

Methods for ultra-broadband correlator development focusing on high-speed digital sampling techniques



Adam Coates
St Cross College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2013

Contents

1	Introduction	3
1.1	Radio instrumentation	3
1.1.1	Interferometry	4
1.2	The Correlator	6
1.2.1	Correlator implementations	8
1.2.2	Real and Complex correlators	10
1.2.3	Double-sideband complex correlators	17
1.3	Analogue and digital correlators	20
1.3.1	Analogue correlators	20
1.3.2	Digital correlators	21
1.4	Current radio-telescopes	22
1.4.1	AMiBA	23
1.4.2	CBI	23
1.4.3	AMI	26
1.4.4	ATCA	27
1.4.5	CARMA	27
1.4.6	EVLA	28
1.4.7	ALMA	28
1.5	The need for broader bandwidth telescopes	29
1.6	Overview of thesis	30
2	Analogue correlation techniques	33
2.1	System overview	33
2.1.1	Prior tests	35
2.2	Analogue correlator bench tests	36
2.2.1	Noise behaviour	36
2.2.2	Single-frequency recovery	38
2.3	Analogue correlator calibration - simulation and implementation . . .	38

2.3.1	Imperfections	38
2.3.2	Calibration methods	41
2.3.3	Simulation of calibration scheme	43
2.3.4	Reduced bandwidth calibration	44
2.3.5	Implemented calibration	47
2.3.6	Broadband calibration	50
2.4	Conclusion	53
3	Digital Sampling techniques: 1-bit sampling	55
3.1	Direct signal input sampling	55
3.1.1	Introduction to direct sampling	55
3.1.2	Hardware platform	56
3.1.3	Transceiver configuration	60
3.1.4	RF ROACH interfacing	62
3.1.5	Initial testing, pseudo random bit sequence and interface development	63
3.1.6	Single-frequency data capture	70
3.2	External comparator sampling	74
3.2.1	Comparator setup	75
3.2.2	Single-frequency capture	76
3.3	Conclusion	80
4	Digital Sampling techniques: Interleaving techniques	82
4.1	Basic principles of interleaving	82
4.2	ROACH implementation of channel alignment	84
4.3	Correlation technique for improving sensitivity	90
4.3.1	Correlated noise	90
4.3.2	Phase switching	94
4.4	Testing interleaving using phase switching	99
4.5	Conclusion	102
5	Digital Sampling techniques: Multiple-bit techniques	103
5.1	Multi-bit techniques	103
5.2	Hysteresis-based multi-level sampling	109
5.3	Reference-based multi-level sampling	115
5.3.1	Low-frequency square wave bias	115
5.3.2	Removal of the AC coupling	116

5.3.3	Re-introducing AC coupling	124
5.3.4	Producing a multi-bit sampler	127
5.4	Use of bit-depth and interleaving	133
5.5	Conclusion	136
6	Hardware Digital correlator	139
6.1	Correlator architecture	139
6.2	Correlator implementation	140
6.3	Hardware correlation of single-frequencies	145
6.4	Broadband recovery	146
6.5	Broadband recovery with manual path delay added	150
6.6	Single-frequency with uncorrelated noise	150
6.7	Broadband signal with uncorrelated noise	154
6.8	Correlating interleaved signals	155
6.9	Conclusion	158
7	RF and IF systems for broadband interferometry	160
7.1	RF and IF components	161
7.1.1	2-20 GHz amplifier	161
7.1.2	Low pass filters	169
7.1.3	Path compensation	169
7.1.4	Variable attenuation	174
7.1.5	Phase switching	174
7.1.6	Slope compensation	176
7.1.7	Reduced-bandwidth amplifier	187
7.1.8	Low-pass filter	190
7.1.9	Digitally-controlled attenuator	190
7.2	IF chain	194
7.2.1	Basic IF chain construction	194
7.2.2	IF chain complete with S-parameters	195
7.3	Conclusion	199
8	Conclusion	200
8.1	Overview	200
8.1.1	Analogue correlator	200
8.1.2	Digital sampling techniques and digital correlator	200
8.1.3	RF and IF device development	201

8.2	Comparing analogue to digital approach	201
8.3	Future research	202
8.3.1	Analogue correlator	202
8.3.2	Digital sampling and correlation	203
8.3.3	On-sky testing	204
	Bibliography	205

Glossary

AC	Alternating Current
ADC	Analogue-to-Digital Converter
AGC	Active Gain Control
ASIC	Application Specific Integrated Circuit
BPF	Band Pass Filter
CASPER	Collaboration for Astronomy Signal Processing and Electronic Research
CLB	Configurable Logic Block
CMB	Cosmic Microwave Background
CMT	Clock Management Tiles
CW	Carrier Wave
CX4	(Differential digital connector standard)
DC	Direct Current
DRP	Dynamic Reconfiguration Port
DSP	Digital Signal Processing
ENOB	Effective Number of Bits
FET	Field Effect Transistor
FFT	Fast Fourier Transform
FPGA	Field Programmable Gate Array
FT	Fourier Transform
FWHM	Full Width Half Maximum
GPIO	General Purpose Input/Output
GPU	Graphical Processing Unit
GTP	(Xilinx high speed serial interface)
GTX	(Xilinx high speed serial interface)
HDL	Hardware Description Language
HPF	High Pass Filter
IC	Integrated Circuit
IF	Intermediate Frequency
JTAG	Joint Test Action Group
LFSR	Linear Feedback Shift Register

LNA	Low Noise Amplifier
LO	Local Oscillator
LPF	Low Pass Filter
PC	Personal Computer
PCI	Peripheral Component Interconnect
Pk-pk	Peak-to-Peak
PLL	Phase Locked Loop
PRBS	Pseudo Random Bit Sequence
RAM	Random Access Memory
RF	Radio Frequency
RFI	Radio Frequency Interference
ROACH	Reconfigurable Open Architecture Computing Hardware
RX	Receive
SINAD	Signal in Noise and Distortion
SMA	SubMiniature version A
SNR	Signal-to-Noise Ratio
SPDT	Single-Pole Dual-Throw
SPI	Serial Peripheral Interface
SZ	Sunyaev-Zeldovich
TX	Transmit
USB	Universal Serial Bus
VHDL	Very-High-Speed Integrated Circuits Hardware Description Language
VNA	Vector Network Analyser
VVA	Voltage-Variable Attenuator
Z-DOK	(Type of differential digital connector)

Chapter 1

Introduction

Observations of the Universe can be made at many different frequencies. For most of history astronomy has been dominated by observations made at optical frequencies. However, a wealth of information can be obtained about the Universe if different frequencies of light are considered. From the middle of the twentieth century developments in technology have led to new telescopes being constructed that observed in frequency bands spread across the entire electromagnetic spectrum.

This thesis will be concerned with the technology used in radio frequency telescopes, concentrating on technologies used for implementing ultra-broadband correlators.

1.1 Radio instrumentation

The radio frequencies considered in this thesis are in the microwave region (GHz or tens of GHz). At these frequencies radio instrumentation can be divided into two categories. The first of these is single-dish receivers. In this design a single dish is used to focus the radiation onto a feed and hence to a detector system. This can then be scanned across the sky to construct a map of the source. Multiple pixels in an array can be used to speed up the mapping process.

The beam width for a single dish device, and therefore the angular resolution on the sky, is given by:

$$\theta \propto \frac{\lambda}{D}, \quad (1.1)$$

where θ is the angular resolution or beam width, λ is the wavelength of light and D is the dish diameter. In the gigahertz frequency range and with typical dish sizes, this results in a very low angular resolution on the sky. The largest fully-steerable

dishes are ≈ 100 m in diameter, however at these sizes practical limitations begin to restrict their operation, such as the pointing accuracy. The increase in dish size has primarily been intended to improve the sensitivity of the telescope, but represents the limit of resolution achievable from a single dish.

To improve the resolution available, the second category of telescopes was invented, namely interferometers. In an interferometer, multiple antennae are linked in a coherent array where the signals received from each dish are correlated, in pairs, with those from all other dishes. This results in the angular resolution being modified to:

$$\theta \propto \frac{\lambda}{B}, \quad (1.2)$$

where D is replaced by B , the longest baseline (separation between two antennae) of the array. High spatial resolution can therefore be generated for the same dish size by using large antenna separations. The dishes of the array form a single synthesised dish with a diameter equal to the longest baseline. However, this synthesised dish is missing large parts of its collecting area (i.e. has a poor filling factor). This can be improved by adding more dishes at intermediate baselines to form a larger and better filled array. For many purposes a sparsely-filled interferometer will still produce perfectly acceptable images.

1.1.1 Interferometry

The basic layout of an interferometer can be seen in the top diagram of Figure 1.1. This shows a two-element (single baseline) interferometer where the signals are received by two antennae and fed to a common point where they are correlated. The correlator output corresponds to one spatial frequency component of the sky within the beam defined by the individual antennae. In larger arrays the schematic is repeated for all combinations of pairs of antennae.

Between the antenna and the correlator RF signal processing is performed, designed to prepare the raw signals being received by the antennae for correlation. The most common components used in this area are amplifiers and filters, which will be discussed later.

One key component that is often used is a heterodyning stage. This is where the signal from each antenna is first mixed with a local oscillator (LO) signal in a mixer. This reduces the frequency of the signals used without destroying any information. If a sky signal $V_{sky} \cos(\omega_{sky}t)$ is mixed with an LO signal $V_{LO} \cos(\omega_{LO}t)$:

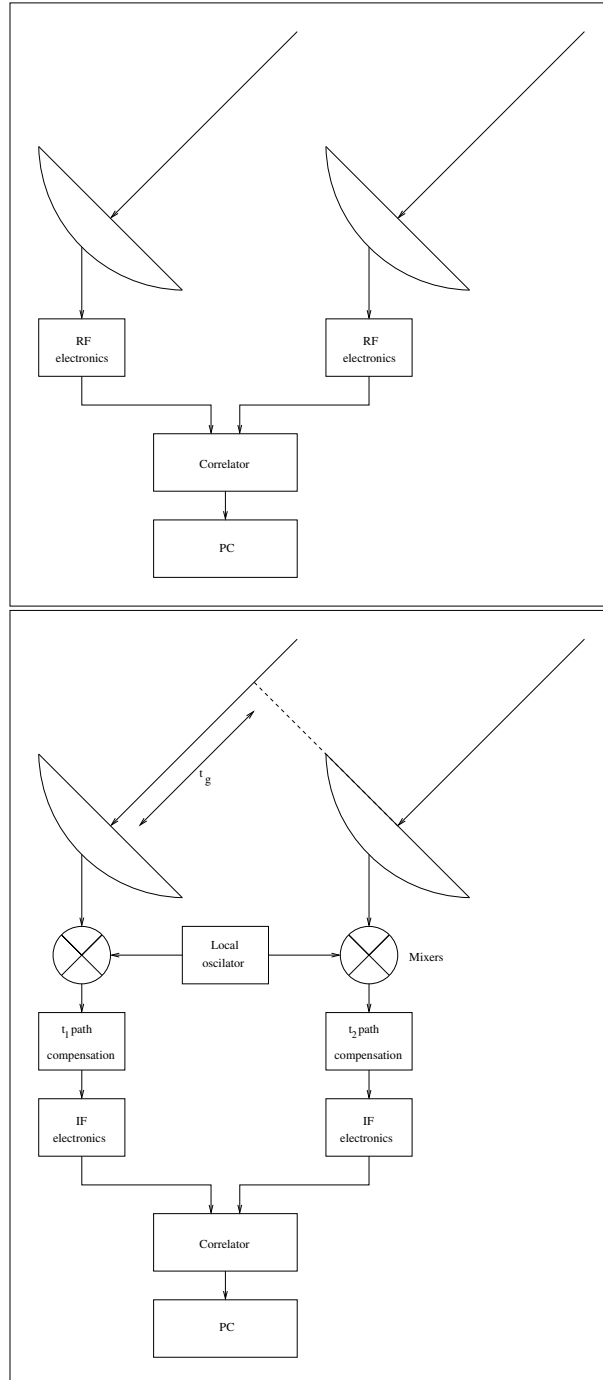


Figure 1.1: Schematic diagram of an interferometer. (Top) Simple setup with no heterodyning, (Bottom) setup with heterodyning and path compensation. t_g represents the geometric path difference for signals entering the two dishes. t_1 and t_2 are path compensation added to the telescope system to combat the geometric delay.

$$\begin{aligned}
V_{out} &= V_{sky} \cos(\omega_{sky}t) \times V_{LO} \cos(\omega_{LO}t) \\
&= V_{sky} V_{LO} \left(\frac{\cos((\omega_{sky} - \omega_{LO})t) + \cos((\omega_{sky} + \omega_{LO})t)}{2} \right) \quad (1.3)
\end{aligned}$$

As can be seen the output is two sideband signals, one at $\omega_{sky} - \omega_{LO}$ and the other at $\omega_{sky} + \omega_{LO}$.

The upper sideband signal can be filtered out using a low-pass filter leaving all the original signal information that existed at ω_{sky} now at the intermediate frequency (IF) frequency. The IF is where most of the signal processing occurs. Heterodyning is usually done since RF components are more complicated to make at higher observation frequencies than at a lower IF frequency. Loss from cables and other devices is also greater at higher frequencies. It also means that one generic IF chain and correlator can be used with many different front ends operating at different sky frequencies. The layout for an interferometer with a heterodyning stage can be seen in the bottom diagram of Figure 1.1.

Also included in this figure is the concept of path compensation, represented by the path delays t_1 and t_2 . When an object is viewed on the sky both antennae will be pointed at it. If, as in the figure, the direction of the source is not perpendicular to the baseline between the two antennae then the propagation delay of the light from the source to each of the dishes is different by an amount t_g . This will result in the signals from the two antennae not arriving at the correlator at the same time (coherently). To compensate for this geometric propagation delay, electronic path delay is added between the antennae and the correlator. This can be electronically controlled to exactly cancel the geometric delay, bringing the signals into the correlator coherently.

To form an interferometer capable of producing images of the sky, multiple dishes are required, arranged at different separations (baselines). If the vector from the telescope to the source in the sky is defined as S_0 , then a set of orthogonal basis vectors, (w, u, v) can be constructed to define the coordinate space. If w acts along S_0 , then u and v define the plane perpendicular to it. The projection of the baseline vectors onto this u-v plane defines the imaging characteristics of the telescope.

1.2 The Correlator

The heart of any interferometer is the correlator. This is the device that takes the signals from each of the antennae and performs the cross-correlation between all of




Figure 1.2: Schematic diagram of a 16-lag correlator. The signal from antenna A enters from the left and from antenna B from the right. Figure from [18].

the pairs. The output from the correlator can be used to gain both spectral and spatial information about the observed source. The correlation function is defined as:

$$R(\tau) = \frac{1}{T} \int_0^{2T} V_A(t)V_B(t - \tau)dt, \quad (1.4)$$

where V_A and V_B are two input signals, τ is the relative propagation delay between the two signals reaching the correlator and $R(\tau)$ is the correlator output for a given propagation delay, with the integral performed over a large timescale which is multiple periods of the signals. The output is the time-averaged correlation of the two signals. If the signals are assumed to be single-frequency wave-like signals of the form $V_A(t) = V_B(t) = A_\nu \sin(2\pi\nu t)$, the above formula can be rewritten as:

$$\begin{aligned} R(\tau) &= \frac{1}{T} \int_0^{2T} A_\nu^2 \sin(2\pi\nu t) \sin(2\pi\nu(t - \tau)) dt \\ &= \frac{1}{2T} \int_0^{2T} A_\nu^2 \left(\cos(2\pi\nu t - 2\pi\nu(t - \tau)) - \cos(2\pi\nu t + 2\pi\nu(t - \tau)) \right) dt \\ &= \frac{1}{2T} \int_0^{2T} A_\nu^2 \left(\cos(2\pi\nu\tau) - \cos(4\pi\nu t - 2\pi\nu\tau) \right) dt. \end{aligned} \quad (1.5)$$

The second cosine term will integrate to zero when the integration is performed over a whole number of periods. This leaves

$$R(\tau) = \frac{1}{2T} \int_0^{2T} A_\nu^2 \cos(2\pi\nu\tau) dt = A_\nu^2 \cos(2\pi\nu\tau). \quad (1.6)$$

This is for a single-frequency signal. If instead a broadband signal is assumed, then the individual frequencies can be treated separately and the total correlator response is given by:

$$R(\tau) = \int A_\nu \cos(2\pi\nu\tau) d\nu, \quad (1.7)$$

where A_ν is the complex magnitude of the signal of frequency ν received on both antennae from the source, and the integral is performed over the frequency band. The above formula can be seen to be the real part of the Fourier transform of the input passband in propagation delay (or lag) space:

$$R(\tau) = \text{Re}[FT(A_\nu)] \quad (1.8)$$

1.2.1 Correlator implementations

Two distinct methods exist for performing correlations in hardware environments, these are XF (or lag) correlators and FX correlators. These terms refer to the order in which the cross multiplication of different antenna data streams (X) and the Fourier transform (F) are performed. Both these methods are equivalent but rely on a different ordering of the steps involved in the correlation.

1.2.1.1 XF or lag correlators

In a lag correlator each pair of data streams is cross-multiplied and summed in an accumulator. This is performed for different relative propagation delays of the RF streams to fully sample the correlation function. The output of the correlator is then passed through a discrete inverse Fourier transform to recover the spectral information of the original signal.

To provide the relative propagation delays within the correlator, different path lengths from each input to the multipliers must be used. This can be seen in Figure 1.2 where the correlations are performed along the diagonal, with the top lag having a short delay from the left and a long delay from the right, etc.

At the frequency range being considered this technique is very feasible as the wavelengths of the signals are of the order of mm. This means any delay length required would be easy to produce on a circuit board using common circuit construction techniques.

1.2.1.2 FX correlators

The FX correlator works by performing an independent Fourier transform on the signals from each antenna. The different spectra are then multiplied together in the correlator to calculate the cross-correlation values for each frequency channel. A

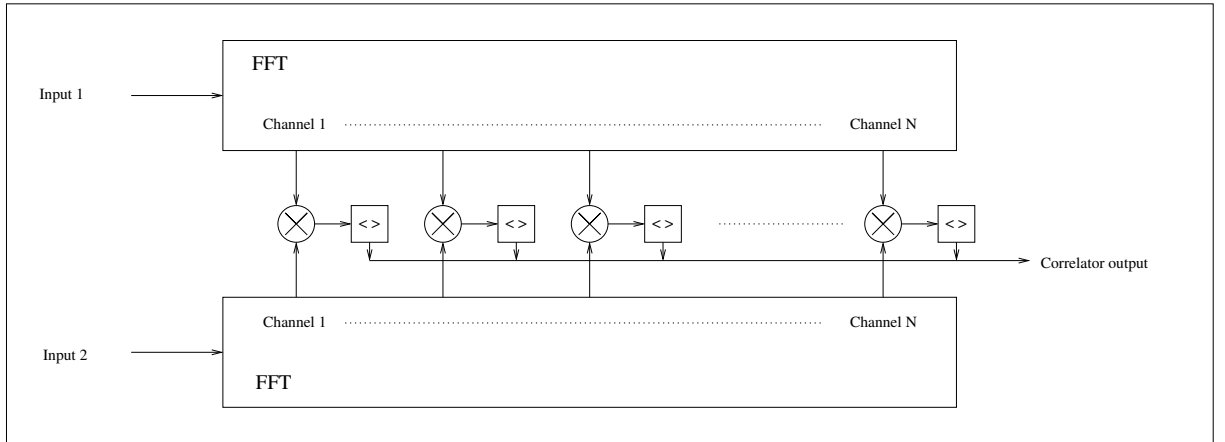


Figure 1.3: Schematic diagram of an FX correlator, with FFT, multiplier (\otimes) and accumulator ($\langle \rangle$) blocks. Figure from [41].

schematic representation of an FX correlator can be seen in Figure 1.3. As can be seen, two different blocks are required, a Fourier transform block and a cross-multiplier and accumulator block.

1.2.1.3 Comparing XF and FX correlators

The two different correlator architectures are equivalent. This can be seen from the convolution theorem, which states that a convolution in normal space is equivalent to a multiplication in Fourier space [40]:

$$FT(a \otimes b) = FT(a) \times FT(b)^*, \quad (1.9)$$

where $FT(a)$ is the Fourier transform of a , \otimes represents the convolution and $*$ is the complex conjugate.

The diagram of Figure 1.4 shows how the two methods follow opposite paths around the same processes to get from the raw data stream to the correlator outputs [40].

The major difference between the methods is the use of system resources. An XF correlator has a system resource requirement defined by:

$$n_s^2 n_t, \quad (1.10)$$

with an FX using:

$$n_s \log n_t + n_s^2 \quad (1.11)$$

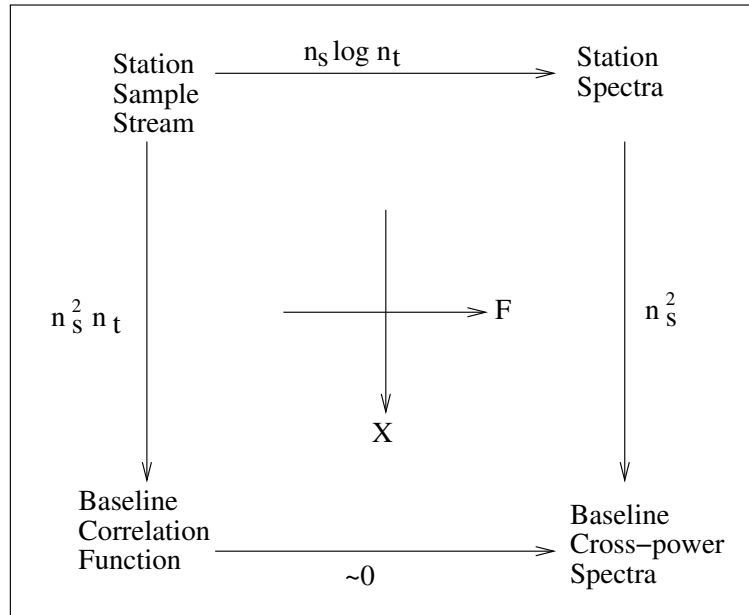


Figure 1.4: Data flow from antenna to output for FX and XF correlators. n_s stations and n_t samples transformed or lags formed. Figure from [40]. This implies that if good spatial and spectral resolution are required, then an FX approach will use less resources.

where n_s is the station number (or antenna number) and n_t is the number of samples transformed or lags measured. For this reason, as more antennae are added it quickly becomes more efficient to use an FX architecture. This is the case with most modern telescope designs where large numbers of antennae are used to produce a good filling factor and image reproduction.

1.2.2 Real and Complex correlators

In the case of the FX correlator, the initial Fourier transform on the real data stream produces a complex output, and subsequent multiplications are necessarily complex. In the XF case, there are several different ways of sampling the data using either real or complex sampling. The sampling needs to be sufficient to fully capture the data contained within the correlation function or the recovered spectrum will not be a true representation of the input signals.

1.2.2.1 Real correlators

In a real correlator the correlation function is sampled at the Nyquist rate, τ_N , for the system, given by:

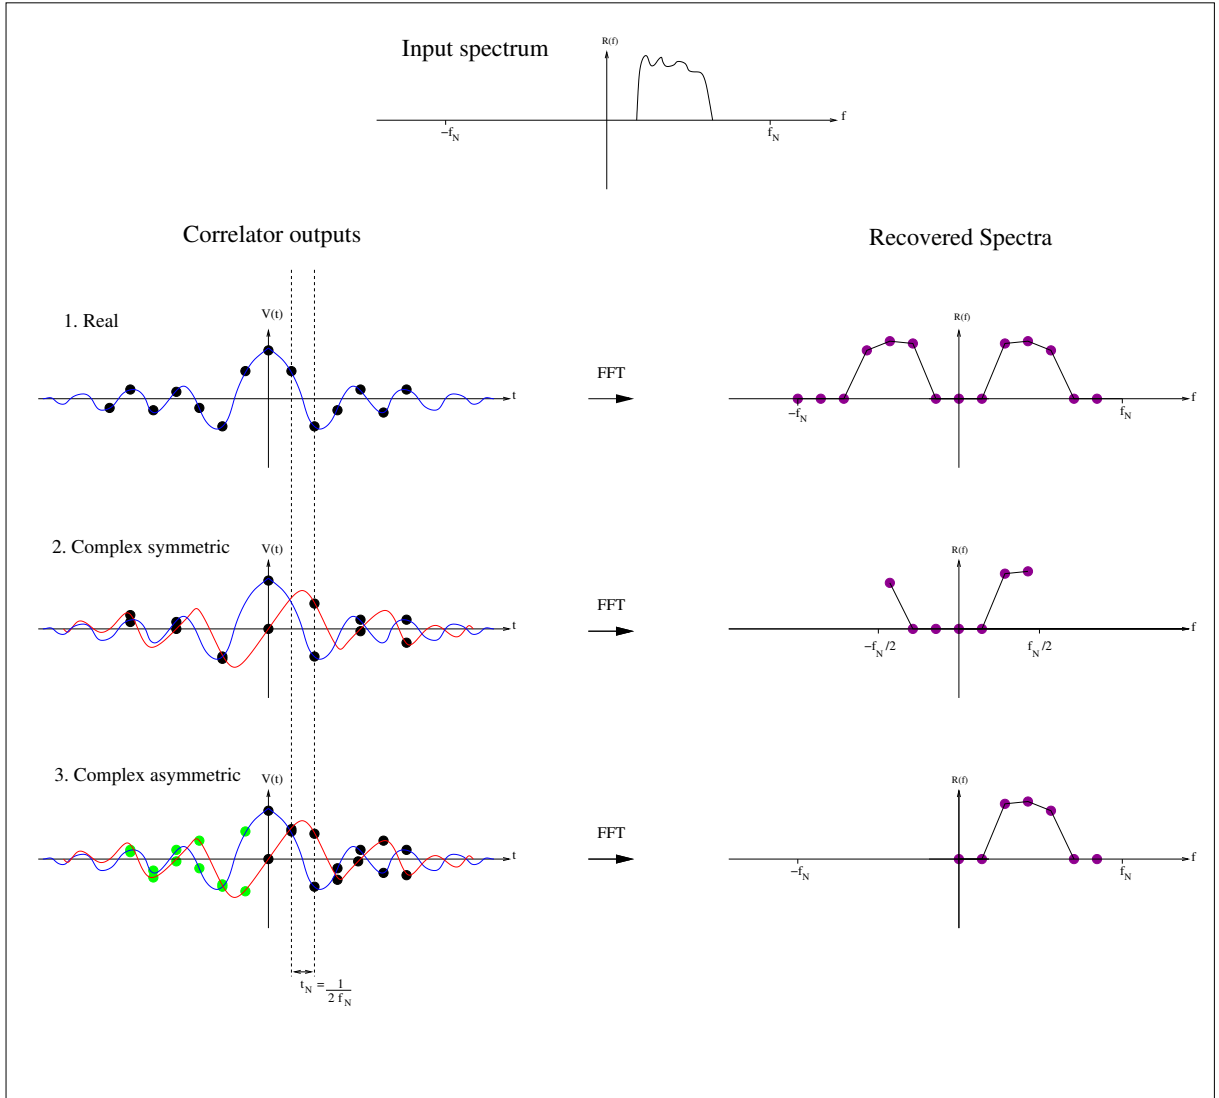


Figure 1.5: (Top) Input RF spectrum, (Below) Correlator outputs (left) and recovered spectra (right) for 1. a real correlator, 2. a complex correlator with symmetric lags and 3. a complex correlator with positive only lags. In each case the blue line shows the real component of the correlation function with the red line the imaginary. Black dots show correlation function sample points. Magenta dots show recovered spectra points. For the complex asymmetric case the negative side of the correlation function is extrapolated based on the positive half. These extrapolated points are shown in green and the recovered spectrum is found by Fourier transforming the combination of the measured points (black) and extrapolated points (green). τ_N is the lag separation for Nyquist sampling.

$$\tau_N = \frac{1}{2\Delta f} \quad (1.12)$$

where Δf is the IF bandwidth. The correlation function is then sampled at this separation, symmetrically about the zero lag.

To recover the spectrum from this output, the correlation function values at each lag are then fed into an Fourier transform. The output will be the same length as the number of lags, however as the input was real, only the positive half of the output spectrum will contain the desired spectrum, with the negative half being a redundant-copy. This means that if $2N$ lag spacings are measured the useful spectrum will contain N points. This can be seen in Figure 1.5.

1.2.2.2 Complex correlators

In a complex correlator both real and imaginary samples are taken at each lag position. The real sample is made as above, with the imaginary sample being produced by applying a 90-degree phase shift to one of the inputs.

As twice as much data is being sampled for every lag spacing, only half the number of lags is required to fully sample the correlation function. This can be done in one of two ways.

The most common approach is to double the separation between the lags but still sample symmetrically about the zero lag. If N complex lag samples are taken then the full spectrum can be recovered simply by again taking the Fourier transform. Unlike the real case, where the negative frequency domain of the output was a symmetric replica of the positive, for a complex correlator the negative side of the spectrum contains the second half of the spectrum. This can be seen in the middle case of Figure 1.5.

The second approach to producing a complex correlator is to maintain Nyquist sampling, but only sample the correlation function for the positive lag spacings. The second half of the correlation function can be filled in by taking the complex conjugate of the positive half. This requires the correlation function to be exactly aligned with the edge of the lag samples taken. The spectrum is recovered by taking the Fourier transform of the full correlation function, this is comprised of both the measured and extrapolated points. As with the real correlator, all the spectral information is contained within the positive half of the output spectrum. The negative half of the spectrum will be zero for all frequencies as it represents the out of bound portion of the input spectrum. This can be seen in the bottom case of Figure 1.5.

1.2.2.3 Finite sampling of the correlation function and bandwidth smearing (chromatic aberration)

It is important to consider the effects of the finite extent of the sampling of the correlation function. In the ideal case the correlation function extends to infinity. However in practice the lag space is only sampled to a finite extent. This means that the ideal correlation function is effectively being multiplied by a top hat window. A multiplication in lag space is equivalent to a convolution in frequency space. This means that the spectrum recovered by a correlator is the ideal spectrum convolved with the Fourier transform of a top hat function (a sinc function). The larger the lag space sampled within the correlator, the broader the top hat function is in lag space, and therefore the narrower the sinc function convolved with the spectrum. This convolution has the effect of smearing power from a sharp feature into the surrounding frequency channels and therefore decreases the frequency resolution of the correlator. It also means that, instead of the ideal rectangular passband for each bin of the recovered spectrum, the actual bin responses is as shown in Figure 1.6. As can be seen, each channel has a sinc function response with the primary peak slightly overlapping the neighbouring channels' central peaks. Also present are large side lobes, where power is scattered away from the central peak into other frequency bins. To reduce the effect of sidelobes, a windowing function can be applied to the output of the correlator. This is a mathematical function that is multiplied with the correlator output before the spectral recovery is performed. This means that the ideal correlator output is convolved with the Fourier transform of the windowing function instead of a Fourier transform of a top hat function. The windowing function can then be chosen to optimise the correlator response.

A common windowing function used is a Hann window, which has the shape shown in Figure 1.7 and the equation:

$$w(n) = \frac{1 - \cos\left(\frac{2\pi n}{N-1}\right)}{2}, \quad (1.13)$$

where $w(n)$ is the Hann window function in vector form, n is the position in the vector and N is the length of the vector. This vector can then be multiplied element-by-element with the correlator output before performing a Fourier transform.

The Hann window is a good general use windowing function, which provides moderate side lobe suppression and peak broadening [15]. The channel response, with the Hann window applied, can be seen in the lower plot of Figure 1.6. The side lobes have




Figure 1.6: Frequency response of the first four output channels out of sixteen in the recovered spectrum. The y-axis is not normalised. Single-frequency inputs from -20 GHz to +20 GHz have been simulated. (Top) No windowing used (Bottom) Hann windowing used

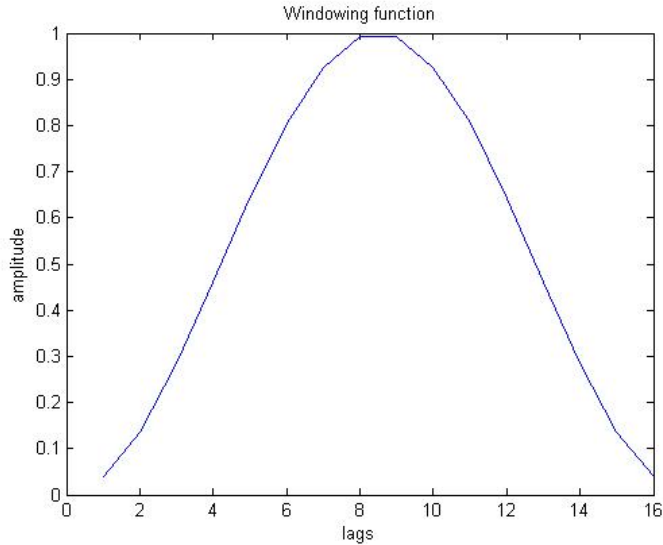


Figure 1.7: Hann window function shape.

been considerably reduced due to this method, but the width of the central peak has increased; This is a trade off made by windowing functions.

For broadband sources most of the power is contained in the central peak of the correlation function. This means that to properly sample the power in the signal the primary peak in the correlation function must fall within the lag space being sampled.

The situation is made more complicated by the beam size of the telescope. An object at the edge of the beam will have its correlation function peak at a different lag to an object in the centre of the beam. To ensure both objects are properly detected, the number of lags has to be selected such that both correlation peaks fall within the sampled lag spacing. This results in a limit on the ratio between the relative dish size to baseline with the fractional resolution bandwidth. If the case of two dishes is considered, both looking at the azimuth, then an object at the edge of the beam will have a net propagation delay between the two receivers. This can be seen in Figure 1.8, where the beam width is given by the angle α , the dish separation is given by D , d is the dish diameter and the propagation delay between the light from the source at the edge of the beam arriving at the two dishes is t_g .

