.10)$$

As an example, if the MAR-6 amplifier discussed in Chapter 1, which has a noise figure of 1.9 dB, is operating with a P_{AVS} of -23 dBm the noise floor is then:

$$-177 + 23 + 1.9 = -152.1 \text{ dBc/Hz} \quad (3.11)$$

The cross-correlation residual noise system was used to measure the MAR-6 at 100 MHz under these operating conditions and the results are shown in Figure 3.15. To confirm that the measured results are above the noise floor of the system itself, a separate measurement was made where the amplifier was replaced with a straight-through connection and the power level was set so that the power into the RF ports of the mixers was the same for both measurements. The results confirm that the measured far-out noise floor of the amplifier, -151.5 dBc/Hz, is very close to the predicted value of -152.1 dBc/Hz and well above the internal noise floor of the system.

As a secondary check, the same measurement was performed using the Rohde & Schwarz FSWP phase noise analyser as a comparison, the results of which are also shown in Figure 3.15. Both measurements were set up to run for around 30 minutes with similar resolution bandwidth and correlation factor settings. Comparisons between the measurements cannot easily be made at frequencies below 1 kHz as the amplifier noise floor is not sufficiently above the noise floor of the FSWP, however the measured far-out noise is very similar for both measurement systems.

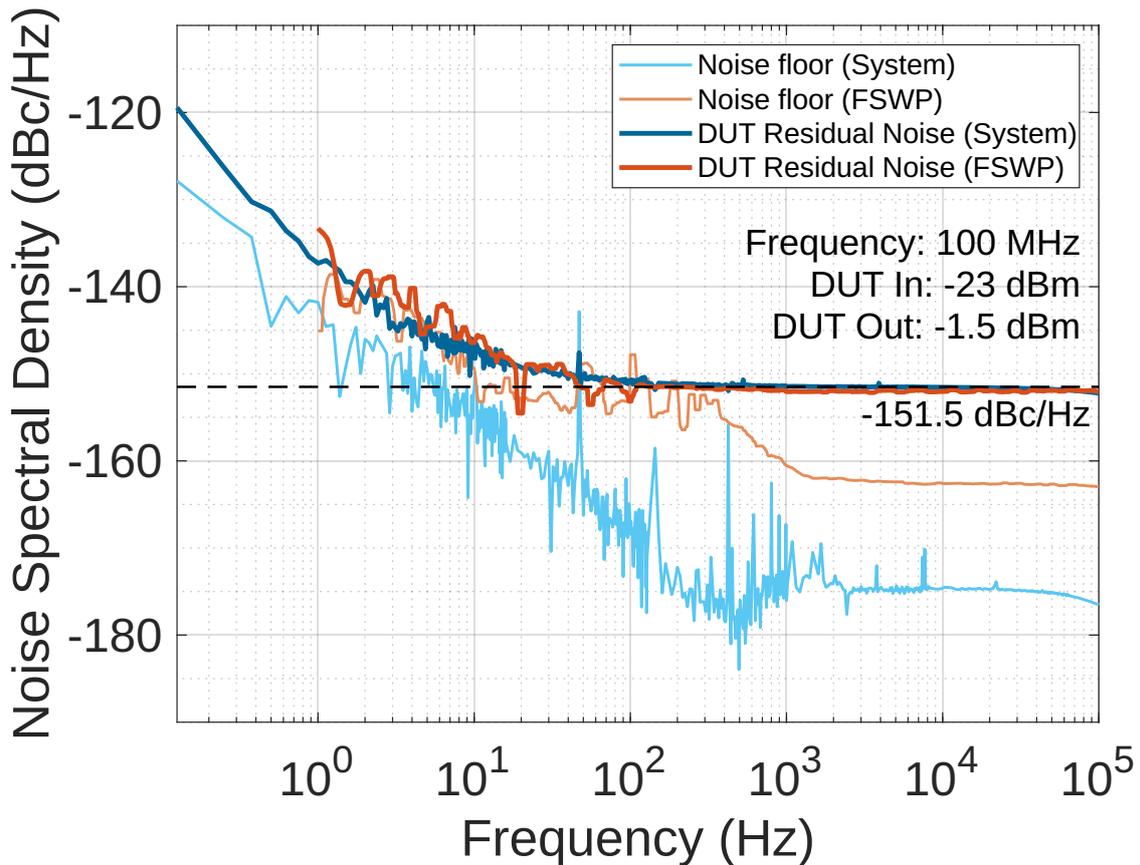


Figure 3.15: Comparison of the residual noise of the MAR-6 amplifier when measured using the system described in this chapter versus the Rohde & Schwarz FSWP Phase Noise Analyzer

As the DUT residual noise measured with this system is in agreement with both the theoretical noise calculation and the noise measured with an established system in the FSWP, confidence can be placed in the accuracy of residual noise measurements made using this system.

Comparisons with Other Systems

It can also be seen in Figure 3.15 that despite its comparatively low cost, the internal noise floor of this system when using a similar resolution bandwidth and measurement time is significantly lower than that of the Rohde & Schwarz FSWP. It was possible to accurately measure the flicker noise corner of the MAR6 amplifier

using this system, which was not possible with the FSWP due to its close-to-carrier internal noise floor.

The main drawback of this system when compared to the FSWP is the relative difficulty in setting up a measurement and performing the calibration; the phase shifters must be set manually to get the required 90° phase shift between the RF and LO ports of the mixers and the calibration procedure involves using external measurement equipment. This is an automated process with the FSWP, which makes measurements significantly easier to perform.

When compared to the system described by Bale et al. [22] [23] based on two HP 11848A units, this system does not exhibit such a low internal noise floor. This is not particularly surprising as this system is designed to operate at much lower power levels. However, since the 11848A units are replaced with inexpensive amplifiers, mixers and filters, the result is a much lower-cost and compact system that can be easily modified if necessary. The number of correlations versus the measurement time is also greatly improved thanks to the digital processing techniques discussed in Section 3.5.

3.5 Digital Signal Processing Techniques for Residual Noise Measurements

This section describes in detail the digital signal processing techniques that are used in the residual phase noise measurement system, focussing primarily on the efficient calculation of Fourier transforms and the optimisation of residual noise measurements through frequency banding techniques.

An efficient variable window size frequency banding algorithm for processing residual noise measurements has been developed, which allows for a measurement with sub-Hz close-to-carrier resolution bandwidth and over 3,000,000 cross-correlations at a 100 kHz offset to be performed in around 13 minutes.

All of the processing algorithms described in this section are to be released as open-source C++ software library.

System Requirements

The primary purpose of the software that handles the digital signal processing for the residual noise system is to convert time-domain domain voltage samples captured at the inputs to a cross-correlated power spectral density of the device under test. This involves performing the digital Fourier transform for each channel and calculating the cross spectrum. This means the software must be able to calculate the digital Fourier transform of a block of samples and calculate the cross spectrum of two sets of frequency-domain samples.

The secondary purpose of the software is to reduce the amount of time it takes to obtain useful measurements by utilising a frequency-banding algorithm that can adjust the resolution bandwidth versus the number of correlations for different frequency ranges of the same measurement, which makes it possible to achieve a high degree of internal noise floor suppression at high offset frequencies from the carrier while maintaining a narrow resolution bandwidth close-to-carrier. This algorithm should be flexible to allow measurements to be optimised for different resolution and

noise floor requirements and it should be efficient so that the processing time does not significantly impact the overall measurement time.

Existing Systems

The residual noise measurement system described by Bale et al. [22] [23] that the work in this chapter is based on does not include any techniques for adjusting the resolution bandwidth versus number of correlations in different frequency bands, and as such, achieving a low instrument noise floor required either long measurement times or very limited resolution bandwidths. This section looks to improve these aspects of the system through digital signal processing techniques.

There exists many commercial systems such as the Rohde & Schwarz FSWP Phase Noise Analyzer that employ frequency banding techniques to improve the instrument noise floor at higher offset frequencies, however the implementation of these techniques is not typically described in detail and there is relatively little scope in adjusting the banding algorithm to suit different measurements. Here, a thorough review of various digital signal processing techniques are presented and the algorithms are described in detail.

There does exist in the literature an open-source direct digital phase noise measurement system described by Holme [29]. Although this system is designed for measuring absolute phase noise, similarly to the techniques described in this section it performs digital signal processing on time-domain voltage samples to produce a frequency-banded power spectral density measurement. The main difference being that Holmes' system uses an FPGA-based hardware solution for this processing, whereas a purely software-based approach is described in this section. Additionally, Holmes' system uses a decimation-in-time approach for the frequency banding algorithm in contrast to the variable window size method that is explored here. The software libraries that have been written as part of the work described in this section are designed to be much more flexible in terms of how its frequency banding is employed and can be run directly on a PC, so that they can be used with a wide

variety of systems or simply to process raw captured data.

3.5.1 Variable Window Size Frequency Banding Algorithm

When performing a residual noise measurement the sampling rate and the number of samples captured in each data block determine both the resolution bandwidth and the overall frequency range of the output data along with the number of measurements that can be performed in a given time frame, which affects by how much the internal system noise can be suppressed in a cross-correlation system. If the sampling rate remains fixed, then capturing more points in each data block will improve the resolution bandwidth, however each measurement will take longer meaning that fewer cross-correlations can be performed over the same-time frame, so the system noise floor will be higher.

It is often desirable to have a narrow resolution bandwidth when measuring the close-to-carrier noise of a device under test; this is so that it is possible to accurately measure the flicker noise of the device. The $1/f$ noise curve means that the measured noise close-to-carrier is often much greater than the far-out white noise and so having a large number of cross-correlations to suppress the system noise floor is comparatively less critical at low offset frequencies. A much wider resolution bandwidth can be used to measure the far-out noise floor of the device, but as many correlations should be performed as possible so that this measurement is not affected by the internal noise floor of the system. The idea behind splitting the measurement into different frequency bands is so that the resolution bandwidth versus the number of correlations can be optimised for each band.

Variable Window Size Method

One method by which it is possible to process captured data samples to produce a residual noise measurement where the resolution bandwidth versus number of cross-correlations are varied for different frequency bands is the variable window size method. Here, a captured set of data samples, the data block, is split evenly

into a number of sets of consecutive data points, which effectively allows more measurements to be performed in the same time frame by reducing the window size of the input data to the Fourier transforms.

When a Fourier transform is performed on data captured using a rectangular sampling window of size N and a sampling rate F_s , the resolution bandwidth of the output data is:

$$R_{bw} = \frac{F_s}{N} \quad (3.12)$$

Therefore if the sampling rate remains fixed, i.e. the data block is split into sets with consecutive data points, reducing the window size will result in a wider resolution bandwidth.

This means that the same data block can either be split into many sets of a small window size to produce many measurements (and hence many cross-correlations), but with a wide resolution bandwidth; or into a few sets of a large window size that results in a very narrow resolution bandwidth, but with reduced system noise suppression.

Examples

Figure 3.16 demonstrates how frequency banding can be used to improve residual phase noise measurements. Figure 3.16(a) shows a hypothetical measurement of a device using a large sampling window without any frequency banding that has been run long enough to perform 10 cross-correlations. A Fourier transform is performed on each data block without splitting the block into sets of a smaller window size. It can be seen that the narrow resolution bandwidth makes it very easy to discern the $1/f$ flicker noise of the device, however the low number of cross-correlations means that the system noise has not been suppressed enough and so the noise floor of the device under test cannot be determined.

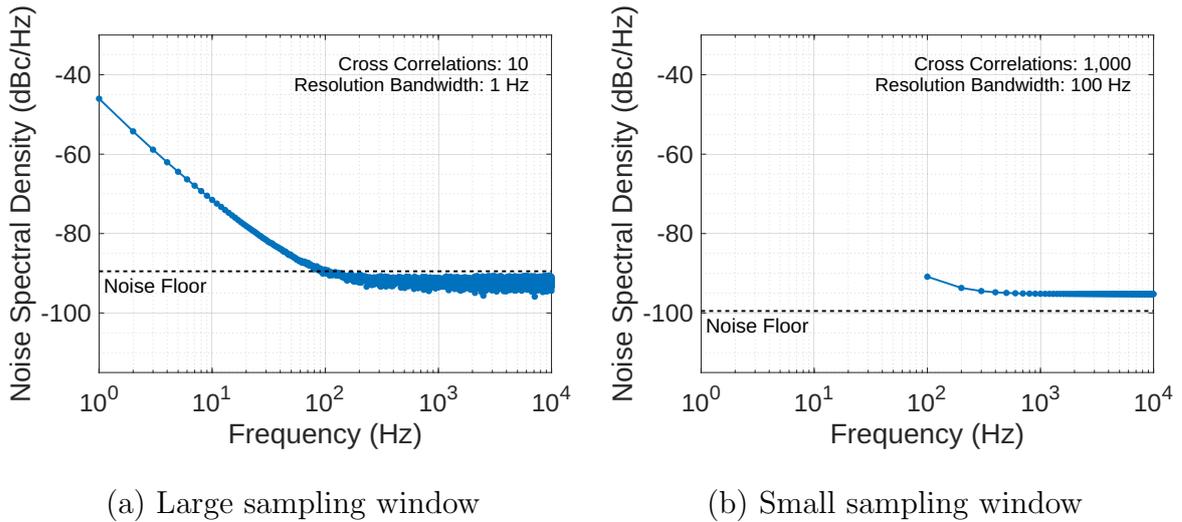
In contrast, Figure 3.16(b) shows the same measurement, but where each captured data block is split into 100 sets before performing the Fourier transforms.

For a measurement in the same time frame the number of cross-correlations has increased by a factor of 100 and so the suppression of the system noise floor has improved by 10 dB. This is sufficient to allow the noise floor of the device under test to be determined, however the resolution bandwidth of the measurement has also increased by a factor of 100, so the data points are far too sparse at close-to-carrier frequencies to see the flicker noise of the device.

Figure 3.16(c) shows how frequency banding can be used to improve the measurement in this case. Here, the measurement is split into 4 bands, each representing a frequency decade. The results in the lowest frequency band are calculated as for (a), i.e. the whole data block is used as the input for the Fourier transforms. For each subsequent band the results are obtained by first dividing the size of the data sets by 10, which means both the resolution bandwidth and the number of correlations that can be performed also increase by 10. It can be seen that this produces a plot that combines the desirable elements of (a) and (b); narrower resolution bandwidths close-to-carrier, so the shape of the flicker noise corner can be easily seen and greater suppression of the system noise floor at far-out frequencies, so that noise floor of the device can be accurately measured.

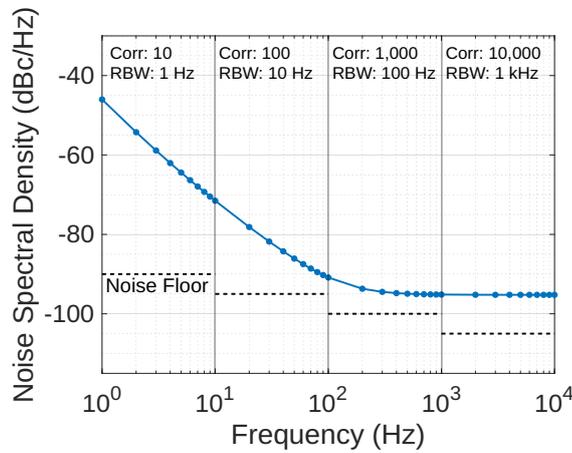
There are a number of ways in which this method can be adjusted to suit the device that is to be measured. For example, the sizes of the frequency bands used do not need to be uniform; a larger band spanning multiple decades could be used if a certain suppression level or resolution bandwidth is required to be maintained over a wider range. Such an option is incorporated into this system in the form of a band 'shift' operation.