From Figure 1.8, it can geometrically be seen that:

$$\lambda_g = D \sin \alpha, \quad (1.14)$$

where λ_g is the propagation delay. Substituting for the beam width:




Figure 1.8: Bandwidth smearing geometry. The beam width is given by the angle α , the dish separation is given by D , d is the dish diameter and the propagation delay between the light from the source at the edge of the beam arriving at the two dishes is t_g .

$$\lambda_g \approx D \times \frac{\lambda_{RF}}{d} = \frac{D}{d} \frac{1}{\nu_{RF}}, \quad (1.15)$$

and re-expressing in terms of t_g , the propagation delay,

$$ct_g = \frac{D}{d} \frac{1}{\nu_{RF}}, \quad (1.16)$$

where λ_{RF} and ν_{RF} are the wavelength and frequency of the light being received. To detect the power from the object at the edge of the beam the correlator must measure far enough into lag space to sample the peak of the correlation function at t_g . The extent of the correlation function sampled by a correlator, $t_{correlation}$, is given by:

$$t_{correlation} = \frac{1}{2\Delta\nu}, \quad (1.17)$$

where $\Delta\nu$ is the frequency resolution of the system.

$$\Rightarrow t_g \ll t_{correlation}, \quad (1.18)$$

therefore,

$$\frac{\Delta\nu}{\nu_{RF}} \frac{D}{d} \ll 1, \quad (1.19)$$

This effect is known as bandwidth smearing or chromatic aberration and must be considered when choosing the frequency resolution of any correlator constructed.

When the complex correlator with only positive lags is considered, the above constraint is further limited by the fact that the correlation peak has to be exactly aligned with the edge of the lag space sampled. This must be the case for the source anywhere in the beam and means that the chromatic aberration is further reduced to:

$$\frac{2\Delta\nu_{full}}{\nu_{RF}} \frac{D}{d} \ll 1, \quad (1.20)$$

where $\Delta\nu_{full}$ is the full bandwidth of the system instead of the frequency resolution.

1.2.3 Double-sideband complex correlators

The discussion so far has considered only positive frequency signals. However, due to the heterodyning used in most interferometers, negative frequencies will also be generated. In some systems these negative frequencies are filtered out using a single-sideband mixer, in which case the above discussion applies. The left-hand plot of Figure 1.9 shows a diagram of a theoretical sky RF signal, along with the frequency of the local oscillator (LO). As can be seen, the LO frequency is within the band of the sky signal. This means that the lower sideband IF output from the mixer is as shown in the right-hand plot. Both positive and negative frequencies are present and represent different parts of the RF spectrum. Filtering in the IF band (shown as a blue box) is used to select the correct frequency band to be fed to the correlator. Negative frequencies can be differentiated from positive frequencies as they have the opposite phase response to the addition of propagation delay, which enables the correlator to separate the signals. Twice as much information is now contained in the same physical frequency band, which results in two complex correlators being required to reconstruct the input spectrum with full signal-to-noise. These four correlators cover all possible signal combinations, i.e. with a 90 degree phase shift added to neither, one or both of the inputs, as shown in Figure 1.10. Although only one in-phase and one quadrature correlator is needed to generate the full frequency spectrum and separate the sidebands, the full set of four correlators is needed to cancel out cross-sideband noise correlations. This is shown below.



Figure 1.9: Sketch of (left) the input signal spectrum into the telescope with the position of the LO marked and (right) IF band. The blue square shows the area selected by a bandpass filter in the IF.




Figure 1.10: Diagram of correlator layout for a four-element, double-sideband, complex-correlator. The mixers and IF components have been omitted for clarity.

To analyse the workings of the system in Figure 1.10, the following convention is adopted. The signal from the antennae is subdivided into the upper (positive frequencies) and lower (negative frequencies) sidebands. F is then used to represent the Fourier transform of the component (i.e. F_{1u} is the Fourier transform of the upper component of the signal from antenna 1). The correlation of two signals is then given by:

$$R = Re[(F_{1u} + F_{1l})(F_{2u} + F_{2l})^*], \quad (1.21)$$

where the star is used to denote the complex conjugate of the signal. The four correlators are labelled by cc, cs, sc and ss where c represents an in-phase input, i.e. a cosine, and s represents a quadrature input, i.e. sine. The four correlators shown give the outputs:

$$R_{cc} = Re[F_{1u}F_{2u}^* + F_{1l}F_{2u}^* + F_{1u}F_{2l}^* + F_{1l}F_{2l}^*] \quad (1.22)$$

$$\begin{aligned} R_{cs} &= Re[F_{1u}(iF_{2u})^* + F_{1l}(iF_{2u})^* + F_{1u}(-iF_{2l})^* + F_{1l}(-iF_{2l})^*] \\ &= Re[-iF_{1u}F_{2u}^* - iF_{1l}F_{2u}^* + iF_{1u}F_{2l}^* + iF_{1l}F_{2l}^*] \end{aligned} \quad (1.23)$$

$$R_{sc} = Re[iF_{1u}F_{2u}^* - iF_{1l}F_{2u}^* + iF_{1u}F_{2l}^* - iF_{1l}F_{2l}^*] \quad (1.24)$$

$$R_{ss} = Re[F_{1u}F_{2u}^* - F_{1l}F_{2u}^* - F_{1u}F_{2l}^* + F_{1l}F_{2l}^*]. \quad (1.25)$$

These four outputs can then be combined to produce a single complex correlation function response:

$$\begin{aligned} R_{total} &= \frac{R_{cc} + R_{ss} + i(R_{cs} - R_{sc})}{2} \\ &= Re\left(\frac{2F_{1u}F_{2u}^* + 2F_{1l}F_{2l}^* + i(2F_{1u}F_{2u}^* - 2F_{1l}F_{2l}^*)}{2}\right) \\ &= F_{Tu}^* + F_{Tl}, \end{aligned} \quad (1.26)$$

where F_{Tu} is the real and imaginary combined parts of the upper sideband Fourier transform. When this is then inverse Fourier transformed, the upper and lower sidebands both appear separately and with full signal-to-noise. The addition of the two extra correlators acts to cancel out the cross-terms of the upper sideband correlated with the lower sideband and of the lower sideband correlated with the upper sideband. These terms contain no useful signal (since they represent the cross-correlation of different parts of the RF spectrum) but would add extra noise if not cancelled out by the linear combination of the raw correlations as shown.

1.3 Analogue and digital correlators

The above sections highlight the different architectures that can be used to produce correlators. How these are implemented in hardware depends primarily on whether an analogue or digital approach is adopted. The signals received by the antennae are analogue RF signals. In modern astrophysics all data are ultimately processed in the digital domain. It is therefore necessary to perform an analogue-to-digital conversion at some stage within the telescope. In the case of an analogue correlator, the correlation is performed on the analogue signals and the integration values are digitised before being read into a computer. The digital correlator takes the raw RF streams from the antennae and feeds them through analogue-to-digital converters. Signal processing and correlation is then performed in the digital domain.

Analogue correlators are generally used when broadband continuum observations are desired with low spectral resolution requirements. In all other applications digital correlators are usually the optimum solution.

1.3.1 Analogue correlators

Analogue correlators can be constructed in either the XF or FX configuration. To produce an XF correlator the signals are combined in multiplier chips at different relative propagation delays, as in Figure 1.2. These propagation delays are often constructed using physical path lengths. The multiplier outputs are integrated and then Fourier transformed to give the output power spectrum.

An FX construction can be formed using a filter-bank array. In this configuration the signals from the antennae are split into multiple frequency bands. This can be done in two ways. Either a bank of different filters is used to isolate each frequency band or a bank of mixers with different LOs is used to down-convert each frequency channel to IF. In the latter case, only one set of identical filters is then required to isolate the IF frequency channels. The bands are then cross-multiplied to obtain the correlated signal at each frequency. The output from this device is a correlated spectrum of the two input signals.

In an analogue XF correlator, the output of each of the multiplier devices has to be calibrated for both gain and lag position in order to be able to perform the inverse Fourier transform and accurately recover the spectrum. In the FX case consideration has to be given to the exact responses of all the filters used. Overlaps between the filter passbands will result in cross-talk between the channels of the output spectrum. Also

the fact that the different filter and mixer bands will have different loss characteristics means that an overall gain calibration is required to produce an accurate spectrum.

Since the bandwidth of the individual components of a filter-bank correlator is smaller, it is normally easier to implement this for larger bandwidth systems. This assumes that a calibration can be performed on the results. The bandwidth of a lag correlator is limited by that of the multipliers used. However, several very broad-bandwidth multipliers have now been created that offer consistent gain characteristics. This has led to the possibility of constructing ultra-broadband lag correlators. The specifics of these multiplier chips and the bandwidths obtainable are discussed in Section 1.4.

1.3.2 Digital correlators

Digital correlators can also be constructed using either the XF or FX architectures. The choice of architecture used is defined by the system resources available for the digital processing, as well as the bandwidth, frequency resolution and number of antennae required. As discussed in Section 1.2.1.3, in most cases the FX correlator becomes the most efficient design. Also, since the Fourier transform is done digitally, an FFT algorithm can be used, which allows a very efficient method of finding the spectral content of a signal.

As well as the correlation being carried out in the digital domain, other processing can also be performed, including the filtering of the correct band and the correction for path delays.

To perform the digital correlation described above, a large amount of digital signal processing is required. This can be performed in either hardware or software. In a hardware approach, FPGAs (Field Programmable Gate Array) or ASICs (Application-Specific Integrated Circuit) are used as these contain a high density of digital logic as well as other required digital features. An FPGA is a digital signal processing device where reconfigurable blocks of logic, RAM and input/output (i/o) are present on a single IC. These function blocks are linked by reconfigurable signal routing sections, which allow a large range of operations to be implemented by different interconnections of the function blocks. A hardware description language (HDL) is used to write the desired functionality. This is then compiled into a logic implementation and used to configure the device. In this way an FPGA can be programmed to perform many different tasks, with different sections of the FPGA being capable of performing entirely independent processes simultaneously.

Two main manufacturers exist for FPGAs, Xilinx and Altera. Both these companies produce a wide selection of devices capable of catering for different specific needs. FPGAs are utilised in radio telescope systems when reconfigurability is an important aspect, or where the scale of production is small.

An ASIC, like an FPGA, contains a large amount of digital logic as well as memory and interfacing. However, unlike an FPGA, the ASIC is fabricated with a single logic architecture and can not be reprogrammed. The advantages of ASICs over FPGAs are that they offer lower power and heat dissipation requirements for a given operation design. If the product is being mass produced, ASICs also offer a more cost effective solution. For this reason many large-scale telescopes, such as ALMA, use ASICs in their correlators after an initial design is developed and tested using FPGAs.

Software digital correlators also exist. These operate by capturing the data into a computer-based architecture and then performing all the digital signal processing in software. This can be done on a single processor device but is normally done on a parallel computing cluster or GPU based system. This is done as the task of correlation is massively parallel and therefore parallelization can be used to provide the necessary processing rates. The main limitations to performing digital correlations in software are the data input rates into the system and the associated data processing rates. An increased power consumption is also seen compared to FPGA or ASIC designs. The use of software correlators is currently undergoing development in the field to establish their full potential [5].

1.4 Current radio-telescopes

To gain a full understanding of the current developments in radio instrumentation it is important to consider the current generation of telescopes being constructed and in use. Discussed below is a selection of radio telescopes that use similar technologies to those discussed in this thesis. The instruments are listed in Table 1.1 along with their key specifications, and images of each are shown in Figure 1.11. This list encompasses telescopes with a wide range of operational and scientific goals, from general-purpose, common-user instruments to those with specific science targets. Of relevance to this thesis is the hardware used in each case, and specifically the limiting factors for the IF bandwidth, i.e. the analogue-to-digital converter for digital correlators and the multiplier chips for analogue correlators. Of the telescopes listed AMiBA, CBI and AMI have analogue correlators, with the others having digital systems. Table 1.2

shows the specifications for the multiplier chips used in the analogue correlators, with Table 1.3 detailing the analogue-to-digital converters used in the digital systems.

1.4.1 AMiBA

The AMiBA telescope [28] was designed to observe the CMB (Cosmic Microwave Background) at 86-102 GHz. It has three primary science goals[28]: surveying high redshift clusters using the Sunyaev-Zeldovich (SZ) effect, searching for missing baryons in large scale structures using the SZ effect, and measuring the polarization of the CMB. These science goals all require continuum observations over large bandwidth with low spectral resolution. The small baselines (≈ 2 m) relative to the dish size (1.2 m) also means that bandwidth smearing is not an issue. This led to an analogue correlator being chosen.

As can be seen in Table 1.2, the AMiBA correlator is based around a lag correlator constructed from four double-balanced diode mixers per baseline [26]. The nominal bandwidth for this system is 16 GHz, representing one of the largest bandwidth correlators currently constructed. With a four lag correlator, a frequency resolution of 8 GHz is achieved with just two frequency bins. The telescope is designed to observe continuum structure, and not sharp spectral features, and this bandwidth results in the achieved resolution being sufficient to eliminate bandwidth smearing ($\Delta\nu/\nu < 10\%$). When the AMiBA project was originally proposed in 2001 [28] the system was to have a bandwidth of 20 GHz (1-21 GHz), however it was found that this bandwidth was too large to be achievable using available RF components. The system was then redesigned to have the current bandwidth of 16 GHz. From observations made of Jupiter [27] it was then found that only an effective bandwidth of 10 GHz was being achieved. This shows a dramatic reduction in the desired capability of the system and highlights the difficulty in producing an ultra-broadband correlator.

1.4.2 CBI

The CBI telescope [30] was also designed to observe the CMB, at a frequency of 26-36 GHz. The key science goal of CBI was to observe the structure of the CMB on scales of 5 arcminutes to 1 degree. It was also designed to observe the SZ effect and foreground emission from our Galaxy. As with AMiBA, these are continuum observations with low spectral requirements driven by the need to reduce bandwidth smearing. An analogue correlator with a bandwidth of 1 GHz per channel was therefore used ($\Delta\nu/\nu \approx 3\%$).



Figure 1.11: (Top to Bottom, Left to Right) The EVLA and ATCA. AMiBA and CBI. AMI and CARMA. ALMA.

Telescope name	Number of dishes	Dish size	Max baseline	Observing frequency	IF bandwidth	frequency resolution
EVLA (Extended Very Large Array)	27	25 m	36 km	1-50 GHz	8 GHz	500 kHz-1 Hz
ATCA (Australia Telescope Compact Array)	6	22m	6 km	1.1-105 GHz	2 GHz	0.5 kHz-64 MHz
AMiBA (Array for Microwave Background Anisotropy)	13	1.2 m	2 m	86-102 GHz	10 GHz	8 GHz
CBI (Cosmic Background Imager)	13	0.9 m	5.51 m	26-36 GHz	10 GHz	1 GHz
AMI (The Arcminute Microkelvin Imager)	18	3.7 m and 12.8 m	110 m	13.5-18 GHz	6 GHz	0.75 GHz
CARMA (Combined Array for Research in Millimeter-wave Astronomy)	23	6 m and 10 m	2 km	27-270 GHz	4 GHz	Variable
ALMA (Atacama Large Millimeter Array)	66	7 m and 12 m	14 km	31-950 GHz	2 GHz sub-bands	Variable

Table 1.1: Specifications of a selection of the current generation radio telescopes. [33][43][28][42][7][3][4]

Telescope name	Correlator architecture	Multiplier name	Multiplier bandwidth	Achieved bandwidth of system
AMiBA	4 lag correlator	double-balanced diode mixers	16 GHz	10 GHz
CBI	filter bank correlator	Gilbert cell multipliers	1 GHz	10 GHz
AMI	16 lag correlator	Schottky diodes	6 GHz	6 GHz

Table 1.2: Multiplier chips used in the analogue correlators.[26][30][7][19]

Telescope name	ADC name	ADC sample rate	ADC bit depth
ATCA	high speed buffer amplifier/comparator	2.048 GS/s	2 bit 3 level
CARMA		20 GS/s	4 bit 10 level
EVLA	Atmel TS83102G0B	2 GS/s	8 bit
	Hittite and Teledyne	4 GS/s	3 bit
ALMA		4 GS/s	3 bit

Table 1.3: Analogue-to-digital converters used in the digital correlators.[39][8][37][38][29][6]

The CBI correlator is an analogue FX design, taking the 10 GHz band and filtering it into ten separate 1 GHz sub-bands. Each 1 GHz sub-band is then cross-correlated with its corresponding sub-band from another antenna using a Gilbert-cell multiplier. This is implemented in a complex correlator and produces the desired 10 GHz of bandwidth with 1 GHz resolution. The system requires regular calibration to account for the different filter responses as well as the multiplier gains. To achieve this, a hardware calibration system was built into the telescope which injected a correlated noise source with a flat passband into all the antenna signals prior to the sub-band filters.

1.4.3 AMI

AMI [18] is another instrument designed to study the CMB and the SZ effect. It operates between 13.9-18.2 GHz. The telescope was designed to observe both known SZ clusters as well as to carry out blind SZ surveys. The minimum frequency resolution for the AMI system to limit bandwidth smearing was 0.75GHz, and as very broadband continuum observations with little need for spectral resolution were required an analogue correlator was again selected.

The AMI correlator is a lag correlator with sixteen real correlations over an IF bandwidth of 6-12 GHz. This gives eight complex channels of 0.75 GHz. The multiplying elements are zero-bias Schottky diodes[19].

1.4.4 ATCA

Unlike the above telescopes, ATCA [1] was designed with very broad scientific goals. It is a general-purpose telescope where astronomers can apply for observing time. The system is designed to have a broad range of observing modes from emission lines to continuum observations. For this reason a digital correlator approach was selected, as high spectral resolution was required and also greater system reconfigurability to provide different observing modes at different times. High spectral resolution was required for both the science goals and to counter bandwidth smearing due to the large baseline-to-dish ratio .

The ATCA digital system [39] uses 2.048 GS/s 2-bit comparators to directly sample the 1 GHz bandwidth signal from the antennae [8]. The digital signal is then demultiplexed into a sixteen sample-wide parallel-data bus, operating at $1/16^{th}$ the sampling frequency. This signal is then fed into the digital signal processing block. The serial-to-parallel conversion is performed to account for the lower clock speed of the digital electronics compared to the sampler. The digital correlator is implemented using FPGAs.

1.4.5 CARMA

Like ATCA, CARMA [37] is a general purpose telescope. Examples of its key science goals include [3]: observations of the first generation of galaxies, determining the density structure of the dark matter in the centres of dwarf galaxies and measuring the evolution of large scale structure in the Universe using the SZ effect. These science goals and the effects of bandwidth smearing require a digital correlator to provide the necessary spectral resolution and reconfigurability for different observing modes.

CARMA has 10 GHz bandwidth and uses 20 GS/s 4-bit ADCs [37]. These are linked directly to the high speed interfaces on Altera Stratix IV FPGAs via a 2:1 demux serial-to-parallel converter. This results in a very high data-rate output into the digital signal processing block. Multiple stages of FPGAs are then used to down convert and distribute the signals into the correlator architecture.

1.4.6 EVLA

The EVLA [34] has many science goals, examples of which are [34]: the measurement of the strength and topology of cosmic magnetic fields, and enabling unbiased surveys and imaging of dust-shrouded objects that are obscured at other wavelengths. As well as these science goals it was required that observations of both continuum and spectral line emissions be made with different observing modes possible. For this reason a digital correlator was again selected.

The system used for the EVLA was developed in collaboration with the eMERLIN telescope[9]. With both telescopes sharing a common design for the correlator and samplers.

The design established has two sets of ADC converters. The first is an 8-bit pair of 2 GS/s samplers, providing 2×1 GHz bandwidth, with the second being 4×2 GHz bandwidth provided by four 3-bit 4 GS/s ADCs [38][29]. Two systems exist since different sensitivities and RFI immunity are required for different bands in the system. The 8-bit ADC system then uses a 1:4 demultiplexer chip to parallelize the data stream for input into the digital signal processing logic.

In the case of the EVLA, two different 3-bit samplers are used in the system. Most of the antennae are connected with Teledyne samplers, however a small number of the antennae have Hittite samplers. Originally the system was designed entirely around the Teledyne 3-bit samplers. When the system was constructed and tested it was found that these did not meet the required signal-to-noise performance[32]. To rectify the issue the Hittite samplers were investigated and found to have a much better response. Hittite samplers have not been used to entirely replace the Teledyne system, however, as considerable funds have already been spent on the Teledyne devices. The EVLA experience illustrates that at the level of 4 GS/s, high-speed ADCs are still a cutting-edge technology.

1.4.7 ALMA

ALMA [4] is another common-user instrument with a very broad range of science goals. Two of the main science goals for ALMA are the detection of spectral line emission from CO or CII in a normal galaxy like the Milky Way at a redshift of $z = 3$ and the imaging of the gas kinematics in protostars and in protoplanetary disks around young Sun-like stars at a distance of the nearest star-forming clouds [4]. These goals again call for a high spectral resolution and therefore a digital correlator approach.

The ALMA telescope uses eight 4 GS/s 3-bit ADCs per antenna [6][2]. These sample the 8 GHz IF band which is split into four 2 GHz sub-bands. A 1:16 demultiplexer is used to parallelize the data and reduce its data rate before being fed into the digital signal processing (DSP).

1.5 The need for broader bandwidth telescopes

This thesis focuses on developing new ultra-broadband devices, using both existing and new techniques. In radio instrumentation a drive has always existed towards extending the bandwidth capabilities of the devices being constructed. Broader bandwidth means more data can be acquired in a single observation. This means that, for example, multiple features in a spectrum can be observed in a single sampling instead of multiple captures. Also, if continuum information is required, broader bandwidth observation gives greater sensitivity in a fixed observing time.

An example of the need for broadband continuum observations is the development of the GUBBINS [12] (220 GHz Ultra-BroadBand INterferometer for SZ) telescope for observing the SZ effect. GUBBINS is a single-baseline prototype ultra-broadband interferometer that is being developed by the University of Oxford. As shown in Figure 1.13, the system is a small telescope with a short 0.5 m baseline. This results in a beam of about 7' FWHM at 220 GHz, with entire clusters appearing in a single pixel. The system is designed to observe the thermal SZ modification of the CMB spectrum in the direction of galaxy clusters. The thermal SZ effect occurs when photons of the CMB pass through clusters of galaxies, causing an increase in the energy of the CMB photons. This is due to a small fraction of the photons interacting with the hot gas in the galaxy cluster and being inverse-Compton scattered to higher energies. This is seen as a shift in the black body spectrum as in Figure 1.12.

If the change in CMB spectrum is plotted it can be seen that at low frequencies the change in intensity is negative and at high frequencies the change is positive. This means that there is a zero crossing in the relative spectral brightness at a frequency around 220 GHz. The exact frequency of crossing shows a dependence on the temperature of the gas that caused the spectral change, given by:

$$f = (217 + 0.45T)\text{GHz}, \quad (1.27)$$

where T is the temperature of the cluster in keV [12]. Values of T for typical clusters range from 5 - 15 keV, leading to frequency changes of 2.5 - 7.5 GHz. This means that if the CMB spectrum is observed along a line of sight of a cluster then it is




Figure 1.12: (left) Sketch of spectrum of CMB before and after SZ modification, (right) Sketch of difference between unmodified and modified CMB spectrum. Arbitrary intensity units used.

possible to determine the cluster temperature from the spectral modifications and zero crossing frequency. To accurately determine the crossing point of this spectrum it is therefore necessary to take measurements across a wide enough range of frequencies to allow the crossing frequency to be determined whatever the temperature of the cluster. This leads to a system requirement for a bandwidth of at least 10 GHz.

The GUBBINS telescope was designed to operate at this 220 GHz crossing with the minimum required bandwidth of 10 GHz in each sideband.

The bandwidth of the GUBBINS system was chosen as it represented the limitations of the technology being used. If a greater bandwidth were achievable, this would allow a stronger constraint to be placed on the interpolation of the zero crossing point and therefore the cluster temperature. This shows the continued need for larger bandwidth telescopes.

1.6 Overview of thesis

From the above descriptions of current radio telescopes it can be seen that a selection of different design ideas have been implemented, with no generally accepted 'best in all cases' solution. It can also be seen that problems have been encountered in attempting to push the limits of available bandwidth, most notably the AMiBA telescope only achieving half the originally designed bandwidth and the EVLA having to go through different ADC manufacturers to achieved the desired performance.

Frequency	185-275 GHz
Antenna aperture	0.4 m
Baseline	0.5-0.6 m
Primary beam	11' @ 220 GHz
Spatial resolution	7'-11' @ 220 GHz
IF band	2-20 GHz
Instantaneous bandwidth	2x 18 GHz
Spectral resolution	1.25 GHz
Target system temperature	50 K
Brightness sensitivity	350 $\mu\text{K/s}^{1/2}$




Figure 1.13: A schematic of the GUBBINS telescope with primary details listed[12].

This thesis will present two approaches to the problem of ultra-broad bandwidth correlators. The first of these is a 2-20 GHz analogue correlator developed prior to the work of this thesis [20]. This system is tested and a calibration method simulated. The second, and major part of this thesis, investigates a new approach to digital correlator development with an integrated sampler and signal processing unit utilising commercially developed high-speed serial interfaces as direct RF samplers.

The last chapter of the thesis discusses the associated ultra-broadband RF devices that are required for ultra-broadband instruments. The structure of this thesis is as follows:

Chapter 2: Analogue correlation techniques

Chapter 2 discusses the characterisation and calibration of a 2-20 GHz analogue correlator constructed for the GUBBINS project. Single-frequency tests are carried out, along with the simulation and implementation of a calibration technique.

Chapter 3: Digital Sampling techniques: 1-bit sampling

Two different digital sampling techniques are discussed, focusing on extending the bandwidth of current digital systems towards that obtainable by the analogue devices discussed. For both direct sampling and external comparator methods, single-frequency measurements are made. The signal-to-noise ratios of the systems are

analysed using an effective number of bits measurement.

Chapter 4: Digital Sampling techniques: Interleaving techniques

The direct sampling method is then extended using interleaving techniques to provide faster sampling. A large amount of investigation is done into the methods for aligning and anti-aligning samplers. Measurements are made using a simple two interleaved sampler setup.

Chapter 5: Digital Sampling techniques: Multiple-bit techniques

The alignment method produced in the interleaving chapter is then used to align samplers for use as a multi-bit sampler. Two techniques are identified and investigated, with one being used to form a 1.5-bit system. Interleaving and multi-bit techniques are then combined to form an effective 1.5-bit 6.25 GS/s sampler.

Chapter 6: Hardware Digital correlator

The 1-bit 3.125 GS/s direct sampling method is used to construct a six-baseline lag correlator within the FPGA architecture. Characterisation tests are carried out on the system with both single-frequency and broadband measurements being made.

Chapter 7: RF and IF systems for broadband interferometry

The ultra-broadband RF devices used throughout the experimentation are then discussed in full with characterisation shown. The development of several new RF devices required for a full IF system is also documented. Lastly, simulated IF chains are designed for the two correlators discussed.

Chapter 8: Conclusions

Finally, the results of the thesis are summarised, with areas for future research identified.

Chapter 2

Analogue correlation techniques

In Section 1.5 the GUBBINS telescope was discussed. This is an ultra-broadband single-baseline telescope designed to make continuum observations of the SZ effect. The telescope was designed by the University of Oxford experimental cosmology group and is being constructed in-house.

The correlator boards for the system were designed, fabricated and initially tested prior to this thesis. The basic parameters of the correlator were tested but correlations were never performed to establish signal recovery ability.

In this chapter further investigation is carried out into the noise performance of the correlator cards. Signal recovery is then simulated and measured, with finally a calibration method described, simulated and implemented.

2.1 System overview

The GUBBINS correlator is an ultra-broadband, analogue, double-sideband, complex lag correlator[20]. It is constructed using four correlator cards, shown in Figure 2.1. Each correlator card implements sixteen multipliers with lag separated by 7.5 mm of free space distance, which corresponds to Nyquist sampling at 20 GHz:

$$\Delta l = \frac{c}{2f} = \frac{c}{2 \times 20 \text{ GHz}} = 7.5 \text{ mm.} \quad (2.1)$$

The sixteen multiplier chips are placed along the diagonal line between the two inputs, in order to provide the correct propagation delay to each multiplier. Two 16-way splitter trees are used to distribute the incoming RF signals into the different delay lines. The splitter trees use Wilkinson dividers, each with seven quarter-wave matching elements, which provide the broadband matching required.




Figure 2.1: (Clockwise from top left) A photograph of one of the four GUBBINS correlator cards with coin for scale, a photograph of the inside of one of the correlator cards, a close up of one of the multiplier chips used and a detailed photograph showing the Wilkinson divider. Photos taken from [20].

The outputs from the 64 multipliers (4×16 lag correlator cards) are read into a computer via a set of filter cards paired with ADC sampling boards. These communicate to a host PC using USB interfaces. The ADC cards also act as digital accumulators, lowering the output data rate being fed into the computer. A schematic of the readout system for a single correlator card can be seen in Figure 2.2.

The filter cards act to amplify the signals coming from the correlators and filter them ready for digitisation. This filtering includes both anti-aliasing for the digitisation and blocking of the DC biasing used on the multiplier chips. The filter card also acts to convert the output of the correlator into a differential input signal for the ADCs to sample.

The DC block means that any constant output from the correlator is not sampled. As sky signals only vary slowly over time they would appear as a near-DC signal on the correlator output and therefore be filtered out. To overcome this, one of the signals fed into the correlator undergoes 180-degree phase switching. When one of the signals is phase shifted it acts to flip the sign of the correlation. The phase switching modulates the correlated signal so that it passes through the DC block. In the readout FPGA the modulation is removed before the signals are accumulated. Any slowly-varying offsets in the data are therefore cancelled out in the accumulated

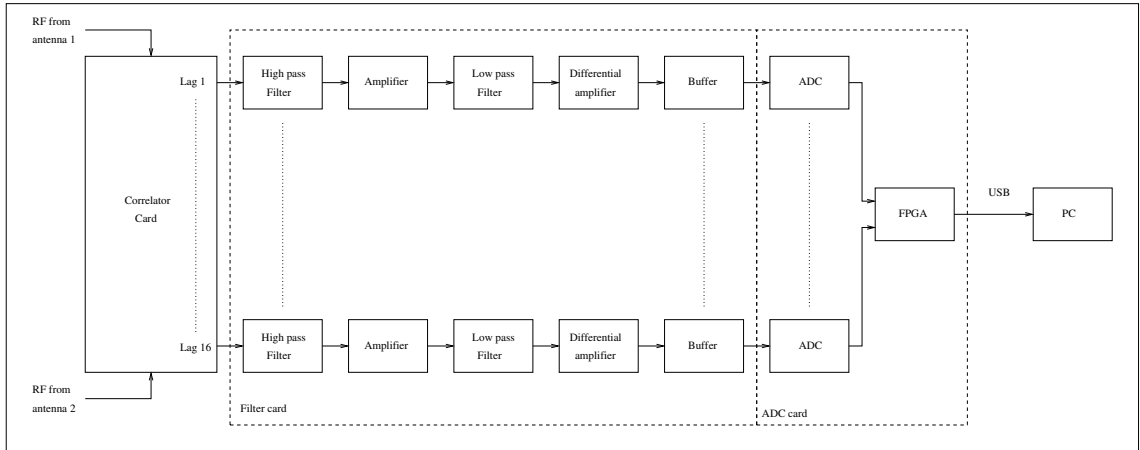


Figure 2.2: Schematic of the GUBBINS correlator readout setup.

data.

The ADCs used are 2.8 MS/s differentially-fed 14-bit converters. Sixteen ADCs are placed on a single ADC card, so that all of the multiplier outputs can be read into a single FPGA. The FPGA then performs the demodulation of the phase switching, as well as accumulation and data packeting ready for transfer over the USB interface to the PC.

2.1.1 Prior tests

Prior to this thesis, measurements were made of the following properties of the correlator cards [20]:

- Return loss
- Lag spacing
- Frequency response
- Signal power response
- Video bandwidth
- Noise behaviour

The noise measurement made showed a level of $\approx 10^{-6.5}\text{V}/\text{Hz}^{1/2}$ over a frequency range of 10 kHz to 100 kHz [21]. This measurement was made using a spectrum analyser, with the lower frequency limit being set by the channel width of the spectrum analyser. However, it is important to measure the noise power spectrum at

much lower frequencies than this since any astronomical signal will be present in the data at relatively low frequency. A key experiment performed in this thesis was the extension of this noise measurement using the direct readout system to improve the lower frequency threshold. This is described in Section 2.2.1.

The lag spacing for the correlator cards is determined by the geometric layout of the tracks on the correlator. When the lag spacings were measured, position errors of around 10% in the 7.5 mm expected lag spacing were seen [20]. In this thesis the impact of these errors is investigated, with particular reference to the correlator calibration scheme.

2.2 Analogue correlator bench tests

The following section describes the new tests performed on the correlator.

2.2.1 Noise behaviour

The first consideration when testing the correlator was a re-measurement of the level of noise in the output. This was done to investigate at what frequency an optimum phase-switching system should operate, as the correlator output noise level should be dominated by thermal noise from the RF system rather than noise generated in the correlator. A measurement of the noise behaviour of the correlator was taken prior to this thesis, by connecting a spectrum analyser to the outputs of the correlator card and measuring the power received. The lower frequency limit of this noise spectrum was set by the 10 kHz channel width of the spectrum analyser. The connection method for the spectrum analyser was not ideal as clip-on probes were used. To improve on this measurement the experiment was repeated using the full readout system described in Section 2.1. Direct data were captured by the FPGA with no accumulation performed. The raw ADC samples were then read into a computer and used to plot a spectrum. To convert the time series outputs of the correlator, $V(t)$, into a power spectrum, $Y(\nu)$, the following formula was used:

$$Y(\nu) = \frac{2 \times |\text{FFT}[V(t)]|^2}{F_s \times L}. \quad (2.2)$$

To convert this to a voltage spectrum the square root is taken. $V(t)$ is the voltage measured on the output of the correlator. This is found by taking the ADC count and scaling into volts, accounting for the $\times 5$ gain present in the filter board. F_s is




Figure 2.3: Noise spectrum for the correlator. Two cases plotted, that of open inputs (system noise) and with 0.7 dBm of uncorrelated noise input (signal noise). Figure from [20]

the sampling rate of the ADCs (2.8 MS/s), L is the number of samples taken and FFT is the fast Fourier transform function.

The result can be seen in Figure 2.3. The noise of the correlator remains well below the thermal noise down to a frequency of ≈ 100 Hz, where the correlator noise power begins to rise as $1/f$. This means that phase switching should be implemented at a frequency above this value. This measurement shows a slight increase in the noise level compared to that of the original measurement. This is caused by the more rigorous nature of the experiment performed.

Another important consideration in the testing of the correlator was to determine the amount of the correlator output noise generated by the ADC sampling and the filter boards. To investigate this, the correlator was removed and the noise was sampled for the two cases of no input signal into the ADC card alone, and the filter card feeding the ADC card with no correlator board attached. It was found that the level of noise was $\approx 1\%$ and $\approx 2\%$ of the correlator noise, for the ADC alone and ADC with filter card respectively. This means that almost all of the noise seen on the correlator output originates within the multiplier chips and is not added by the readout system.

2.2.2 Single-frequency recovery

After the correlator noise was analysed and the required phase-switching frequency identified the next step was to test the correlator response to RF signals.

Initial tests on the correlator were performed using a bench signal generator producing single-frequency outputs, which were fed into one of the four correlator cards. The setup can be seen in Figure 2.4, with the simulated output for a 3 GHz input signal shown in Figure 2.5. The response obtained, shown in Figure 2.6, is as predicted by the simulation, with the output being a cosine function of lag delay with the correct period (≈ 13 lags ≈ 330 ps). The simulation output used a different normalisation to the correlator so the amplitudes could not be compared. As seen, the observed output of the correlator is not as smooth as the simulation, due to imperfections in the system that require calibration. This will be discussed in Section 2.3.

To further test the response of the correlator, the relative propagation delay for the RF signal to the two inputs of the correlator was varied using the path compensation. To do this the top setup of Figure 2.4 was modified to that of the bottom. This was done by the addition of seven path compensator devices which, when combined, can generate any path delay from 2.5 mm to 317.5 mm in the arm of the IF. These devices will be discussed in Section 7.1.3. When a delay equivalent to one lag spacing is added to the path on one of the arms then the correlator output should move by one spacing. This can then be repeated to move the signal across the lag outputs. The plot of this data can be seen in the right-hand graph of Figure 2.6. This shows the correlation pattern moving as expected, but with some distortion of the pattern due to errors in the lag positions and gain responses.

2.3 Analogue correlator calibration - simulation and implementation

2.3.1 Imperfections

The previous section outlines the basic measurements made with the analogue correlator. As was shown, the results obtained are not exactly as predicted by the simulations. These discrepancies are caused by imperfections in the correlator. The two main sources of imperfection are the relative positions of the lags and the gain behaviour of the different multiplier chips.

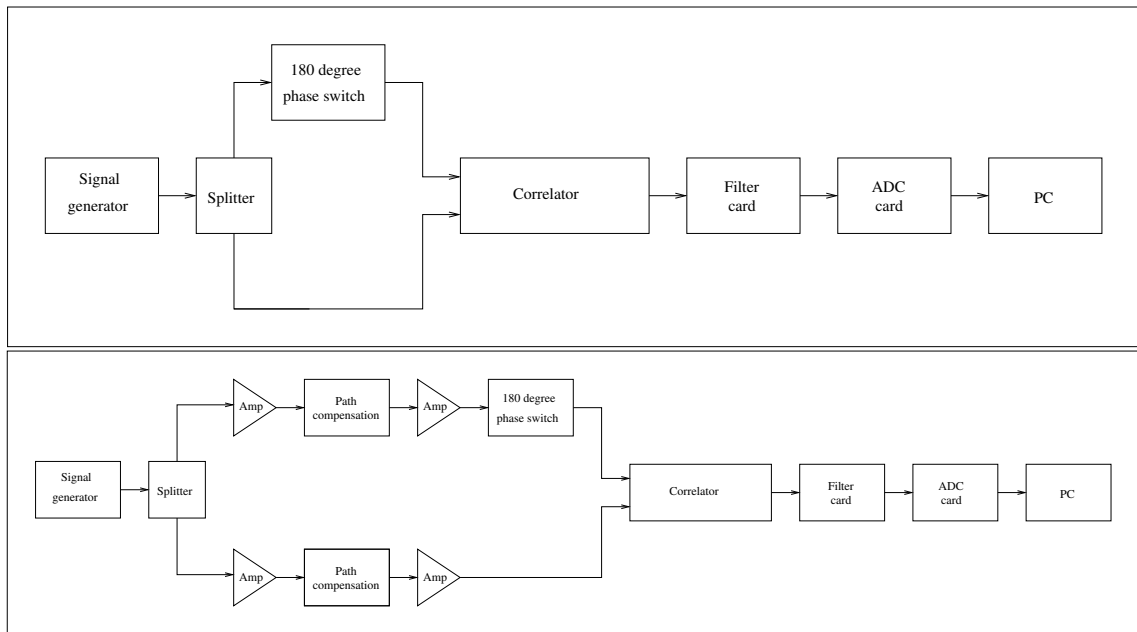


Figure 2.4: RF setup for single frequency correlator tests. (Top) Simple IF configuration with no path compensation present. (Bottom) More complicated IF setup with path compensation and the required amplification included.

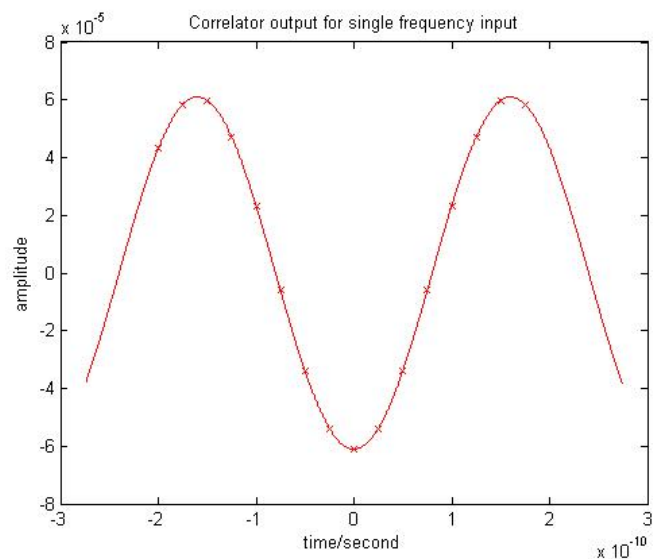


Figure 2.5: Simulation of correlator response for a 3 GHz input signal. The crosses on the line represent the lag values measured.

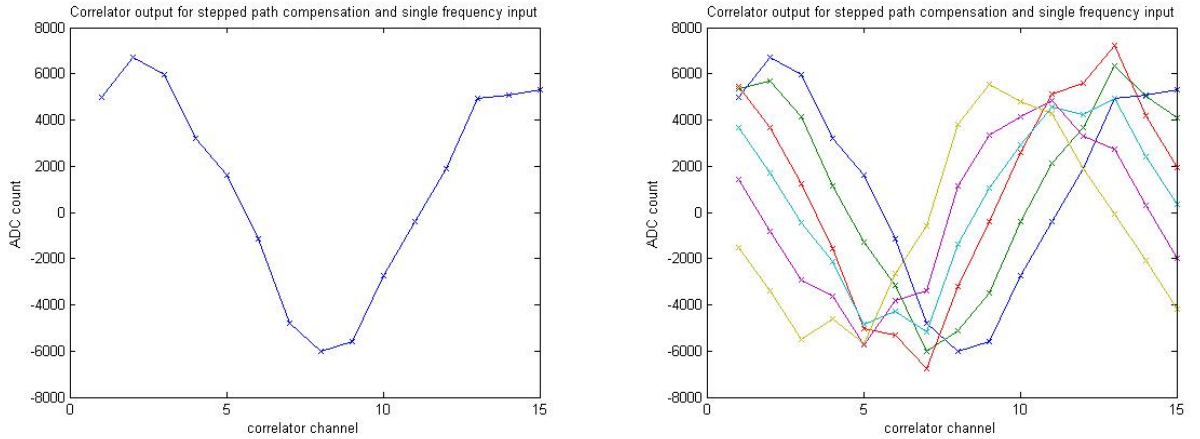


Figure 2.6: (left) Actual correlator response for a 3 GHz input signal. (right) Actual correlator response for a 3 GHz input signal stepped across lag channels using path compensation. Each line represents a 7.5 mm path difference from its neighbours. In all plots, the crosses correspond to the same individual lag positions.

2.3.1.1 Lag discrepancies

Each lag output samples the correlation function at a different relative propagation delay. These lag values should form a uniform sampling of the correlation function, which enables a standard inverse Fourier transform to be used to recover the original spectrum. If the lag positions are not distributed evenly this will cause corruption of the recovered spectrum. In the GUBBINS correlator, the lag positions are determined by the relative track lengths of the circuit. They are therefore affected by the accuracy of the construction of the board, by variations in the dielectric constant of the board material and also by the exact bonding of the multiplier chips. Prior to this thesis, the actual lag positions for the correlator cards were measured. The result for one of the correlator cards can be seen in Figure 2.7 [20]. For the GUBBINS system the optimum spacing of the lags is 7.5 mm, but the plot shows that the actual spacings differ from this by up to 10%. The primary cause for this discrepancy will be variations in the dielectric constant of the board material, this will cause stray capacitance due to the impedance matching, affecting the propagation time of the signals.

2.3.1.2 Multiplier gain response

The multiplier chips used to calculate the cross-correlation and the accumulation are custom made Gilbert-cell multiplier chips [10][13][14]. As each lag correlation is calculated on a separate chip and the different devices have different gain responses, errors in the amplitude of the individual lag measurements need to be considered.




Figure 2.7: Measured lag spacing in actual correlator. Figure taken from [20]. The nominal value is 7.5 mm.

2.3.1.3 Dispersion

As well as the two main imperfections listed above, dispersion in the microstrip transmission lines in the correlator will result in the distance between lags varying across the band. This is particularly important given the large bandwidth of the correlator. This means that any calibration adopted will need to have a frequency dependent correction to the lag position.

2.3.1.4 Band slope

The response of the correlator also shows a frequency dependence [20], due to the gain of the multiplier chips reducing at higher frequencies, and also the loss in the microstrip being greater at higher frequencies. This is a common problem in broadband RF devices and will need to be factored into any calibration. The band slope may also vary between different multiplier chips. This effect will lower the sensitivity of the system to higher frequency signals, and is calibrated using a frequency dependent gain correction.

2.3.2 Calibration methods

To account for the above imperfections in the system, a calibration routine is required that is a modification of a standard inverse Fourier transform. The calibration method builds upon one identified for the AMI correlator [18] and relies on measuring the

output when a set of known inputs are applied to the correlator. Two options exist for the input signals used in the calibration, both providing an equivalent calibration of the system. The first option is to use a set of single-frequency RF signals. To completely constrain the system the number of single frequencies used needs to match the number of lags in the correlator, with the frequencies chosen to be the centre frequencies of the bins of the recovered spectrum. The alternative to using single frequencies is to apply a broadband noise source. The relative propagation delay for the noise source into the two channels of the correlator can then be controlled to step the correlation function, such that the peak of the sinc response is centred on each of the lags in turn. This results in the same number of measurements being made as for the single-frequency case. The correlator characterisation measurements described below use the single-frequency method which is most useful for laboratory testing, while the broadband method would most likely be used during astronomical observations.

In the single-frequency case, the output of the correlator for the i^{th} given input frequency is defined as a vector \mathbf{v}_i . If the expected ideal recovered spectrum is \mathbf{r}_i , then the calibration method assumes that there is a transformation matrix \mathbf{F}_m to convert between the two such that

$$\mathbf{v}_i \mathbf{F}_m = \mathbf{r}_i. \quad (2.3)$$

When multiple frequency inputs are used, they can be arranged in a matrix such that

$$\mathbf{V} \mathbf{F}_m = \mathbf{R}, \quad (2.4)$$

where \mathbf{V} and \mathbf{R} are now matrices with the rows corresponding to the vectors \mathbf{v}_i and \mathbf{r}_i respectively. The frequencies labelled by i are the centre frequencies of the outputs of the Fourier transform over the lag spacing, i.e. $f_i = i\Delta f$ where $\Delta f = c/(N_{\text{lags}}\Delta l)$ and Δl is the lag spacing. In the ideal case of no lag or gain errors, the matrix \mathbf{F}_m is the inverse Fourier transform matrix, \mathbf{F} . However, for the non-ideal case, a correction matrix, \mathbf{C} , must also be applied, such that,

$$\mathbf{F}_m = \mathbf{C} \mathbf{F} \quad (2.5)$$

From Equation 2.4 it then follows that

$$\mathbf{C} = \mathbf{V}^{-1} \mathbf{R} (\mathbf{F})^{-1}. \quad (2.6)$$

In the ideal case the calibration should produce the identity matrix for \mathbf{C} .

2.3.3 Simulation of calibration scheme

The correlator system was simulated in MATLAB to understand both the impact of the imperfections identified and the success of the calibration scheme in correcting for them. The simulation also allowed the imperfections to be studied in isolation, so as to ascertain their individual contributions to the output corruption.

The simulation takes an input spectrum and produces the corresponding correlator outputs. This can then be processed in an attempt to recover the original spectrum, simulating all aspects of the correlator and data recovery. To test the program, the parameters of the simulation were set to match an 8-lag double-sideband complex correlator, with positive-only lags. The lag spacing was set to Nyquist sample the correlation function. This gives a frequency resolution of 2.5 GHz, matching the resolution used in the testing of the single GUBBINS correlator cards in Section 2.3.5. As discussed in Section 1.2.2.2, to recover the spectrum from a complex correlator with positive lags only, the negative half of the correlation function is calculated by taking the complex conjugate of the positive half. In the simulation this second half of the spectrum is implicitly calculated, as FFTs are being used throughout that naturally produce the required data. This results in sixteen lags appearing in all the plots shown for the simulation in this chapter.

When the simulation was run with a single-frequency input the results shown in Figure 2.8 were obtained. The input spectrum, the correlator output and the recovered spectrum with no calibration can all be seen. These show how the correlator produces the expected discrete sampled Fourier transformed spectrum output as identified in the introduction section. These outputs represent the ideal case for the correlator and as a result, a standard inverse Fourier transform has been used to recover the spectrum from the correlator outputs.

The first imperfection that was investigated to characterise the calibration scheme was the introduction of lag errors. This was introduced into the simulation by allowing the positions of each of the lags to be set manually. The actual correlators have a small scatter on all the lags. However, to allow the results to be analysed, the imperfections were first simplified so that only one lag was moved away from its ideal position. To investigate the effect of this imperfection, the correlator output for the sixteen different single-frequency calibration inputs was simulated. From this the \mathbf{C} matrix, shown in Figure 2.9, was derived using equation 2.6. In the ideal case of no lag errors the \mathbf{C} matrix is the identity matrix. Figure 2.9 shows that only the

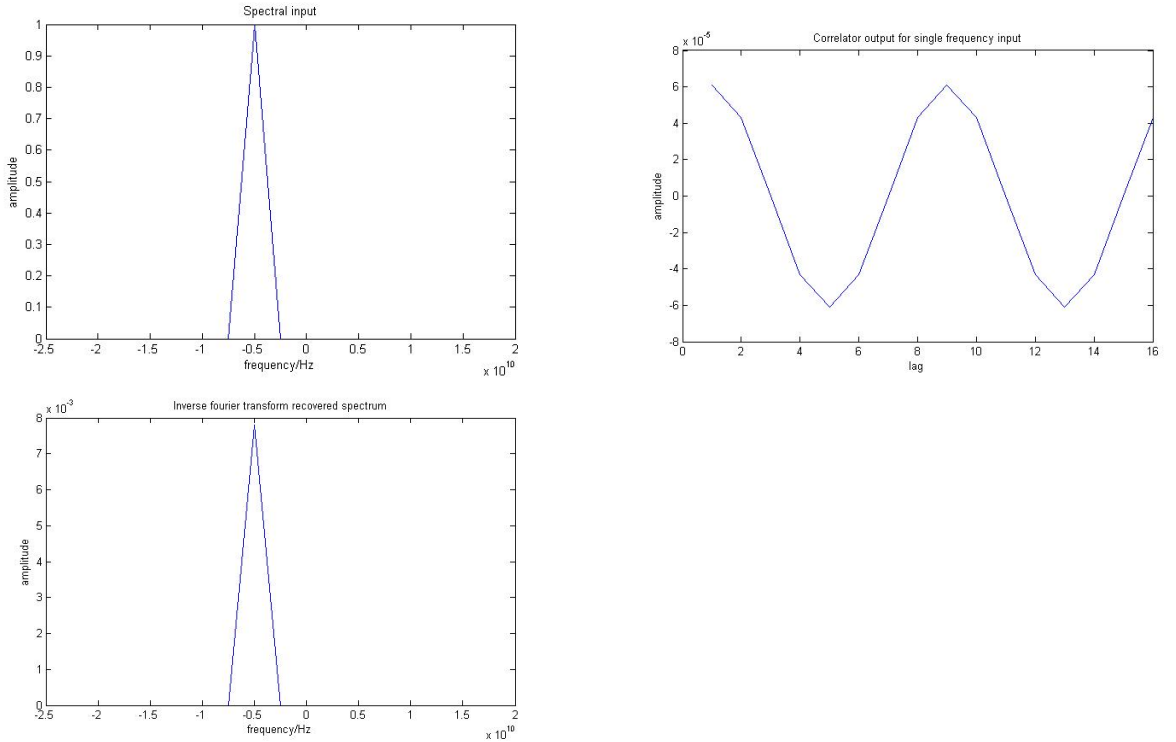


Figure 2.8: (top left) Input spectrum into ideal correlator (top right) output of ideal correlator (bottom left) recovered spectrum from ideal correlator output using inverse Fourier transform.

row corresponding to the altered lag displays different behaviour. The result shows how the calibration method has effectively altered the measured value of the moved lag by mixing in components of all the other lags. If a broadband signal filling the band is considered, then the response of the correlator is the Fourier transform of the passband, namely a sinc function. The ideal lag positions will lie on the central peak and all the zero crossings. When a lag error is introduced, the zero crossing will not exactly align with the lag position and a non-zero value will be measured for a broadband input at any relative propagation delay. The calibration method interpolates the values from all the other lags in order to correct the misplaced lag.

2.3.4 Reduced bandwidth calibration

When the system is observing real sky signals, a possible problem can arise due to the presence of radio frequency interference (RFI) from terrestrial sources. This effect is especially pronounced at low frequencies where some of the initial tests will be carried out. To compensate for RFI, it may be necessary to add extra filters to block the parts of the spectrum affected by the interference. The simulation can be used to

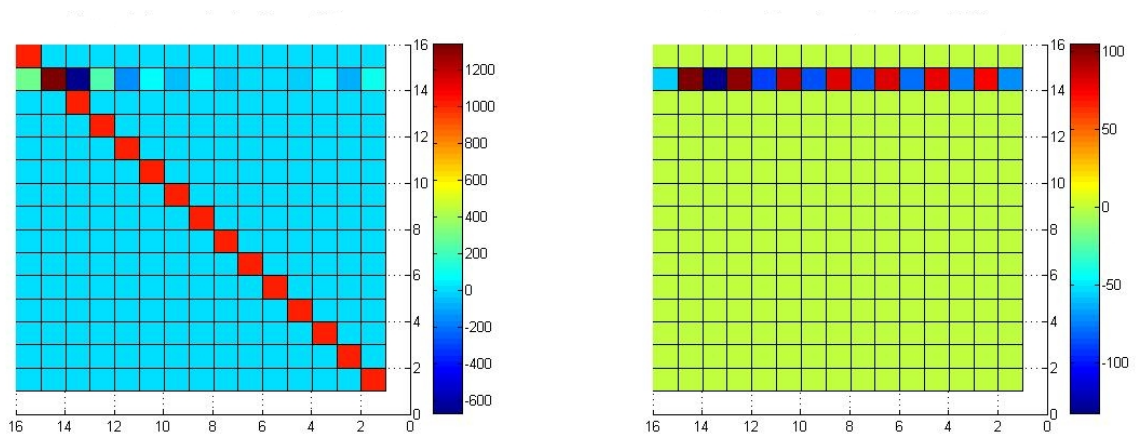


Figure 2.9: Real (left) and imaginary (right) parts of correction matrix generated from calibration with a single imperfect lag spacing. The matrix differs from that of the identity matrix (the situation for no imperfections) by altering the row associated with the lag error. As this lag is out of place, it is sampling the wrong point in the correlation function and therefore must have the other lag values mixed in to correct for the error.

predict the result of this reduced bandwidth on the system output.

The test performed included running the simulation over a reduced bandwidth and checking that signal recovery was not corrupted. The signals used in the single-frequency calibration correspond to the centre frequencies of the bins in the recovered spectrum. For the system described these are at 2.5 GHz increments. The first reduced bandwidth considered was less than 2.5 GHz wide, and spanned from 4 GHz to 6 GHz. This meant that only one calibration frequency existed in the range (5 GHz). After the system was calibrated, single-frequency signals with a frequency within the reduced bandwidth, from 4 to 6 GHz, were input and recovered using the calibration. It was found that the recovered spectra consisted of a single peak at 5 GHz, with varying amplitude according to the input frequency. This was expected as only one of the frequency bins had been calibrated. No leakage was seen into other bins.

To gain an understanding of reduced bandwidth with some spectral resolution, the bandwidth was expanded to encompass two calibration frequencies. A frequency range of 5 GHz to 7.5 GHz was selected, which includes both the 5 GHz and 7.5 GHz calibration frequencies. After calibration, single-frequency signals in this frequency range were again input and the calibration used to recover their spectra. The results from this can be seen in the five plots of Figure 2.10. The recovered spectra are shown in red, with a superimposed sinc response in blue. As discussed in Section 1.2.2.3, the

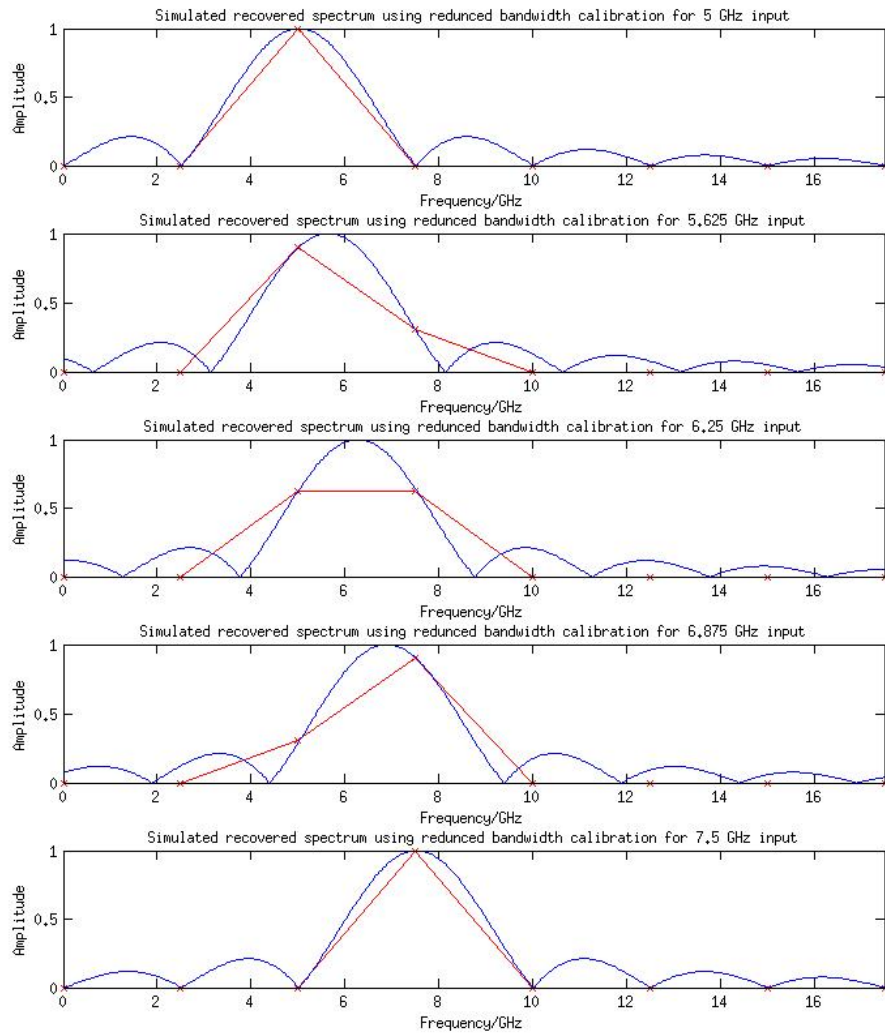


Figure 2.10: Five Recovered signals in the band of calibration for a two lag-frequency calibration.

measured spectrum for a given input is the ideal spectrum, a delta-function for single frequency input, convolved with a sinc response generated by the finite sampling of the correlation function. When the single-frequency input is selected to coincide with a centre frequency for one of the FFT bins then the central peak of the sinc function lies on this bin with the zero crossings of the function aligning with the other bins. This means the measured spectrum is still a delta-function at the correct frequency. This can be seen in the cases of the 5 GHz and 7.5 GHz input signals. When the frequency of the signal does not align with the centre of a bin, then the zero crossings of the sinc response also do not align with the frequency bins, and power is seen across the entire spectrum. The central peak measured also spans two bins, with the relative amplitudes giving the exact position of the single-frequency signal between the bins. In these cases, to recover the exact frequency of the input signal a sinc interpolation can be used. In Figure 2.10 the middle three plots show cases where the single-frequency does not align with the frequency bin centres. As can be seen the peak is two bins wide and matches the sinc response. Power is not seen to be scattered across the rest of the frequency range from the side lobes of the sinc pattern, as the calibration is zeroing the other frequency- channel values. This is a result of the reduced bandwidth calibration. These results show that the calibration can be applied to a subset of the entire frequency band. This will be used to implement a calibration with the real correlator cards.

2.3.5 Implemented calibration

The bench setup used for the initial testing of the correlator worked with a reduced bandwidth of 3-13 GHz. This frequency range was selected as the initial testing of the GUBBINS system was also going to be performed over this range, matching the IF bandwidth of the mixers to be used. This setup was used to calibrate the output of a single correlator card with a simple IF chain comprising some amplification, band pass filtering and path compensation, as shown in Figure 2.11. This represents a simplified version of the GUBBINS IF system, with all seven path compensator elements used in each arm and enough amplification to counter the loss they produce. The calibration signals were produced by a signal generator and were therefore single frequencies. As the bandwidth being used in the experiment is 3-13 GHz only, the 5 GHz, 7.5 GHz, 10 GHz and 12.5 GHz calibration frequencies were used. The correlator response for these input signals can be seen in Figure 2.12, but the 12.5 GHz signal has not been plotted as it is too small to see. This figure shows three sinusoidal correlator responses, but with a severe reduction in amplitude as the frequency is increased.

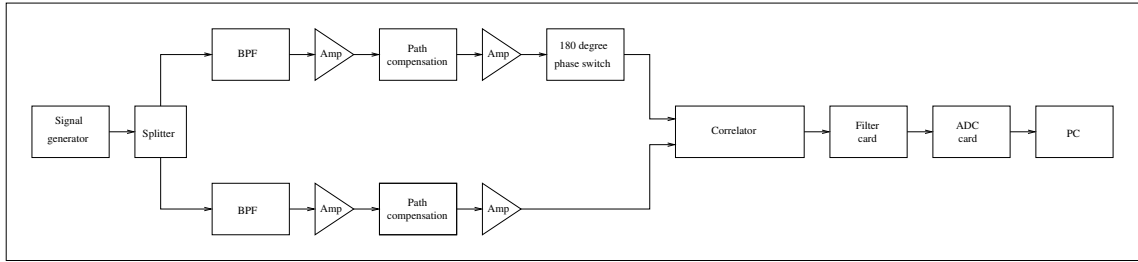


Figure 2.11: Initial setup for bench calibration test of the correlator.

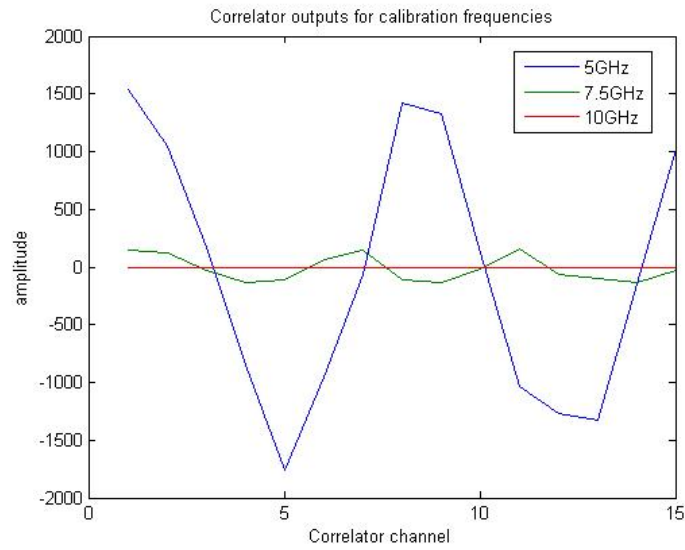


Figure 2.12: Correlator outputs for the calibration frequencies used. The 10 GHz output shows wave features with an amplitude of order 10 counts.

This reduction in amplitude is caused by the band slope in the devices used in the IF for the test setup as well as in the correlator itself. Methods for dealing with this band slope will be presented in Section 7.1.6, however at the time of the correlator testing these were not available. The presence of the band slope resulted in the calibration not working effectively, as the error in the measurement of the 5 GHz signal is of the same order as the signal for the 10 GHz case and the 12.5 GHz signal was not visible at all.

To improve on the results obtained, the majority of the IF chain was removed from the setup. This can be done as no dynamic path compensation was required to centre the correlated output, and no bandpass filtering was required to isolate the IF frequencies. This results in the much simplified setup of Figure 2.13. The lack of path compensators results in the band slope being dramatically reduced. However the correlator itself also generates a certain amount of slope, and this was manually

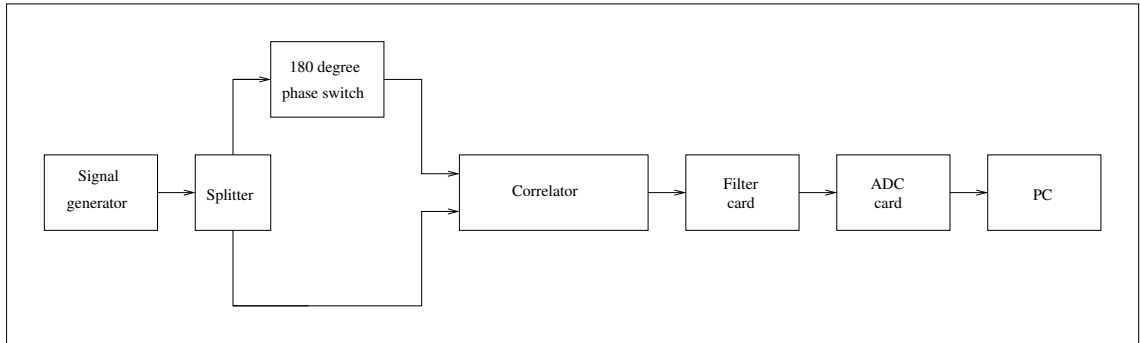


Figure 2.13: Reduced correlator setup used for second calibration attempt.

corrected for by increasing the power output of the signal generator as the frequency was increased. For the purpose of this calibration it was empirically found that a signal generator output power, P , given by the following formula dramatically reduced the slope:

$$P = f + 1, \quad (2.7)$$

where P is the output power measured in dBm and f is the frequency in GHz.

The frequency range used in the calibration was also altered to 2.5-10 GHz, which encompasses four of the calibration frequencies (2.5 GHz, 5 GHz, 7.5 GHz and 10 GHz). When the setup was implemented it was found that one of the ADC channels was malfunctioning, meaning it could not be used. To combat this the channel value was linearly interpolated from the values on either side. This will add corruption to the recovered spectra, but has a less corrupting effect than removing the channel from the data or leaving it in as a zero value. The correlator outputs for the different frequency inputs, including both the calibration frequencies and non-calibration frequencies, can be seen in Figure 2.14. The slope across the band has not entirely been corrected for as the high frequency signals are of a lower amplitude than the low frequencies, however this is a dramatic improvement on Figure 2.12.

The correlator outputs for the calibration frequencies (2.5 GHz, 5 GHz, 7.5 GHz and 10 GHz) were then used to generate a calibration recovery matrix using the method outlined in Section 2.3.2. Figure 2.15 shows the recovered spectra for the calibration frequencies using both a standard inverse Fourier transform and also the calibration matrix. As expected, when the calibration matrix is used a perfect delta-function response is seen at the calibration frequencies. The inverse Fourier transform recovery shows a series of strongly peaked functions with the band slope clearly visible. Power can also be seen to have been scattered out of the primary peak into the rest of

the spectrum. The inverse Fourier transform recovery plots are symmetric as a real correlator is being used.

The main test of an implemented calibration method is to analyse if an improvement is seen in the recovery of non-calibration-frequency signals. To do this the other single-frequency inputs (3 GHz, 6 GHz and 9 GHz) were studied. Figure 2.16 shows the comparison between the recovered spectra using both an inverse Fourier transform and the calibration matrix. To produce this plot, the inverse Fourier transform recovered spectra were normalised in the y -axis based on the peak amplitude of the 2.5 GHz recovered spectrum. This value was selected as, due to the band slope, it is the largest peak output. As the non-calibration frequencies in these plots do not perfectly align with a bin in the output spectra, the results are ideally expected to be 2 bins wide peaks. A sinc function, given by the Fourier transform of the lag spacing window function, would fit these data when centred at the correct frequency.

The recovered spectra for the 6 GHz and 9 GHz input frequencies show a clear improvement going from an inverse Fourier transform recovery to a calibration matrix recovery. Less power can be seen away from the main peak and the band slope can be seen to have been corrected for.

The recovered spectra for the 3 GHz input does not show as clear an improvement. The peak width is reduced when using the calibration recovery, but a large secondary peak is seen at 10 GHz. This secondary peak is in part caused by the band slope still being present, as high frequency components in the correlation function are given higher gain in the calibration to account for the slope.

The results show a clear ability of the calibration function to improve on the signal quality for more than half of the band. The rest of the band would also see an improvement if the band slope in the system was fully corrected for.

2.3.6 Broadband calibration

The simulated calibrations from Section 2.3.3, and the implemented calibration from Section 2.3.5, have been achieved using a series of single-frequency signals and show the feasibility of the calibration method described. Single-frequency calibration signals represent the simplest choice for performing the calibration in a lab using a bench setup.