Examples of how this works are given in Figure 3.17. If the specified number of bands to use is fewer than will fill up the entire frequency range of the measurement, the system will default to setting the higher frequency bands as normal and then extending the lowest frequency band to fill the space, as in Figure 3.17(b). This has the effect of increasing the noise floor suppression at the lowest frequencies, but without the narrow resolution bandwidth close-to-carrier that would otherwise be



(a) Large sampling window

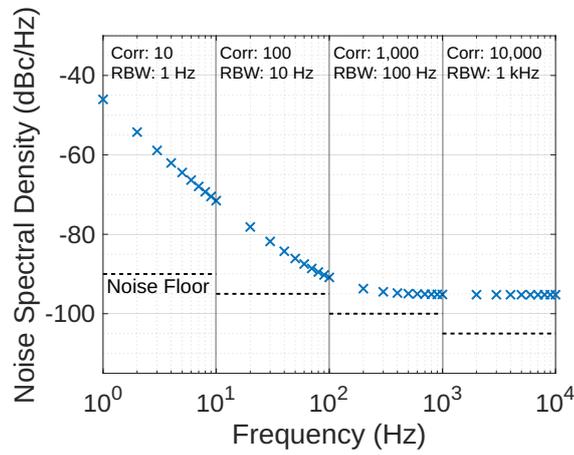
(b) Small sampling window



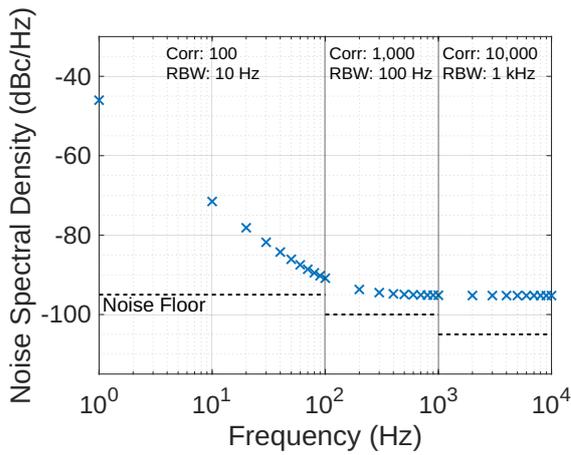
(c) Variable sampling window

Figure 3.16: Demonstration showing a comparison of variable and fixed sampling window sizes for a theoretical 10 sec cross-correlation measurement at 200 kHz

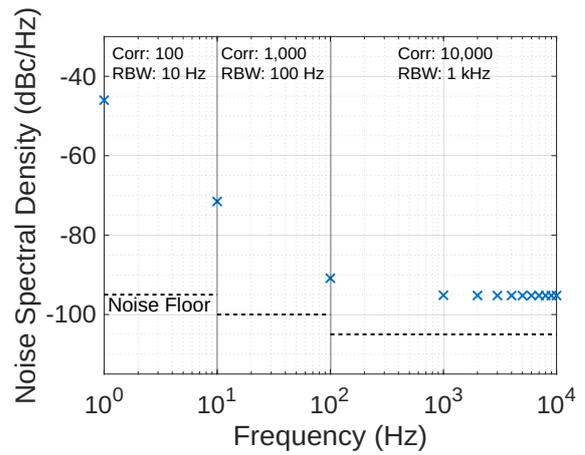
achieved by using more frequency bands. If the placement of the bands is shifted left, as in Figure 3.17(c), this allows for the noise floor suppression to be prioritised over the resolution bandwidth for intermediate frequencies. This method allows measurements to be adjusted to produce useful results even in situations where it is not practical to measure very narrow resolution bandwidths or the system noise floor needs to be as low as possible over a wide frequency range.



(a) 4 bands - no shift



(b) 3 bands - no shift



(c) 3 bands - shifted left

Figure 3.17: Demonstration showing a comparison of different frequency banding styles for a theoretical 10 sec cross-correlation measurement at 200 kHz

Banding Algorithm

During a measurement, the data block of captured voltage samples is split up into separate smaller data sets, the number of which differs for each band, and allows for additional cross-correlations to be performed at the expense of resolution bandwidth. The number of these sets is largest in the final band (the highest offset frequency) and is limited by the minimum allowable sampling window size, *MinWindowSize*, for producing Fourier transforms of the captured data. For each subsequent band going down in frequency, the number of sets is then divided by the band correlation ratio, *BandRatio*, which has the effect of reducing the number of cross-correlations that can be performed in the same time frame, but improves the resolution bandwidth due to the increased number of samples in each measurement.

The number of sets per band, b , can be calculated as:

$$NumSets_b = \frac{DataSize}{MinWindowSize \cdot BandRatio^{b-1}} \quad (3.13)$$

The sampling window size, N , per band for a captured data block of size *DataSize* is then:

$$N_b = \frac{DataSize}{NumSets_b} \quad (3.14)$$

The resolution bandwidth, *ResBW*, is equal to the sampling rate divided by the sampling window size when using a rectangular window function, so the resolution bandwidth for each band is then:

$$ResBW_b = \frac{SampleRate}{N_b} \quad (3.15)$$

An FFT function is applied individually to each set in the band, and since only positive frequencies are considered, the size of the FFT outputs, k , for each band is:

$$K_b = \frac{N_b}{2} + 1 \quad (3.16)$$

Each band covers only a limited frequency range, so not all of the FFT output data points are needed. To save computation time and to be able to collate the separate data from each band into the final cross-correlated output data, it is necessary to know the start of the range of useful FFT output data, $kStart$, for each band as well the end, $kStop$. $kStart$ is the FFT output data point that corresponds to the lower frequency range of the band and $kStop$ is one before the FFT output data point that corresponds to the upper frequency range of the band (or the final data point if this is the highest frequency band). This algorithm also takes into account the ability to shift the position of the frequency bands by multiples of $BandRatio$ in order to prioritise resolution bandwidth over number of correlations and vice versa. The range for each band can be calculated as follows:

$$kStart_{b=1} = 0 \tag{3.17}$$

$$kStart_{b=2 \rightarrow B} = \frac{MinWindowSize}{2 \cdot BandRatio^{BandShift+2}}$$

$$kStop_{b=1 \rightarrow (B-1)} = \frac{MinWindowSize}{2 \cdot BandRatio^{BandShift+1}} - 1 \tag{3.18}$$

$$kStop_{b=B} = \frac{MinWindowSize}{2}$$

Cross-correlation of the FFT data for each channel is performed over this range and the size of the overall output data, S for this band is then:

$$S_b = kStop_b - kStart_b + 1 \tag{3.19}$$

It is also necessary to know where this band's data should be placed in the final correlated output data. This point, $sStart$ is calculated by simply adding the data size, S , of the previous band to its start position as follows:

$$sStart_{b=1} = 0 \tag{3.20}$$

$$sStart_{b=2 \rightarrow B} = sStart_{b-1} + S_{b-1}$$

3.5.2 Fast Fourier Transform vs. Discrete Fourier Transform

Since these residual phase noise measurements involve converting digitally-stored time-domain voltage samples into frequency-domain power spectral density, it is necessary to apply a discrete Fourier transform. Additionally, since the cross-correlation technique requires a large number of measurements to be made to suppress the system noise floor, it is necessary to perform a large number of Fourier transforms. Considering that a digital Fourier transform is a complex operation, it is essential that an efficient algorithm be implemented so that the data processing time does not limit the noise floor suppression that can be achieved with this system. In this section, the suitability of two different methods of implementing a digital Fourier transform in a residual phase noise measurement system are explored.

Discrete Fourier Transform (DFT)

The discrete Fourier transform (DFT) is the most simple way of applying a Fourier transform in a digital system. The DFT for a set of time-domain samples, x , of size N produces a set of frequency-domain samples, X , also of size N , and is defined as:

$$X_k = \sum_{n=0}^{N-1} x_n \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k} \quad (3.21)$$

Where k is an integer in the range $0 \rightarrow (N - 1)$.

From this it can be seen that to calculate a DFT of size N would require N^2 calculations of the result $x_n \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k}$. Therefore, the time taken to compute the DFT increases rapidly with the size of the input data. Even with modern computer systems, implementing a Fourier transform in this way quickly becomes impractical for large data sets due to the processing time. For a low-cost cross-correlation residual noise measurement system to be able to perform a large number of Fourier transforms for large sampling window sizes on modest hardware it is necessary to implement an algorithm to reduce the number of calculations that

must be performed.

Cooley-Tukey Fast Fourier Transform (FFT) Algorithm

The most popular method of implementing a fast Fourier transform algorithm was described by Cooley and Tukey [30] and enables the number of computations to be reduced by a factor of $\frac{N}{\log_2(N)}$ without affecting the accuracy of the output data. The principle behind this algorithm is that a DFT of size N can be expressed as a combination of two interleaved DFTs of size $N/2$. This property can then be applied recursively such that each of the smaller DFTs can be expressed as a combination two DFTs of size $N/4$ and so on.

An implementation of this algorithm for input data x of size $N=8$ is shown in Figure 3.18. The combination function, W , is defined as:

$$W_N^k = E + e^{-\frac{2\pi j}{N} \cdot k} \cdot O \quad (3.22)$$

Where E represents the value from the even-indexed input sequences and O represents the value from the odd-indexed.

The combination function is in a similar format to that seen in equation 3.21, $(A \cdot e^{-\frac{2\pi j}{B} \cdot C})$, and we can use this to compare the complexity (and hence the processing time) of the DFT and FFT. We already know that the DFT has a complexity of N^2 ; by using the recursive combination of smaller DFTs as shown in Figure 3.18, the radix-2 FFT reduces the complexity to $N \cdot \log_2(N)$.

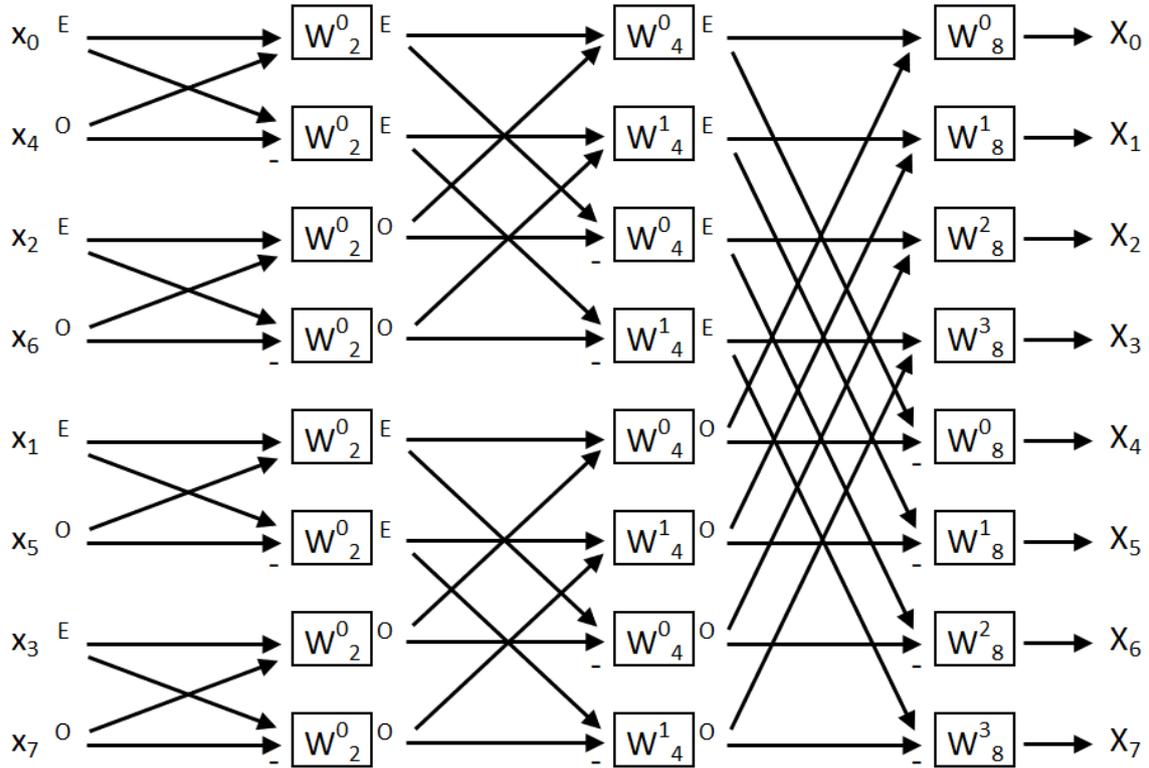


Figure 3.18: Diagram showing a radix-2 Cooley-Tukey FFT for N=8

Comparison

Despite this drastic improvement to the overall Fourier transform computation time, there is a potential advantage in using the DFT. This is because the frequency banding technique used in this system and discussed in Section 3.5.1 means that potentially most of the output data for each Fourier transform is effectively unused. When applying the standard DFT, only the outputs for the required values of k need to be calculated, and the rest can be ignored, whereas a standard Cooley-Tukey FFT algorithm does not operate in this way; the full transform is performed to get the required output data.

As an example, if we had a data size, $N = 128$, a band correlation ratio, $BandRatio = 8$, and we wanted to calculate the DFT outputs, X_k for the lowest frequency band (assuming no band shifting), we can see from equation 3.18 the number of values of X_k that need to be calculated is just 8. We can write this out

more generally below and compare the computation of a DFT and an FFT in this scenario for S required output values:

$$\begin{aligned}Q_{DFT} &= N \cdot \frac{N}{2 \cdot S} = 128 \cdot \frac{128}{2 \cdot 16} = 512 \\Q_{FFT} &= N \cdot \log_2(N) = 128 \cdot \log_2(128) = 896\end{aligned}\tag{3.23}$$

Where Q is the number of calculations in the format $(A \cdot e^{-\frac{2\pi j}{B} \cdot C})$ that need to be performed for the required output data.

Although this is a very specific scenario, it shows how using a full FFT algorithm will not always be the most efficient method depending on the output data required. That being said, the processing time of this DFT implementation will still increase much more rapidly with N compared to an FFT. In fact, the DFT will only be more efficient when the condition $(\frac{N}{2 \cdot S} < \log_2(N))$ is satisfied.

Considering the requirements for the residual noise system, it is necessary to be able to compute Fourier transforms for large sampling window sizes in excess of 100,000 in order to achieve high resolution bandwidth at low offset frequencies whilst still being able to measure at high offset frequencies. Therefore, implementing an FFT algorithm will help to greatly reduce processing compared to methods utilising a standard DFT.

3.5.3 Optimising the FFT for Frequency Banded Data

The frequency banding method described in section 3.5.1 involves the calculation of Fourier transforms where the number of required output data points is potentially significantly less than the size of the input data, N . It was discussed in the previous section that even when only part of the output data is required, an FFT algorithm is often much faster than a standard DFT, but ideally the FFT algorithm would be further optimised so that any unnecessary calculations can be ignored. In this section it will be explained how a partial FFT algorithm can be implemented so that the overall processing time when using the variable window size frequency banding method is comparable to performing a single Cooley-Tukey FFT of the full input

data set.

The Partial Fast Fourier Transform (PFFT)

A partial fast Fourier transform (PFFT) is a term used to describe a fast Fourier transform algorithm that calculates only a subset of the output data points calculated by a conventional FFT. There exists various descriptions of partial fast Fourier transform algorithms in the literature [31] [32]; in this section a PFFT algorithm specific to the requirements of a frequency-banded cross-correlation residual noise system is presented.

We can see from Figure 3.19(a) that if we only want to calculate a single output data point of a Fourier transform from a set of N input samples, the number of computations of W required is $2^0 + 2^1 + \dots + 2^{\log_2(N-1)}$ or more simply, $N - 1$.