However, when the telescope is constructed it is advantageous to achieve calibration using broadband sky signals, as this will allow the entire signal path of the telescope, and any errors associated with it, to be calibrated. The correlator output for a flat broadband spectrum occupying all the bandwidth is a sinc function, where

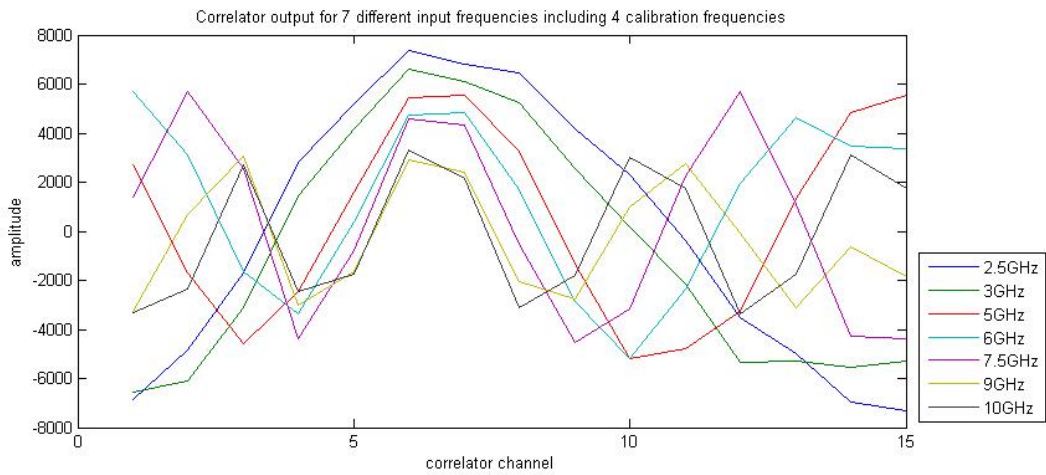


Figure 2.14: Correlator outputs for the second calibration attempt using no IF chain and forced slope over the band.

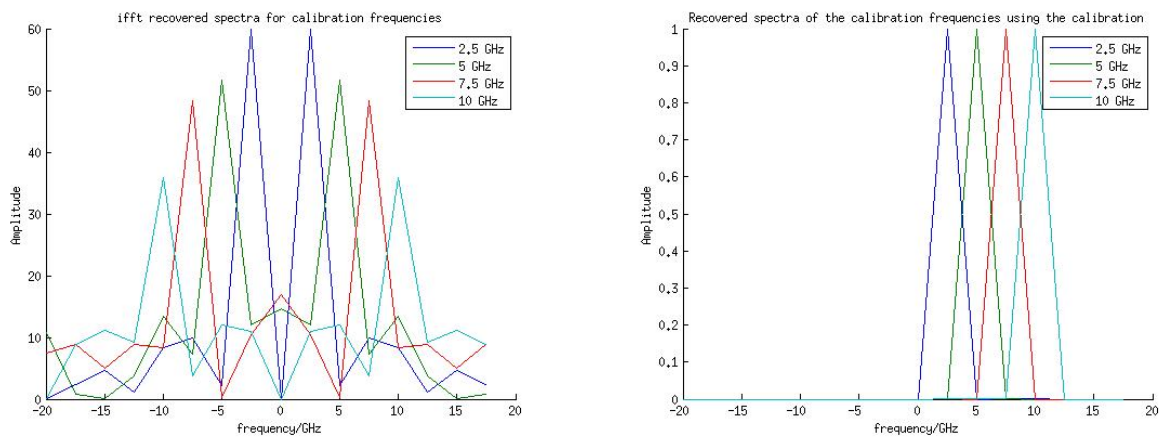


Figure 2.15: Recovered spectra for input calibration frequencies of 2.5 GHz, 5 GHz, 7.5 GHz and 10 GHz. Using (Left) a standard inverse Fourier transform and (Right) the calibration matrix generated by these inputs.

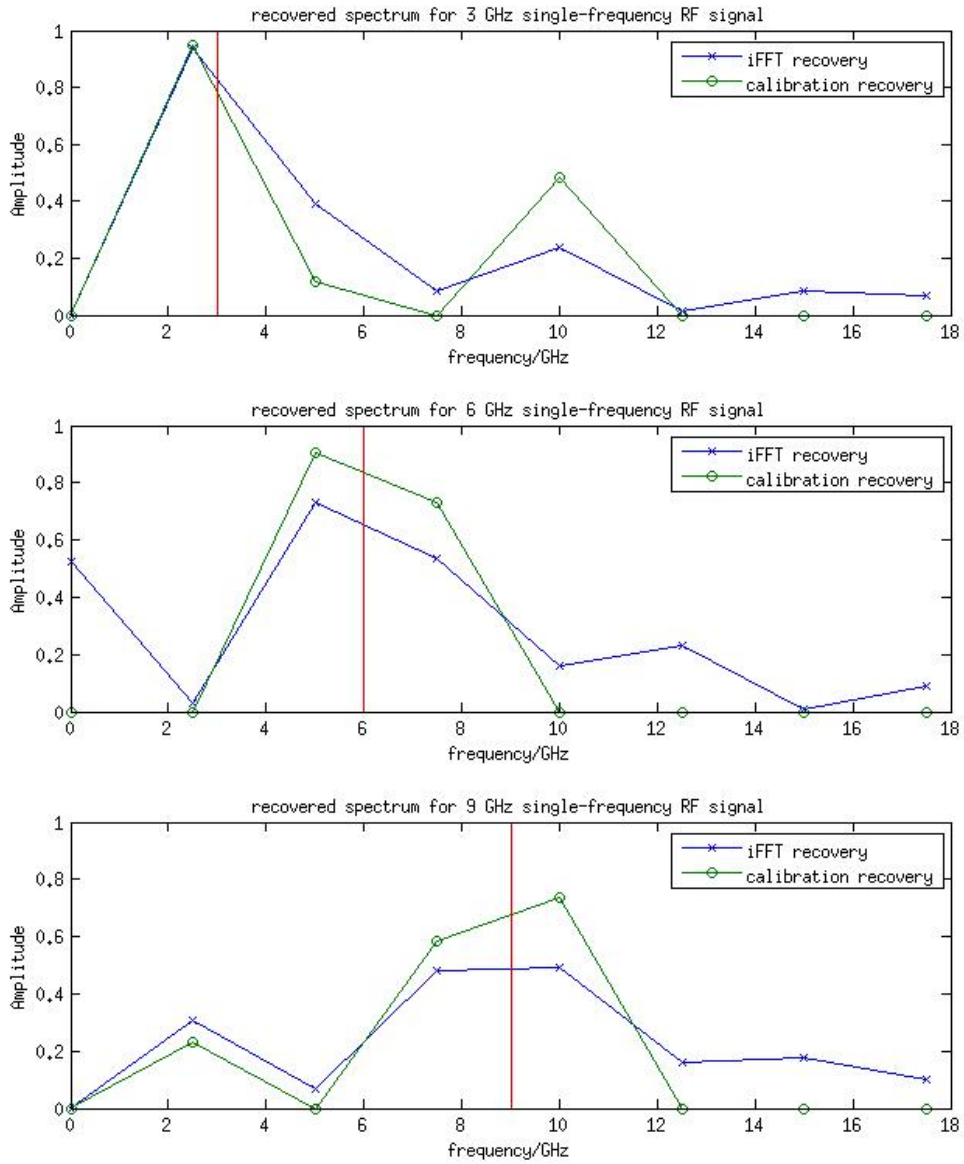


Figure 2.16: Comparisons between a standard iFFT recovery and the calibration recovery for three frequencies in-between the calibration frequencies used. The crosses and circles show the data points. The red line shows the input frequency.

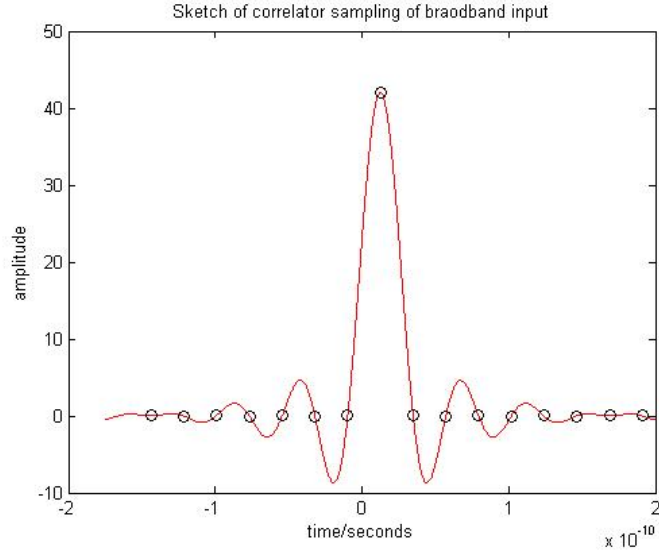


Figure 2.17: Sketch of correlator samples when given a 20 GHz broadband input signal.

the discrete samples made by the correlator lie at the central maximum and all the zeros lie in between the side lobe peaks, as shown in Figure 2.17. To perform a calibration the path compensation is used to step this sinc function across the correlator, centralising the primary peak on each of the different lags in turn. The correlator outputs can then be used to produce the \mathbf{V} matrix described in Equation 2.4. The \mathbf{R} matrix is formed using the expected amplitude spectrum for the broadband source, modified by a complex phase term used to describe the path compensation added. This phase term has the form:

$$p(l,\nu) = e^{2\pi il\nu} \quad (2.8)$$

where $p(l,\nu)$ is the phase term multiplied to the frequency bin ν of matrix \mathbf{R} , for a path delay of length l added to the input.

This calibration method is equivalent to the single-frequency case and therefore gives the same \mathbf{C} matrix for a given setup. Once the calibration is achieved, the windowing described in Section 1.2.2.3, can be used to further improve the response of the calibrated system.

2.4 Conclusion

This chapter describes the initial testing performed on the 2-20GHz analogue correlator developed for the GUBBINS system. Noise measurements of the correlator

and readout system were made showing a correlator noise floor of $\approx 10^{-6}\text{V}/\text{Hz}^{1/2}$ and a $1/f$ knee frequency of ≈ 100 Hz. This excess low-frequency noise can be effectively removed by phase-switching the signals at a frequency significantly bigger than 100 Hz.

Single-frequency RF signals were then correlated to test the output, with manual path delay being added to step the correlation function through the lags. This showed the imperfections in the system due to gain variations and lag errors, as the correlation function changed as it was moved through the lags.

A calibration method was then outlined, with simulated results. An implementation of the calibration method was performed with good signal recovery. The analogue correlator has not yet been commissioned into a working system due to delays in the production of the mm-wave RF front end. However these studies show that it should be able to function as required and the effect of the correlation imperfections can, in principle, be calibrated out.

Chapter 3

Digital Sampling techniques: 1-bit sampling

In most modern digital correlators one of the primary limiting factors of the bandwidth of the system is the analogue-to-digital converter. When designing the next generation of digital correlators it is therefore necessary to put significant consideration into the approach of analogue-to-digital conversion. As discussed in the introduction, the current generation of radio telescopes are pushing this bandwidth boundary by moving to higher speed analogue-to-digital conversion chips which offer increased data rates at reduced bit depth. These ADC chips are, however, relatively expensive and consume a lot of power.

This chapter discusses two different approaches to the task of digital sampling that can potentially avoid the need for dedicated ADC chips. The first is a direct sampling method using an FPGA, with the second being the use of external high-speed comparators.

3.1 Direct signal input sampling

3.1.1 Introduction to direct sampling

The first method of digital sampling discussed is that of direct sampling by the digital processing device.

When using high-speed digital sampling devices, a high-speed serial link is used to pass the sampled data stream into the digital processing hardware, usually an FPGA (Field Programmable Gate Array) or custom ASIC (Application Specific Integrated Circuit).

This work investigates whether the ADC chip in a data acquisition system could be dispensed with and the RF directly fed into the high-speed digital serial interfaces

on the digital processing hardware. These interfaces would then act as high-speed 1-bit RF samplers without the need for any external circuitry. This presents the advantage of greatly reducing the complexity of the system and, assuming a low bit-depth sampling is acceptable, allows multiple RF signals to be sampled by a single digital processing board, reducing the interconnect problem often associated with digital correlator development. The use of the high-speed digital interfaces as ADCs also removes the requirement to independently develop faster converter circuits. The rapid development cycle being performed by the electronics industry can then be utilised to advance to greater sampling speeds. 3.125 Gb/s serial interfaces are now commonplace, and are used, for example, as bonded channels to implement 10 Gb/s Ethernet. Much faster interfaces are now also becoming available, with speeds of up to 28 Gb/s now implemented on some devices.

A 1-bit ADC will generate a much lower signal-to-noise level than a high-bit-count ADC, and consequently decreases observing efficiency and lowers RFI rejection. From this it is clear that this technique is not suitable for all applications. However, if the decrease in complexity and increase in data rates can be achieved certain situations, such as high frequency broadband radio telescopes, could greatly benefit from its development.

The aim of this research is to investigate whether an entire high-bandwidth correlator with multiple RF input streams operating at multiple GS/s can be incorporated onto a single chip, where all data processing is being performed on the same device.

3.1.2 Hardware platform

To develop the direct sampling technique and ascertain its feasibility, a development platform was selected. The key requirement was that multiple high-speed serial interfaces could be incorporated easily into an accessible DSP (Digital Signal Processing) package. For this reason an FPGA-based solution was adopted, based on Xilinx hardware. Xilinx FPGAs incorporate multiple high-speed transceivers into a reconfigurable architecture, and Xilinx also offer good product support to academic institutions, meaning greater accessibility to the hardware. Alternatives that were considered were Altera based FPGAs or ASIC devices. ASICs do not offer the required reconfigurability and Altera devices were not readily available.

The Xilinx FPGA being used is part of a CASPER¹ (Collaboration for Astronomy Signal Processing and Electronics Research) signal processing board called a ROACH

¹<https://casper.berkeley.edu/>

Device	XC5VSX95T
CLB Array	160 x 46
CLB slices	14720
CLB Max Distributed RAM	1520
DSP48E slices	640
18 KB black RAM blocks	488
36 KB block RAM blocks	244
Max Block RAM	8784
CMTs	6
PowerPC blocks	N/A
Endpoint blocks for PCI Express	1
Ethernet MACs	4
Max RocketIO GTP transceivers	16
MAX RocketIO GTX transceivers	N/A
Total I/O Banks	19
Max user I/O	640

Table 3.1: Specification for the Xilinx Virtex 5 FPGA used in the ROACH development board[44]. CLBs are Configurable Logic Blocks and CMTs are Clock Management Tiles

(Reconfigurable Open Architecture Computing Hardware). This incorporates a Xilinx Virtex 5 XC5VSX95T FPGA that contains the logic, memory and interfaces listed in Table 3.1[44].

The most important feature of this device is the RocketIO GTP transceivers. These are the high-speed digital interfaces on the FPGA that can operate up to data rates of 3.3 Gb/s per transceiver. Sixteen of these transceivers exist on the XC5VSX95T device. These are divided into eight tiles, each containing two transceivers. The schematic layout of the ROACH board can be seen in Figure 3.1, detailing all the interfaces available.

3.1.2.1 The CASPER approach to digital correlators using ROACH

In standard operation, connectivity with the ROACH is provided by 10 Gb/s Ethernet connections, which are made using sets of four transceivers mounted to a single CX4 header. RF sampling and other peripheral activities are then provided by add-on hardware modules that connect to the Z-DOK connectors, which interface to banks of the GPIO (general purpose input output) pins on the FPGA. The Z-DOK connectors are differential high-speed digital interface connectors capable of handling data rates up to 6.25 Gb/s ².

²Z-DOK (<http://www.te.com/en/home.html>)

A standard ROACH correlator configuration can be seen in Figure 3.2. This design highlights the scalability principle of the CASPER architecture. The signals are distributed by a standard Ethernet network switch and therefore the system can be expanded to accommodate more baselines by simply connecting more ROACH boards into the network. The design shows two antennae in an FX correlator configuration. The initial FFT is performed by a different ROACH board (F-Engines) for each antenna. The cross-multiplication of these spectra is performed on a set of ROACH boards labelled X-Engines. The distribution of the signals is entirely performed by the 10 Gb/s Ethernet switch and the data are read out on a PC.

The CASPER system incorporates a Simulink-based FPGA code design environment where complicated architectures can quickly be constructed without any knowledge of the HDL (Hardware Description Language) operating underneath. This environment allows interfacing to external CASPER hardware boards as well as providing other functions such as Ethernet connectivity via the use of inbuilt function blocks called “yellow blocks”, which can be inserted into the Simulink design. This method of programming also incorporates features for using the on-board Power PC as both a configuration store for the FPGA code and as an interface method for external computers.

3.1.2.2 The approach to digital correlators for this thesis

For the purposes of this project the configuration of the system had to be dramatically redesigned so that the high-speed serial interfaces could be used to sample the RF data, with the GPIO used for digital communication with the FPGA.

The Simulink environment described above was not adopted, as the overhead involved in writing new hardware interfaces is relatively high. It is possible to add new yellow blocks which would provide the required connectivity, however this was an unnecessary level of complication for this project. All the operation code for the FPGA used has been written in VHDL (VHSIC (Very-High-Speed Integrated Circuits) Hardware Description Language) and programmed directly into the FPGA using the JTAG (Joint Test Action Group) interface as opposed to the via PowerPC. This method of interaction with the FPGA does bring about certain constraints, as none of the existing CASPER tools could be used to simplify the design. The primary effect of this was the complication of pulling data off the ROACH. Normally this would either be done using a Power PC interface or an Ethernet module, but these are both accessed through the Simulink environment and are otherwise inaccessible without extensive code development.

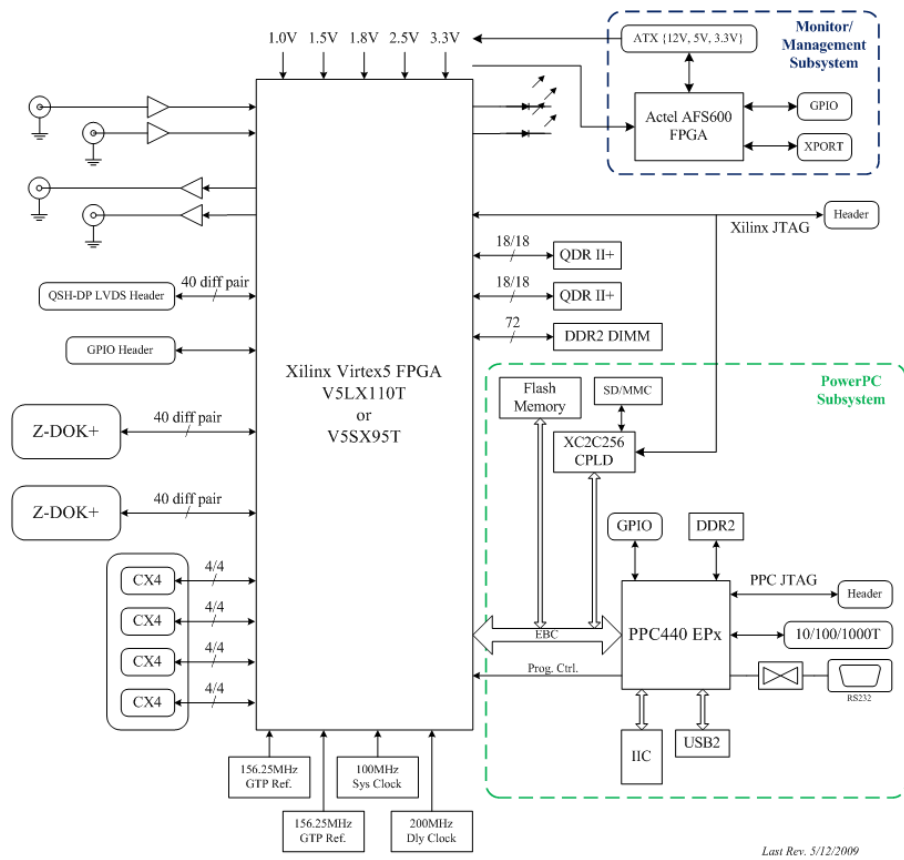


Figure 3.1: Schematic of a ROACH board.

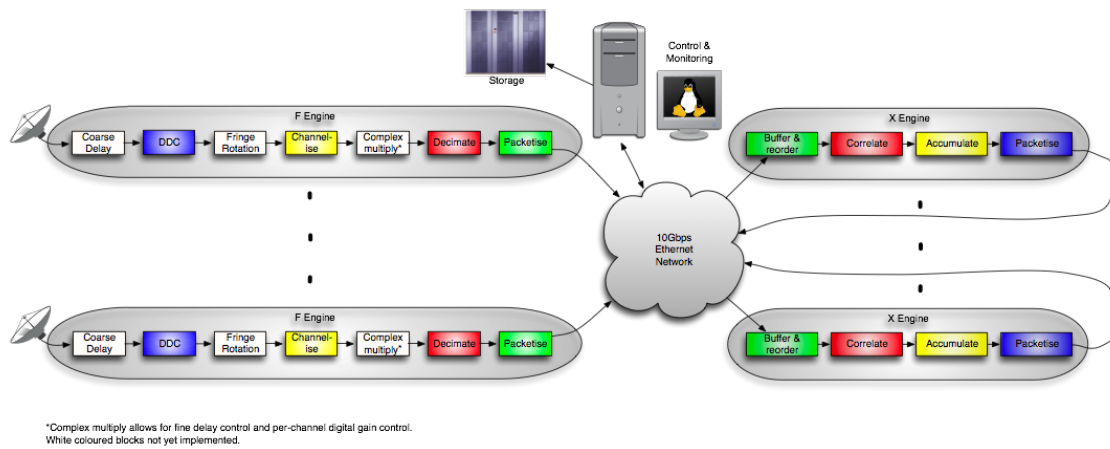


Figure 3.2: CASPER packetized correlator architecture[31].

The alternative approach adopted was to use the FPGA GPIO ports which are available on the ROACH board as an exposed header. These were linked to an Arduino microprocessor board via an SPI (Serial Peripheral Interface), with the Arduino being the master. The Arduino was also linked to a PC via a RS232 serial interface. In this manner, commands from the computer could be relayed to the FPGA via the Arduino, with data being transmitted back the other way. Both SPI and serial interfaces are supported with built-in libraries for the Arduino and were therefore straightforward to implement. This technique allowed adequate bandwidth for both device setup and data capture. The ease of developing this interface allowed the majority of the firmware development time to be placed on the transceiver interfaces.

The FPGA was programmed, as seen in Figure 3.3, with a central SPI control block. This handles access to the FPGA via the Arduino link and is connected with all other parts of the FPGA code. System components are accessed from the SPI block by reading from or writing to a distinct address. These addresses are used to refer to all aspects of the FPGA system, from control registers to direct access to the memory used to store captured data. If the PC requires a control variable to be altered, it writes the desired value to the correct address location. If data retrieval is required, the computer sends a read request addressed to the destination register, and the FPGA then issues the data as a single packet response.

As shown in the schematic for the code system, Figure 3.3, as well as allowing access to the control registers and direct sampling memory, the addressing structure also allows access to the accumulated lag value outputs from an in-built hardware correlator. This is set up by writing values to control registers, and provides accumulated cross-correlation values for pairings of RF inputs. This will be discussed in more detail in Section 6.2.

3.1.3 Transceiver configuration

One of the important aspects of the ROACH system described above is the ability to configure the high-speed transceivers so that they can be used to capture the RF data. The Virtex 5 offers a large amount of control over the operation of the transceivers with many accessible and tunable features[45]. In normal operation as a digital serial link, these functions are used to tune the link characteristics to maximise the channel throughput and reliability. They include, but are not limited to, the ability to vary the termination impedance, switch between AC and DC coupling, and vary the sampling point within the data eye.

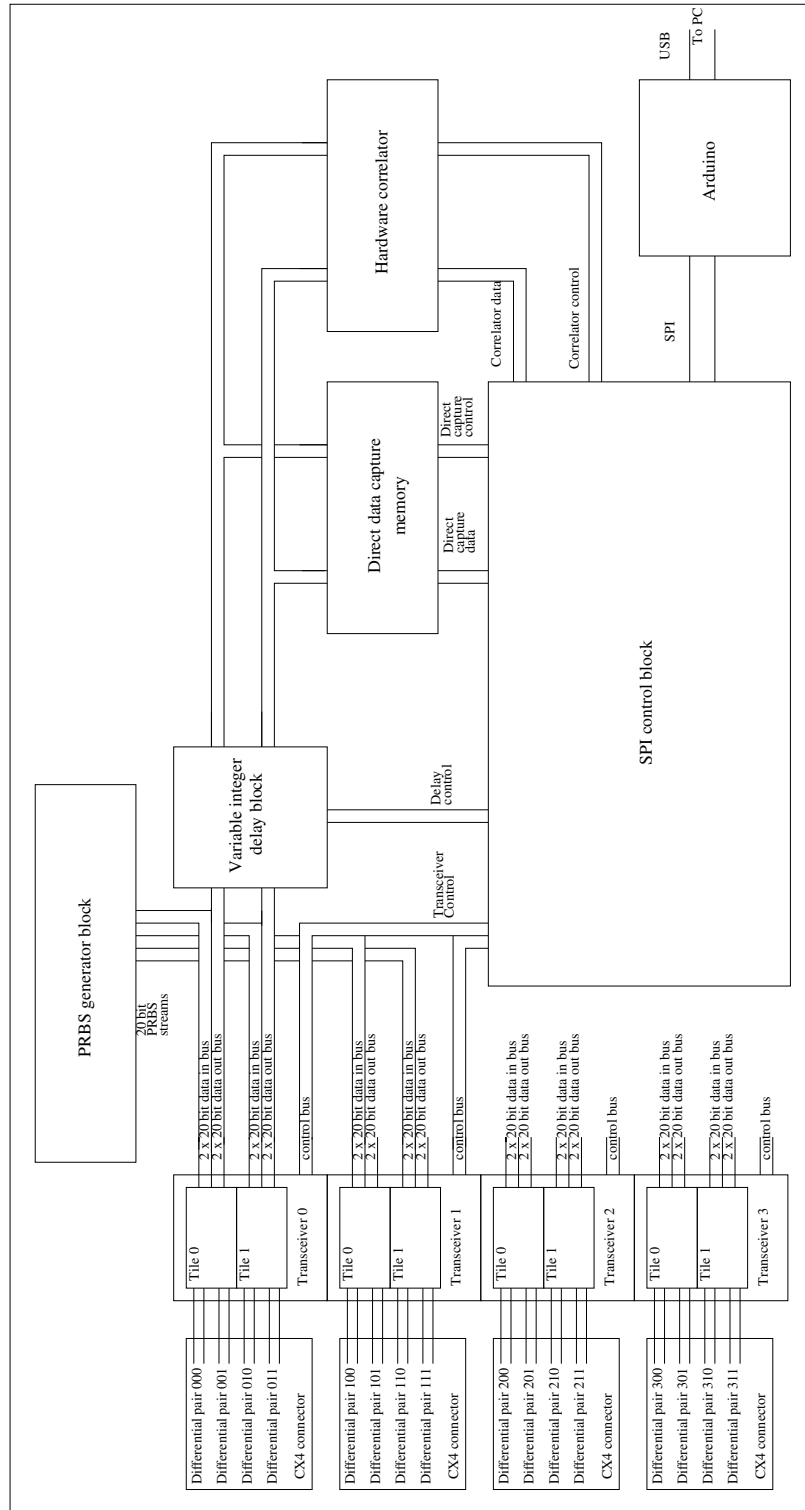


Figure 3.3: Schematic of the ROACH internal setup used in this thesis.

The configuration of the transceivers is accessed using two methods, one for settings that can be changed in real time, and another for those requiring a reset to reconfigure the change.

Settings that can be changed in real time include control over in-built transceiver functionality such as comma detection and elastic buffering. These are controlled by port mappings within the FPGA, meaning that control signals can directly be routed to the setting being altered.

For more fundamental aspects of the transceivers that can not be altered during operation, such as impedance matching or clock division, a DRP (Dynamic Reconfiguration Port) is provided. This is a block of memory with synchronous parallel access. The memory locations are mapped to different control registers within the transceivers. This means that when values are written to or read from the DRP memory, the state of the transceiver can be altered or read.

In the ROACH design shown in Figure 3.3, both these configuration methods are accessed. This allows the settings of the transceivers to be modified during testing to reconfigure the device into the desired settings.

3.1.4 RF ROACH interfacing

The above section describes the PC interface method and FPGA code environment developed to allow testing of the direct sampling method. The final piece of the test system required was a method for injecting RF signals into the transceivers. The ROACH board offers four CX4 connectors as an interface to the high-speed GTP transceivers. CX4 is one of the standard copper physical interfaces for 10 Gb Ethernet, and terminates four thin-coaxial differential lines in one screened rectangular connector. It is not ideally suited to connecting RF signals, however it is not feasible at this stage to remove or bypass this interface. An adaptor was therefore made that could be fed by a SMA RF connector and output over CX4 cable. This is made more complicated by the fact that the CX4 cable and ROACH transceivers use differential signals, whereas an SMA connector is single-ended. This meant that two SMA connectors were required for every transceiver. The first prototype of this board can be seen in Figure 3.4, where two SMAs feed a differential input and two more SMAs are fed by a differential output. The SMA connectors feed the signals into the board using coplanar wave guide, whereas the CX4 connector is fed by striplines. This means that a transition was required. This was achieved using a simple taper structure where the signal lines from the two coplanar sections narrow together into the stripline, with the associated ground planes tapering away. This design was simulated in Ansoft




Figure 3.4: First prototype SMA to CX4 interface board. This allows RF signals to be directly fed into the high speed serial transceivers. Note: One of the SMA connectors is not connected.

Designer to test the signal transmission levels. It was found that this design provided better than -10 dB reflection across the entire band.

3.1.5 Initial testing, pseudo random bit sequence and interface development

3.1.5.1 Pseudo random bit sequence testing

With the above system it was possible to begin testing of the basic principles of direct sampling. The first experiment was a verification of the FPGA code suite developed, along with the computer interface and python scripts, all without the complication of RF injection. To do this the configuration shown in Figure 3.5 was set up, where the output from one of the CX4 connectors was directly linked to another via a CX4 cable. This allowed the receive transceivers to be fed a digital data stream generated within the FPGA and output through a separate transmit transceiver. The data stream chosen for this task was required to be uniquely identifiable, to spot bit errors, as well as having lots of transitions to thoroughly test the system. For these reasons a PRBS (Pseudo Random Bit Sequence) was used. This is a digital signal where the bit stream possesses the statistics of a random sequence but is entirely deterministic, with the next state being calculable from the previous. A PRBS also has a periodicity determined by the base bit length of the pattern. This means that the stream will

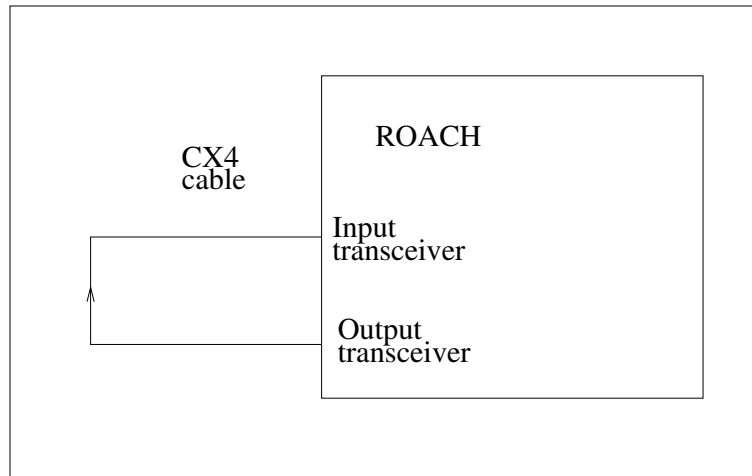


Figure 3.5: Initial setup of hardware where the input transceiver was directly fed via a single CX4 cable from a digital transmitter.

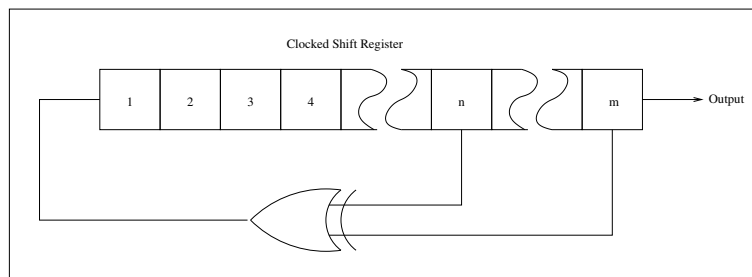


Figure 3.6: A linear feedback shift register[23].

have peaks in the autocorrelation function at regular intervals, instead of just at the zero lag as with a perfectly random sequence.

To generate the PRBS, a linear feedback shift register (LFSR) was used. This circuit, shown in Figure 3.6, is a shift register where the input is fed from a logical expression generated from the shift register bits. These shift registers can be used in many applications, however with the correct choice of logical operation for the feedback, a PRBS generator is formed.

When implementing a PRBS generator within the ROACH setup, the rate of data generation had to be considered. The transceivers operate from a 156.25 MHz clock that is up-scaled in frequency by a factor of 20 to produce the data clock at 3.125 GHz. All the transceiver logic also operates at the lower 156.25 MHz clock rate, with data being parallel loaded in 20-bit chunks into the data buffer. This means that the PRBS generator had to be able to produce the next 20-bits of the sequence at each clock cycle and therefore could not be constructed as a simple LFSR. The implementation adopted was to re-express the function of the LFSR as a matrix manipulation of

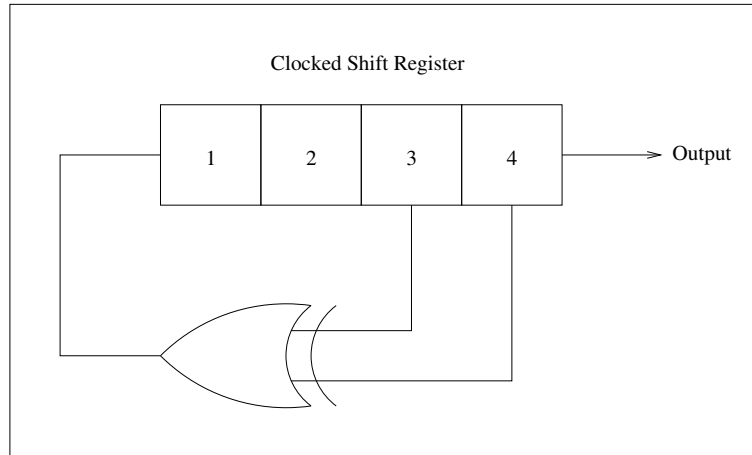


Figure 3.7: A 4 bit LFSR[23]

the current register states[16]. If each multiplication of the register values by the transformation matrix generates the next register states, and therefore one more bit of the sequence, the multiplication must be performed 20 times per clock. This is achievable as the transformation matrix is constant for a given LFSR and therefore multiplying by the matrix 20 times is equivalent to multiplying once by the matrix raised to the power of 20. This can be seen in the following example.

If a 4-bit long LFSR is made, it will look like Figure 3.7 and, given the seed of 1111, the shift register states will be as in Table 3.2. If the initial state is called S and the next state is S' , then the above can be represented as:

$$S' = AS \quad (3.1)$$

where A is the matrix

$$A = \begin{pmatrix} 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (3.2)$$

If the seed vector is represented as S_0 , then the n^{th} state of the shift registers can be written as

$$S_n = A^n S_0 \quad (3.3)$$

state number	state	output bit
1	1111	1
2	0111	0
3	0011	0
4	0001	0
5	1000	1
6	0100	0
7	0010	0
8	1001	1
9	1100	1
10	0110	0
11	1011	1
12	0101	0
13	1010	1
14	1101	1
15	1110	1

Table 3.2: LFSR output states[23].

where

$$A^2 = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \quad (3.4)$$

and similarly for subsequent outputs and higher powers of A .

This method can therefore be used to generate the state of the shift register for multiple clocks in the future, without the need for working out the intermediate states. This method was then implemented in the FPGA. For each slow clock cycle, the current state of the PRBS generator is multiplied by the correct transform matrix and the next 20 states are generated and outputted in 20-bit parallel form.

The main advantage of using a PRBS at this stage of testing was that it was easy to analyse the bit stream to determine whether or not there was any data corruption occurring since, given any state, the next state can be predicted and checked against. It is also preferable over, for example, a square wave pattern as the rate of bit change is higher and therefore the pattern is more distinguishable. When the experiment was performed the output and input data streams were viewed in a real time internal FPGA logic analyser as well as capturing the data within the PC. These three PRBS streams were seen to be the same showing that the data capture functionality was working successfully. However, the received data stream was seen to run behind the transmitted stream. This delay (or lag) was expected and was caused by the time taken for the signals to propagate through the output transceiver, along the CX4

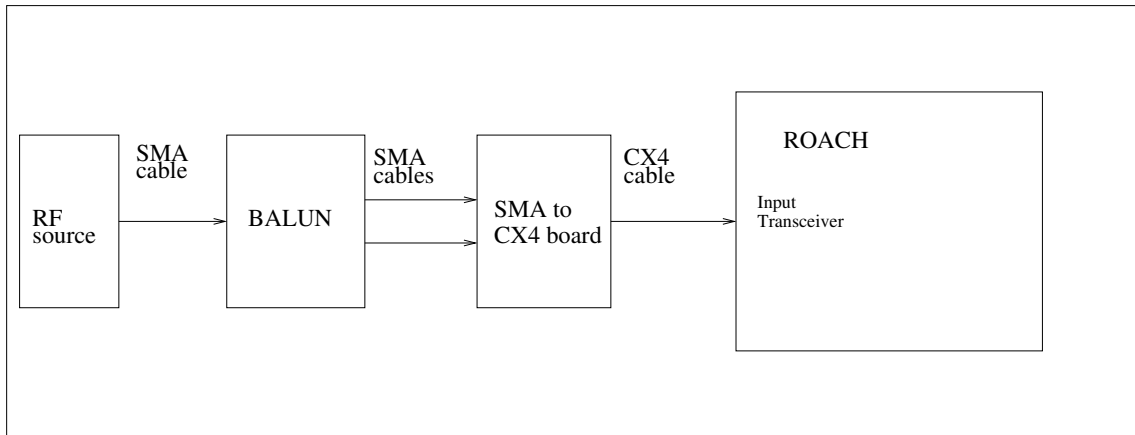


Figure 3.8: RF setup for basic single frequency data capture using an RF balun.

cable and back through the receive transceiver. The clock rate for the system is 3.125 GS/s. This means that a clock cycle occurs every 320 ps, which is equivalent to approximately 100 mm of free-space delay. This means that for every 100 mm of signal path, the received signal will lag behind the transmitted by another clock cycle.

The experiment demonstrated that the FPGA code and Arduino access were both working properly, with direct data capture achieved.

3.1.5.2 Initial RF testing

The next stage of testing was to inject a known RF signal into the SMA inputs on the CX4 interface board and read these into the FPGA. A schematic is shown in Figure 3.8. As the inputs to the CX4 board are designed to take differential signals and a single-ended signal generator was being used, a balun was employed to generate the necessary signal conversion. The balun used was a passive RF device capable of taking unbalanced (ground referenced) RF signals up to a frequency of 6 GHz and output a balanced (differential) signal on two SMA connectors. The RF frequency was set at 10 MHz, well below the Nyquist frequency (1.5625 GHz) for the system, and the power level slowly increased until data was seen on the input. A 1000-point data capture was then acquired on the system. A plot of the raw data gathered is shown in Figure 3.9. This plot shows the expected square wave representation of the 10 MHz RF signal injected with no signs of data corruption and confirms the feasibility of using the GTP transceivers to directly acquire RF data. The next stage of testing was to ascertain whether the system could be simplified by the removal of the balun from the RF signal path.




Figure 3.9: Single ended injection of a 10 MHz signal into a balun feed of the transceiver.

The transceiver inputs are designed to operate in conjunction with a differentially-fed signal. However, assuming the input acts like a comparator (as shown in the data sheet schematic[45]), then this differential signal can be replaced with a single-ended feed with the second line being connected to ground. In principle, this should produce the same sampling ability without the complication or expense of separate baluns for every transceiver. To investigate this, the RF signal was fed directly into one of the SMA inputs with the second input shorted to ground. Data were then captured as before with a direct sample of 1000-points. Figure 3.10 shows these data, which also shows a clean square wave at the input frequency of 100 MHz. From these experiments it was decided to adopt the data input approach without a balun for all future testing as it dramatically simplifies the RF injection method.

Before any further testing, a new interface board was produced based on the findings of the previous experiments. This board was designed to be fed with single-ended signals and to provide direct connectivity to all four of the transceivers on a single CX4 connector. This is shown in Figure 3.11. Also present is a multi-channel digital-to-analogue converter that will be used in later experiments.




Figure 3.10: Single ended injection of a 100MHz signal directly into transceiver.




Figure 3.11: New interface board capable of receiving four RF signals and feeding them, single endedly, into one CX4 connector.

3.1.6 Single-frequency data capture

Once single-frequency capture and processing using the ROACH implementation was established, the next tests were a more thorough analysis of the signal quality being achieved. The first stage of this testing was to compare the spectra of the single-frequency data being captured relative to theoretical models. To achieve this, the setup from the previous experiment was used but with data being collected for a range of single-frequencies. A simulation was then constructed to produce perfect 1-bit sampled data of the same frequency signals, with an FFT (Fast Fourier Transform) being performed on both data sets to compare their spectral components. For the case of a perfect many-bit sampler, the spectral response of a single frequency data capture will show a single peak at the frequency of the original signal with a low noise floor existing elsewhere. When a perfect 1-bit signal is considered, the sine wave of the RF is effectively being sampled as a square wave. When the FFT is performed on this, an infinite set of harmonic peaks of reducing amplitude will be seen. These will be folded back into the Nyquist spectral range resulting in the observed spectrum consisting of a number of peaks, the structure depending on the relationship between the signal frequency and the sampling frequency.

The plots of Figure 3.12 show the simulated and measured results for this experiment with different RF frequencies. As can be seen, the peak position and relative amplitudes for the measured data show good agreement with those of the simulation. The primary difference between the plots is the level of the noise floor. This is due to the noise floor on the simulated system being entirely generated by quantisation error. In the measured response, RF noise is inherent in all stages of the system, with further noise being generated by imperfections in the sampling, such as clock jitter.

As well as showing the effectiveness of the sampling technique, the above experiment highlighted another key feature of the sampling setup being used. Throughout the experiment the RF signal source had to be manually configured in both frequency and power level. The power level was selected by slowly ramping up from a minimal value until stable triggering was achieved. It was noted that the amount of power required by the system grew steadily with frequency. This was to be expected since all components in the signal path have loss, which increases with frequency. However, the extent of the slope was more severe than anticipated. The signal source being used had a maximum power output of 30 dBm which was found to be insufficient to trigger the samplers near the top of the first Nyquist zone (above ≈ 1 GHz).

To investigate this further, the RF chain was isolated and its S-parameters explicitly measured. The S-parameters for a system describe the energy propagation

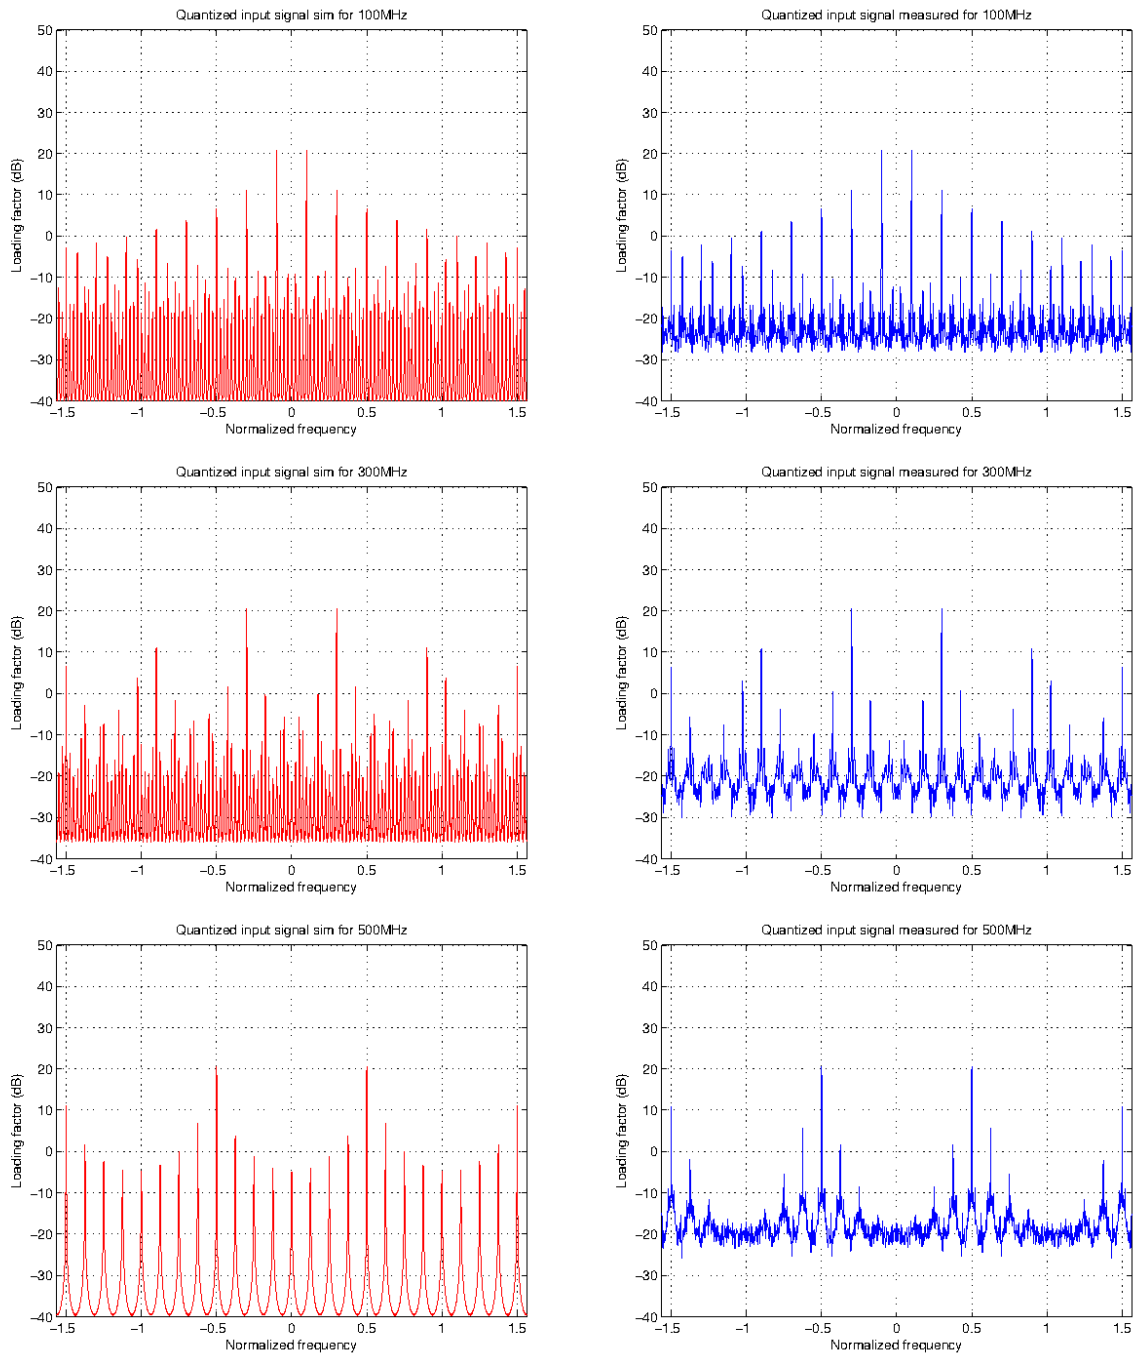


Figure 3.12: (Left) Simulated and (Right) measured spectra from a 3.125 GS/s 1-bit sampler with (Top) 100MHz, (Middle) 300 MHz and (Bottom) 500 MHz RF signal injected

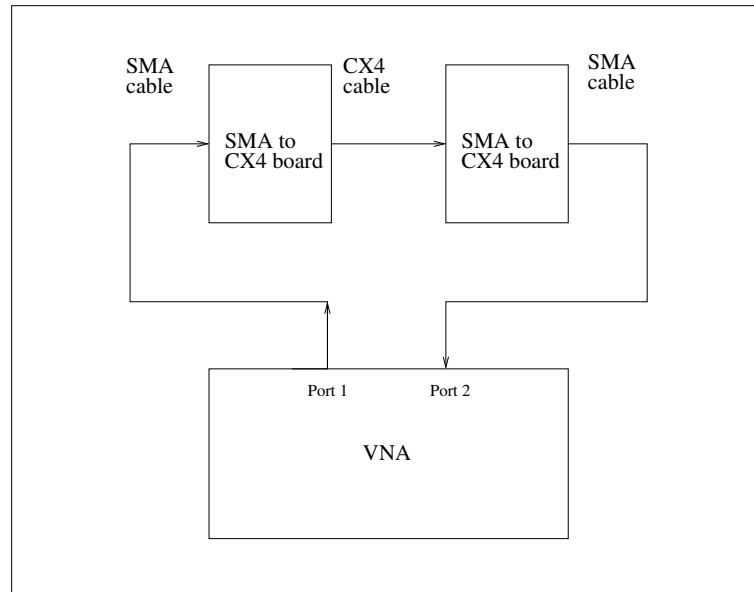


Figure 3.13: VNA setup for measuring the loss in the CX4 cable and connectors.

through an RF network as a function of frequency. Each S-parameter is denoted by two numbers, the first of which is the destination port in the RF network, with the second being the source port. In this manner, for a given single port being excited, the amount of power received on all ports can be expressed. This also allows the amount of reflected power from an RF network to be seen.

To measure the S-parameters for the boards, two CX4-to-SMA boards were connected to either end of a CX4 cable to measure an approximate value of the loss due to the cable and CX4 connector. The test setup for this process can be seen in Figure 3.13, and Figure 3.14 shows the reflection and transmission coefficients for two different length CX4 cables (1 m and 3 m). The graphs quantify the slope observed in the previous test, showing between 10 and 15 dB of slope over the first 1.5 GHz. It is also clear from the responses of the two cable lengths that the CX4 cable provides a considerable contribution towards the overall slope. This is unsurprising since the CX4 cable is a low-cost cable only designed to transmit signals of this frequency over short distances.

The above tests did not provide any indication of the losses caused by the RF signal path on the ROACH board, which was designed for high-speed digital signals but not with RF in mind. If a high-speed, direct-sampling FPGA-based system were to be produced for a final application, a custom-designed FPGA board would also be manufactured, allowing more direct and cleaner injection of the RF signals into the system.




Figure 3.14: Measured (left) reflection and (right) transmission coefficients for 3 m and 1 m CX4 cables connecting two SMA connector boards. ≈ 5 dB per metre loss seen from CX4 cable.

The results so far provide a qualitative understanding of the effectiveness of the sampling technique. However, to gain a fuller understanding, a quantitative approach was adopted that allows a measure of the signal-to-noise in the received data stream to be expressed in terms of the bit depth of an ideal sampler required to produce the same signal-to-noise level. This measure, called the effective number of bits (ENOB), makes the assumption that all the noise in the system is caused by quantisation effects in the sampler. The expected signal-to-noise ratio, $\frac{S}{N}$, for a sampler with bit depth n is given by[11]:

$$\frac{S}{N} = \frac{2^n \sqrt{3}}{\sqrt{2}}. \quad (3.5)$$

This can be expressed to give the signal-to-noise ratio in units of dB as:

$$\frac{S}{N}_{dB} \approx 6.02n + 1.76, \quad (3.6)$$

where n is the number of bits of the system. If a signal-to-noise ratio is measured for a system, this formula can then be inverted to find the effective bit depth of a perfect sampler required to generate the same signal-to-noise value:

$$n = \frac{S/N_{dB} - 1.76}{6.02}. \quad (3.7)$$

At this stage it is important to clarify what is meant by the signal-to-noise ratio of the system, as two distinct measures exist: signal-to-noise ratio (SNR), and sinusoidal signal power to total noise and distortion (SINAD). The difference between these two




Figure 3.15: Simulated and measured ENOB as a function of frequency for a 1-bit sampler

measures is that the SNR just takes into account the sinusoid peak and the noise floor where as the SINAD also incorporates any harmonics into the noise term. As seen in the previous spectral plots, a large amount of the spectral power exists in harmonic peaks caused by the 1-bit nature of the sampler. It is therefore important to use the SINAD value of signal-to-noise as this fully accounts for the noise contribution from these secondary peaks. Since the amount of noise in the system depends on frequency, it is necessary to plot ENOB as a function of frequency to gain a full understanding of the samplers' limitations. The plots of the ENOB for both the simulated case and actual 1-bit samplers are shown in Figure 3.15.

The real sampler shows a lower ENOB than the simulation, with a greater difference seen at higher frequency. This is expected, as imperfections in the system exist that limit the overall performance. The shape of both curves is due to the frequencies being chosen not coinciding with the centres of the frequency bins produced by the FFT. This affects the amount of power aliased from higher Nyquist zones into each frequency bin. These results show, however, that the direct sampling technique can function as a 1-bit sampler with close to ideal performance, providing sufficient power can be delivered to the sampler.

3.2 External comparator sampling

The previous experiment shows the use of the high-speed digital transceivers on the Xilinx FPGA as 1-bit ADCs. However, there exist dedicated 1-bit comparator chips which can be used as 1-bit samplers. These will be considered here as they represent




Figure 3.16: Hittite HMC875LC3C chip layout. The differential input signal is presented to pins 2 and 3, which are terminated through internal resistors to pins 1 and 4.

an effective comparison to the direct sampling method described above. By comparing the performance of the FPGA direct-sampling with that of chips designed to do precisely the task required, it is possible to assess whether the FPGA sampling performance is adequate.

The 1-bit comparator chips used are the Hittite HMC875LC3C³, which are 20 GS/s clocked 1-bit comparator chips with differential inputs, output and clock lines, as seen in the schematic of Figure 3.16.

3.2.1 Comparator setup

To capture and process the digital output from the comparator, the previous ROACH setup was again used. The differential output from the comparator was fed directly into one of the differential inputs in the ROACH board via the SMA-to-CX4 board. The comparator was clocked from an external 3.125 GHz signal generator fed through a differential balun connector. As the ROACH interface and the comparator were clocked from different signals, the clock recovery options were used within the GTP interface to synchronize the data capture. The hardware setup for this experiment is shown in Figure 3.17.

³<http://www.hittite.com/>

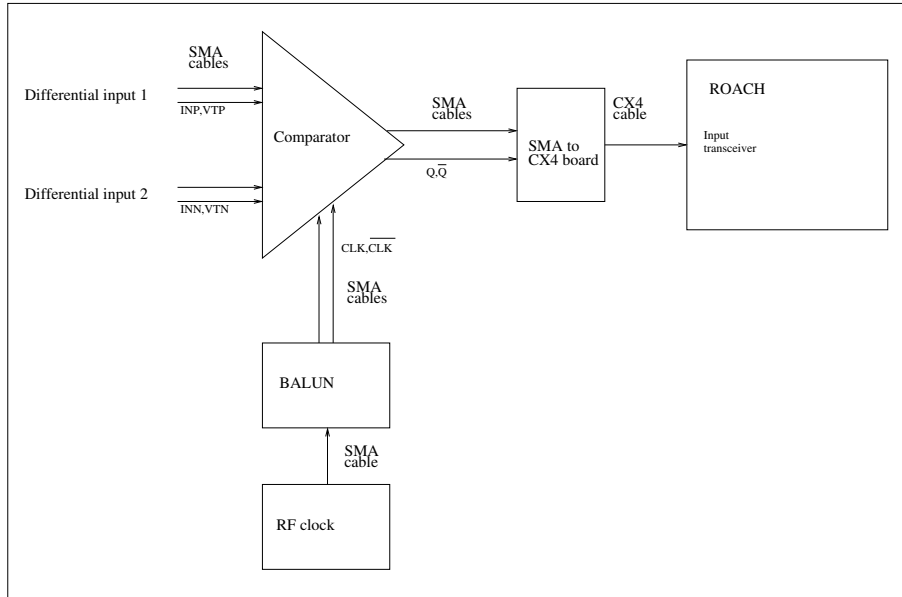


Figure 3.17: RF setup for external comparator.

3.2.2 Single-frequency capture

As with the case of direct sampling, the first test that was performed on the external sampling setup was the acquisition of single-frequency RF signals. These were fed into the differential input of the comparator as a ground-referenced single-ended signal. Direct data were then captured by the ROACH and stored on a computer for processing. The plots of Figure 3.18 show simulated and measured data spectra for different frequency inputs. As can be seen, the data closely resemble the simulated cases with the primary difference being an increase in the noise floor. As before, this is expected as the only noise present in the simulation is the quantisation error, whereas the measured data has RF noise present as well.

A limitation was quickly identified in the system, as for input frequencies above 400 MHz the ROACH interface began to struggle to detect the incoming data stream from the comparator. This is a similar problem to that experienced in the direct sampling case and again comes from the poor RF signal routing present in the setup. It can also be seen in Figure 3.18 that for frequencies above 200 MHz the noise floor in the measured case is lower than that of the simulated case. This suggests the data loss in the connection is already corrupting the conversion.

To counteract this issue an amplification stage was added between the comparator and the ROACH transceiver. This was done to boost the signal level so that it could again be detected by the receiver. The setup was made more complicated by the

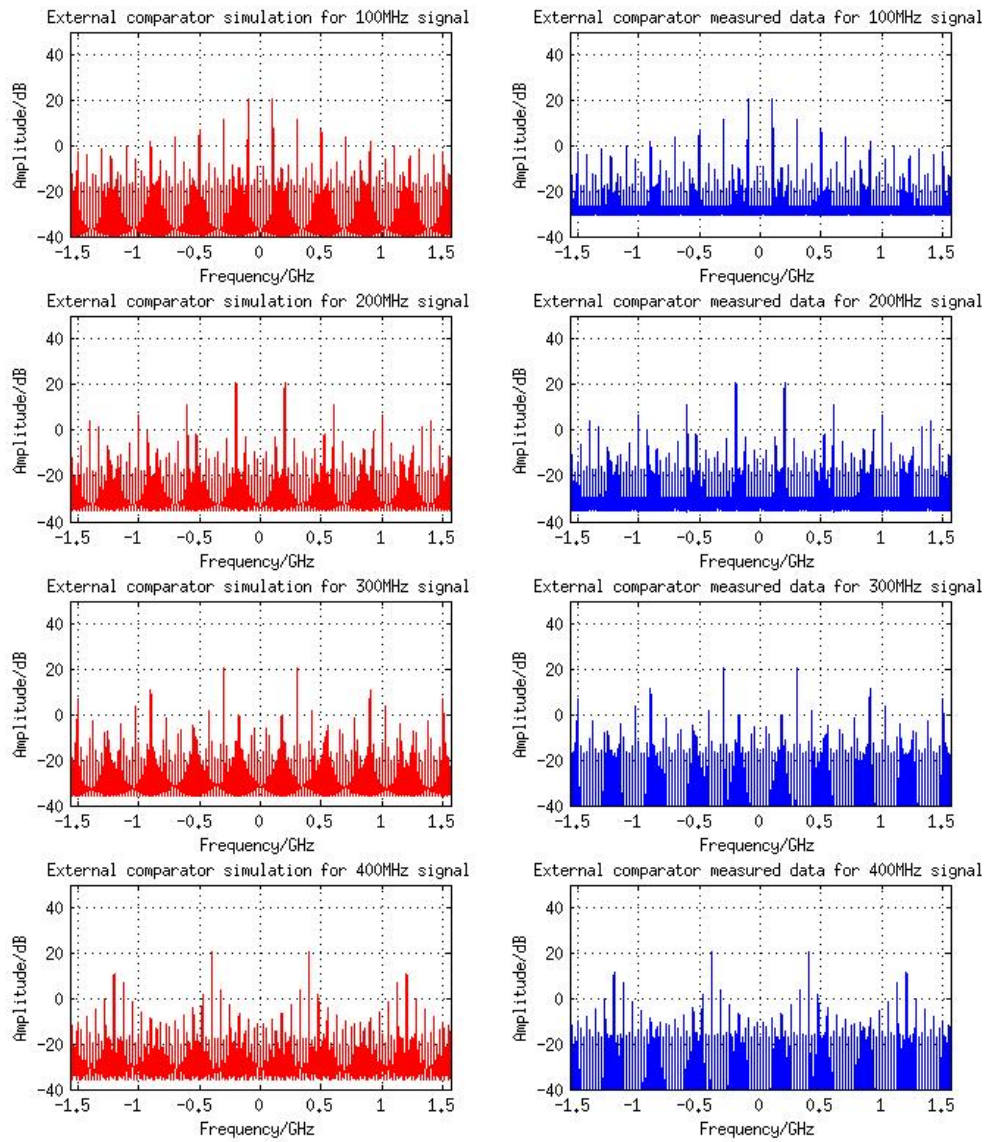


Figure 3.18: (Left) Simulated and (Right) measured spectra from an external 1-bit comparator with different single-frequency RF signals injected.

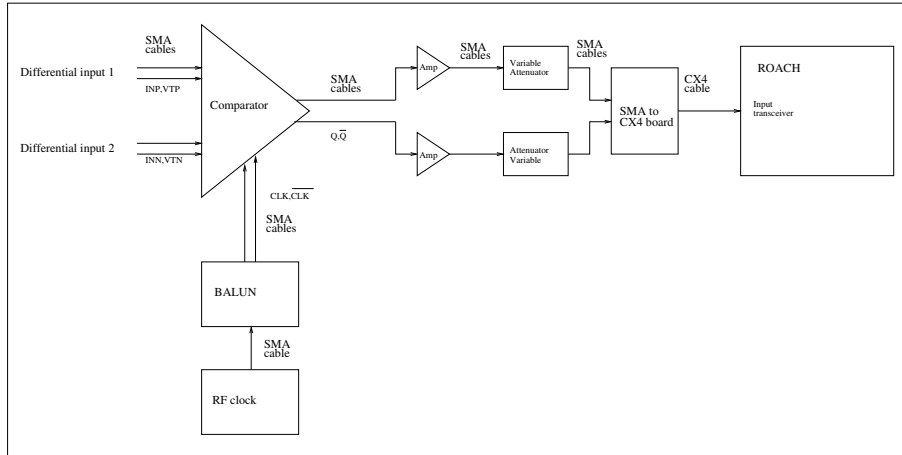


Figure 3.19: New RF setup for external comparator with two amplifiers added between comparator and ROACH.

fact that the transmission between the comparator and the ROACH was differential, and only single-ended amplifiers were available. The setup adopted, shown in Figure 3.19, was to use two DC-3 GHz amplifiers running in parallel. As the amplifiers are identical, and the signals being considered are digital, it was believed that this would maintain the required signal quality. Also included were a pair of digital attenuator devices, added to allow the signal power level to be tuned to match the needs of the receiver for any given frequency.

The first test with the new setup was to repeat the above experiment with single-frequency inputs. The equivalent plots of Figure 3.18 for the new setup can be seen in Figures 3.21. As can be seen, the signal quality observed in the low frequency cases matches that of the original experiment.

It was also found with this new setup that the full first Nyquist zone could be sampled by the device. This then enabled the ENOB calculation, previously performed on the direct sampling method, to again be carried out, establishing the efficiency of the comparators and setup being used.

The results for this ENOB measurement along with the simulated case can be seen in Figure 3.22. The results show a match in performance for the simulated and real cases up to 1100 MHz. After this point the measured ENOB is significantly worse than that of the simulation. Since the comparators are working well within their published frequency range, this is probably due to the RF quality of the signal link into the ROACH board, despite using the additional amplifiers.

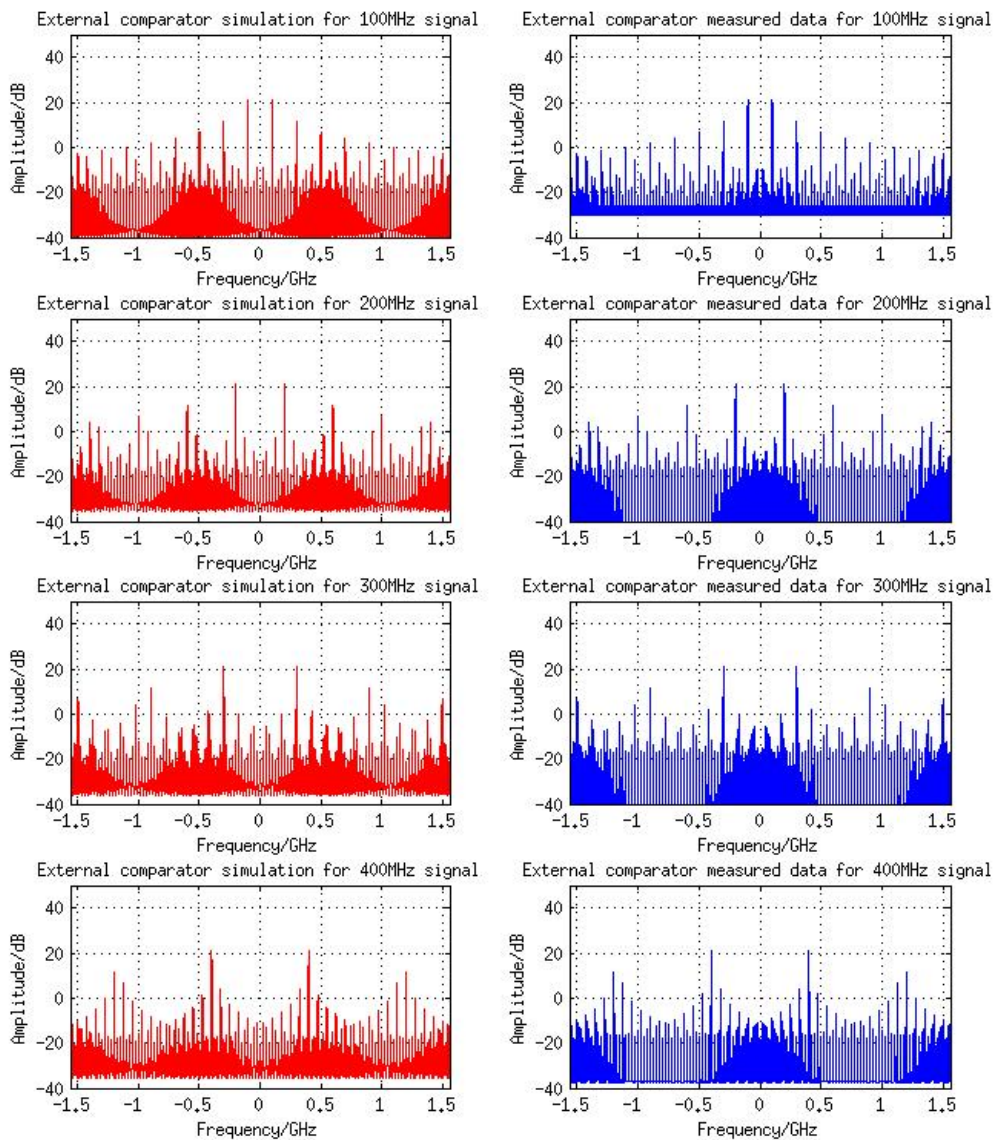


Figure 3.20: Single frequency data captures from the external comparator with amplification added between comparator and ROACH.

Figure 3.21: (Left) Simulated and (Right) measured spectra from a external 1-bit comparator with added amplification and different single-frequency RF signals injected.




Figure 3.22: Simulated and measured ENOB data for the external comparator.

3.3 Conclusion

In this chapter two methods for high-speed digital sampling have been outlined. The first method was the use of the high-speed serial transceivers on a Xilinx FPGA as 1-bit direct samplers. The second high-speed sampling method explored was the use of high-speed 1-bit external samplers from Hittite.

The single-frequency data capture results and ENOB of Figure 3.23 for the two systems show the results were consistent with each other and with the simulation. Both setups were able to generate an effective number of bits around 0.8 at 500 MHz. The external sampling method showed slightly higher signal-to-noise. This could in part be caused by the cleaner RF injection, as well as higher stability on the clock used for the external sampler. Both systems experienced problems with higher frequency RF signals due to the quality of the RF input into the ROACH.

Although the external comparators showed near-ideal performance over the first 1 GHz, the direct sampling technique also achieved a very high fraction of the ideal signal-to-noise, with a much simpler system. It was therefore decided to continue to investigate the behaviour of the direct sampling method in more complex setups, namely interleaving, and multi-level sampling.




Figure 3.23: Comparison between external and direct sampling techniques.

Chapter 4

Digital Sampling techniques: Interleaving techniques

It has now been shown that 1-bit sampling at the 3.125 GS/s clock frequency can be obtained by a simple interface system using the direct sampling technique. The extension of this is to attempt to increase the sampling data rate, either by interleaving multiple comparators together to obtain a finer sampled signal, or using multiple comparators to achieve better than 1-bit sampling. This chapter focuses on interleaving.