We can also consider the computation of partial FFTs for other output data sizes by using the variable R to represent the ratio of the input data size, N , to the desired output size, M :

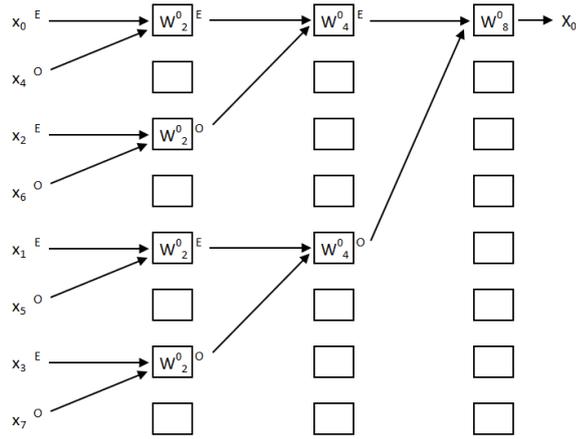
$$R = \frac{N}{M}$$

Figure 3.19(b) and (c) show the processes required for calculating the FFTs for different values of R . For the sake of simplicity, we are only considering scenarios where the output data is consecutive and whose size is a power of 2. The number of computations required to perform a partial FFT for a given R can be calculated by subtracting the unnecessary computations from the $N \cdot \log_2(N)$ computations required for a standard FFT as follows:

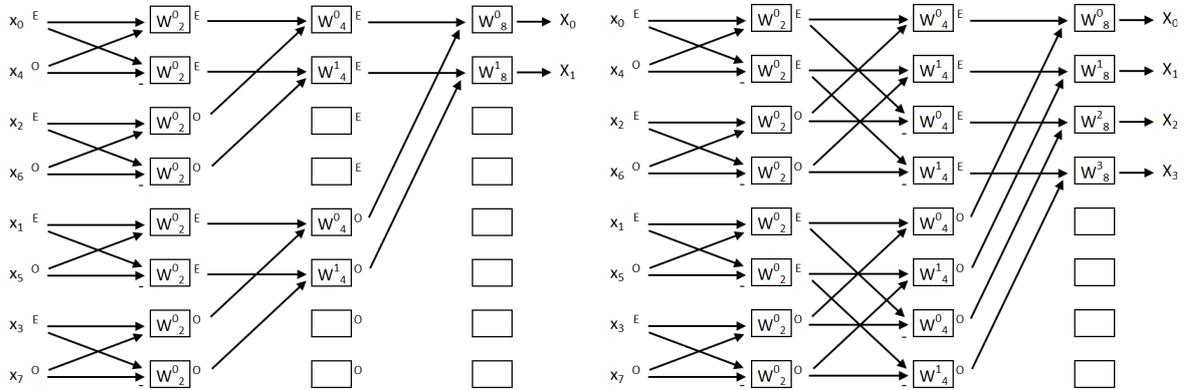
$$\begin{aligned} Q_{PFFT} &= N \cdot (\log_2(N) - \log_2(R) + \frac{R-1}{R}) \\ &\simeq N \cdot (\log_2(N) - \log_2(R) + 1), \quad (R \gg 1) \end{aligned} \tag{3.24}$$

Or expressed in terms of the desired output size, M :

$$Q_{PFFT} \simeq N \cdot (\log_2(M) + 1), \quad (N \gg M) \tag{3.25}$$



(a) Single output ($R = 8$)



(b) 2 outputs ($R = 4$)

(c) 4 outputs ($R = 2$)

Figure 3.19: Diagram showing the calculation of partial FFTs for $N=8$

This means that by using a partial FFT algorithm in situations where not all of the FFT output data is required, the processing time can be decreased compared to a standard Cooley-Tukey algorithm by a factor of roughly:

$$\frac{Q_{PFFT}}{Q_{FFT}} \simeq \frac{\log_2(N)}{\log_2(M) + 1}, \quad (N \gg M) \quad (3.26)$$

Single-Point Frequency Bands

If we take the idea of splitting the output data into different frequency bands to its extreme, we could split the data such that every output data point is in a different band. In other words $\log_2(N)$ FFTs are performed to produce $\log_2(N)$ data points,

where for every successive point along the frequency axis the distance between each point doubles and the required FFT window size halves.

We now know that to calculate a single FFT output from N input data points requires $N - 1$ computations of W , which if we assume that $N \gg 1$ can be approximated to N computations. The first output data point then requires N computations. The next data point requires $N/2$, as the window size has halved, but this leaves half the available data unused. As we want to maximise the number of measurements for each data point so that more cross-correlations can be performed, the computations are repeated for the other half of the data, which means the total number of computations at this data point will also be N . This process repeats for all the subsequent points; the window size halves again and twice as many FFTs can be performed for N computations, so the number of correlations increases exponentially for each output data point with the trade-off being that the resolution bandwidth also increases exponentially. It can be easily seen that this results in a total of approximately $N \cdot \log_2(N)$ computations for the whole data set, the same as performing a single FFT of the data using a fixed resolution bandwidth.

Multiple-Point Frequency Bands

The problem with using single-point frequency bands is that it does not provide any flexibility in terms of the resolution bandwidth, and hence the total number of output data points. It is often useful not only to split the data into larger frequency bands, but to use different sized bands within the same measurement, such as the examples given in Figure 3.16. This allows the resolution bandwidth versus number of correlations to be adjusted for different frequency ranges as required.

We can consider a more general case than the single-point frequency band example where each frequency band can have M_b output data points (as long as M_b is a power of 2). We still want to maximise the number of measurements for each data point, and the number of correlations that can be performed in each band, CC_b , for a given data set of size N is:

$$CC_b = \frac{N}{N_b}$$

Where N_b is the FFT input size used for each band.

Equation 3.25 can be altered to give the total number of computations to be performed in each band:

$$\begin{aligned} Q_b &\simeq N_b \cdot CC_b \cdot (\log_2(M_b) + 1) \\ &\simeq N \cdot (\log_2(M_b) + 1) \end{aligned} \tag{3.27}$$

Therefore calculating the Fourier transforms for bands that have the same number of output data points will require roughly the same number of computations regardless of how many cross-correlations are to be performed.

Comparison of Computation Time

It was previously shown that to calculate a number of consecutive bands with a single output data point ($M = 1$) and where the resolution bandwidth doubles each time requires approximately N computations for each band. If i is the number of bands and Q is the total number of computations required, we can write:

$$Q \simeq N \cdot i$$

If we were to replace these exponentially spaced bands with a single band that has a fixed resolution bandwidth, the number of points in the band would then be 2^{i-1} . If we apply this to Equation 3.27 to find the number of computations required for this band:

$$\begin{aligned} Q_b &\simeq N \cdot (\log_2(M_b) + 1) \\ &\simeq N \cdot (\log_2(2^{i-1}) + 1) \\ &\simeq N \cdot (i - 1 + 1) \\ &\simeq N \cdot i \end{aligned} \tag{3.28}$$

We can see that the number of computations is approximately the same meaning that, assuming the size of the frequency windows used are a power of 2 and do not overlap, the size of the windows used does not significantly affect the overall computation time. Using larger window sizes will maintain a fixed resolution bandwidth over a wider frequency range, resulting in more output data points, at the expense of fewer cross correlations at the higher frequencies in each band than could be achieved by using smaller bands.

The conclusion to this is that by applying an optimised partial FFT algorithm a residual phase noise measurement can be split into different frequency bands, where more measurements are performed for each data point at higher frequencies at the expense of resolution bandwidth, without significantly increasing the computation time for a sufficiently large N .

3.5.4 Twiddle Factors

A 'twiddle factor' is a term commonly used to describe a multiplication by a fixed value used when calculating a DFT (or FFT). Results that appear frequently during the DFT can be calculated beforehand and stored so that the results can be looked up when needed. This is especially efficient when a large number of Fourier transforms of the same size are required to be performed, as the twiddle factors will be the same. In this section a set efficient equations for calculating the Fourier transform in digital system and associated twiddle factors are derived.

Derivation

To understand how Fourier transform calculations are performed in a digital system, we can apply Euler's formula ($e^{jx} = \cos x + j \cdot \sin x$) and rewrite (3.21) as:

$$X_k = \sum_{n=0}^{N-1} x_n \cdot \left(\cos\left(-\frac{2\pi}{N} \cdot n \cdot k\right) + j \cdot \sin\left(-\frac{2\pi}{N} \cdot n \cdot k\right) \right) \quad (3.29)$$

The complex frequency-domain output of the Fourier transform can then be split into its real and imaginary parts:

$$\begin{aligned}
 \text{Re}(X_k) &= \sum_{n=0}^{N-1} \text{Re}(x_n) \cdot \cos\left(-\frac{2\pi}{N} \cdot n \cdot k\right) - \text{Im}(x_n) \cdot \sin\left(-\frac{2\pi}{N} \cdot n \cdot k\right) \\
 \text{Im}(X_k) &= \sum_{n=0}^{N-1} \text{Im}(x_n) \cdot \cos\left(-\frac{2\pi}{N} \cdot n \cdot k\right) + \text{Re}(x_n) \cdot \sin\left(-\frac{2\pi}{N} \cdot n \cdot k\right)
 \end{aligned} \tag{3.30}$$

Generally speaking, trigonometric functions require significantly more computer processing time than simple floating-point operations, such as addition, multiplication, etc. To calculate just the real parts of a DFT would require N^2 cosine operations and N^2 sine operations to be processed (or $N \log(N)$ in the case of an FFT) and the same again for the imaginary part. However, it can be seen that for a given data size, N , the sine and cosine factors that are applied at each step in the calculation when iterating through n and k will always be the same regardless of the input data, x_n .

Since this system is required to perform a large number of Fourier transforms on data that use the same sampling window size, it makes sense to pre-calculate the sine and cosine values for each possible value of $n \cdot k$ and place them in a lookup table, so that the time-expensive trigonometric functions can be replaced with a simple value lookup.

We can rewrite (3.30) incorporating these lookup tables, where A represents the cosine coefficients, and B represents the sine coefficients:

$$\begin{aligned}
 \text{Re}(X_k) &= \sum_{n=0}^{N-1} \text{Re}(x_n) \cdot A_{n \cdot k} - \text{Im}(x_n) \cdot B_{n \cdot k} \\
 \text{Im}(X_k) &= \sum_{n=0}^{N-1} \text{Im}(x_n) \cdot A_{n \cdot k} + \text{Re}(x_n) \cdot B_{n \cdot k}
 \end{aligned} \tag{3.31}$$

In addition to this, it is not necessary to calculate and store results for all the values of $n \cdot k$, as it can be seen from the trigonometric functions in (3.30) that these values will repeat for every integer multiple of $\frac{n \cdot k}{N}$. Only the first N possible values of these functions, (i.e. for $(n \cdot k)$ in the range $0 \rightarrow (N-1)$) need to be calculated. When calculating the Fourier transform coefficients, the correct lookup table positions can

be found by taking the modulus of $(n \cdot k)$ in the base, N . Assuming that N is a power of 2, calculating the modulus is equivalent to a simple bitwise operation, but it means the size of the lookup table can be reduced by a factor of N , which is essential in situations where a very large sampling window is used or where available memory is limited.

Finally, a set of equations for efficient calculation of DFT coefficients in a digital system:

$$\begin{aligned}
 A_{m=0 \rightarrow (N-1)} &= \cos\left(-\frac{2\pi}{N} \cdot m\right) \\
 B_{m=0 \rightarrow (N-1)} &= \sin\left(-\frac{2\pi}{N} \cdot m\right)
 \end{aligned}
 \tag{3.32}$$

$$\begin{aligned}
 Re(X_k) &= \sum_{n=0}^{N-1} Re(x_n) \cdot A_{n \cdot k(modN)} - Im(x_n) \cdot B_{n \cdot k(modN)} \\
 Im(X_k) &= \sum_{n=0}^{N-1} Im(x_n) \cdot A_{n \cdot k(modN)} + Re(x_n) \cdot B_{n \cdot k(modN)}
 \end{aligned}$$

These equations can be equally applied to an FFT algorithm, such as that described in the previous section. The implementation of the Fourier transform in this system combines these methods to optimise this process in terms of processing-time as well as memory requirements.

3.5.5 Simultaneous Fourier Transforms of Real Data

A Fourier transform operates on a series of complex input data points to produce a series of complex output data points. In the case of this system, the time-domain voltage samples that are used as the input must always be real-valued, so the imaginary parts of the input data to the transform are effectively set to 0, and due to the symmetry property of Fourier transforms of real data the FFT outputs for positive frequencies will be duplicated for negative frequencies. Applying the FFT in this way for a cross-correlation system means that two separate FFTs would need to be performed on the real-valued input data for each channel and the negative frequency

data is simply ignored.

E. Chu and A. George [33] describe in their book a method by which the separate Fourier transforms of two distinct sets of real-valued data can be computed simultaneously by applying a single complex Fourier transform that uses one set of the real-valued input data as the real part of the input data to the transform and the other set of data as the imaginary part.

Derivation

To show how this method is derived, we can write the desired DFTs of the of the real input data for the two channels, F_k and $G(k)$ as:

$$F_k = \sum_{n=0}^{N-1} f_n \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k} \quad \text{and} \quad G_k = \sum_{n=0}^{N-1} g_n \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k} \quad (3.33)$$

And the complex DFT that uses for its input the time-domain data samples for each channel, f_n and g_n , as the real and imaginary parts respectively:

$$\begin{aligned} X_k &= \sum_{n=0}^{N-1} x_n \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k} \\ &= \sum_{n=0}^{N-1} (f_n + jg_n) \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k} \\ &= \sum_{n=0}^{N-1} f_n \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k} + j \sum_{n=0}^{N-1} g_n \cdot e^{-\frac{2\pi j}{N} \cdot n \cdot k} \\ &= F_k + jG_k \end{aligned} \quad (3.34)$$

Therefore, the separate DFTs for each channel, F_n and G_n , can be obtained from X_n if a suitable method can be used to recover them. Chu and George go on to describe how this can be done by using the symmetry property of real-valued DFTs: F_k is equal to the complex conjugate of F_{N-k} .

Applying the symmetry property gives:

$$F_k = \overline{F_{N-k}} \quad \text{and} \quad G_k = \overline{G_{N-k}} \quad (3.35)$$

Using (3.34) we can then express the complex conjugate of X_{N-k} in terms of F_k and G_k :

$$\begin{aligned}\bar{X}_{N-k} &= \bar{F}_{N-k} - j\bar{G}_{N-k} \\ &= F_k - jG_k\end{aligned}\tag{3.36}$$

Finally, we can rearrange these and substitute them back into (3.34) to find the desired DFTs of the real input data for the two channels, F_k and G_k , in terms of the single complex DFT, X_k , that uses the two data sets as its real and imaginary parts:

$$\begin{aligned}F_k &= \frac{1}{2}(X_k + \bar{X}_{N-k}) \\ G_k &= \frac{j}{2}(\bar{X}_{N-k} - X_k)\end{aligned}\tag{3.37}$$

Implementation

These equations can be most efficiently implemented into a digital system by splitting them into their real and imaginary parts, as for (3.32). It is also worth bearing in the mind that the factors of $\frac{1}{2}$ present in these equations do not need to be performed for every Fourier transform in a cross-correlation system, as they can be taken into account using a single multiplication after the cross-correlation and summation of all the transforms have been performed:

$$\begin{aligned}Re(F_k) &= \frac{1}{2}(Re(X_k) + Re(X_{N-k})) \\ Im(F_k) &= \frac{1}{2}(Im(X_k) - Im(X_{N-k})) \\ Re(G_k) &= \frac{1}{2}(Im(X_k) + Im(X_{N-k})) \\ Im(G_k) &= \frac{1}{2}(-Re(X_k) + Re(X_{N-k}))\end{aligned}\tag{3.38}$$

The benefits of using this single Fourier transform method in a digital system as opposed to performing two separate transform can be seen quite simply. The required system memory is reduced by a whole transform size (N , in this case), which is especially critical when using large sampling window sizes. Furthermore,

since the only additional processing required for each transform are some simple additions/subtractions, the processing time is roughly halved for an N of any appreciable size.

3.5.6 Cross-Correlation and Post-Processing

Although the complex nature of performing Fourier transforms and the importance of using frequency banding techniques to improve phase noise measurements mean that optimising and refining these aspects of the signal processing should be prioritised, efficient algorithms for the other aspects, such as post-processing and performing cross-correlations should also be implemented and are detailed in this section.

Cross-Correlation

To achieve suppression of the system noise floor, the results of every Fourier transform performed are cross-correlated between the two channels. The cross-correlation of two complex values, A and B, is found by taking the product of B with the complex conjugate of A as follows:

$$(A * B) = \bar{A} \cdot B \quad (3.39)$$

Seeing as a set of equations for the real and imaginary parts of the Fourier transforms of the two channels has already been established in (3.38), the cross-correlation of the respective points for each value of the output data points, k, can be very easily calculated as follows:

$$\begin{aligned} Re(F_k * G_k) &= (Re(F_k) \cdot Re(G_k)) - (Im(F_k) \cdot -Im(G_k)) \\ Im(F_k * G_k) &= (Re(F_k) \cdot -Im(G_k)) + (Im(F_k) \cdot Re(G_k)) \end{aligned} \quad (3.40)$$

As the cross-correlation can be calculated from a relatively simple arithmetic process, it will not have much effect on the processing time assuming that the Fourier transform window sizes are kept sufficiently large.

Conversion to Power Spectral Density

Once the cross-correlations have been performed for each measurement, the complex-valued results at each frequency point are summed together, the absolute value is calculated and this summation is divided by the total number of measurements at each point to give the frequency components of the measured noise in terms of mean squared voltage. Residual noise is typically measured in noise power spectral density: the noise density per unit bandwidth. This is because when performing a Fourier transform of a noise source, any frequencies present in the signal inside one of the frequency bins (between two data points) will contribute to the amplitude measurement for that bin, so if a larger resolution bandwidth is used, the measured noise power at each data point will be greater. To calculate noise power spectral density, this value is divided by the width of the band, so that the magnitude of the data points will be the same regardless of the resolution bandwidth used for the measurement.

The noise power spectral density of the data points, Y_k , in each band, b , can be calculated as follows:

$$Y_k^b = \frac{10}{CC_b} \cdot \log_{10}\left(\frac{|Z_k^b|}{ResBW_b}\right) \quad (3.41)$$

Where Z_k is the summation of the cross-correlations at each data point; CC_b is the number of cross-correlations performed in that band; and $ResBW_b$ is the resolution bandwidth of the band.

This produces results in the unit of dBV/Hz, which is useful when measuring absolute noise levels, however for the purposes of measuring residual phase noise it is useful to express the noise power relative to the carrier signal, i.e. in dBc/Hz. A calibration procedure is required to for this, which was discussed in 3.3.2. The calculation of the calibrated result can easily be added to (3.41) so that the results are given in dBc/Hz:

$$Y_k^b = \frac{10}{CC_b} \cdot \log_{10}\left(\frac{|Z_k^b|}{ResBW_b}\right) + (Calib_{tone} - Calib_{RF} - Calib_{out} - 6 \text{ dB}) \quad (3.42)$$

This only needs to be calculated once for each of the output data points of the measurement at the point when the data is needed to be displayed and so has minimal effect of the overall processing time.

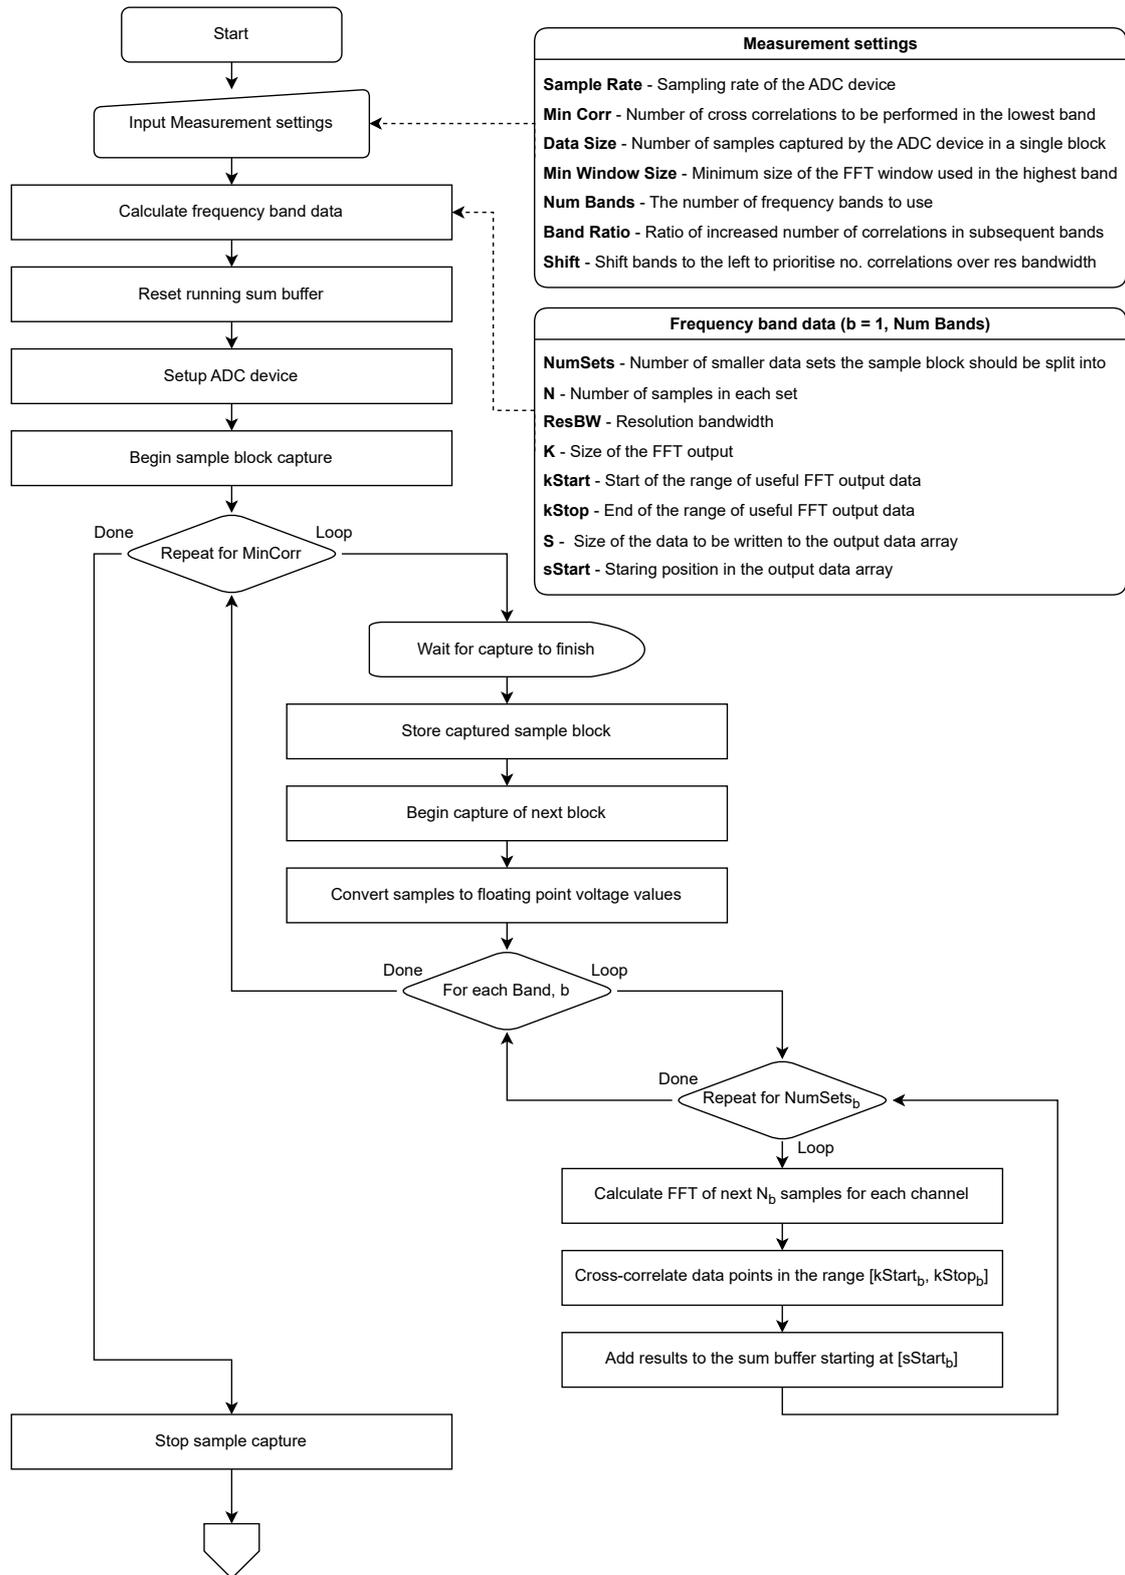
Calculation of Frequency Data

In order to plot the output data points of the residual noise measurement it is necessary to calculate the corresponding frequency values for each of the measured power spectral density values. These only needs to be calculated once and can be done so quite simply if the resolution bandwidth and the start and end points for the useful output data of the Fourier transforms for that band are known, the calculations for which are shown in 3.5.1.

The frequency points increment by the resolution bandwidth for each subsequent point along the frequency axis, so the frequency points, X_k^b , for each band are:

$$X_{k=0 \rightarrow (S_b-1)}^b = (k + sStart_b) \cdot ResBW_b \quad (3.43)$$

3.5.7 Complete System Flowchart



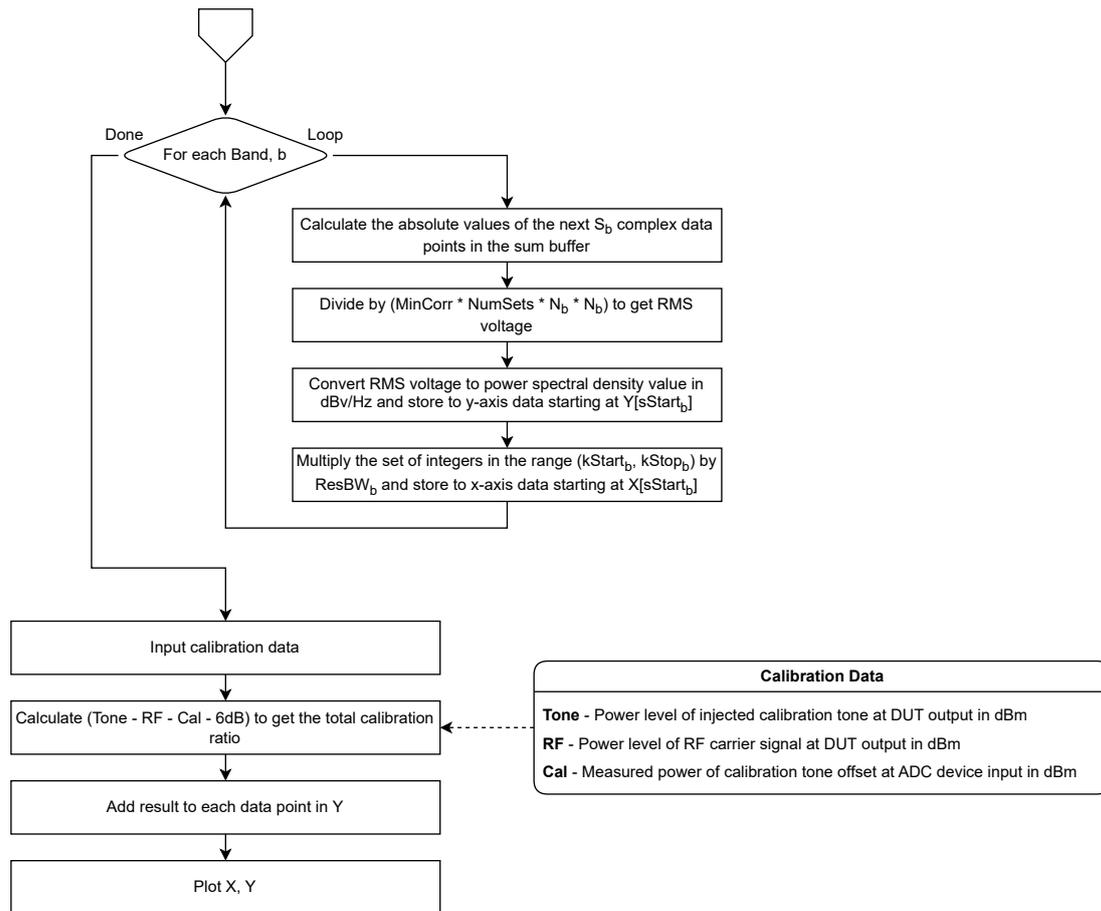


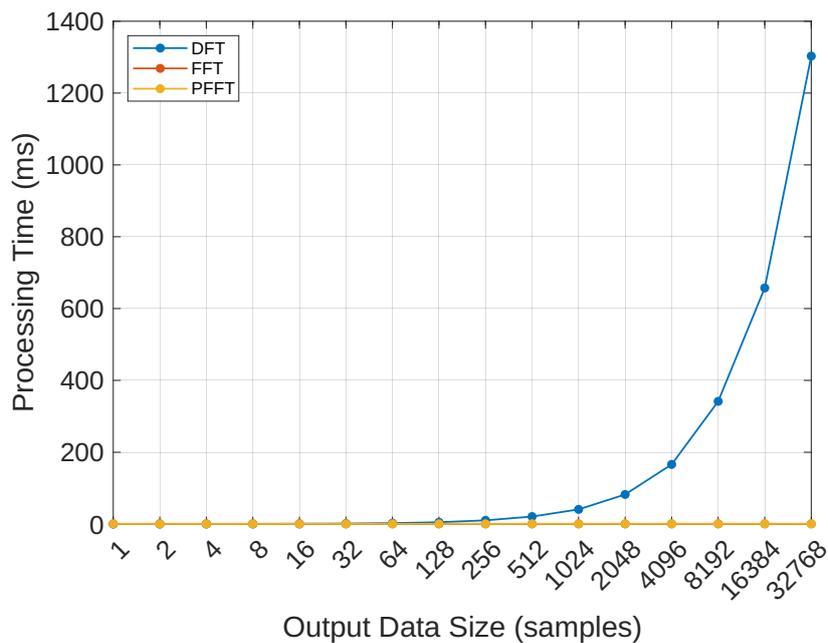
Figure 3.20: Complete flowchart of signal processing for a cross-correlation residual phase noise measurement

3.5.8 Results

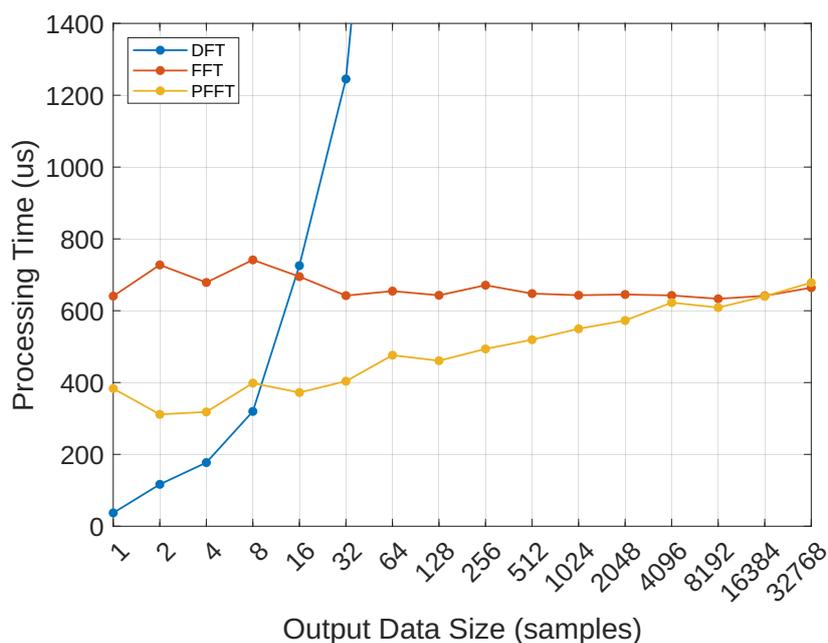
Comparison of the Processing Time of Fourier Transform Algorithms

To demonstrate the difference in processing time between the various Fourier transform algorithms outlined in this section (the DFT, FFT and PFFT) a test program was created that repeatedly performs a Fourier transform on a set of 32,768 random samples to calculate a specified number of output data points and records the average time taken to perform a single transform. This process is repeated for different output sizes, M , to see the relative performance of each of the Fourier transform algorithms as the desired number of output data points increases.

The results are shown in Figure 3.21. It can be seen in (a) that, as expected, the processing time for the DFT increases rapidly with M and is much greater than that for either the FFT or DFT when M is any considerable size. (b) shows a zoomed-in view that allows the different algorithms to be compared more easily. The FFT results should be a flat line as the algorithm is the same regardless of how many output data points are required, however there is some deviation due to the fact that this is running on a PC with other processes running concurrently. In all cases the PFFT was either faster or took roughly the same time as the FFT, and was considerably faster when the output data size was less than 1024. For very small output data sizes the improvement in processing time wasn't as much as was predicted by Equation 3.26, which is likely due to the additional processing involved in the actual implementation of the PFFT, such as setup etc., that has to be performed as well as the calculations that were considered in Section 3.5.3. It can also be seen that there are a few instances where the DFT is the fastest algorithm, specifically when the M is less than eight, but this would not be very common scenario. For the vast majority of situations the PFFT algorithm is the most efficient option, and so this is the algorithm that was selected to be implemented into the residual noise measurement system.



(a) Zoomed-out view showing the rapid increase in DFT computation time for high values of M



(b) Zoomed-in view showing the relative efficiencies of each method for lower values of M

Figure 3.21: Plots showing the average processing time versus the size of the required output data, M, when using different methods to calculate a single Fourier transform with an input data size of 32,768 samples

Processing Time of Simultaneous Fourier Transforms of Real Data

Another test program was created to compare the average processing time to calculate the Fourier transform of a set of two random sets of real data using an FFT algorithm for various input data sizes and the results are shown in Figure 3.22. In one case the two FFTs are performed separately, and in the other, the algorithm for the simultaneous transforms of real data described in Section 3.5.5 is used. It was estimated that simultaneous transform algorithm would be able to reduce the processing time by a factor of two, and it can be seen that holds roughly true across the full range of the measurement.