4.1 Basic principles of interleaving

Interleaving is the technique of using multiple slower samplers at staggered delays to emulate a single faster sampler. A diagrammatic representation of this can be seen in Figure 4.1, where two samplers clocked 180 degrees out of phase sample the same 1 GHz RF signal. When the two data streams are combined together a single stream is constructed that is equivalent to a single sampler working at twice the original sampling frequency. This process can theoretically be repeated for an unlimited number of individual samplers interleaved together with appropriate clock phases to generate whatever effective sample rate is required. In reality this process will be limited by how accurately the samplers can be aligned and how much the sampling position jitters. This technique is used in commercial samplers to achieve higher sampling rates. An example of this is the Fujitsu CHAIS demonstration chip, which uses 80 interleaved 8-bit ADCs to achieve a sampling rate of 56 GS/s¹.

The most important aspect of any interleaving system is the method used to stagger the sampling positions. Two approaches are possible: either the same RF

¹https://www.fujitsu.com/downloads/MICRO/fma/pdf/56G_ADC_FactSheet.pdf




Figure 4.1: Simulated example of how two transceivers can be combined to capture an RF signal at twice their individual sampling rates

signal is fed in phase to each sampler and the clocks are staggered by a constant phase shift (180 degrees in the case of two samplers being interleaved) or the clocks are fed in phase and the RF signal into one of the samplers is delayed such that it arrives half a clock behind the other. In the latter case the phase delay added to the RF line is not constant with frequency and therefore can not easily be added as a phase delay circuit. Instead a propagation delay must be added to the signal path to achieve the desired time offset. This propagation delay can simply be provided by the lengthening of the transmission line going into the sampler. Staggering the clock arrivals can be achieved using commercially available chips, these can perform advanced clock functions such as phase shift and propagation delay. One example of these devices is the Hittite microwave HMC988LP3E², an ultra low-noise clock divider capable of stepping down the frequency of an input clock by a factor of 1/2/4/8/16/32, as well as adding a phase delay and a propagation delay. The device can also be programmed using an SPI interface so offers dynamic reconfigurability. If this device were added to the clock input to the samplers it could be set to maintain the frequency but add the required propagation delay to the clock line to ensure the clocks arrived out of phase.

4.2 ROACH implementation of channel alignment

No guarantee can be made that the signal paths in the ROACH board that link the CX4 connector to the transceivers on the FPGA are matched in propagation delay. This means that the actual offset requirements to achieve interleaving will either first need to be measured, or the system constructed will need to have a dynamically reconfigurable setup to compensate for the intrinsic offsets found. This makes it difficult to implement the propagation delays in the RF lines using path lengths, as no reconfiguration will be available. Reconfigurability is an important aspect as it will allow all the different tests to be performed on the system without constantly switching the physical devices being used. It was therefore decided to adopt a clock delay approach to the interleaving. The ROACH board has all clock lines already routed for the device with internal oscillators providing the required clock frequencies, which means that external clock management devices can not be used. Fortunately, as the GTP transceivers within the FPGA are designed to allow a huge amount of customisation to the sampling environment, clock management settings are built into the device. The feature that can be used within the FPGA to allow clock offsetting

²<http://www.hittite.com/>




Figure 4.2: Diagram of a data eye in a standard high speed serial digital data link[45]. Shown are the relative phase offsets for the sample with respect to the recovered clock.

is the horizontal sampling-point shift within the RX transceiver clock data recovery. In normal transceiver operations, where a serial digital data stream is being fed into the transceiver, the transceiver generates a recovered clock by phase-locking to the incoming data stream. The zero phase of the recovered clock is aligned with one of the data transitions and the sampling point is set relative to this. This sampling-point shift is used to align the sampling point with the centre of the data-packet eye to guarantee clean data recovery. A data eye is a representation of the periods within a digital data stream where transitions between bit states do not occur. Figure 4.2[45], shows a typical representation of a data eye. The Figure also shows the options for the sampling offset between the recovered clock (with clock transitions synchronised to data transitions) and the point where the data is actually sampled. In this example 90 degrees would be selected, for the transceiver to sample the desired data and not the transition. The horizontal sampling-point shift function of the transceiver allows 128 positions to be selected between 0 and 180 degrees (corresponding to the separation between data transitions). To determine the correct value, each position would be selected with a known input data stream, normally a PRBS. The bit error rate for each position would be measured and the shift value that minimised errors would be selected. This results in the sampling point being aligned with the correct position in the data stream.

For the purposes of this research the horizontal shift can also be used to vary the phase of the sample clock. The sample clock is derived directly from the on-board clock and is not recovered by the transceiver from the incoming data stream. Therefore, any shift is done in absolute terms with respect to this stable reference. From this, the relative sampling positions of multiple transceivers can be controlled to allow interleaving. To access and control this feature within the FPGA, the DRP is used. This was discussed earlier in Section 3.1.3 and allows access to most of the

base functionality inherent in the FPGA. The horizontal shift is accessed through the last 8 bits of PLL_RXDIVSEL.OUT, a data registry within the DRP. This is independently controlled for each of the two transceivers in a tile to allow different alignments to be achieved.

The functionality present within the system to provide the offsetting of transceiver sampling and therefore interleaving has now been identified. However to use this the correct configuration values must be found. A method was identified that allowed the relative alignment of transceivers to be measured and the correct horizontal shift value to be selected to provide the desired offset. The method relies on the fact that a maximum in the cross-correlation value between two identical signals exists at the zero lag (corresponding to a zero difference in the sampling points on the signal). This means that if the relative sampling point of the centre lag is shifted from $-1/2$ lag to $+1/2$ lag, a maximum should be seen in the centre corresponding to the two channels being perfectly aligned. To do this in the system with two channels, one transceiver is kept at a constant horizontal shift value while the other is incremented through all 128 options. For each of these positions the cross-correlation is calculated and the value at the centre lag taken. This approach effectively samples the peak of the correlation function at an accuracy of $1/128$ of the lag separation. As the zero lag will always show a maximum for any input signal, a single frequency can be chosen as this is the simplest form of input. This will show a periodic maximum peak on the correlation function at the period of the signal, so the approximate alignment of the system must be known to obtain the zero lag. A simulation of this effect can be seen in Figure 4.3, where an infinite-bit-depth correlator has been modelled. The plot shows a periodicity of 128 shift positions, caused by the zero lag moving from the centre position to the next lag over.

The above explanation and example were for infinite bit depth; the situation is more complicated when the 1-bit case is considered. If the frequency of the test signal is picked to be an exact multiple of the sampling frequency, then the left plot of Figure 4.4 is obtained, where all information about the position of the peak is lost. If the test frequency is chosen to be a non-integer multiple of the sampling frequency, a peaked function, as shown on the right plot from Figure 4.4, can again be obtained. To simplify the frequency selection and remove any ambiguity, a broadband noise source can be used as the RF signal. This has two advantages, the first being that as all frequencies in the first Nyquist zone are present, a highly peaked correlation function is produced as the horizontal shift is altered. The second advantage is that a broadband noise source only has a single correlation peak around the zero lag and




Figure 4.3: Simulation of the variation in correlation peak as the horizontal shift is changed for a 100MHz sampled RF signal.




Figure 4.4: Simulation of the variation in correlation peak as the horizontal shift is changed for a (left) 312.5 MHz and (right) 279.5795 MHz 1-bit sampled RF signal. The left case shows no peak feature due to the frequency chosen being an exact division of the clock frequency.

not the repeat pattern seen when using single frequencies. This means the alignment will always find the zero lag and not centre on another position. The shape of the alignment plot seen from simulations of the broadband case also differs from those of Figure 4.4. This is because the correlation of a broadband noise source gives a smooth sinc response, even for the 1-bit sampled case. This means the alignment function resembles that of Figure 4.3.

As described, the zero lag position coincides with the peak of the above plots: this is a perfect alignment of the sampling points of the two channels. To anti-align the channels the minimum position is selected.

When this above experiment was repeated using the ROACH sampler, a problem was quickly identified due to an ambiguity in the boot cycles of the samplers. When the transceiver is first powered up, either after reset or when the board is first turned on, the samplers need to configure themselves and fill their 20-bit multiplexing buffers. This can take a different amount of time for the different transceivers, resulting in the samplers being an unknown number of samples apart. To perform the change in sampling position, configured through the DRP, a reset must be sent to trigger the transceiver to boot in the new configuration. As the entire transceiver is rebooting this integer sampling offset will occur, changing every time a new horizontal shift value is selected. The system was designed to cope with this integer delay problem by having a user-defined variable integer delay manually added to each data stream, which can be manipulated to bring the streams back into alignment. This is one of the reasons why a broadband noise source was selected for calibration instead of a single frequency, as it allows this alignment to be performed. The process does result in one remaining ambiguity however, if the sampling point is moved half a sample away from the zero lag maximum, the peak of the correlation function will be evenly spread over the two centre lags. This makes definition of the centre lag difficult, as two options exist. This will mean that when two channels are anti-aligned it will have to be manually checked to see which sampler is triggering first.

After all the above simulations were performed, and the expected result determined, it was necessary to run the actual experiment using the ROACH sampler and the broadband noise source. The setup for this test can be seen in Figure 4.5. The horizontal shift was swept through its 128 positions to produce the graph of Figure 4.6. As can be seen this accurately reproduces the required shape where the aligned and anti-aligned positions can easily be identified.

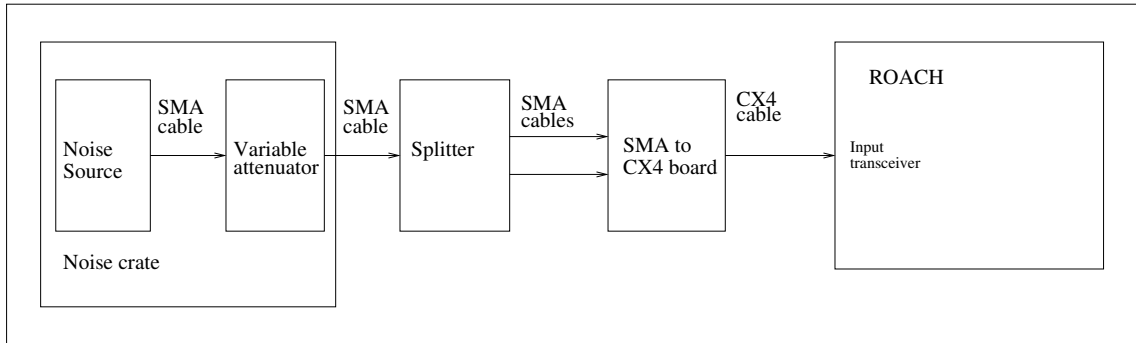


Figure 4.5: RF setup for channel alignment using a broadband noise source fed into two transceivers and the horizontal sampling position altered. An anti-aliasing filter is used to limit the noise source to the first Nyquist zone.




Figure 4.6: Fractional alignment routine. This is the variation of the peak of the correlation function as the fractional delay between the two channels is changed. Fractional delay position ≈ 50 is aligned and position ≈ 115 is anti-aligned.

4.3 Correlation technique for improving sensitivity

Now that the channels can be set in anti-alignment ready for interleaving it is necessary to test the full implementation. With multi-bit samplers this is easy to achieve with a sinusoidal applied signal, as the interleaved samples can clearly be seen to fit between each other to produce a fuller representation of the applied signal. In the case of a 1-bit sampler, however, the same approach does not yield the same ability to see the increased resolution. Instead the fact that the Nyquist frequency has been doubled must be utilised. This means that an interleaved system will be able to distinguish between two signals at f and $2f_n - f$, where f_n is the Nyquist frequency for a single sampler, which a non-interleaved system would be unable to do.

However, in order to test interleaving at these high frequencies, it was necessary to improve the systems ability to detect high-frequency signals. When looking at the frequency response of the system it was found that a large amount of slope exists across the band. This limits the maximum frequency that can be sampled by the transceivers, as the signal generator's maximum output power is 30 dBm. This limit is reached in normal setup around the end of the first Nyquist zone and therefore will not allow signals at higher frequencies than this to be sampled. To overcome this issue a correlation technique can be adopted. As outlined in the introduction, correlators can be used to isolate a low-level signal from within a noisy environment. This means that a weak high frequency signal can be detected by the system if it is fed coherently into two inputs along with uncorrelated noise. The noise can then be set at sufficient power to trigger the samplers. When the correlation between the two inputs is performed with a long integration, the noise level will be reduced and the desired weak single-frequency signal will be clearly detected.

To do this, independent uncorrelated noise signals were fed into each sampler, covering a frequency range of 0-1.5 GHz with sufficient power to trigger the samplers. The desired single-frequency signal was then added coherently to both channels. As the samplers are being triggered by the noise sources no requirement is put on the power level of the single-frequency signal.

4.3.1 Correlated noise

In any real system, spurious correlated components will exist within the noise which do not integrate down, and therefore need to be of lower level than the desired signal being measured. To measure the quantity of unwanted correlated components

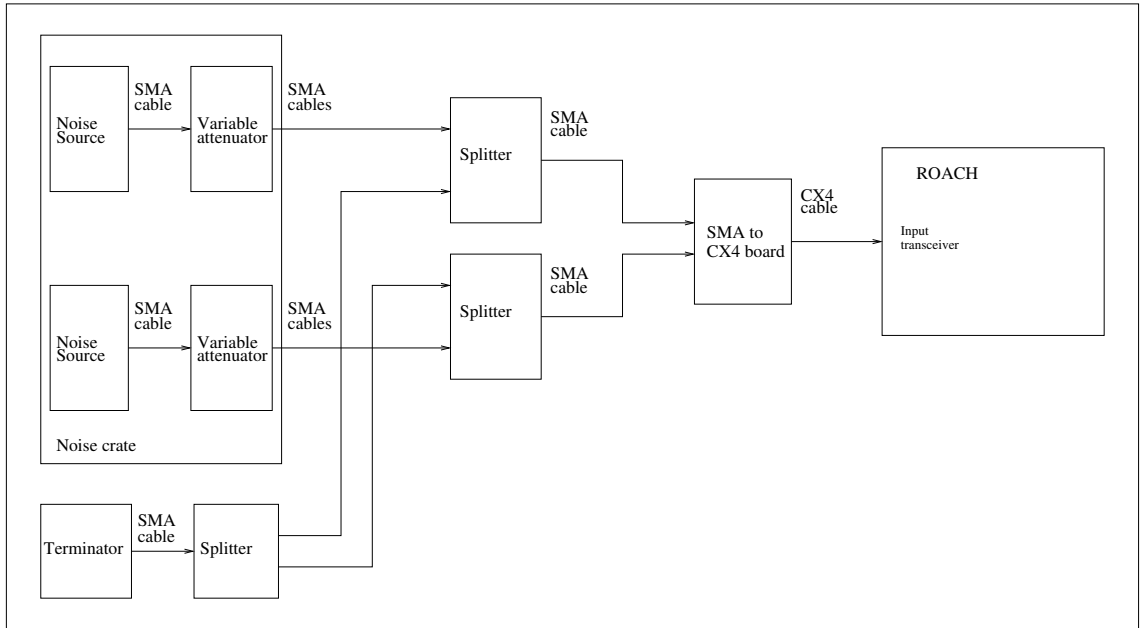


Figure 4.7: RF setup for independent noise triggered sampling. This was used to measure the correlated component of the two noise sources.

the correlation of the noise sources without a signal present was checked. This measurement also allows the maximum length of integration to be ascertained, as once the correlated component of the noise dominates over the uncorrelated component, further integration does not achieve a better signal-to-noise ratio. The setup used can be seen in Figure 4.7. Although no single-frequency signal is being added, the point of injection has still been included. This was done since a primary cause of coherent noise is part of the noise generated on one channel passing backwards over this link and into the second sampler. The test was performed by taking correlations with increasing integration time. The expected result is to see the level of correlation between the two channels decrease as the integration time increases until the limit of correlated noise is reached at which point the level remains constant. The graph of Figure 4.8 shows the result that was obtained. This differs from the expected, showing very little reduction in the correlation level as integration time is increased. This shows that a large amount of correlation must be present between the noise sources.

As correlation between the two sources was seen it was important to gain an understanding of the underlying cause and attempt to reduce its effects. Three areas were identified as being possible causes for the large correlation:

1. The cross link between the two arms used to inject the single frequency signal
2. The noise sources themselves, probably due to shared power lines




Figure 4.8: Correlator outputs for independent noise fed into two channels with increasing integration period. The correlation amplitude axis is plotted logarithmically.

3. Cross talk in the CX4 cable and on the FPGA board

To test the effect of the cross link on the correlation between the two arms the setup was altered to that seen in Figure 4.9, where the noise sources are directly linked to the transceivers and no RF path exists between the two arms. This was done since if a high correlation is still seen it will prove that the cross link is not a dominant contributor. When this experiment was performed little reduction was seen in the levels of cross-correlation.

To test the amount of cross-correlation present within the noise sources, the setup seen in Figure 4.10 was then constructed, where four transceivers were sampled with the connections shown in Table 4.1. In the ideal case where the noise sources are independent and no other cross-talk exists, only the correlation between channels 0 and 1 will show a strong signal, with all other correlations showing zero. If cross talk exists in the noise crate between the two independent noise sources then correlation should also be seen in the outputs from the correlation of pairs 0 and 1, and 0 and 2. The relative strength of these two correlations will give a measure of the extent of the internal cross talk. If a correlation is seen with transceiver 3 and any other transceiver then cross talk must be occurring in the CX4 cable and on the FPGA

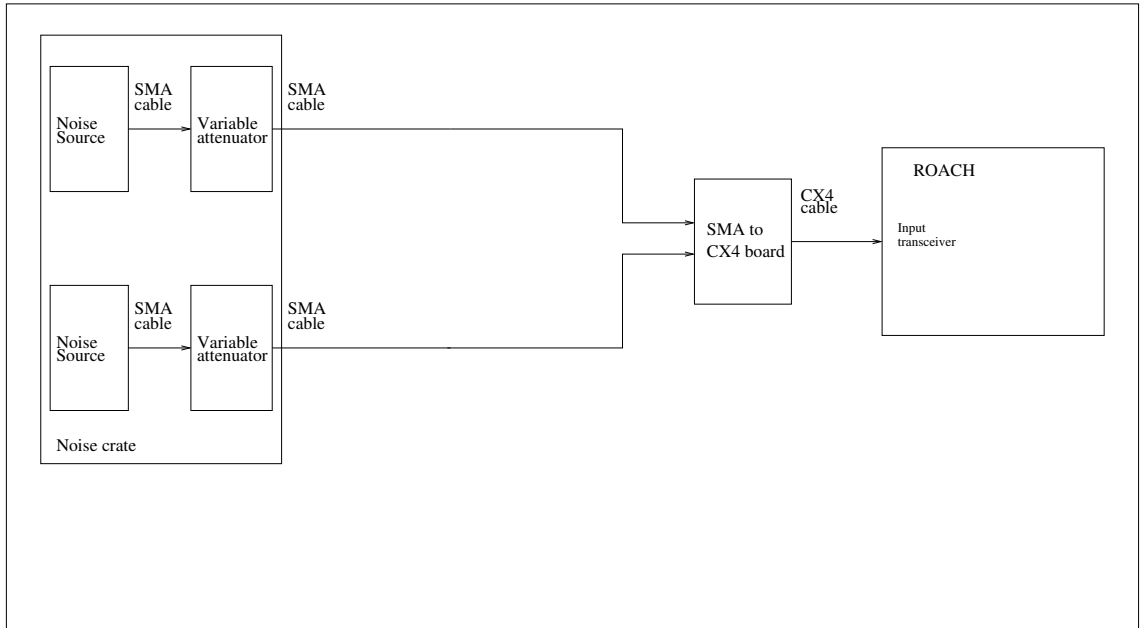


Figure 4.9: Modified RF setup for independent noise triggered sampling where no direct RF link exists between the two samplers. This tests the correlated component of the noise sources and any cross talk effects.

Transceiver	signal
0	noise a
1	noise a
2	noise b
3	carrier wave signal

Table 4.1: Signal connections for the four transceivers in Figure 4.10

board, while if the level of this correlation is comparable to that of the 0-2 pair then this effect must be dominating over any cross talk within the noise crate itself.

It is therefore possible to discern the different possible causes of cross talk by analysing the amounts of correlation between the different channel pairs. The correlations seen show a high level of cross talk on all pairs of transceivers, which is indicative of crosstalk occurring on the CX4 connector and the FPGA board. No difference can be seen between the levels of correlation for the pairs including the CW signal and the pairings of the two noise crates, this suggests that any cross talk in the noise crate is much lower. Once the crosstalk in the CX4 cable and board is corrected for, other effects might become apparent that are currently being hidden. These were considered at a later stage.

The cross talk that is seen in the system is most likely being caused by the physical proximity of the signal wires as the microstrips approach the CX4 connector, within

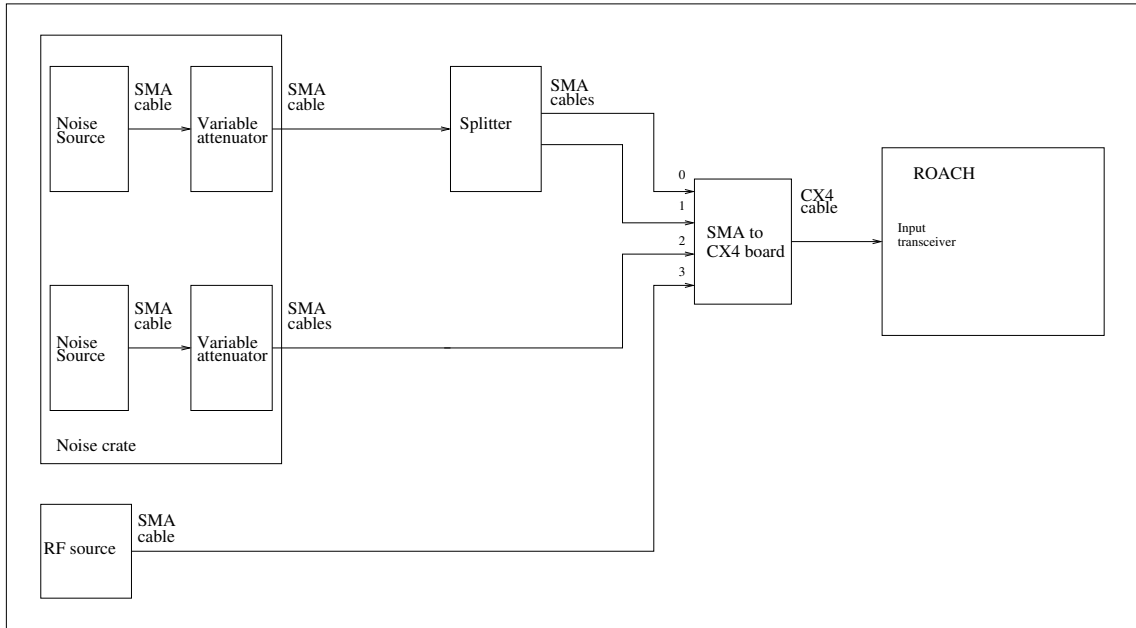


Figure 4.10: RF setup to locate the source of the correlated noise term.

the CX4 cable and once they are on the board. An obvious solution to this would be to increase the physical separation of the signals and add shielding where possible. This approach would require drastic redesign of the board and signal input method which is unfeasible at this stage of development.

4.3.2 Phase switching

Instead of the physical approach described above the electronic solution of phase switching was adopted. Phase switching works by periodically adding and removing a 180 degree phase shift to the inputs before they pass through an area of high cross talk. The switching is controlled by the sampling hardware, which then removes the added phase switching in software, restoring all signals to their original forms. This means that no effect is seen on the desired signals and their correlations. Any cross-coupled signals will have occurred within the phase switching region and will therefore still be phase switched after the software switching within the device. If the phase switching periods are selected appropriately this results in all correlations having no net contribution from the cross talk.

To explain this further, the example of the two channel device shown in Figure 4.11 is considered. In this example only the signal from source a is being phase switched with the signal from source b left unchanged. The graphs of Figure 4.12 show the two signals at the different stages of the system where a 100 MHz sine wave

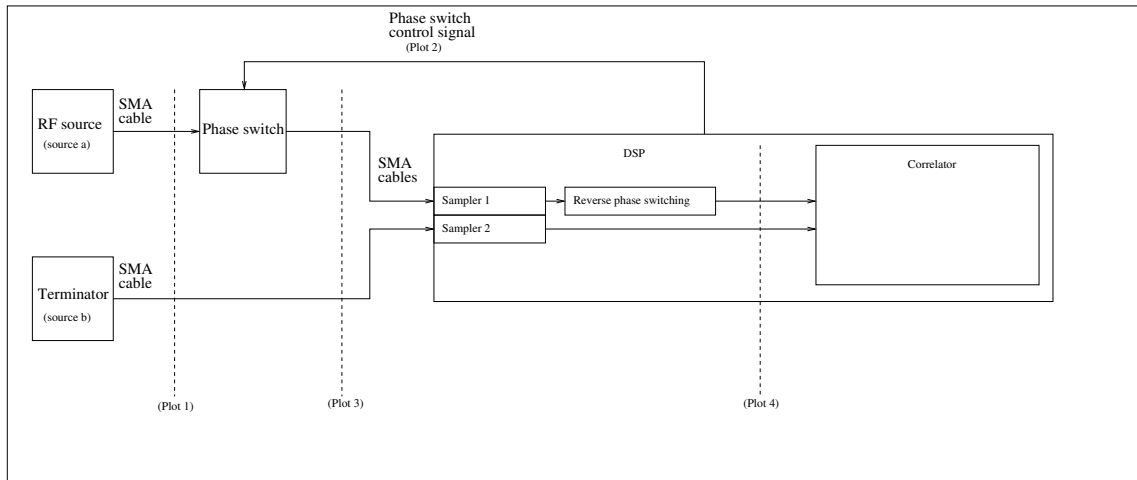


Figure 4.11: Example setup of a phase switched two channel system. The signals used in the plot of Figure 4.12 are shown.

is supplied by source a and no signal is applied by source b. After sampling and phase switch demodulation has been performed, sampler 2's output is entirely comprised of a cross-coupled term, which alternates between correlating and anti-correlating with the signal from sampler 1. This means that when the integration is taken over a whole period of the phase switching signal, no net correlation will be seen. The above applies for any selection of input signals on either channel with no coupled term correlating consistently with the signal that generated it.

When phase switching is required for more than two inputs, the selection of the switching signals becomes important. The signals usually used are a set of functions called Walsh functions. These are an orthogonal set of two-state signals, the first four of which are shown in Figure 4.13. If each input is phase switched using a different Walsh function then all cross talk terms between all pairs will cancel.

When phase switching was introduced into the system, the previous setup of Figure 4.11 was modified to that of the top layout of Figure 4.14 with phase switching only being used on one arm, using the first two Walsh functions. When the correlation of the noise sources was again measured a substantial reduction in the level of cross-correlation was seen, however, a small component still existed. This was shown to be coming from the RF link between the two arms, as the signal sensitivity is now such that even this dramatically lowered signal can be detected. To finish the setup and completely isolate the two arms the setup was modified again to that shown in the middle setup of Figure 4.14 where amplifiers and attenuators have been added to the RF link. These were used as they show a large attenuation for reverse signals while amplifying the forward transmission.




Figure 4.12: Example of how phase switching can work on a two channel system where one channel is fed with an RF signal. Note in the last plot how the cross coupled signal switches between correlating and anti correlating with its forming signal. This means after integrations of whole numbers of phase switching periods, no net correlation will be seen.




Figure 4.13: The first four Walsh functions.

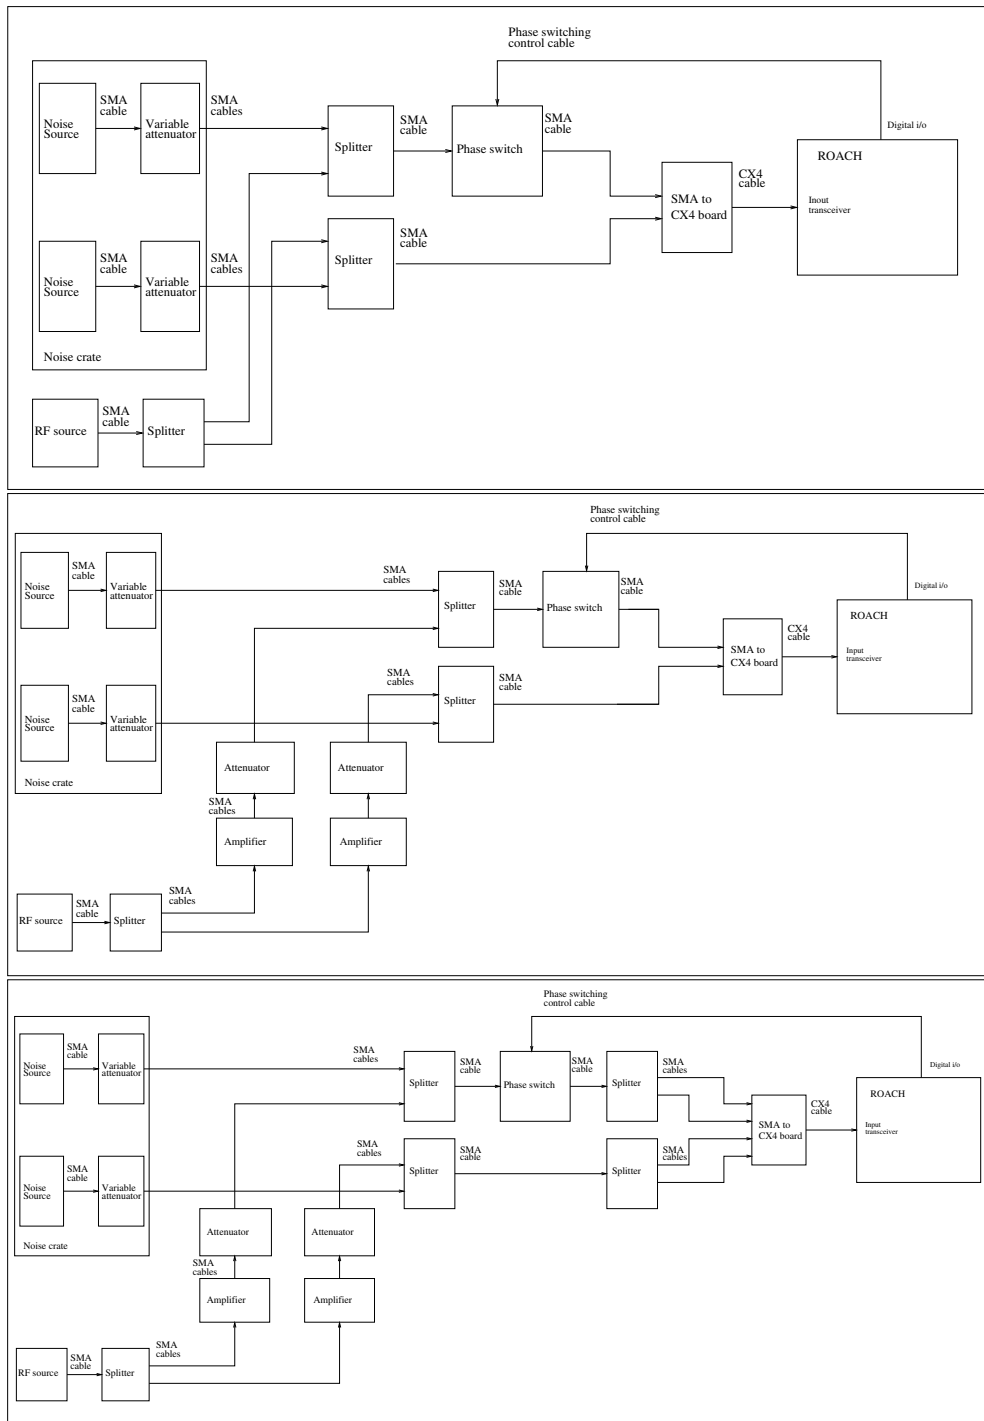


Figure 4.14: The evolution of the RF setup used to detect a low level RF source signal from within a noisy environment using correlation techniques. Unwanted correlations were experienced between the noise signals, these were cancelled in the top setup using phase switching. It was then found that the RF source injection was allowing noise leakage between the two arms. This was corrected in the middle setup using amplifiers to isolate the two RF paths. The bottom setup shows the additional splitting of the signals to drive 4 transceiver inputs ready for interleaving.

4.4 Testing interleaving using phase switching

After the unwanted correlating components were removed the system was tested using single-frequency input signals. The frequency was slowly increased with the signal power remaining constant, the correlator output was taken and an FFT performed to obtain the spectrum. It was found that signals could now be detected into the third Nyquist zone ie up to 4.6875 GHz. This shows a considerable improvement over the original single-frequency experiments where it was not possible to detect signals above about 1 GHz. This highlights the ability of correlation-based techniques to greatly improve the sensitivity of experiments.

Now that phase switching and the associated correlator response were well understood it was possible to revisit the problem of interleaving. The bottom setup of Figure 4.14 was constructed that employs interleaving techniques as well as phase switching. The input pairs were anti-aligned using the previously described method and data captured for a range of frequencies of input signals. Signals at 100 MHz, 200 MHz, 2925 MHz and 3025 MHz were injected into the samplers, with all cross-correlation pairs being calculated by the digital lag-correlator within the FPGA. The non-interleaved spectra were then recovered using an inverse Fourier transform of the cross-correlation between the two signal arms. The cross-correlation outputs were also combined, as described in Section 6.8, to generate a single interleaved-correlator output, with an inverse Fourier transform again being used to recover the spectra.

The plots of Figure 4.16 show the obtained results. Peaks can clearly be seen in the spectra, corresponding to the input-signal frequencies. Secondary peaks are present, and can be seen in both the non-interleaved and interleaved spectra. These are therefore being caused by spurious correlations within the system, and are not by the interleaving method discussed. The experiment was designed to test the distinction between signals in the first and second Nyquist zones for the interleaved and non-interleaved systems. The simulation of this effect can be seen in Figure 4.15 for two samplers, one at 3.125 GS/s and the other interleaved up to 6.25 GS/s. If sampling at 3.125 GS/s is considered, ie with Nyquist frequency $f_N = 1.5625$ GHz, then signals at f (eg 200 MHz) and $2f_N - f$ (eg 2.925 GHz) can not be distinguished from each other. With an interleaved system, with an effective sampling rate of 6.25 GS/s and $f_N = 3.125$ GHz, both are detected at their correct frequency.

The plots of Figure 4.16 show the second Nyquist zone frequencies being clearly distinguishable from their corresponding first Nyquist zone frequencies when interleaving is used. This shows that the interleaving technique adopted is working.

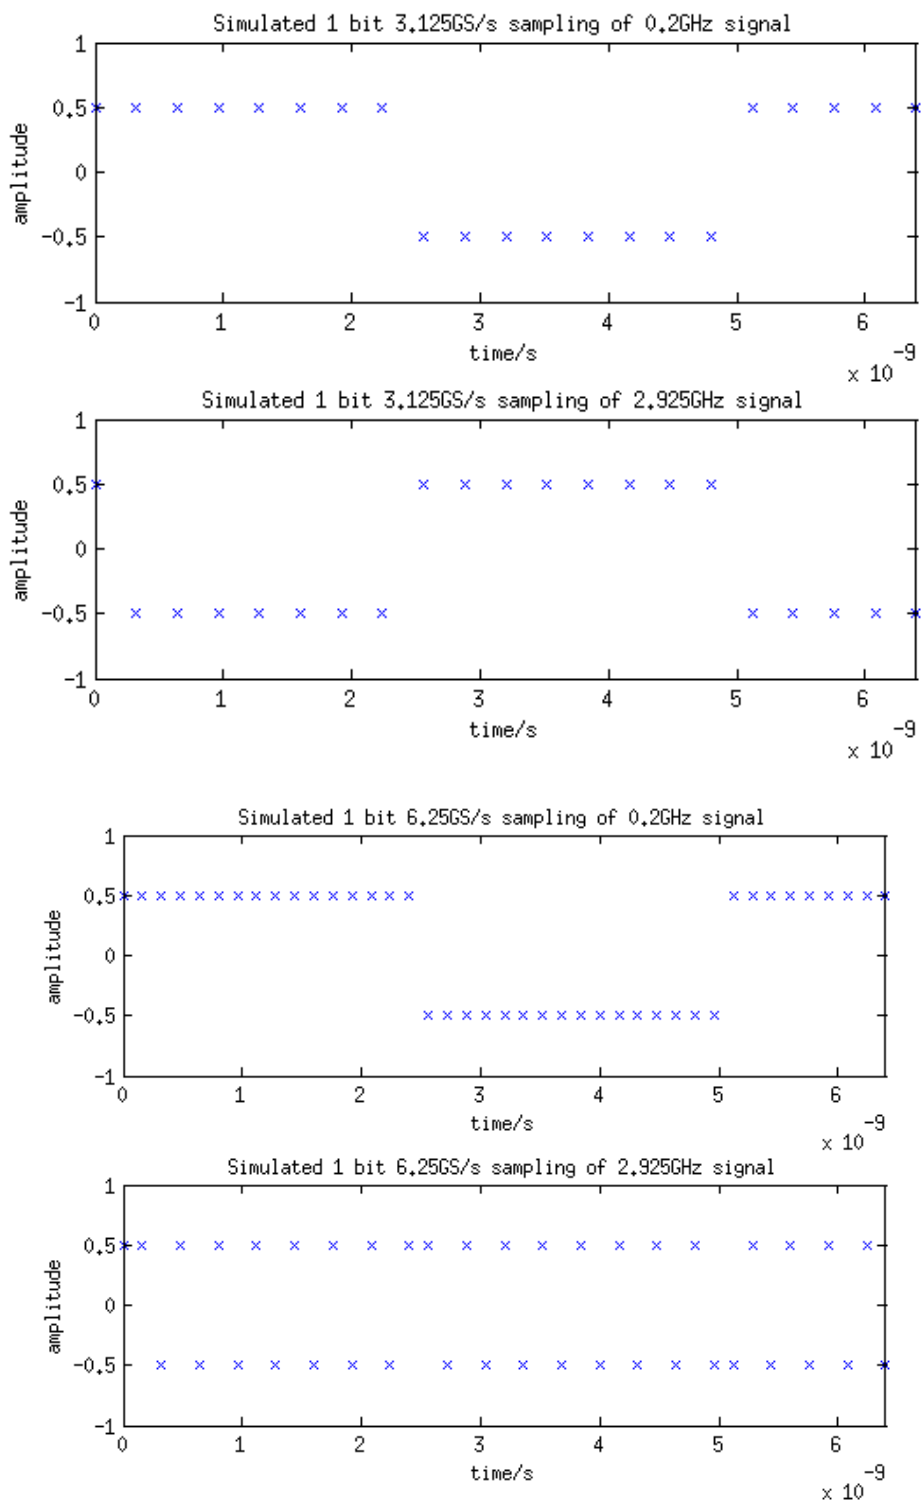


Figure 4.15: Simulation of two RF signals being sampled by a 1-bit (top) 3.125 GS/s sampler and (bottom) 6.25 GS/s sampler. The bottom case can clearly distinguish the two different frequency signals as its sampling rate is double that of the top.

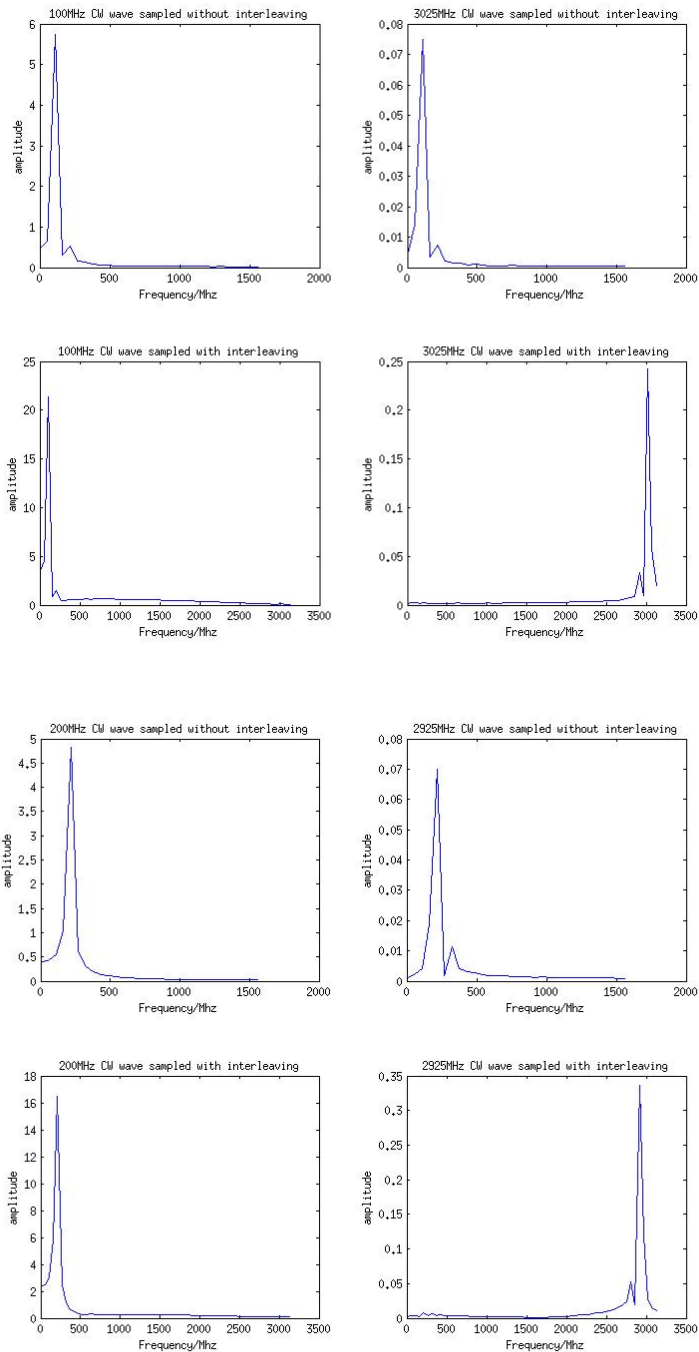


Figure 4.16: Comparison between an interleaved and non-interleaved system when fed (Top) 100 MHz and 3025 MHz signals and (Bottom) 200 MHz and 2925 MHz signals. The interleaved system can distinguish the two signals where as the non-interleaved system can't.

4.5 Conclusion

A technique can be seen to interleave samplers together to sample at a higher rate. The method used shows the ability to measure the relative sampling positions of different comparators and move them into alignment or anti-alignment. This was shown to allow interleaving to be performed, with the identification and differentiation of signals in the first and second Nyquist zone of a single sampler.

Chapter 5

Digital Sampling techniques: Multiple-bit techniques

The investigation in the previous chapter has demonstrated the use of GTP transceivers as ultra-fast 1-bit samplers. However, many radio telescope applications require a multi-bit sampler to be used, to increase RFI immunity and also increase the observing efficiency. A common approach in high-speed sampling to produce a multi-bit sampler is to use multiple 1-bit samplers in parallel[17]. This is the so-called 'flash' architecture, where all the bits of the conversion are obtained simultaneously, in contrast to other ADC architectures that use fewer comparators but taken a finite time to perform the conversion. This flash approach can also be applied to the GTP transceivers being used here and will be discussed in this chapter.

5.1 Multi-bit techniques

Two techniques were identified to allow a multi-level sampler to be produced. Both require multiple transceivers to be fed the same RF signal, with their clocks synchronized. The first method was to shift the reference voltage applied to the differential inputs of the transceivers. The transceiver inputs act as 2-input differential comparators, therefore if the second input, previously tied to ground, is set to a threshold voltage the comparison will now be made with respect to the new level. If multiple transceivers are set up with different threshold levels, as in Figure 5.1, then a multi-level signal can be constructed based on a combination of the outputs. This can be seen in the simulated plots of Figure 5.2, where a 100 MHz RF signal is sampled by three transceivers operating at different thresholds. When the signals are combined a 2-bit representation of the original signal is constructed.

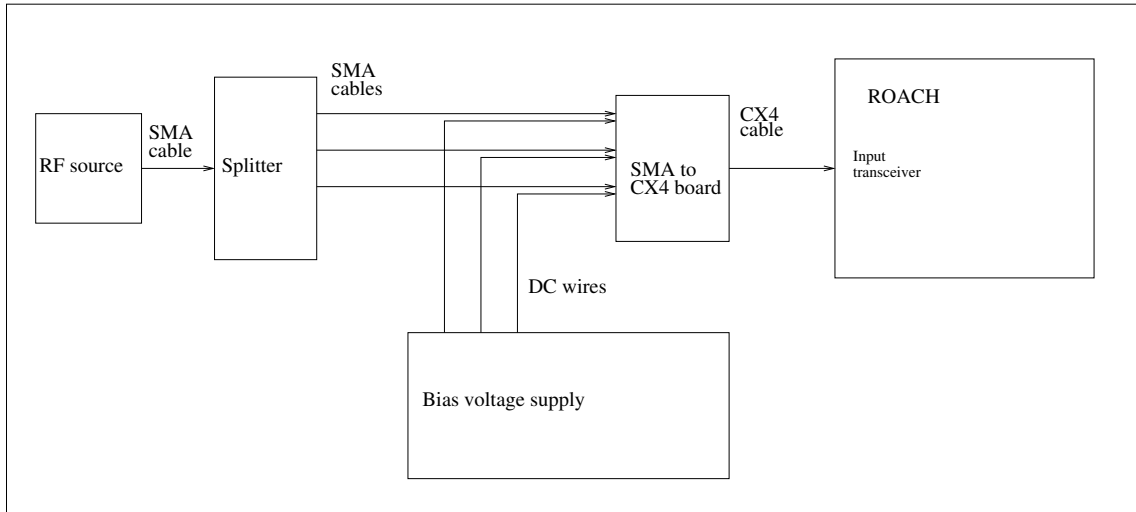


Figure 5.1: RF setup for initial Multi-bit tests. An external bias voltage supply is used to supply different reference voltages to each of the differential transceivers.

This technique can theoretically be cascaded together to perform sampling at any bit accuracy required by the system, at the cost of more transceivers being required.

The second approach considered to multi-bit sampling was the use of hysteresis effects within the samplers. All previous characterisation of the samplers has assumed that the output switch-over point occurs exactly when the two RF inputs are equal. In reality all systems will show a level of hysteresis, which means that the transition point for positive-going output change occurs at a different point to that of a negative-going change. In the case of a CW input signal, the effect of hysteresis in the sampler is simply to shift the phase of the output, as seen in the upper plots of Figure 5.3. If a more complicated signal, such as in the lower plots of Figure 5.3 is used, the differences between the hysteresis and non-hysteresis cases become more apparent. These examples used a threshold level of 0.7 of the peak signal amplitude, which is a considerable fraction of the RF signal level. If the RF power is increased such that the amplitude of the RF is much greater than the threshold, the hysteresis sampler output approaches that of the ideal case, which was the situation in the previous experiments.

To construct a multi-level sampler using the hysteresis effect, different transceivers are fed with the same RF signals but at different signal powers. In the case of a three level sampler, two samplers are configured with synchronised clocks but differing RF signal powers. The signal powers are chosen to be of the same order as the hysteresis threshold level for the system, but with one transceiver receiving a higher signal power. The outputs of the two transceivers are then combined by adding




Figure 5.2: Example of 2 bit sampling of a 100 MHz RF signal using three transceivers with different bias settings.

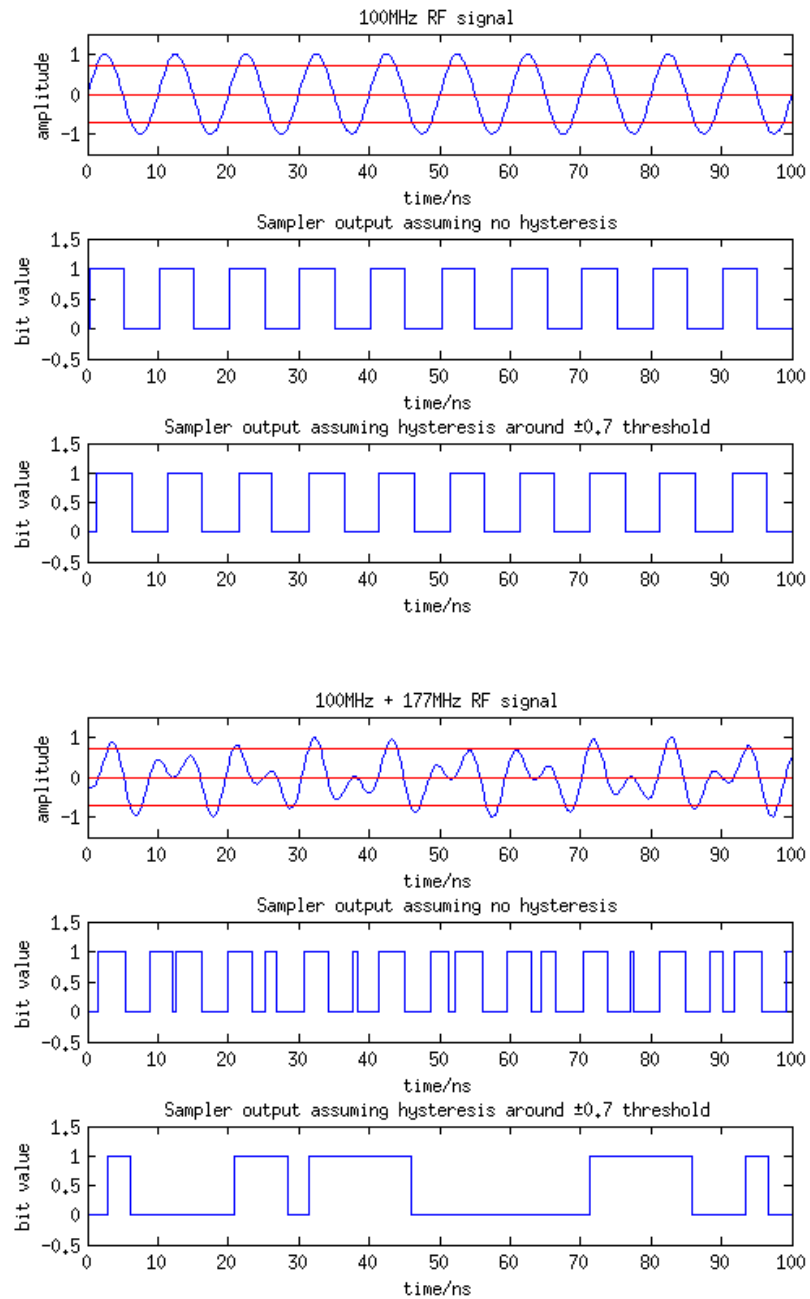


Figure 5.3: Example of theoretical sampler output with and without hysteresis. The top plot shows the outputs for a simple single-frequency input. The bottom plot is for a more complicated input signal. The red lines show the threshold levels used.




Figure 5.4: Example of theoretical sampler output with hysteresis used to generate a 1.5 bit sampler. The red lines show the threshold levels used.




Figure 5.5: Theoretical variation of the ENOB of a 1.5 bit sampler when the thresholds are varied.

the individual transceiver data streams to construct a multi-level output, and then weighting the output so that it is symmetric about zero. This is shown in Figure 5.4. As with the previous method, the above approach to multi-level sampling can be extended using more transceivers and different power levels to generate any bit-depth required.

The two above methods both theoretically allow complete control over the sampling levels. This introduces the question as to what the best values for these levels should be. To check this a parameter space exploration was performed using a 1.5 bit sampler (3 levels), where the values of the threshold level was varied. In a symmetric 3-level system there is only one adjustable parameter, which is the threshold level. The ENOB for the system was calculated as a function of threshold level with the aim of maximising the ENOB. The results of this can be seen in Figure 5.5, showing that it is possible to generate the maximum sampling at the appropriate threshold level, but that it is also possible for the ENOB to fall below that of a single bit sampler with a poor choice of threshold. An easy example of this is if the threshold level is moved above the peak of the CW signal. In this case the samplers do not trigger, and no signal information is contained in the output.

5.2 Hysteresis-based multi-level sampling

The first of the methods for multi-level sampling that was investigated was the hysteresis-based sampling. The most obvious sign of hysteresis in the system discussed so far is the minimum power level required for stable sampling. If the samplers were true differential comparators referenced from 0 V, then any signal level would cause switching in the output. In all previous experiments the effects of this hysteresis have been reduced by ensuring the signal power level well exceeds that required to start triggering, which means that the threshold level approaches the zero crossing on the RF waveform and the hysteresis in the system could be ignored. To test the effects of the hysteresis and develop a hysteresis-based multi-level sampler, the signal levels chosen must be approximately equal to the threshold level of the sampler to ensure hysteresis is a significant effect.

The initial test that was performed to investigate the effects of the hysteresis on the sampling was to sample a broadband noise source with a single transceiver, capturing raw data and plotting a spectrum. The noise power was slowly ramped up from below the sampling threshold to above it, with spectra being plotted at each stage. As the noise power is increased, the largest voltage spikes cross the triggering threshold first, and as these are widely spaced they appear as a low frequency component in the spectrum. As the power level is ramped up more of the spikes cross the threshold and the spectrum increases in power with more high frequency components. If just the average power in the spectrum is considered, as opposed to the entire shape, the simulation result of Figure 5.6 is seen, where a smooth transition occurs between the non-triggering and fully triggering regimes.

To perform this experiment using the ROACH setup the broadband noise sources previously described were used. These have built-in digital attenuators with a resolution of 1 dB, which therefore provides the limit of the granularity of the power axis. The results from this experiment can be seen in the bottom plot of Figure 5.6, which shows a comparable response to that of the simulation. The x axis of the plot is not calibrated in absolute power. However from the experiment it was possible to discern a scale over which hysteresis effects are visible. This was found to be within 6 dB of the sampling threshold. For future experiments it will therefore be necessary to remain in this 6 dB range, as beyond this range hysteresis effects become negligible.

The next stage of testing was to exactly determine the threshold point for triggering of the samplers. To do this a single-frequency signal was injected with the power slowly ramped up, as before. Unlike the case of broadband signals, in the case of




Figure 5.6: (Top) Simulated and (Bottom) measured average spectral power from a sampled noise source as the noise power is increased through the threshold level.




Figure 5.7: Voltage Variable Attenuator board: This is a DC-20 GHz attenuator controlled by a DC voltage input.

a single-frequency input the switch-on point is much sharper, with an ideal sampler having no transition period between the non-triggering and triggering states. Once triggering is reached for a single-frequency signal, any further increase in power only results in a phase shift of the output signal relative to the input, as the triggering point moves from the crests of the sine wave towards the zero crossing. This phase effect will be used for the next experiment; however it was first important to accurately measure the triggering threshold.

In the previous broadband experiment, the in-built digital attenuators were used to provide the power attenuation, providing a resolution of 1 dB. This was too coarse for this experiment as the entire hysteresis effect occurs over just 6 dB and therefore in six steps. To improve on this power resolution a voltage-variable attenuator, constructed from a Hittite HMC346LC3B chip¹ (shown in Figure 5.7) was used. This was developed by a previous member of the group and provides variable attenuation based on an input control voltage for any signals between DC-20 GHz. To control this, a digital-to-analogue converter board with an SPI interface was used, controlled using an Arduino microcontroller and allowing the power level to be set with a resolution of approximately 0.1 dB.

When the single-frequency experiment was performed the result shown in Figure 5.8 was achieved. From this it is clear that the switch-on transition occurs much faster than in the broadband noise case, as expected. The transition is not instantaneous, as imperfections in the sampler and signal source, as well as noise effects, will blur

¹HMC346LC3B (<http://www.hittite.com/>)




Figure 5.8: Plots of average spectral power for input single-frequency signal as the power level is swept through the triggering threshold. The x-axis is uncalibrated and represents the attenuation level. As can be seen, for above ≈ 30 attenuation the samplers are not triggering.

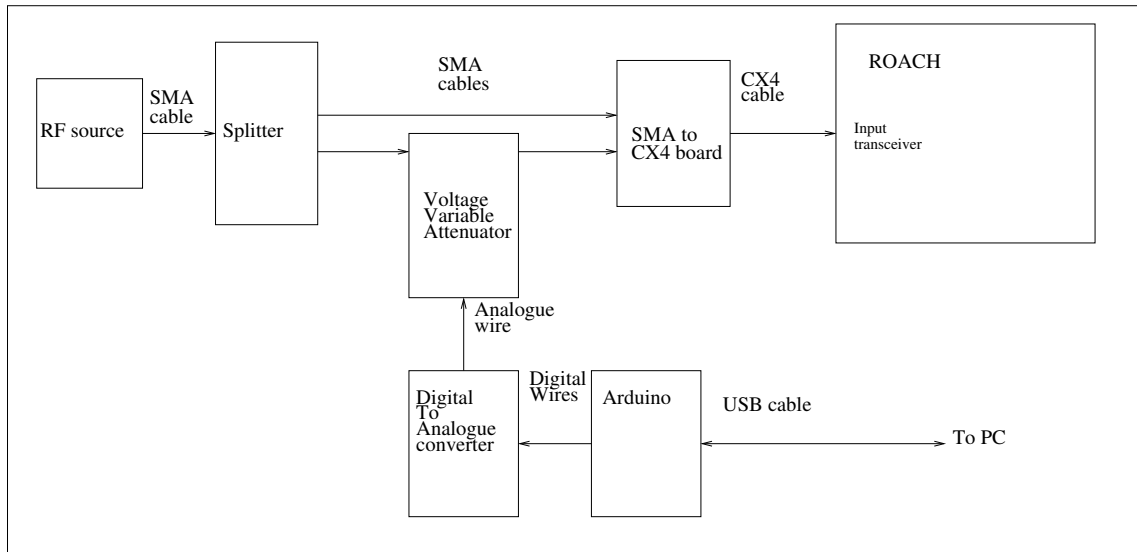


Figure 5.9: RF setup for phase difference tests of the hysteresis multi-bit technique.

the transition.

Now that the transition point for the samplers has been well-established, the phase-variation effect of hysteresis sampling with single-frequencies can be observed. This is the effect that is used to generate multi-level sampling and is therefore very important for the production of any multi-bit samplers. To test this effect the setup of Figure 5.9 was constructed, where two inputs are fed with the same RF signal, one kept at constant power and the other varied. The constant power level is chosen to consistently trigger the sampler and the other power slowly increased through the threshold point. The relative phase between the two sampler outputs should be seen to move through 90 degrees, as the second sampler starts triggering on the crest of the sine wave and moves through to the zero crossing. A simulation of this effect can be seen in the top plot of Figure 5.10, with the actual results obtained shown in the bottom plot. As can be seen the shape of the graphs agrees however the extent of the phase change is much lower in the real experiment. This suggests that the sampler is not behaving like the simulated case, with the hysteresis being less defined. Although there is a minimum power level to activate the samplers, above this level the sampler demonstrates very little hysteresis.

As only a very small phase variation was observed, no further testing into the use of hysteresis as a multi-level technique was carried out. It is clear from simulations that a much larger phase change effect would be required to produce any discernible bit-depth improvement over the 1-bit case.




Figure 5.10: Plot of (Top) simulated and (Bottom) measured phase difference between two samples of the same 100 MHz RF signal, fed into two samplers at different power levels. The power level into one of the samplers is kept constant with the other varied. In the simulation, as soon as the signal begins to trigger the samplers the phase difference is 90° , and this phase difference reduces towards zero as the RF power level becomes large relative to the threshold. In the measured case, there is only a phase shift of a few degrees over the entire power range.

5.3 Reference-based multi-level sampling

As discussed in the introduction to multi-level sampling, the reference voltage on the second of the differential pair pins can be used to generate different sampling levels and therefore a multi-bit sampler. A problem arises when implementing this technique with the current system as the inputs to the ROACH board are AC coupled (a small capacitor is placed in series with the input line to block DC signals and act as a high-pass filter). For the previous experiments this has not presented an issue, as all the signals used have been of sufficiently high frequency to pass through unimpeded. The idea behind referenced-based multi-level sampling is to feed a DC voltage into the second input, but this will obviously be blocked by the coupling capacitor and the differential pair will still see a 0 V reference independent of the voltage applied to the input. To overcome this, the ideal approach would be to modify the hardware such that no AC coupling is present, at which point a DC signal could be applied with the desired affect. As this requires a hardware modification to the RF signal path for the board, an area sensitive to changes in line impedance, it was deemed appropriate to first investigate the feasibility of the overall technique.

5.3.1 Low-frequency square wave bias

Another approach exists that allows the principle behind this sampling method to be tested without any modification to the hardware environment used. This is to apply a low frequency square wave signal to the second input of the differential pair, instead of using a DC reference voltage, such that the frequency of the square wave f_{sq} is much less than that of the RF f_{RF} . It is also necessary to make the square wave frequency larger than the cut-off frequency for the high-pass filter formed by the AC coupling f_{HP} such that

$$f_{HP} \ll f_{sq} \ll f_{RF}. \quad (5.1)$$

This will allow the reference voltage to pass through the AC coupling, but the sampler still see multiple periods of the RF signal with a constant reference voltage.

The first test that was performed with this setup was to input a constant level 100 MHz RF signal into the sampler at a level sufficient to just trigger the sampler. The reference voltage was fed with a square wave, initially at 100 kHz, and with a voltage level between zero and that required to trigger the sampler without the RF input. Multiple reference amplitude steps were taken in this range, with direct data being captured. The data were then analysed to look at the mark-to-space ratio of

the signal received. This is a measure of the relative time for a signal of period P_{total} spent in the high state P_{high} as compared to the low state P_{low} such that:

$$P_{total} = P_{high} + P_{low} \quad (5.2)$$

Figure 5.11 shows the theoretical result of having a square wave reference voltage in addition to a constant frequency RF signal, with diagrams showing how this effect occurs. As can be seen, the mark-to-space ratio of the output follows the reference square wave level, and it is therefore possible to verify that the reference input is being used in the comparison by analysing the mark-to-space ratio results. Figure 5.12 shows the actual results obtained for an increasing amplitude of reference square wave. These show a substantial difference to the theoretical model of Figure 5.11. Note a saturation value of ± 0.5 exists as when the bias wave triggers the samplers entirely the binned sampled output matches the bias input with no RF frequency oscillations. Although there is a step in the recovered data on each transition of the bias state as expected, this is then followed by an exponential decay. This discrepancy is caused by the presence of the AC coupling capacitor; the exponential decay is a capacitor discharge curve and indicates that the 100 kHz input frequency was below the cutoff for the high pass filter.

If the shape of the plots is ignored it can be clearly seen that the desired effect is occurring, as the peak value of the mark-to-space ratio change scales with the amplitude of the square wave reference signal, as would be expected. This indicates that if the correct frequency was used, or if the AC coupling was removed, the sampler would indeed reference off whatever reference input signal was used. To further test this, the frequency of the bias signal was increased to 1 MHz. This is as high as can be produced using the signal generator used in this setup. It also allows multiple bias signal periods to be captured in a single data grab. This allows the full form of the wave to be observed, and permits a better measurement of the variation in mark-to-space ratio with bias signal amplitude. Figure 5.13 shows a binned captured data plot and the maximum mark-to-space ratio against applied signal respectively. These show the expected relation of the mark-to-space ratio to the bias amplitude, up to a value of 7, where the sampler saturates.

5.3.2 Removal of the AC coupling

From the above tests it was decided that the hardware modification of the ROACH board was justified. The AC coupling on the transceivers is provided by small surface-

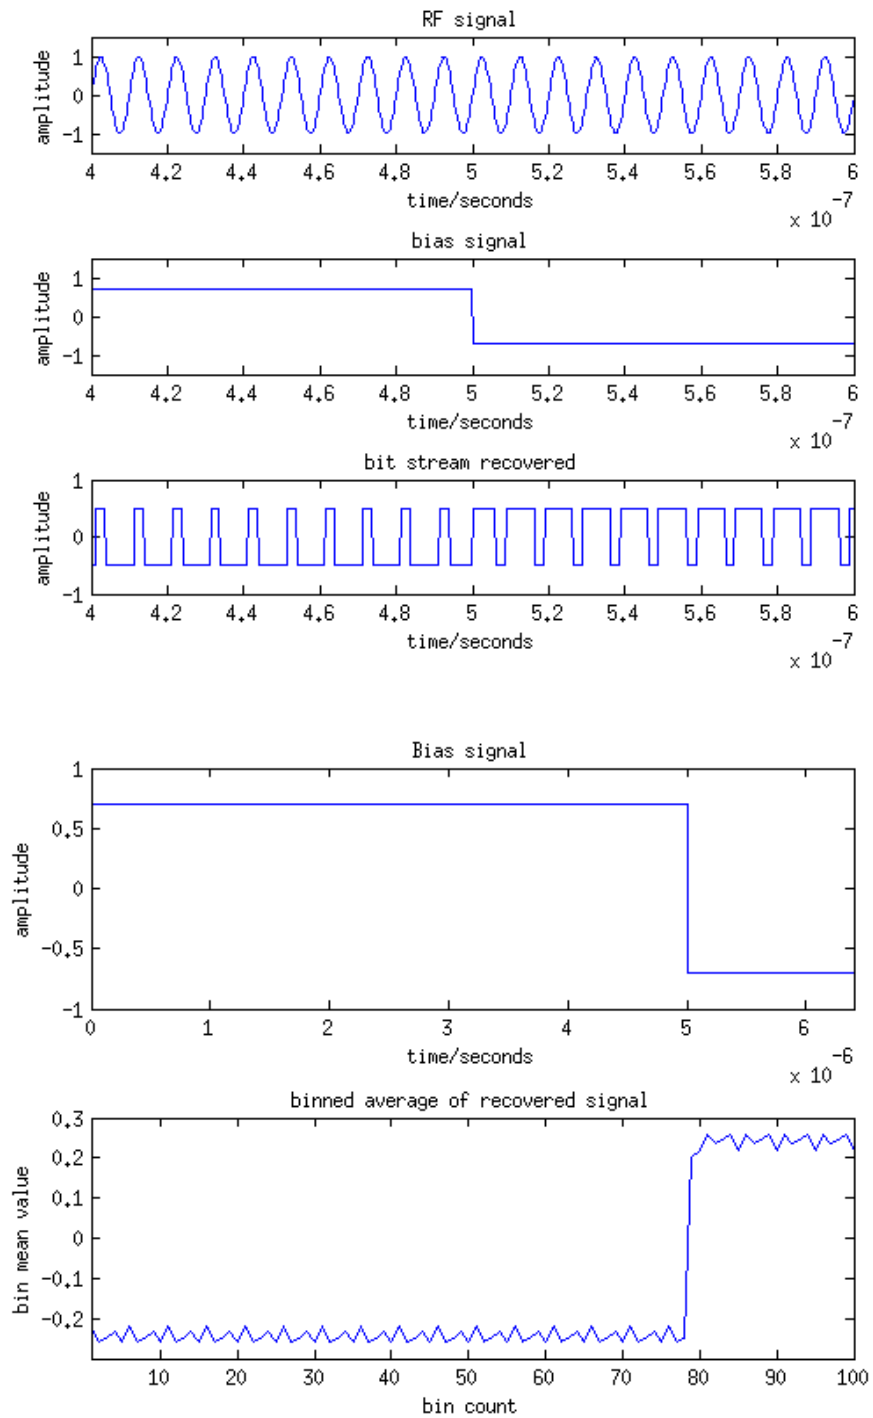


Figure 5.11: (Top) Plot of RF input, bias signal and recovered bit pattern for simulated data. As the bias signal changes sign, the mark-to-space ratio of the recovered bit stream changes. (Bottom) Plot of average value of the binned recovered signal for simulated data. This shows the variation in the mark-to-space ratio as the bias signal changes sign.




Figure 5.12: Plots of average value of binned recovered data, for real data and a low bias signal, for three different bias values. The steps correspond to the switching of the 100 kHz bias signal.


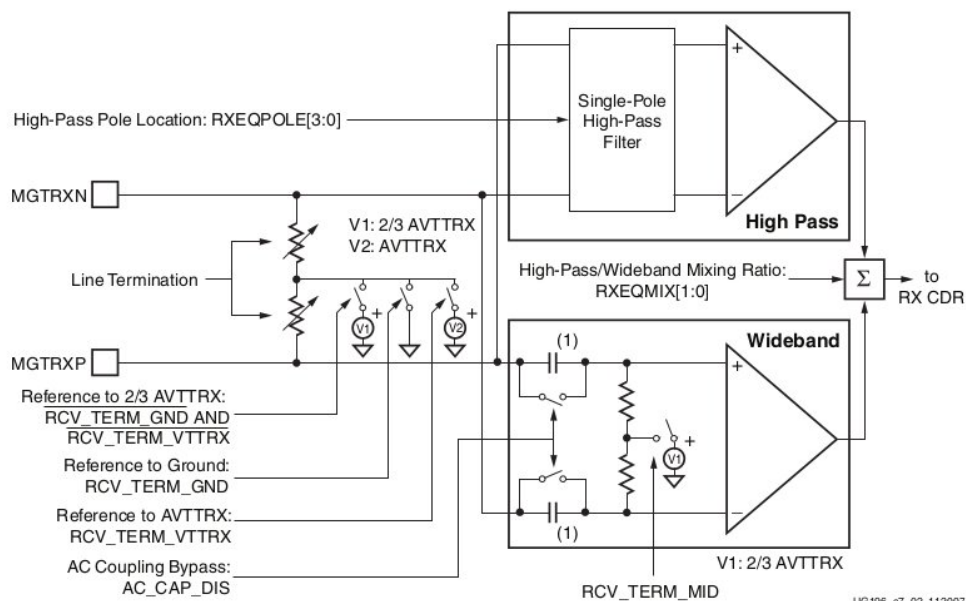


Figure 5.13: (Top) Average value of the binned recovered signal for real data and 1 MHz bias signal frequency. (Bottom) Peak value from binned plots as a function of bias amplitude. The bias amplitude was not calibrated, and therefore no units are given.