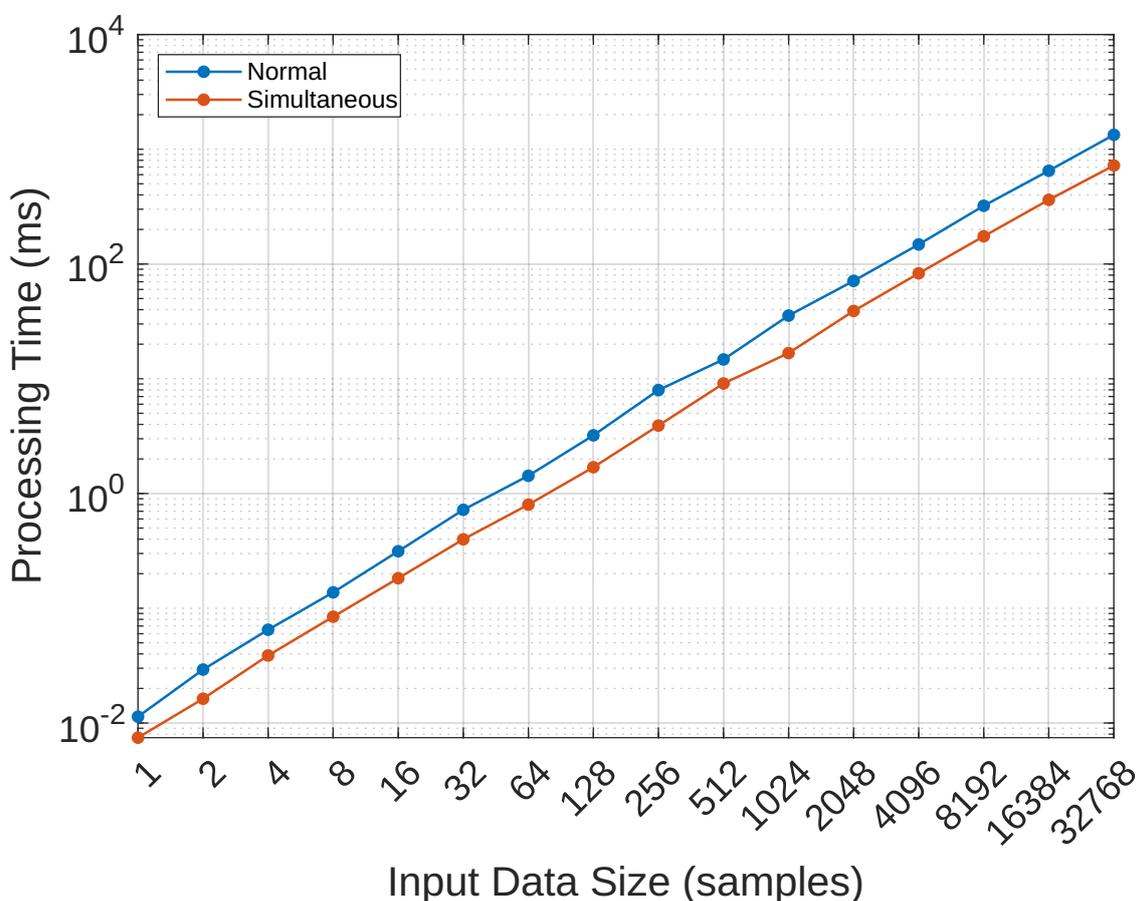


Figure 3.22: Comparison of the average processing time to calculate the Fourier transforms of two sets of real data plus the cross correlation of the results for different input data sizes when using either a standard approach or the simultaneous method

Comparison of Processing Time versus Number of Frequency Bands

It was suggested in Section 3.5.3 that by implementing a PFFT algorithm the measurement could be split into separate frequency bands with an increased number of correlations at higher offset frequencies, as described in Section 3.5.1, without any significant impact on the processing time. To test this a program was made that performs the Fourier transforms and cross-correlations for a set of 16,777,216 random input samples using the frequency banding algorithm to split the results into a specified number of frequency bands. The processing time was then plotted against the number of bands when using both the PFFT algorithm and the FFT algorithm as a point of reference. The results are shown in Figure 3.23 and show that when the number of bands used was between 1 and 8, the processing time is roughly the same when using the PFFT, while it increases linearly with the FFT. When more bands are used the processing time with the PFFT also starts to increase linearly, although less steeply than that for the FFT.

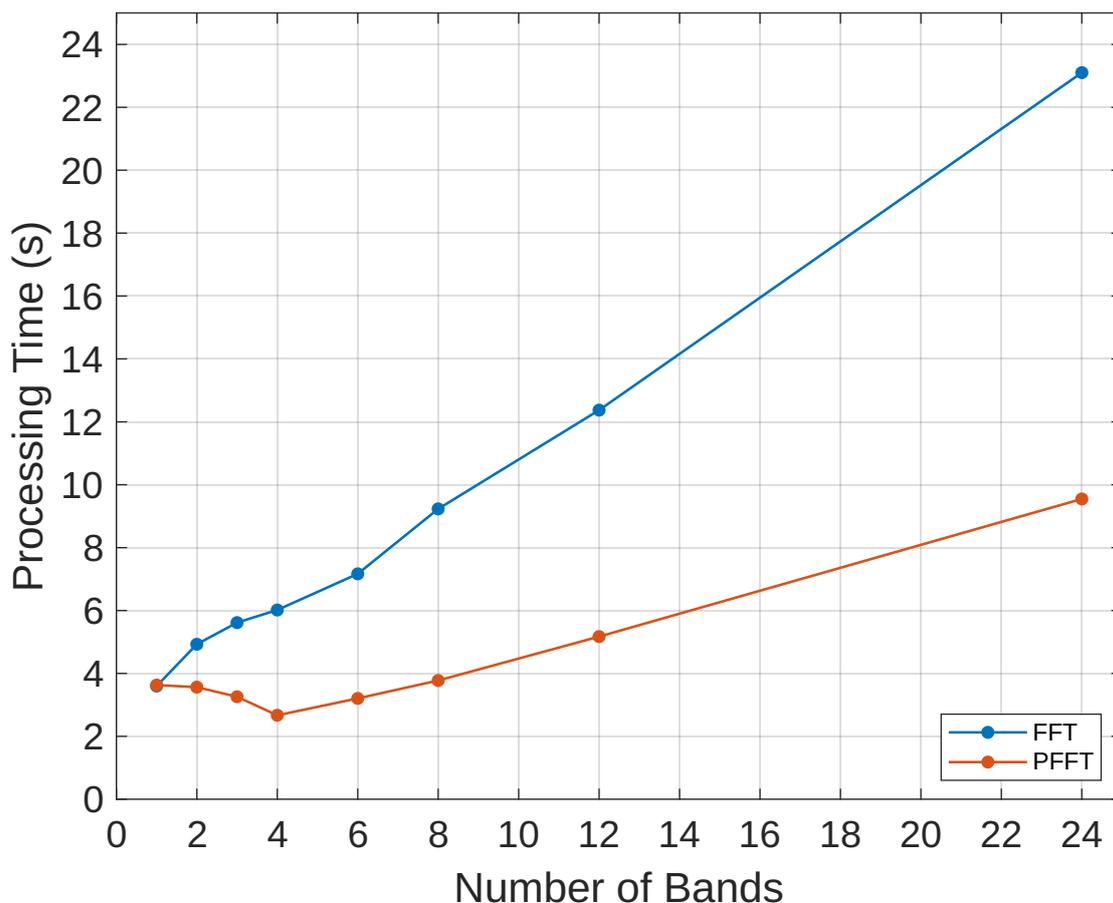


Figure 3.23: Comparison of the processing time for a single data block of 16,777,216 captured samples versus the number of equal-ratio frequency bands that the measurement is split into when using either an FFT or PFFT algorithm

Examples of Measurement Processing Time

Table 3.1 shows various example measurement setups with different sampling rates and frequency bands. The resolution bandwidth and number of correlations in each band are shown for a given minimum measurement time. The minimum measurement time is based on the theoretical minimum time it would take to capture the samples required to produce a measurement with the specified resolution bandwidth and the specified number of correlations. These measurements were performed by capturing samples using the PicoScope and running the full digital signal processing software in real time to produce a frequency-banded power spectral density mea-

surement.

The actual measurement time was recorded so that it could be compared with the minimum theoretical measurement time. For all the measurements with a sampling rate of 2 MHz or less the actual measurement time did not exceed the theoretical minimum, which suggests the capture time is the limiting factor of overall the measurement time and the digital signal processing is able to be completed within this period. In fact, the actual measurement time for these setups was slightly less than the theoretical minimum. This is believed to be due to a mismatch between the requested sampling frequency and the actual frequency that the PicoScope uses, likely due to the limited number of bits used to set the frequency. A nominal sampling frequency of 1.048576 MHz was used for all measurement results in this chapter, which according to the table likely has a frequency error of around 0.75% and so will not have much effect on the accuracy of the results. If a frequency of 2.097152 MHz is used however, the error appears to be as much as 6.5%, which is a significant discrepancy and highlights an issue that should be addressed.

S. Rate (MHz)	Min Res BW (Hz)	No. Correlations							Measurement Time (s)	
		@ 0.125 Hz	@ 2 Hz	@ 16 Hz	@ 128 Hz	@ 1024 Hz	@ 8192 Hz	@ 65536 Hz	Min	Actual
1.048576	0.125	100	100	100	100	100	100	100	800	797
1.048576	0.125	100	100	800	6400	51200	490600	3276800	800	797
1.048576	8	-	-	6400	6400	51200	490600	3276800	800	794
		0.25	4	32	256	2084	16384	131072		
2.097152	0.25	100	100	100	100	100	100	100	400	375
2.097152	0.25	100	100	800	6400	51200	490600	3276800	400	375
2.097152	16	-	-	6400	6400	51200	490600	3276800	400	374
		0.5	8	64	512	4096	32768	262144		
4.194304	0.5	100	100	100	100	100	100	100	200	290
4.194304	0.5	100	100	800	6400	51200	490600	3276800	200	274
4.194304	32	-	-	6400	6400	51200	490600	3276800	200	207

Table 3.1: A comparison of the the total measurement time for a variety of configurations versus the minimum theoretical time possible based on the desired minimum resolution bandwidth

3.6 Conclusions and Further Work

The aim of the work in this chapter was to produce a low-cost residual phase noise measurement system that is capable of measuring the residual noise of amplifiers to be used in ultra-low noise crystal oscillators such as those described in Chapter 1. It was shown in Section 3.4.2 that the system was capable of accurately measuring the residual noise of the MAR-6 amplifier, which was not possible on the commercially-available Rohde & Schwarz FSWP Phase Noise Analyzer due to its internal noise floor. The minimum noise floor of the system described in this chapter was found to be around -185 dBc/Hz, which while not as low as some systems described in the literature or commercially-available systems, was achieved with open-source system designed for low-power measurements that uses low-cost components.

The secondary aim was to review the performance of various digital signal processing techniques specific to residual phase noise measurements and develop a software library that would allow for efficient conversion of captured samples to power spectral density and provide a way to increase the number of correlations at high offset frequencies that can be performed in a given time frame whilst maintaining a narrow resolution bandwidth at close-to-carrier. An open-source software library has been developed and it was shown in Section 3.5.8 that it could be used to perform a measurement with sub-Hz close-to-carrier resolution bandwidth and over 3,000,000 cross-correlations in around 13 minutes. It was also shown that the frequency-banding algorithm that was developed did not significantly increase the processing time over processing without frequency banding.

3.6.1 Further Work

Minor Improvements

One of the issues highlighted was that the actual sampling frequency of the Pico-Scope was often slightly different from the requested frequency, which could lead to a significant percentage error for some sampling frequencies. Although the error

was not significant at 100 MHz, this issue should be addressed for this system to be useful at a wide range of operating frequencies.

The maximum offset frequency that can be measured by the system is currently somewhat limited at 100 kHz, which is mainly due to the limited gain-bandwidth product of the LT6018 amplifiers used. A change to the amplifier and anti-aliasing filter designs could be made to increase this range.

Automation

One of the main drawbacks of this system is the time it takes to set up measurements. Currently, the required 90° phase shift between the DUT and reference arms must be achieved by manually adjusting both the digital and mechanical phase shifters. The setup time could be significantly reduced by automating this process. It would be simple to implement this for the digital phase shifters as they can be easily controlled from a microprocessor connected to the PC. All that would be required is to sweep through the phase shifter settings and see which produces the lowest DC value measured on the PicoScope. The mechanical phase shifter cannot be controlled in this way, but it may be possible to replace this with a high-resolution digital phase shifter.

The calibration procedure can also take a considerable amount of time. Since this involves injecting and measuring the spectrum of RF signals, automating this process is likely outside the scope of a low-cost system. However, it may be that there is no need to perform the calibration procedure more than once for measurements that use the same operating conditions, i.e. the same frequency and DUT output power, but this will require further investigation.

Variable Sample Rate Method

Although the system used for frequency banding of residual noise measurements described in this chapter uses an entirely software-based approach, another method has been considered that could reduce processing time and memory requirements

as well as making the system more adaptable to operating at different frequencies. This system would split the baseband channels into the desired number of frequency bands each with a separate anti-aliasing filter and capture device so that the sampling rate can be different for each of the channels. This essentially means that multiple residual phase noise measurements can be performed concurrently, where each measurement has a different resolution bandwidth versus number of correlations for a given measurement time. The channels with slower sampling rates would have a narrower resolution bandwidth for measuring close-to-carrier frequencies, whereas the channels with faster sampling rates have a wider resolution bandwidth but can perform many correlations for greater instrument noise floor suppression. These measurements can then be combined to form a single frequency-banded plot.

One of the main benefits of this system is that it would be much easier to process the data for close-to-carrier frequencies if a slower sampling rate is used as far fewer samples need to be captured in each block to achieve the same resolution bandwidth. There would then be no need to limit the maximum sampling frequency for far out frequencies in order to keep the required number of samples in each block to a reasonable size. It would not, however, be practical for a low-cost system to have multiple PicoScope devices running concurrently to capture samples in each band, so a more compact and less expensive solution would be needed. Some work has already been carried out towards this. A small, inexpensive, dual-channel sample capture device with onboard amplifiers and anti-aliasing filters has been built using a Raspberry Pi Pico microprocessor, as shown in Figure 3.24, and is capable of capturing and storing up to 4,194,304 pairs of samples at sampling rates up to 500 kHz. It is hoped that a series of these devices would allow such a system to be built.

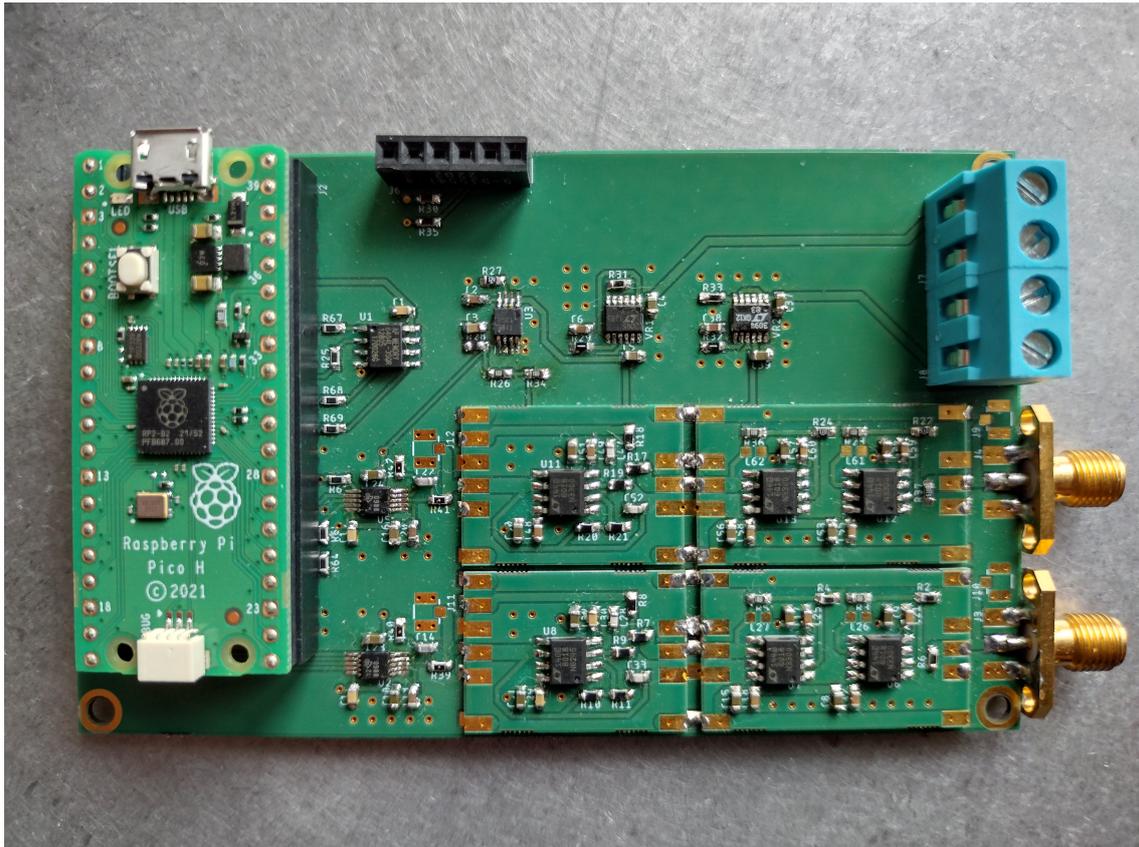


Figure 3.24: Photo of the low-cost Raspberry Pi Pico two-channel ADC capture device

Chapter 4

Overall Conclusions

In this thesis the design for a 100 MHz crystal oscillator with phase noise of around -140 dBc/Hz at a 100 Hz offset from the carrier has been presented. This matches the state-of-the-art performance of commercially-available oscillators and no 100 MHz crystal oscillator with this level of performance has previously been described fully in the literature. As part of this work, a low-flicker-noise differential amplifier using a quad parallel configuration has been designed with a flicker noise corner of around 80 Hz when operating at 4.6 dBm output power, which was measured using a technique that involved placing the amplifier in the loop of an LC oscillator. It has also been shown that the placement of an attenuator between the resonator and the input of the amplifier in an oscillator loop can be used as a simple way to improve the phase noise in situations where the amplifier's flicker noise is sensitive to the source impedance presented to it or its input is not well-matched to the system impedance. A 10 dB improvement in the close-to-carrier phase noise of an LC oscillator was demonstrated using this method. It may be possible to further improve the phase noise performance of the 100 MHz crystal using an alternative amplifier with a higher gain compression point, such as the Mini Circuits MAR-2SM+.

A set of digitally-controlled phase shifters and attenuators for a range of applications have been presented and the complete design procedure for producing these devices with different operating conditions has been described. These devices are

completely passive, so that they exhibit very low residual phase noise, and use off-the-shelf mechanical RF relays, so that they can be produced rapidly and at minimal cost. The devices shown in this thesis include a 100 MHz digital phase shifter with a 180° range and 2.8125° resolution; a 1.5 GHz digital phase shifter with a 360° range and 6° resolution; a digital attenuator with a 6.3 dB range and 0.1 dB resolution and a digital attenuator with a 31.5 dB and 0.5 dB resolution. The maximum non-linearity error of the devices was found to be no worse than $0.55 \cdot \text{LSB}$ in all cases and the residual phase noise was low enough such that it could not be measured with any measurement systems that were available.

Finally, a low-cost cross-correlation residual phase noise measurement system with an internal noise floor of around -177 dBc/Hz when measuring a device operating at 0 dBm output power has been presented. It has been shown that despite using inexpensive, off-the-shelf components, this system has a lower internal noise floor than the commercially-available Rohde & Schwarz FSWP for residual phase noise measurements at 100 MHz. As part of this work, a detailed review of digital signal processing techniques for residual phase noise measurements has been carried out and a C++ software library of efficient algorithms for processing measurement data has been developed and is intended to be released open-source. This software was used to perform a residual phase noise measurement with sub-Hz close-to-carrier resolution bandwidth and over 3,000,000 cross-correlations at high offset frequencies in around 13 minutes. Currently, the system takes a considerable amount of time to set up and calibrate measurements and would benefit from further investigation into automation of these aspects.

References

- [1] Jeremy Everard , Tsvetan Burtichelov, and Keng Ng, "Ultralow Phase Noise 10-MHz Crystal Oscillators" *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Control*, vol. 66, no. 1, January 2019
- [2] J. Everard and K. Ng, "Ultra-low phase noise crystal oscillators," *Proc. IEEE 21st Int. Freq. Control Symp. Joint Eur. Freq. Time Forum*, May/June. 2007, pp. 1246–1250
- [3] M. M. Driscoll, "Low frequency noise quartz crystal oscillator," *IEEE Trans. Instrum. Meas.*, vol. IM-24, no. 1, pp. 21–26, Mar. 1975.
- [4] E. Rubiola and V. Giordano, "On the 1/f frequency noise in ultra-stable quartz oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 1, pp. 15–22, Oct. 2007
- [5] KVG, "O-40-ULPN-100M," O-40-ULPN-100M oven controlled crystal oscillator datasheet
- [6] T. McClelland, C. Stone and M. Bloch, "100 MHz crystal oscillator with extremely low phase noise," *Proc. 1999 Joint Meeting of the European Freq. and Time Forum and the IEEE International Freq. Control Symp.*, pp. 331-334 vol.1
- [7] D. B. Leeson, "A simple model of feedback oscillator noise spectrum," *Proc. IEEE*, 1966-02, Vol.54 (2), p.329-330
- [8] KVG, "O5SC85105A," O5SC85105A crystal resonator datasheet

- [9] T. Baier, "A Small, Simple, USB-Powered Vector Network Analyzer Covering 1 kHz to 1.3 GHz," *QEX*, pp 32-36, Jan/Feb 2009
- [10] J. Everard, *Fundamentals of RF Circuit Design: With Low Noise Oscillators*, Hoboken, NJ, USA: Wiley, 2002.
- [11] K. K. M. Cheng and J. K. A. Everard, "Noise performance degradation in feedback oscillators with nonzero phase error," *Microwave and Optical Technology Letters*, Vol. 4, No. 2, January 20 1991
- [12] Mini Circuits, "Drop-In Monolithic Amplifier," MAR-6+ datasheet
- [13] Rohde & Schwarz, "R&S FSWP Phase Noise Analyzer and VCO Tester," FSWP specifications, Feb 2024
- [14] Infineon, "Low Noise Silicon Bipolar RF Transistor," BFR92P datasheet
- [15] Konstantinos Theodoropoulos, "Residual Phase Noise Modelling of Silicon Bipolar Amplifiers and Ultra Low Phase Noise Ceramic Dielectric Resonator Oscillators," PhD Thesis, *University of York Department of Electronics*, September 2009
- [16] Konstantinos Theodoropoulos and Jeremy Everard, "Residual Phase Noise Modeling of Amplifiers Using Silicon Bipolar Transistors," *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Control*, vol. 57, no. 3, March 2010
- [17] M. M. Driscoll, "Oscillator AM-to-FM noise conversion due to the dynamic frequency-drive sensitivity of the crystal resonator," *Proc. IEEE Int. Freq. Control Symp.*, May 2008, pp. 672–676
- [18] J. Iannacci, G. Resta, A. Bagolini, F. Giacomozzi, E. Bochkova, E. Savin, R. Kirtaev, A. Tsarkov, M. Donelli, "RF-MEMS Monolithic K and Ka Band Multi-State Phase Shifters," *Sensors 2020*, no. 9: 2612. <https://doi.org/10.3390/s20092612>

- [19] E. Hammerstad and O. Jensen, "Accurate Models for Microstrip Computer-Aided Design," *1980 IEEE MTT-S International Microwave symposium Digest*, Washington, DC, USA, 1980, pp. 407-409, doi: 10.1109/MWSYM.1980.1124303.
- [20] E. Hammerstad and F. Bekkadal, "A microstrip handbook, ELAB Report STF44 A74169," *University of Trondheim, Norway, Tech. Rep.*, February 1975
- [21] J. Bahl and D. K. Trivedi, "A Designer's Guide to Microstrip Line", *Microwaves*, pp. 174-182, May 1977
- [22] Simon J. Bale, David Adamson, Brett Wakley, Jeremy Everard, "Cross Correlation Residual Phase Noise Measurements using Two HP3048A Systems and a PC Based Dual Channel FFT Spectrum Analyser," *24th European Frequency and Time Forum*, 13-16 April 2010
- [23] Simon Bale, "Ultra High Q Resonators and Very Low Phase Noise Measurement Systems for Low Noise Oscillators," PhD Thesis, *University of York Department of Electronics*, April 2012
- [24] W. F. Walls, "Cross-Correlation Phase Noise Measurements," *Proc. IEEE Freq. Ctrl. Symp.*, 1992, pp. 257-261
- [25] E. Rubiola and F. Vernotte, "The cross-spectrum experimental method," *arXiv*, 1003.0113v1, 2018.
- [26] Mini Circuits, "Frequency Mixer," ZX05-1L-S+ datasheet
- [27] E. Lamb, "Ultra-Low Noise Floor Phase Noise Measurement Instrumentation for Precision RF/Microwave Component Characterisation," *MEng Project Report*, Department of Electronic Engineering, University of York, 2020-21
- [28] C. W. Nelson, A. Hati, and D. A. Howe, "A collapse of the cross spectral function in phase noise metrology," *Rev. Sci. Instrum.*, vol. 85, pp. 024705–024707, Feb. 2014.

- [29] A. Holme, *Direct digital phase noise measurement* [Online]. Available: <http://www.aholme.co.uk/PhaseNoise/Main.htm>
- [30] J. Cooley, J. Tukey, "An algorithm for the machine calculation of complex Fourier series", *Math. Comp.* 19, 297-301, 1965.
- [31] J. Bowman , Z. Ghoggali, "The Partial Fast Fourier Transform", *Journal of Scientific Computing.* 76, 2018.
- [32] Y. Park and J. Jang, U. Kang, "Fast Partial Fourier Transform", *arXiv*, 2008.12559, 2020.
- [33] E. Chu and A. George, "Chapter 14: FFTs for Real Input," in *INSIDE the FFT BLACK BOX*, Ontario: CRC Press, 2000.

Appendix A

EFTF 2024 Conference Paper: Ultra-Low Phase Noise 100 MHz Crystal Oscillator

Ultra-Low Phase Noise 100 MHz Crystal Oscillator

Luke Dummott
School of Physics, Engineering and
Technology
University of York
York, United Kingdom
luke.dummott@york.ac.uk

Jeremy Everard
School of Physics, Engineering and
Technology
University of York

St. John Gilbert
School of Physics, Engineering and
Technology
University of York

Simon Bale
School of Physics, Engineering and
Technology
University of York

Abstract—This paper describes an ultra-low phase noise 100 MHz crystal oscillator with a phase noise of around -140 dBc/Hz at a 100 Hz offset. The oscillator includes digitally-controlled phase shifter and attenuator elements, which allow for a flexible approach in the oscillator design and a low residual noise quad parallel differential amplifier. The effect of using an attenuator in the oscillator loop to improve phase noise is also explored.

Keywords—oscillators, low phase noise, crystal oscillators

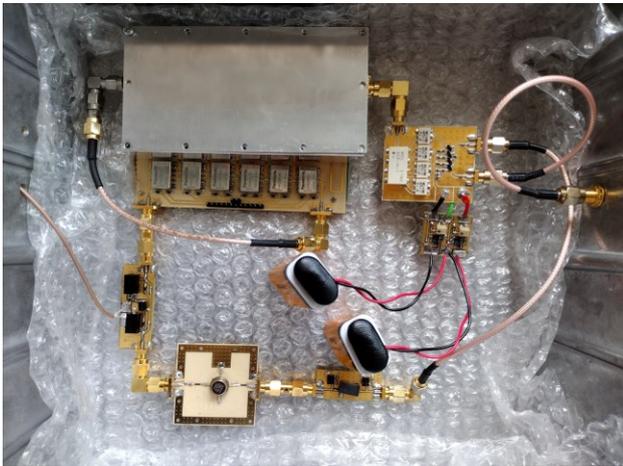


Fig. 1. Photo of the ultra-low phase noise 100 MHz crystal oscillator

I. INTRODUCTION

Minimising phase noise and jitter is an important factor in designing oscillators used to provide a very stable reference frequency required for many modern systems. For example, ultra-low phase noise oscillators operating at 100 MHz have applications in communications & RADAR systems, atomic clocks and particle accelerator systems.

There exist in the literature and as commercially-available products 100 MHz oscillators that exhibit a phase noise performance around -140 dBc/Hz at a 100 Hz offset [1] [2], however their construction and architecture are not described in detail.

This research builds on work done by this group in producing low phase noise crystal oscillators operating at 10 MHz [3] [4].

II. DESIGN OVERVIEW

A block diagram of the oscillator is shown in Fig. 2. and its constituent components are described in detail in the following sections.

A. Crystal Resonator

The resonator is a KVG #O5SC85105A [4] 5th overtone SC-cut quartz crystal with a stated maximum phase noise of

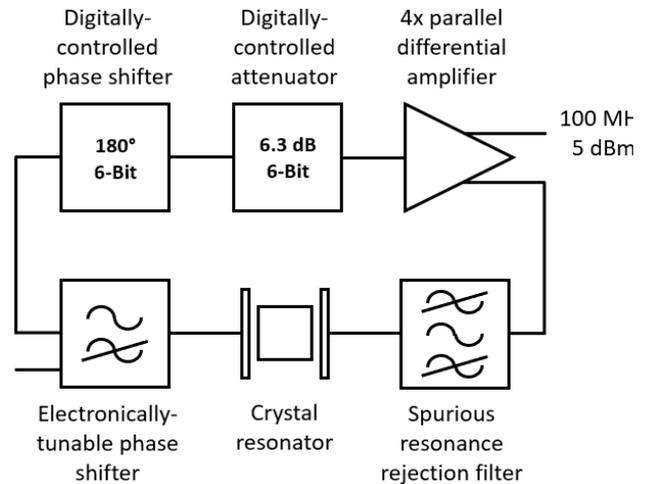


Fig. 2. Oscillator block diagram

-130 dBc/Hz at a 100 Hz offset. It is operated at a drive level of around 5 dBm.

B. Spurious Resonance Rejection Filter

An LC band pass filter is used to prevent oscillation at spurious resonant frequencies of the crystal. It is a 3rd-order Butterworth filter with a 40 MHz bandwidth and centre frequency at 100 MHz. The circuit diagram is shown in Fig.3.

C. Electronic Phase Shifter

The electronic phase shifter is used for narrow-band frequency tuning of the oscillator via an external control voltage. The circuit diagram is shown in Fig.4. It is based on the design described in [3]. The circuit is a 5th-order LC high pass filter with a cutoff frequency at around $0.6 \cdot f_0$ where varactor diodes biased through a voltage applied to the centre inductor are used as tuning elements. In addition to allowing the oscillator to be used in a phase-locked loop configuration, the oscillator loop phase can be adjusted to minimise phase noise degradation due to open-loop phase error as described in [5].

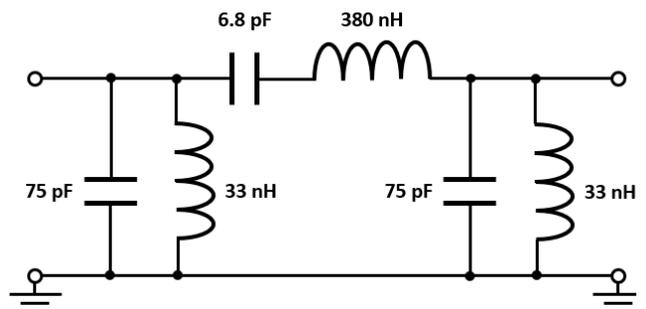


Fig. 3. Circuit diagram of the spurious resonance rejection filter

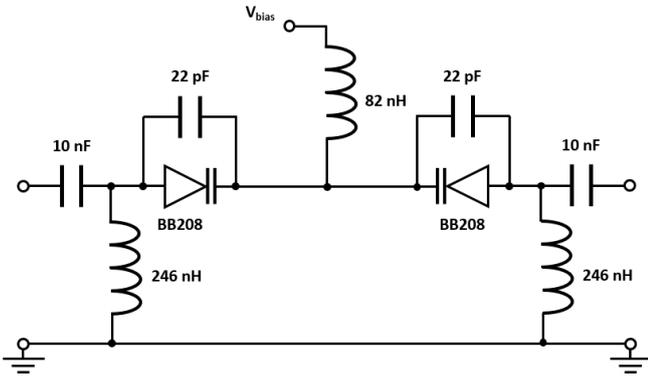


Fig. 4. Circuit diagram of the electronic phase shifter

III. DIGITALLY-CONTROLLED PHASE SHIFTER AND ATTENUATOR

The digitally-controlled phase shifter and attenuator are passive devices that use Panasonic ARA220A12 latching RF relays to switch a series of delay line / attenuator elements into the signal path based on a digital control signal. These devices are designed to provide a quick and simple way to test alternative oscillator configurations and amplifier designs that require different operating conditions. The phase shifter and attenuator share a similar construction and in both cases were found to have an insertion loss of less than 0.5 dB. A passive design was selected to minimise flicker noise and it was found to be low enough that it could not be measured on any residual noise measurement systems that were available, including an R&S FSWP [6].

A. 6-Bit 180° Phase Shifter

The purpose of the digital phase shifter is to set the loop phase of the oscillator to roughly 0° so that it can then be finetuned using the electronic phase shifter. Having a phase shifter that operates over such a wide range eliminates the need to insert additional lengths of coaxial cable into the loop that must be changed whenever the open-loop phase shift of the oscillator is altered during the design and optimisation process.

The device operates by using microstrip transmission lines as delay elements that can be switched into the signal path or bypassed via a straight-through connection. The most significant bit of the digital control signal corresponds the longest delay line and for each subsequent bit the length of the microstrip is roughly halved, plus a fixed length equal to the length of the straight-through path. This allows the electrical length of the device to be varied linearly according to the binary number represented by the digital control code.

B. 6-Bit 6.3 dB Attenuator

The purpose of the digital attenuator is to set the loop power of the oscillator such that the drive level of the crystal is as high as possible without degrading phase noise performance or driving the amplifier too far into saturation. A 0.1 dB resolution gives reasonable control over the overall loop power.

This device operates in a very similar way to the 6-bit phase shifter, although in this case the microstrip delay lines are replaced with resistive attenuator networks. The most

significant bit represents the network with the most attenuation and the attenuation is halved for each subsequent bit.

IV. QUAD PARALLEL DIFFERENTIAL AMPLIFIER

The amplifier uses a parallel configuration of 4 long-tailed pair differential amplifiers that is designed to provide a higher maximum available output power compared to a single long-tailed pair amplifier with a similar gain, whilst reducing the quiescent current through each transistor. The complementary outputs of the amplifier allow for one output to be fed back into the oscillator loop while the other is used as a buffered oscillator output signal.

A. Theory of Operation

The single differential amplifier subcircuit used in the design is shown in Fig. 5. and is based on the design described in [3]. It uses a transformer with 2:1 impedance transformation ratio at its input to produce the differential signal required for driving the inputs to the long-tailed pair. The quiescent current through the transistors is set by the approximate current source created by the combination of a negative bias voltage and fixed 680 Ω resistor.

A block diagram of the quad parallel combination of these amplifiers used in the design is shown in Fig. 6. A 0° power splitter is used at the input to supply the signal to each of the differential amplifier subcircuits and the outputs are then combined by connecting the collectors of the transistors together, effectively summing the collector currents. When this is compared with the single differential amplifier in Fig. 5. it can be seen that the current through each of the 50 Ω collector resistors is increased by a factor of 4, which corresponds to a 12 dB increase in power due to the parallel combination at the output. Since the maximum available output power of the long-tailed pair is set by the quiescent current through each transistor, this is also increased by 12 dB. The gain of the amplifiers can be compared by considering both the increase to the output power and the -6 dB at the input to each differential amplifier subcircuit due to the power splitter. The gain of the quad amplifier compared to the single amplifier is then increased by $12\text{ dB} - 6\text{ dB} = 6\text{ dB}$.

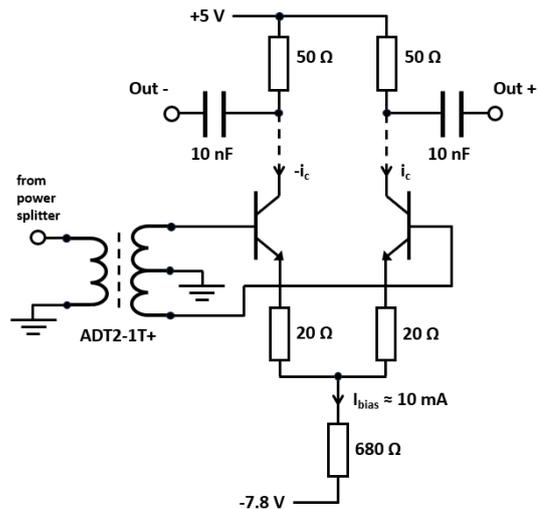


Fig. 5. A single long-tailed pair differential amplifier subcircuit used in the quad parallel differential amplifier

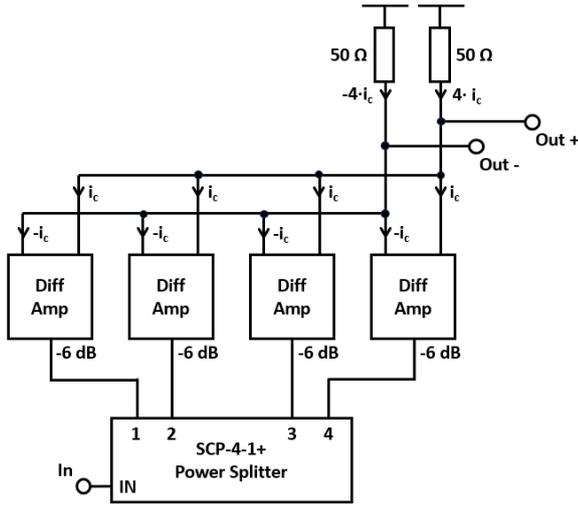


Fig. 6. Block diagram of the quad parallel differential amplifier

B. Measurements

The amplifier has a small-signal gain of 7.6 dB and the 1 dB gain compression point occurs for an output power of around 3 dBm. The noise figure was measured to be 3.8 dB using an R&S FSWP.

Since the close-to-carrier residual phase noise of the amplifier was found not to be sufficiently above the internal noise floor of available residual noise measurement systems, the flicker noise corner of the amplifier, f_c was determined by measuring the absolute phase noise of the oscillator where the crystal was replaced with a comparatively low Q LC resonator. The $1/f^3$ to $1/f^2$ transition point of the oscillator phase noise measurement corresponds to the flicker noise corner of the amplifier [7]. This method was used to determine that the f_c of the amplifier for an output power of 5 dBm is around 80 Hz.

V. ATTENUATOR PLACEMENT

The placement of the attenuator at the input of the amplifier was selected to mitigate against issues relating to the impedances presented to the crystal and the amplifier when in the oscillator loop.

The ratio of the loaded Q to the unloaded Q of the crystal, Q_L/Q_0 , is an important factor in designing low phase noise oscillators. It is shown in [8] that minimum phase noise occurs when Q_L/Q_0 is equal to approximately 0.5 and can be calculated as follows:

$$\frac{Q_L}{Q_0} = \frac{R_{loss}}{R_{in} + R_{out} + R_{loss}} \quad (1)$$

Where R_{loss} is the effective series resistance of the crystal, R_{in} is the source impedance at the input and R_{out} is the load impedance at the output.

If a crystal is designed for optimal Q_L/Q_0 in a 50 Ω system, then a deviation in the load impedance could result in degradation of the phase noise.

The residual phase noise of an amplifier is also typically characterised for use in a 50 Ω system. The noise performance

of an amplifier is often dependent on the source impedance that is presented to it, so if an amplifier with low measured residual phase noise is placed in an oscillator loop where the source impedance is different, the oscillator phase noise could be worse than expected.

A. System Impedances

Fig.7. a) shows a simplified block diagram of the oscillator including only the resonator, amplifier and an attenuator matched to a 50 Ω system at the output of the amplifier. The amplifier is assumed to have an output impedance of 50 Ω, which in the case of the quad parallel differential amplifier is set by the collector resistor, and the output impedance of the attenuator is therefore also 50 Ω. The output impedance of the resonator is represented by Z_R and in this configuration, this is also the source impedance that is presented to the amplifier input, which could affect the noise performance of the amplifier. Similarly, the input impedance of the amplifier, represented by Z_A , is the load impedance presented to the resonator. If Z_A deviates from 50 Ω, which is true in the case of the quad parallel differential amplifier, this could degrade the phase noise by affecting the ratio of Q_L/Q_0 .

Fig.7. b) shows the configuration where the attenuator is instead placed at the input to the amplifier. Here the load impedance presented to the resonator is instead the effective input impedance of the attenuator, Z_{in} , and the source impedance presented to the amplifier is the effective output impedance of the attenuator, Z_{out} . It can be shown that Z_{in} and Z_{out} will become closer to 50 Ω as the amount of attenuation increases.

B. Attenuator Impedance Transformation

A resistive attenuator can be modelled as a pi-pad network as shown in Fig.8, where the values of the resistors R_A and R_B can be calculated as follows:

$$R_a = Z_0 \cdot \frac{1 + K}{1 - K} \quad (2)$$

$$R_b = Z_0 \cdot \frac{1 - K^2}{2 \cdot K} \quad (3)$$

Where K is the attenuation of the device in terms of V_{out}/V_{in} .

The effective input impedance of the attenuator when terminated with a load resistance, R_L , can then be calculated by considering the resulting resistor network as follows:

$$Z_{in} = R_a \parallel (R_b + (R_a \parallel R_L)) \quad (4)$$

If Z_0 is 50 Ω then it can be seen from (2) that as K decreases (i.e. for a higher level of attenuation), the value of R_A will tend towards 50 Ω, and from (3) that the value R_B will become much greater than 50 Ω. Applying these to (4) it can be seen that as the level of attenuation increases, the effective input impedance of the attenuator will become closer to 50 Ω even if R_L is mismatched. Since the device is symmetrical the inverse is also true, the effective output impedance of the

attenuator will become closer to 50Ω even if R_S is mismatched.

Therefore it is believed that the placement of the attenuator at the amplifier input can help to mitigate against degradation of phase noise due to impedance mismatch between the resonator and the amplifier.

VI. OSCILLATOR MEASUREMENTS

The absolute phase noise of the oscillator is shown in Fig.9. It was measured using an R&S FSWP and a smoothing function has been applied so that frequency spurs, particularly that at 100 Hz due to mains interference, can be ignored. This can be seen overlayed on the raw measurement.

The phase noise at 100 Hz can be seen to be just below -140 dBc/Hz and the far-out noise floor is around -168 dBc/Hz.

VII. CONCLUSIONS

A 100 MHz crystal oscillator has been developed that demonstrates phase noise performance comparable to state-of-the-art commercially-available oscillators. Each of the constituent parts of the oscillator has been described in detail and a flexible approach in the construction of the oscillator has been demonstrated.

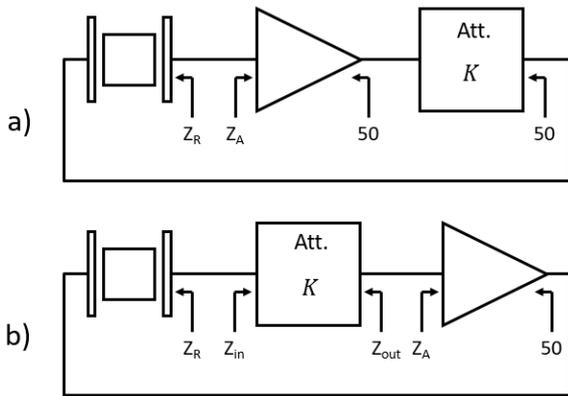


Fig. 7. Simplified block diagram of the oscillator to demonstrate the effects of pre-amplifier and post-amplifier attenuation

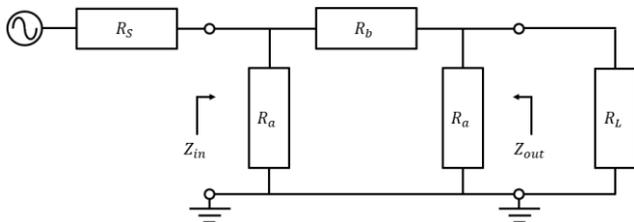


Fig. 8. Equivalent circuit of a resistive attenuator placed between source and load impedances

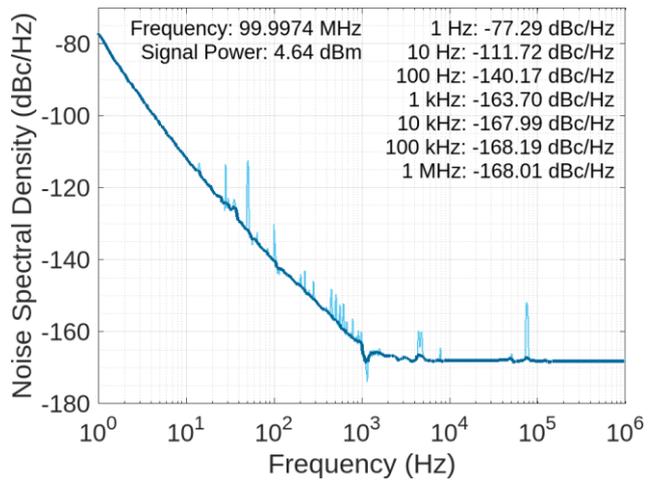


Fig. 9. Oscillator phase noise measured using the R&S FSWP

ACKNOWLEDGEMENTS

The authors would like to thank Agilent Technologies Inc., Santa Clara, California & The UK Engineering and Physical Sciences Research Council (EPSRC) for their funding and support.

REFERENCES

- [1] T. McClelland, C. Stone and M. Bloch, "100 MHz crystal oscillator with extremely low phase noise," Proceedings of the 1999 Joint Meeting of the European Frequency and Time Forum and the IEEE International Frequency Control Symposium (Cat. No.99CH36313), Besancon, France, 1999, pp. 331-334 vol.1
- [2] KVG, "O-40-ULPN-100M," O-40-ULPN-100M oven controlled crystal oscillator datasheet
- [3] J. Everard, T. Burtichelov, and K. Ng, "Ultralow Phase Noise 10-MHz Crystal Oscillators", IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Control, vol. 66, no. 1, January 2019
- [4] J. Everard and K. Ng, "Ultra-Low Phase Noise Crystal Oscillators," 2007 IEEE International Frequency Control Symposium Joint with the 21st European Frequency and Time Forum, Geneva, Switzerland, 2007, pp. 1246-1250
- [5] K. Cheng and J. Everard, "Noise performance degradation in feedback oscillators with nonzero phase error," Microw. Opt. Technol. Lett., 4: 64-66.
- [6] Rohde & Schwarz, "R&S FSWP Phase Noise Analyzer and VCO Tester," FSWP specifications, Feb 2024
- [7] D. B. Leeson, "A simple model of feedback oscillator noise spectrum," in Proceedings of the IEEE, vol. 54, no. 2, pp. 329-330, Feb. 1966
- [8] J. Everard, "Fundamentals of RF Circuit Design: With Low Noise Oscillators", Hoboken, NJ, USA: Wiley, 2002.

Appendix B

EFTF 2024 Conference Paper:

A Low-Cost Residual Phase Noise

Measurement System & Efficient

Digital Signal Processing

Techniques

A Low-Cost Cross-Correlation Residual Phase Noise Measurement System & Efficient Digital Signal Processing Techniques

Luke Dummott
School of Physics, Engineering and
Technology
University of York
York, United Kingdom
luke.dummott@york.ac.uk

Simon Bale
School of Physics, Engineering and
Technology
University of York

Jeremy Everard
School of Physics, Engineering and
Technology
University of York

Abstract— This paper describes a residual phase noise measurement system using cross-correlation techniques that is built from relatively low-cost and readily available components and exhibits a system noise floor of around -177 dBc/Hz when measuring a device operating at 0 dBm output power. A software library for the digital signal processing of residual noise measurements has also been produced, which will be provided open-source.

Keywords— cross correlation, residual phase noise, phase noise measurement



Fig. 1. Photo of the cross-correlation residual noise measurement system

I. INTRODUCTION

When designing electronic devices to be used in the production of oscillators with exceptional phase noise performance, it is important to be able accurately measure their residual phase noise characteristics. This becomes increasingly difficult, especially at lower operating power levels, due to the internal noise limitations of residual phase noise measurement systems

Cross-correlation techniques can be used as a way to suppress the internal noise floor of residual phase noise measurement systems [1]. When combined with efficient digital signal processing algorithms, it becomes possible to perform accurate residual phase noise measurements of very low noise devices in a reasonable amount of time.

This research is based on previous work done at York to build a cross-correlation residual noise measurement system [2], but here the system is designed to be low-cost and operate at 100 MHz with a DUT output power of around 0 dBm for the purpose of measuring amplifiers to be used in an ultra-low

phase noise 100 MHz crystal oscillator. The digital signal processing aspects of the system have also been greatly expanded upon to reduce measurement times.

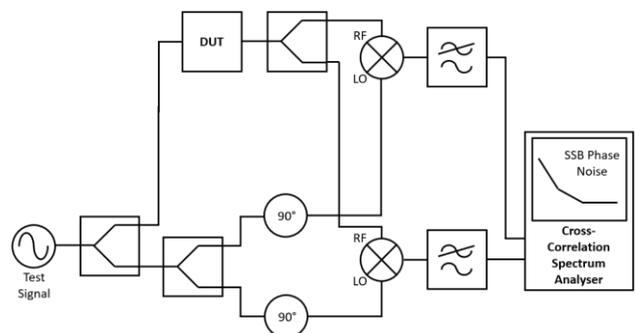


Fig. 2. Simplified block diagram of a cross-correlation residual phase noise measurement

II. THE CROSS-CORRELATION METHOD

A simplified block diagram of a cross-correlation residual phase noise measurement system is shown in Fig.2. Residual phase noise of a DUT (device under test) can be determined by measuring the spectrum of phase variation between a reference signal and the same signal after it has been passed through the DUT.

To measure the phase variation, a frequency mixer with RF and LO ports in quadrature is used as a phase detector. The baseband part of the output of the mixer corresponds to the variation in phase between the RF and LO ports and the high frequency part is removed using a suitable low pass filter.

A cross-correlation system splits the signals from the reference signal and the DUT output so that the measurement is performed in two separate channels. The phase noise of the DUT will be present as correlated noise between both channels. A cross-correlation spectrum analyser can then be used to determine then single-sideband residual phase noise.

A. Derivation

In the model described by Rubiola & Vernotte [3] the time domain signals captured at the input to the cross-correlation spectrum analyser, $x(t)$ and $y(t)$, can be considered to consist of a correlated component, representing the noise due to the DUT; and an uncorrelated component, representing the system noise in each channel. This can be written as:

$$\begin{aligned}x(t) &= a(t) + c(t) \\y(t) &= b(t) + c(t)\end{aligned}\quad (1)$$

Where $x(t)$ and $y(t)$ represent the uncorrelated system noise and $x(t)$ represents the correlated DUT noise.

The Fourier transform can be applied to give the frequency spectrum of each channel:

$$\begin{aligned}X(f) &= A(f) + C(f) \\Y(f) &= B(f) + C(f)\end{aligned}\quad (2)$$

The cross-spectrum of the two channels can then be calculated by performing the cross product of the frequency spectrum of one channel with the complex conjugate of the other and taking the average over M measurements. This is converted to power spectral density by multiplying by the measurement time, T , and dividing by the sampling frequency, f_s :

$$\begin{aligned}S_{XY} &= \frac{T}{f_s} \cdot \frac{1}{M} \sum_{m=1}^M [X_m Y_m^*] \\&= \frac{T}{f_s} \cdot \frac{1}{M} \sum_{m=1}^M [(A_m + C_m) \times (B_m + C_m)^*] \\&= \frac{T}{f_s} \cdot \frac{1}{M} \sum_{m=1}^M [(A_m B_m^*) + (A_m C_m^*) \\&\quad + (C_m B_m^*) + (C_m C_m^*)]\end{aligned}\quad (3)$$

The cross-spectrum of two perfectly correlated signals will be unchanged, while the cross-spectrum of two uncorrelated signals will be reduced in magnitude. In (3) it can be seen that the uncorrelated terms will decrease as the number of measurements increases, leaving just the $(C_m C_m^*)$ term, i.e. the equation becomes approximately equal to the power spectral density of the DUT noise.

It is shown in [3] that the uncorrelated terms which represent the system noise are suppressed by a ratio of \sqrt{m} . In other words, the cross-correlation method can be used to suppress the noise floor of the system along with the number of correlations performed at a rate of 5 dB per multiple of 10 correlations.

III. SYSTEM OVERVIEW

A block diagram of the hardware part of the cross-correlation residual phase noise measurement system is shown in Fig. 2. and the various aspects of the system are described in detail in the following sections.

A. Reference Oscillator

It is important that the reference oscillator has as low AM and PM noise as possible, so as not to limit the internal noise floor of the system. A KVG O-40-ULPN-100M oven-controlled crystal oscillator [4] was selected for the reference oscillator as it exhibits very low phase noise: around -138

dBc/Hz at a 100 Hz offset with a far-out noise floor of around -185 dBc/Hz. It is also capable of providing enough output power to saturate the LO ports of the mixers.

B. Quadrature Phase Shift

The required 90° phase shift between the RF and LO ports is achieved using a combination of 6-bit 180° digitally-controlled phase shifter and a 9° mechanical phase shifter.

The digitally-controlled phase shifter was originally built for use in ultra-low phase noise 100 MHz crystal oscillators. It uses latching RF relays to switch different lengths of microstrip transmission line into the signal path to create a phase shifter that can be adjusted linearly according to a digital control code. It is a completely passive device that was specifically designed for very low residual noise.

Once the digitally-controlled phase shifter has been used to set the phase shift to roughly 90°, the mechanical phase shifter is used to fine tune to the phase shift so that DC output of the mixers is as close to 0 V as possible.

C. DUT Power Level

A digitally-controlled attenuator with a range of 31.5 dB and a resolution of 0.5 dB is used to set the source power to the DUT, so that residual phase noise of the device can be measured at the desired operating conditions. This is a similar design to the digitally-controlled phase shifter, but here the microstrip delay line elements are replaced with resistive attenuator networks.

D. Mixers

To optimise the system for the desired DUT output level of 0 dBm, a low-power mixer was selected: the Mini Circuits ZX05-1L-S+ [5]. This mixer is designed to operate with an LO power of 3 dBm. Using such a low power mixer allows the RF and LO ports to be as close to saturation as possible to achieve the best performance even when measuring low-power device and also removes the need for additional amplifiers at the inputs to the mixer.

E. Signal Capture

A 16-bit PicoScope 4262 USB oscilloscope [6] is used to capture the time-domain signals so that they can be transferred to a PC for processing using the techniques described in later sections. It has two input channels that can capture samples simultaneously and is operated at a sampling frequency of 1 MHz.

F. Signal Conditioning

To provide the required baseband phase noise signals to the PicoScope, the high frequency part of the mixer output signals, i.e. the sum of the RF and LO frequencies, must be removed. This is achieved using a 10 MHz diplexer with high frequency output terminated in 50 Ω, so that it can be removed from the signal without unwanted reflections, and low frequency output connected to the rest of the baseband circuitry.

The amplitude of the output from each mixer is too low to be precisely measured directly with the PicoScope, so two-stage LT6018 operational amplifier circuits with a total voltage gain of 900 (around 59 dB) are used to increase the signal level.

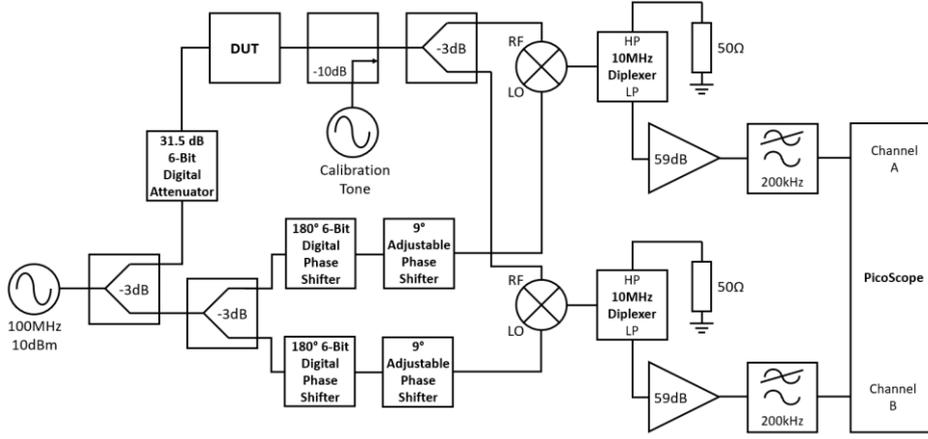


Fig. 3. Block diagram of the cross-correlation residual phase noise measurement system

An anti-aliasing filter is used at the input to the PicoScope to filter out any frequencies above the Nyquist frequency (half the sampling rate) prior to being sampled by the PicoScope, as these frequencies will be aliased into the measured noise spectral density. It is a 5th-order Butterworth low-pass LC filter with a cutoff frequency of around 250 kHz and was designed as part of previous work at York on phase noise measurement instrumentation [7].

G. Calibration

It is necessary to calibrate the system so that the noise power spectral density measured at the input to the PicoScope can be referenced to the carrier signal power to produce a residual phase noise measurement in dBc/Hz. This is achieved by injecting a tone into the system through a 10 dB coupler at the output of the DUT of a known offset frequency and power level in relation to the carrier signal. By observing the amplitude of the peak due to the calibration tone on the PicoScope, a reference can be obtained to convert the absolute noise spectral density measured on the spectrum analyser in dBV/Hz to dBc/Hz. As shown in [2], the calibrated single-sideband phase noise, $L(f)$, can be calculated from:

$$L(f) = S_{\phi} + K_1 - K_2 - 6 \text{ dB} \quad (4)$$

Where S_{ϕ} is the measured double-sideband power spectral density, K_1 is the power ratio of the carrier signal to the injected spur and K_2 is the amplitude of the injected spur.

IV. DIGITAL SIGNAL PROCESSING TECHNIQUES

The primary purpose of the software that handles the DSP (digital signal processing) for the system is to convert captured time-domain domain voltage samples at the inputs to the capture device to a cross-correlated power spectral density of the DUT residual phase noise. This involves performing a FFT (fast Fourier transform) for each channel and calculating the cross spectrum.

The secondary purpose of the software is to reduce the amount of time it takes to obtain useful measurements by utilising a frequency-banding algorithm that adjusts the resolution bandwidth versus number of correlations for different ranges, which allows for a higher degree of internal noise floor suppression at high offset frequencies whilst maintaining a narrow resolution bandwidth close-to-carrier. Frequency banding techniques are used in commercially-

available residual phase noise measurement systems, such as the R&S FSWP, although their implementation is not described in detail.

To allow residual phase noise measurements to be performed in real-time and on inexpensive hardware, a C++ library of efficient DSP algorithms has been produced, which will be provided open-source and explained in detail in further work. A brief description of the techniques used and their effect on processing time is given in the following sections.

A. Frequency Banding Using Variable FFT Window Size

Frequency banding is performed by splitting a block of captured samples into sets of different sizes for each frequency band. An FFT is performed on each set and the number of correlations performed at each frequency of a measurement is equal to the number of sets in the sample block. At lower offset frequencies, the data is split into fewer sets with a larger FFT window, so the results can have a narrow resolution bandwidth at the expense of reduced noise floor suppression due to fewer correlations. For higher offset frequencies, the data is split into more sets of a smaller window size, so a lower noise floor can be achieved at the expense of resolution bandwidth. Fig.4. shows a residual phase noise measurement of the baseband front end part of the system (amplifier, filter & PicoScope). The effect of the frequency banding can be seen in the cross-correlation measurement; the resolution bandwidth decreases with each subsequent band, but the noise floor suppression increases.

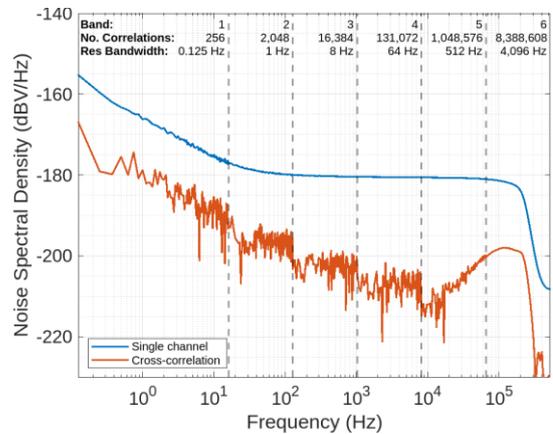


Fig. 4. Comparison of system front end noise floor (only amplifier, filter & PicoScope) for single channel and cross-correlation measurements

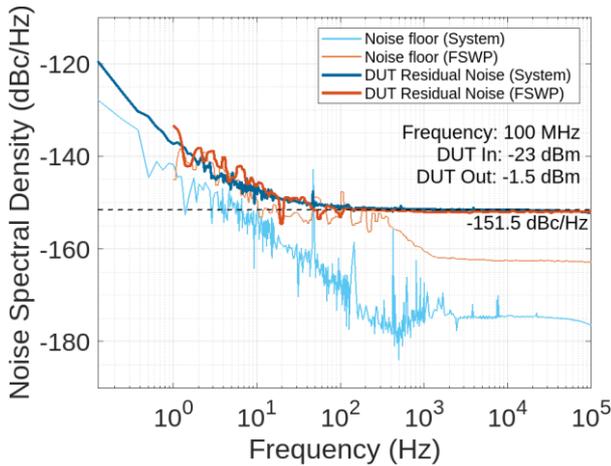


Fig. 5. Comparison the residual noise of a MAR-6+ amplifier when measured using this system and an R&S FSWP

B. Optimising the FFT for Frequency Banded Data

The implementation of the frequency banding means that much of the output data from each Fourier transform performed in a block when using a standard FFT algorithm remains unused, i.e. only a certain frequency range of frequency points are required for each band, which could be much less than the size of the FFT. An optimised FFT algorithm is implemented that is much more efficient when the required output data size is less than the input data size.

C. Processing Time

It has been shown that using the described techniques, the processing time of frequency banded residual noise measurements is comparable to conventional fixed resolution bandwidth techniques even on inexpensive hardware. These techniques have been used to perform a measurement with sub-Hz close-to-carrier resolution bandwidth and over 3,000,000 cross-correlations at high offset frequencies in less than 15 minutes.

V. MEASUREMENTS

A. Verification and Comparison to Existing Systems

To confirm the accuracy of the system, a measurement was made using the Mini Circuits MAR-6+ amplifier as the DUT. This amplifier has a gain of around 22 dB at 100 MHz and it was operated at an input power of -23 dBm. The same measurement was performed using a commercially-available R&S FSWP [8] measurement system and the results are compared in Fig.5. The internal noise floor of each system was also measured by replacing the amplifier with a straight through connection and setting the power level at this point to the output power of the amplifier. Each system was set to run for around 40 minutes with similar resolution bandwidths, though the difference in implementation of frequency banding means data points are not directly comparable.

It can be seen the measured residual phase noise was almost identical for each system, though comparisons cannot easily be made at frequencies below 1 kHz as the amplifier noise is not sufficiently above the noise floor of the FSWP. Despite its comparatively low cost, at these operating conditions the internal noise floor of this system is significantly lower than that of the FSWP; around 12 dB at high offset frequencies.

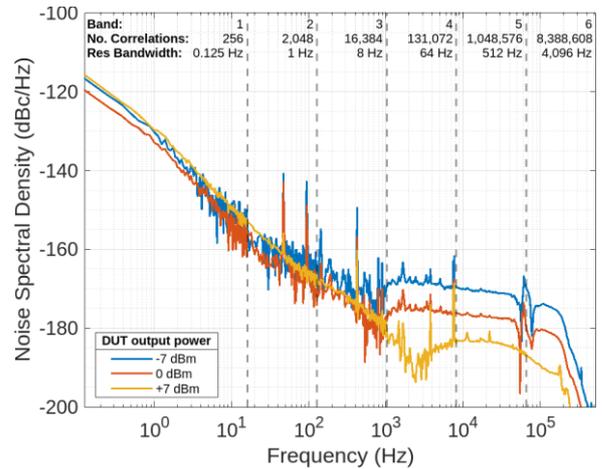


Fig. 6. Comparison of the system noise floor for various DUT output power levels

B. System Noise Floor

The internal noise floor of the system was measured for various DUT output powers as shown in Fig.6. The minimum noise floor was measured for a DUT output power of 7 dBm, when both ports of the mixers were saturated at around 3 dBm. Here, the far-out noise floor was less than -180 dBc/Hz.

VI. CONCLUSIONS

A low-cost cross-correlation residual noise measure operating at low power levels has been produced that has an internal noise floor lower than that of the commercially-available R&S FSWP. A set of digital signal processing algorithms for residual phase noise measurements has been produced that allows for measurements to be performed efficiently, in real-time and on inexpensive hardware.

ACKNOWLEDGEMENTS

The authors would like to thank Agilent Technologies Inc., Santa Clara, California & The UK Engineering and Physical Sciences Research Council (EPSRC) for their funding and support.

REFERENCES

- [1] W. F. Walls, "Cross-correlation phase noise measurements", *Proc. IEEE Frequency Control Symposium*, 1992, pp. 257–261.
- [2] S. J. Bale, D. Adamson, B. Wakley, J. Everard, "Cross Correlation Residual Phase Noise Measurements using Two HP3048A Systems and a PC Based Dual Channel FFT Spectrum Analyser", *24th European Frequency and Time Forum*, 13-16 April 2010
- [3] E. Rubiola and F. Vernotte, "The cross-spectrum experimental method," *arXiv*, 1003.0113v1, 2018.
- [4] KVG, "O-40-ULPN-100M," O-40-ULPN-100M oven controlled crystal oscillator datasheet
- [5] Mini Circuits, "Frequency Mixer," ZX05-1L-S+ datasheet
- [6] Pico Technology, "PicoScope 4262 High-Resolution USB Oscilloscope", PicoScope 4262 datasheet
- [7] E. Lamb, "Ultra-Low Noise Floor Phase Noise Measurement Instrumentation for Precision RF/Microwave Component Characterisation," *MEng Project Report*, Department of Electronic Engineering, University of York, 2020-21
- [8] Rohde & Schwarz, "R&S FSWP Phase Noise Analyzer and VCO Tester," FSWP specifications, Feb 2024