Notes:

1. The internal AC coupling capacitor is not a replacement for offchip coupling capacitors. It is an isolation capacitor for nonstandard termination voltages.
2. These port names have the prefix MGT to identify them easily in a pad file that is very often used to create symbols for board design schematics. In this document, the MGT prefix was removed from those names. However, names with and without the MGT prefix are synonymous to each other.

Figure 5.14: Schematic of the receiver configuration [45]

mount capacitors, where a separate capacitor is used for each of the differential inputs into the transceivers. The setup adopted is to feed RF into one of the differential lines with a DC bias fed into the other. The bias signal requires DC coupling however the RF signal can be input with either. This means that only half of the AC coupling capacitors are required to be removed to provide the necessary functionality. It was decided, however, that the system would be made more simple if all the coupling capacitors were removed.

Before any tests on the new setup were performed the configuration of the transceiver, shown in Figure 5.14, was first considered. With the external AC coupling present, as in all previous experiments, the values for the differential line reference and internal AC coupling were not important. However now that the external coupling capacitors have been removed, the DC coupling characteristics of the board depend entirely on these settings.

The transceivers have two separate receivers in order to provide different frequency responses. One is a wideband receiver, the other has a high-pass filter, whose pole frequency is controllable using the RXEQPOLE control word. The outputs of these two receivers are added together in variable proportions, with the mixing ratio con-

trolled by RXEQMIX. This addition of different frequency responses allows the overall frequency response to be varied to equalise band slope introduced by the cables.

There are two sets of line termination resistors, one which applies to both receivers, and one only to the wideband receiver, both with optional termination at the centre tap. The common resistors are connected to either ground or one of two reference voltages. If RCV_TERM_GND is set, the resistor centre-tap is grounded; if RCV_TERM_VTTRX is set, it is connected to a reference voltage AVTTRX; and if both are not set then it is connected to $2/3$ AVTTRX. The termination resistor on the wideband receiver is controlled by RCV_TERM_MID: if set, the centre tap is connected to $2/3$ AVTTRX, if not set the centre tap is left floating.

Finally, there are AC coupling capacitors on the inputs to the wideband receiver that can be bypassed by asserting AC_CAP_DIS. There are no similar capacitors on the high-pass receiver as it is permanently AC coupled.

As DC coupling is required, the internal AC coupling capacitors were first disabled, and the differential pair reference voltage was then set to 0 V as this matches the input signal references. Once the system was configured in this state the effects of DC biasing could be investigated. This was done by inputting an RF signal into one input of the differential pair. To ensure that no DC offset existed on the RF an external AC coupling SMA connector was used. The second transceiver input was directly connected to a DC voltage source with the output value manually altered. Data were then captured, with the DC bias signal swept through a range of voltages from -500 mV to $+500$ mV. This range was selected as it covered the extremes found in the previous experiment for a bias square wave. The average value of the data capture was then taken and plotted against the input bias value. This plot, shown in Figure 5.15, shows the variation in the net DC offset of the received data stream. This is directly proportional to the mark-to-space ratio and shows the sampling reference variation. The expected result for this plot was to see a symmetric graph about $x = 0$ with a non-linear response. Instead, a highly asymmetric shape was seen. This suggests that the biasing of the transceivers is not functioning as expected. This was thought to be due to improper setting of the transceiver reference levels.

To investigate the effects of the different transceiver reference voltage settings, the ROACH was configured in each of the available options, and the voltage drop between a transceiver line to ground measured. When no input was applied it was expected that the transceiver input would be pulled up to the middle reference point of the sampler by the termination impedance, but when a DC signal was then applied over this, the voltage drop between line and ground should mirror the input




Figure 5.15: Plot of how recovered signal average value changed with applied DC bias voltage after the removal of the AC coupling capacitors.

bias voltage. All sixteen combinations of the configuration bits (RCV_TERM_GND, RCV_TERM_VTTRX, RCV_TERM_MID and AC_CAP_DIS) were tried.

The effects of the AC coupling capacitor were considered, as this allows isolation of the two reference voltage selectors, as seen in Figure 5.14. When the AC coupling is enabled the transceiver input is isolated from the effects of the RCV_TERM_MID setting by the internal AC coupling. Table 5.1 was produced using all of the above settings and measuring the voltage at the transceiver output line with a voltmeter.

These data seem to agree with the expected results, showing the different reference levels and how they are affected by the changing of transceiver settings. During this testing it was however discovered that if an external DC source was added to the bias line then the measured voltage drop did not behave as expected. It was described above that the external assertion of a bias voltage should over-ride the internally-set value. However when this was done, it was found that a minimum voltage existed, different for each transceiver configuration, that needed to be exceeded before the applied voltage was measured on the transceiver input. This shows that external DC biasing disrupts the existing biasing within the transceiver. It is therefore concluded that the circuit diagram of Figure 5.14 for the transceiver shows a idealised circuit. In reality the implementation of the resistances and the biasing of the comparators will

AC_CAP _DIS	RCV_TERM		Measured voltage V	Input resistors connected to/V	Wideband resistors connected to/V
	_GND	_VTTRX			
0	0	0	0.781	2.2	N/C
0	0	0	0.781	2.2	2.2
0	0	1	0.000	3.3	N/C
0	0	1	0.000	3.3	2.2
0	1	0	1.185	0	N/C
0	1	0	1.185	0	2.2
0	1	1	0.086	?	N/C
0	1	1	0.086	?	2.2
1	0	0	0.781	2.2	N/C
1	0	0	0.781	2.2	2.2
1	0	1	0.000	3.3	N/C
1	0	1	0.000	3.3	2.2
1	1	0	1.185	0	N/C
1	1	0	1.185	0	2.2
1	1	1	0.086	?	N/C
1	1	1	0.220	?	2.2

Table 5.1: Measurements of the voltage level at the differential output with respect to ground. The transceiver setting is shown. N/C means not connected and a ? is used to denote not defined.

take a more complicated form, which means that setting the DC biasing externally will not work as expected.

5.3.3 Re-introducing AC coupling

From the previous experiment it was clear that complications exist when using DC biasing into the transceivers. It was therefore necessary to develop an alternate approach to the reference-based multi-level sampling problem. From this the method that developed was to hybridise the DC biasing approach with the square wave approach used in the initial tests. To do this, external AC coupling was re-introduced into the system but at a much higher capacitance value, which had the effect of altering the cutoff frequency for the high-pass AC coupling filter to allow lower-frequency biasing square waves to be used within the system.

It was found that the original low frequency cut off generated by the AC coupling that was on the board was around 1 MHz. With a line impedance of 50Ω this allows the original capacitor value to be calculated as:

$$f_{break} = \frac{1}{2\pi RC} \quad (5.3)$$

$$\Rightarrow C = \frac{1}{2\pi R f_{break}} = \frac{1}{2\pi 50 \times 1 \times 10^6} \approx 3.2 \text{ nF} \quad (5.4)$$

To allow the injection of the bias square wave to be achieved by basic circuitry without RF considerations, and to allow standard analogue circuit components to be used, it was decided that a cutoff frequency in the order of tens of kHz would be required. If the cut off frequency was to be made 1 kHz then the capacitor value would be given by:

$$C_{new} = \frac{1}{2\pi R f_{break}} = \frac{1}{2\pi 50 \times 1 \times 10^3} \approx 3.2 \mu F \approx 3.3 \mu F \quad (5.5)$$

Now that a square wave bias is again being used it means that a sampler will spend half its time at one threshold level and then switch to the negative of the same threshold. To use this system in a multi-bit sampler, a second transceiver must be paired which is fed the same amplitude bias wave at a 180 degree phase shift. This results in at any one time one sampler working at a positive threshold and the other at a negative threshold. To allow the multi-bit stream to be produced from the transceiver outputs, the FPGA must know which transceiver has the positive threshold and which has the negative. This was achieved by feeding the bias square

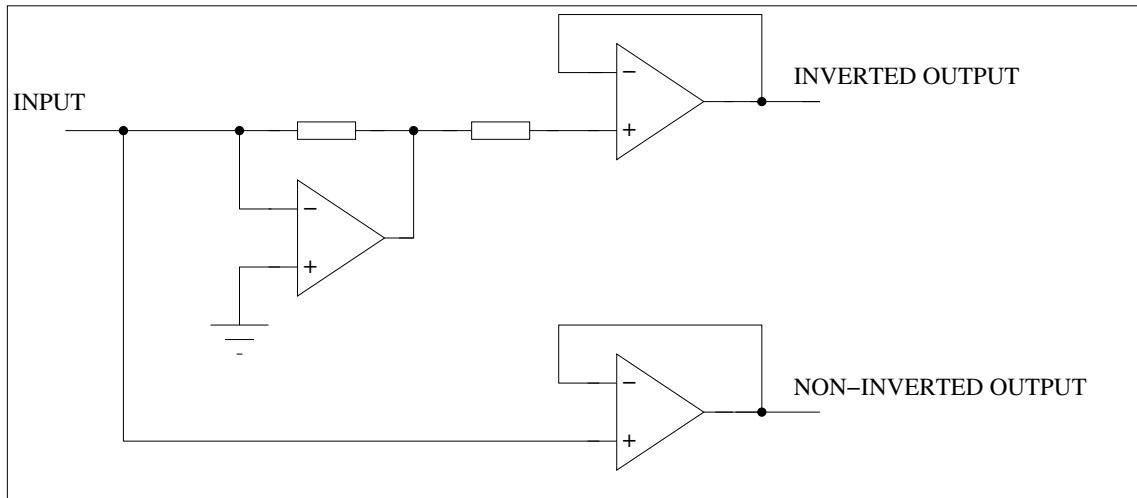


Figure 5.16: Circuit diagram for the bias signal inverter circuit. The op-amps used were contained in a single LM348N package.

wave into a digital input pin on the FPGA. This signal was then used to digitally switch the streams, so that after the switching code the data streams appear to have a constant threshold value.

The circuitry required to take a single square-wave signal and generate two 180-degree phase-shifted outputs is shown in Figure 5.16. A buffering stage was also added to the outputs to isolate the high-pass filter generated by the AC coupling from the effects of the phase-inverting amplifier.

The first test performed with the new system was a repeat of the measurement of a single-frequency RF signal, sampled with a bias square wave applied, and the internal switching of the FPGA turned off. This enabled the testing of the effectiveness of the bias square wave and allowed any capacitive effects to be seen. As with previous tests, the mark-to-space ratio of the output was analysed by looking at the mean value of the binned signal. The results from this can be seen in Figure 5.17, which shows the binned values with and without the bias inverting circuit. The plots show a prominent square wave shape, but significant transition regions can be seen when the bias inverter is used. This will have to be factored into future calculations with the introduction of a null period around switching built into the correlator. The transition effect seen is caused by bandwidth limitations in the bias inverter limiting the upper frequency allowed and removing high frequency components from the square wave.

Now that the basic operation of the bias square wave has been established, the next step was to take two inputs, synchronize their bit streams, and apply the switching technique to establish how effective the FPGA compensation is for the bias square




Figure 5.17: Binned recovered signal (Top) without and (Bottom) with bias inverter circuit after AC coupling removed.




Figure 5.18: Binned recovered signal with bias inverter circuit after AC coupling removed and FPGA signal switching

wave used. A direct-sample capture of a single bit stream after FPGA switching that has been put through the previous binned averaging is shown in Figure 5.18. The transition period can still be seen in the plot as no effort was made to null the output during this time. Also visible is the fact that, even though the two samplers are referencing from the same bias voltages, their responses differ slightly. This will add a low frequency component to all spectra produced that should be easy to remove.

5.3.4 Producing a multi-bit sampler

The next stage of testing was to take the two samplers that have been matched using the previously discussed method and feed a common RF signal into them. This was then directly and simultaneously sampled on the two transceivers and the resulting data used to produce a multi-bit sampled signal. Before this could be done consideration had to be given as to how the signals were combined and the effect of the bias amplitude on the output. To do this a simulation was constructed that was able to generate two sampled bit streams that were biased with plus and minus a given threshold, represented as a fraction of the RF sine wave amplitude. To combine the two signals the truth table of Table 5.2 was generated.

This data combination method is also shown in the simulated data plot of Figure




Figure 5.19: Simulated data showing how an RF signal is sampled by two samplers and the signals combined to generate a multi bit sampler

data in 1	data in 2	multi-bit data out
-0.5	-0.5	-1
-0.5	0.5	0
0.5	-0.5	0
0.5	0.5	1

Table 5.2: Truth table for how to combine two data streams into a single multi-bit stream.

5.19, which demonstrates how the two inputs are used to produce a single three-state output. This three-state output is equivalent to a 1.5-bit sampled signal. However the effective number of bits obtained will depend on the specific value of the threshold generated by the bias signal, as described in Section 5.1.

This can be seen in the limiting cases where the bias signal is set to 0 V, where the sampler works as a 1-bit sampler, and the case where the bias signal exceeds the RF signal, and therefore the output will constantly be zero. From this it is clear that the ENOB function will be peaked in the middle and lower at the extremes. To generate a better understanding of this, the simulation program was run and allowed to sweep through all the bias levels from 0 to 1, calculating the effective number of bits in each case. This resulted in the data shown in the top plot of Figure 5.20, which has a clear peak at a threshold value of 0.4, which is the fractional amplitude of the bias signal with respect to the input RF signal.

This test was then carried out using the real system, with the amplitude of the bias square wave being swept from 0 V to an upper limit determined by removing the RF signal and seeing at what input level the sampler was triggered entirely from the bias square wave. The results from this experiment can be seen in Figure 5.20. The shape of the plots matches that of the simulated plot in Figure 5.20, with a peak being seen around the 0.3 V pk-pk level. In the case of the low frequencies the peak is also seen to be very close to the theoretical maximum of 1.5 effective bits. This shows that the samplers can be run in parallel to generate a greater than 1-bit signal. Note that in practice a 1-bit sampler achieves an ENOB of ≈ 0.7 bits.

Now that a peak in the ENOB has been seen around 0.3 V pk-pk a full frequency sweep was performed with the bias set to this value, allowing the frequency variation of the bit efficiency to be shown. This can be affected by frequency-dependent noise terms as well as clock effects such as jitter. Before the actual frequency sweep was performed, a simulation was again run to determine the expected result. When this simulation was first executed a large amount of frequency dependence was seen with




Figure 5.20: (Top) Simulated and (Bottom) Real data for a single-frequency RF signal sampled by two transceivers working as a 1.5 bit sampler with the threshold level swept.




Figure 5.21: Simulated spectrum for a 1.5 bit sampled RF signal

a minimum ENOB value seen at around 700 MHz. After further investigation it was found that this frequency dependence was actually an aliasing effect with the data points being scattered between a maximum and minimum value. This was found to be due to the frequencies being used by the simulation not aligning with the FFT frequencies of the bins. It was initially thought that this issue would not arise as the ENOB calculation takes a large number of bins on either side of the peak position and defines all the associated power as being the signal. This should mean that if the signal is sitting between two bins, and has its power shared, the ENOB calculation should remain unchanged. When the frequency spectra were analysed the plot in Figure 5.21 was found. This shows how a very broad peak is seen with large tails, caused by the quantisation of the signals. To combat this broadening effect, and include all the signal power from the primary peak, the width of the signal bin used for the ENOB calculation had to be increased. This had the unwanted side-effect of incorporating more of the noise into the signal bin and therefore showing an artificially high ENOB value as well as a larger variance in the result obtained. It does, however, dramatically reduce the variance seen across the frequency plot of the ENOB for the simulated data. The simulated results for two different bin widths are shown in Figure 5.22. These have error bars included which represent the 1-sigma variance in the result value for repeat experiments.






Figure 5.22: Simulated ENOB frequency variation when the primary signal power is contained within (Top) 9 bins and (Middle) 61 bins. (Bottom) Real data ENOB frequency variation when the primary signal power is contained within 9 bins

	3.125 GS/s	6.25 GS/s	9.375 GS/s	12.5 GS/s
1-bit	1	2	3	4
1.5 bit	2	4	5	8
2 bit	3	6	9	12
2.25 bit	4	8	12	16

Table 5.3: Transceiver usage to achieve different sampling bit-depths and interleaving.

The above discussion of bin widths was considered when analysing the real system data taken. It was found, however, that widening the signal bin width produced a larger variation in the results than that seen in the simulated case. This is due to the increased noise level in the real system. The bottom plot of Figure 5.22 shows the result obtained. As in the simulated case, error bars have been added to show the 1-sigma variation in the result obtained. As with the previous experiments, the data used were automatically selected to avoid the bias switching point of the signal. A larger error value is seen compared to the simulated data, which represents the larger variance in the signal levels as well as the increased noise in the system.

5.4 Use of bit-depth and interleaving

In this and the previous chapter the practical use of multiple bit-depth sampling and interleaving have been shown. Both these approaches utilise multiple transceivers working in parallel to perform the given task. These two methods can theoretically be combined to generate a multi-bit sampled signal at higher than the Nyquist frequency for an individual sampler. This quickly becomes very costly in use of transceivers as each of the interleaved bits requires the bit-depth specified number of transceivers. Table 5.3 lists the number of transceivers needed to make a sampler of a given bit-depth and interleaved sampling speed.

For the purposes of testing an interleaving multiple-bit sampler the simplest form will be constructed. This is a 1.5-bit 6.25 GS/s sampler using four transceivers in the configuration of the top of Figure 5.23. When this system was tested an interesting complication was discovered. In all experiments conducted since the removal of the AC coupling capacitors, external AC coupling was provided for the RF lines using a separate SMA DC blocks. This was done as no DC bias effects were desired in the signal output, and it was considered important to ensure that no biasing effects were being caused by the RF circuits. In Figure 5.23 the DC block was again added, however it was used before the 4-way splitter. This splitter is a passive circuit that consists of a split track with matching circuits to minimise reflections and improve

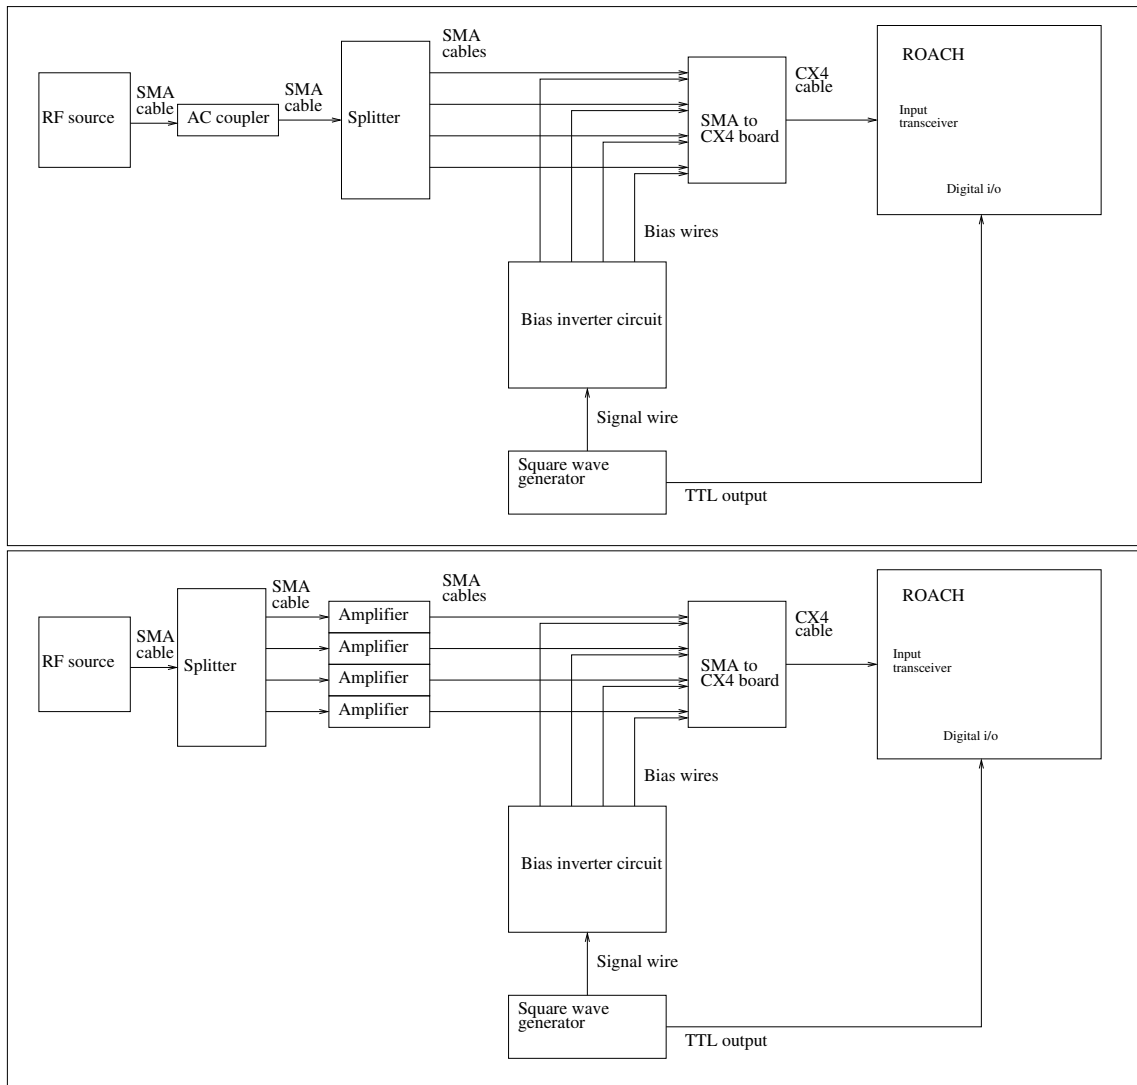


Figure 5.23: (Top) Initial and (Bottom) final RF setup for interleaving and multi-bit sampling.

impedance matching. It was therefore assumed that no DC effects would be added and the transceivers would operate as expected. After the setup was constructed it quickly became apparent that the splitter was providing a DC signal path between the transceiver inputs causing a complete obstruction to RF sampling, with the transceivers outputting a constant stream of either 0's or 1's. This may have been caused by the fact that no consideration had been made as to whether the positive or negative of the differential lines are being used, which means that opposite differential lines could have been connected. To combat this issue the system layout was modified to that shown in the bottom of Figure 5.23, in which amplifiers are used as DC blocks after the splitter. This approach was adopted as only two SMA DC blocks were available for use at this time which would not allow all the channels to be isolated as desired. The low-frequency amplifiers used have AC coupling inbuilt into both the input and output, to allow the DC biasing of the amplifier chip not to spread through the RF chain, which can also be used in this case to stop any DC biasing of the transceiver RF line from propagating through the RF splitter. Once this setup was adopted normal sampling resumed and the experiment could continue.

The first stage of the experiment was to perform an alignment of the channels to get the sampling points correct. This was done with a broadband noise source being injected into the splitter, and the biasing square waves turned off so all the transceivers are sampling with respect to the same thresholds. A python script was then run, similar to that explained in the previous sections on alignment, where channels 0 and 1 are first aligned by looking for a peak in the correlator output. Channels 0 and 2 are then put in anti-alignment with finally channel 3 being aligned to channel 2. This results in channels 0 and 1 sampling at the same point and 2 and 3 also acting as a pair. Once this alignment was complete the noise source was disconnected and the RF signal source attached and set to 100 MHz with the minimum power level to cause consistent triggering. At this stage the bias signals were attached. Channels 0 and 1 were fed with opposite sign bias square waves, as were channels 2 and 3. The relative selection of biasing between pairs 0 and 1 and 2 and 3 was not important, as when the multi-bit wave forms are generated, importance is only placed on whether the two channels being combined are the same or different states. As with the previous experiments on multi-bit sampling, the bias signal was set to 0.3 V pk-pk and a frequency of 25 kHz. The RF signal generator was then swept from 100 MHz to 1 GHz in 100 MHz steps, with the signal power level being manually rebalanced at each stage to stabilise triggering. From the data obtained the ENOB as a function of frequency was first calculated for the non-interleaved pairs of channels 0 and 1, and

channels 2 and 3. The results from this should coincide with those obtained by the straight multi-bit experiment previously described. The results from this can be seen in Figure 5.24 and show a slightly lower response than previously seen. This will be discussed below. The signals of all four channels were then interleaved together and the ENOB for the full system calculated. This is shown in the bottom plot of Figure 5.24, where a slightly lower ENOB is seen at each frequency than the individual channel pairs. This is being caused by an increase in signal corruption caused by the interleaved streams not matching perfectly, a trade off that will be seen if higher sampling rates are desired.

As can be seen the ENOB response for the pairs of channels set as 3.125 GS/s 1.5 bit samplers is much lower than the response seen in the previous section on multi-bit sampling. To trace the source of this error further measurements were made using the same setup of Figure 5.23, however samples at each frequency were taken for the four transceivers with the biasing voltages both on and off. This enables the sampling outputs to also be paired together to make 1-bit 6.25 GS/s interleaved signals as well as straight 3.125 GS/s 1-bit signals. The ENOB for each of these were calculated and it was found that the ENOB for the individual samplers capturing the signals was much lower than the original tests made of the sampling setup. Two possible reasons exist for this reduction in signal-to-noise, the first of these is the removal of the AC coupling capacitors in the RF circuit. The capacitors were replaced with straight 0Ω connections. This was not believed to have a detrimental effect of the signal performance but may be having an effect on the impedance matching of the input. The second possibility is the corruption due to cross talk. It has already previously been established that cross talk is a prominent effect in the system. When the original ENOB measurements were made for the individual samplers only one sampler was being fed RF. This means that no cross talk will of been present in the signal. In the latest multi-level interleaving experiment all four transceivers were being fed an RF signal, this could be greatly reducing the ENOB of the system. Both these situations are limited by the current hardware platform and would be dramatically reduced in any custom built FPGA board.

5.5 Conclusion

The use of the direct samplers as multi-bit samplers has been shown. It was found that an ENOB greater than that of a 1-bit sampler was achievable using a referenced-based sampling method.




Figure 5.24: Measured ENOB as a function of frequency for a 1.5 bit sampler made out of (Top) channels 0 and 1, (Middle) channels 2 and 3 and (Bottom) all four channels. The overall signal-to-noise is worse than expected, but little degradation is seen due to the interleaving process.

Finally, interleaving and multi-level sampling was combined to show a full system. This was found to have a lower signal-to-noise ratio than that expected from the individual interleaved and multi-bit measurements. The reason for this is believed to be the simple RF setup used throughout the experimentation adding cross talk and loss.

Chapter 6

Hardware Digital correlator

The basic concept of the digital correlator constructed has already been discussed in the context of Chapter 4, however the specifics of its creation and operation are detailed below, along with further investigation of its response.

6.1 Correlator architecture

The basic principles of digital correlator design were outlined in the introduction, and it was clear that the two different approaches, XF and FX, are optimal in different extremes of correlator architecture. The FX case applies for the requirements of fine spectral resolution over many baselines, while XF is optimal in the regime of coarse spectral resolution and small number of baselines. Another factor that needs to be considered when designing a correlator is the bit depth of the signals being used. In the case of a 1-bit system, the multiplication used by an XF correlator can be simplified to use an XNOR operation. This means that simple logic gates can be used in place of DSP48 multiplier blocks within the FPGA. The use of XNOR operations results in two consequences. The first is that the interpretation of the data streams from the samplers can be simplified. This is because XNORs use the binary values being output in the correlation whereas multiplications require the correct weights ($\pm\frac{1}{2}$) to be used. FPGA boards also incorporate many more logic blocks than multipliers[44], which means that more lags can be constructed. Finally, the use of direct sampling also introduces the ability to implement a many-baseline correlator on a single FPGA device.

The design specification for the correlator described in this thesis was to produce a system that could coarsely sample the spectrum over a large bandwidth. Given the relative merits of XF over FX, an XF correlator design was selected as it would be the most efficient use of FPGA resources. To test the basic operation of direct sampling

described in the previous chapter and to characterize the correlation technique, four inputs were required. This led to a six-baseline correlator being constructed so that each pair could be simultaneously correlated. The frequency resolution of a lag correlator is determined by the number of lags calculated per baseline. This lag count was manually explored such that the correlator being constructed filled most of the device and a value of 60 lags per baseline achieved. If more baselines were to be required this could easily be achieved by reducing the number of lags calculated per baseline, sacrificing frequency resolution of the system.

The hardware platform used for this correlator was the same ROACH board incorporating the high speed serial interfaces used for direct sampling. This device was described in Section 3.1.2.

6.2 Correlator implementation

In an XF correlator, each lag value is calculated by multiplying each bit from the first incoming data stream with a bit from the second data stream, taking into account the offset of the lag:

$$R(\tau) = \Sigma \left(D_1(n) \times D_2(n + \tau) \right) \quad (6.1)$$

where $R(\tau)$ is the accumulated correlation value for offset τ , D_1 is the first data stream and D_2 is the second. The above process is made more complicated by the fact that 20 bits of data are produced per clock, this means that 20 multiplications and additions must occur per lag per clock. To achieve this an adder tree was used. This is a cascading system of additions where for each clock, each stage of the tree performs a single set of additions before passing the results on to the level below. In this fashion the output from the 20-bit addition can be performed at the correct data rate, but with some latency between the input of the data and output of the result. The adder tree constructed for a single correlator lag can be seen in Figure 6.1. The flow of a single 20-bit data correlation, shown in cyan, through the tree can be seen in Figures 6.2 and 6.3. This shows how it takes eight clocks for the data to be output from the transceiver and added to the accumulator. The multiplier boxes are constructed from an XNOR operation as this is equivalent to the correlation between the two 1-bit signals, as shown in Table 6.1.

The integration period for the correlations runs as a separate clock, and is set in code via the SPI interface discussed in Section 3.1.2.2. Once the integration time has expired, the lag values are dumped to memory and a data-ready flag is asserted. The

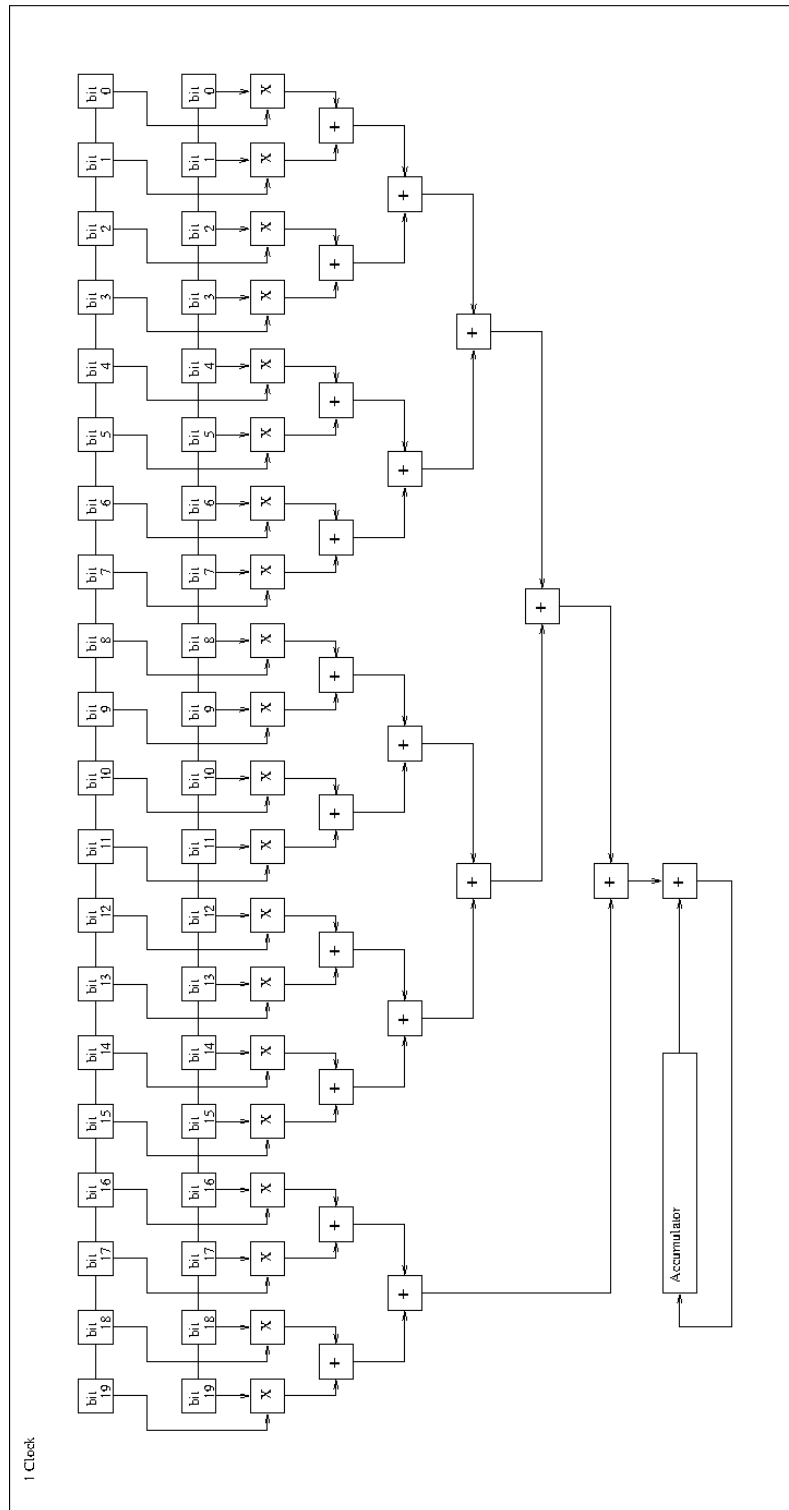


Figure 6.1: Correlator architecture: This is the operation of one lag of one of the correlators. 20-bit parallel data streams are received from the transceiver outputs, multiplied using XNOR operations, and added to the accumulator using an adder tree.

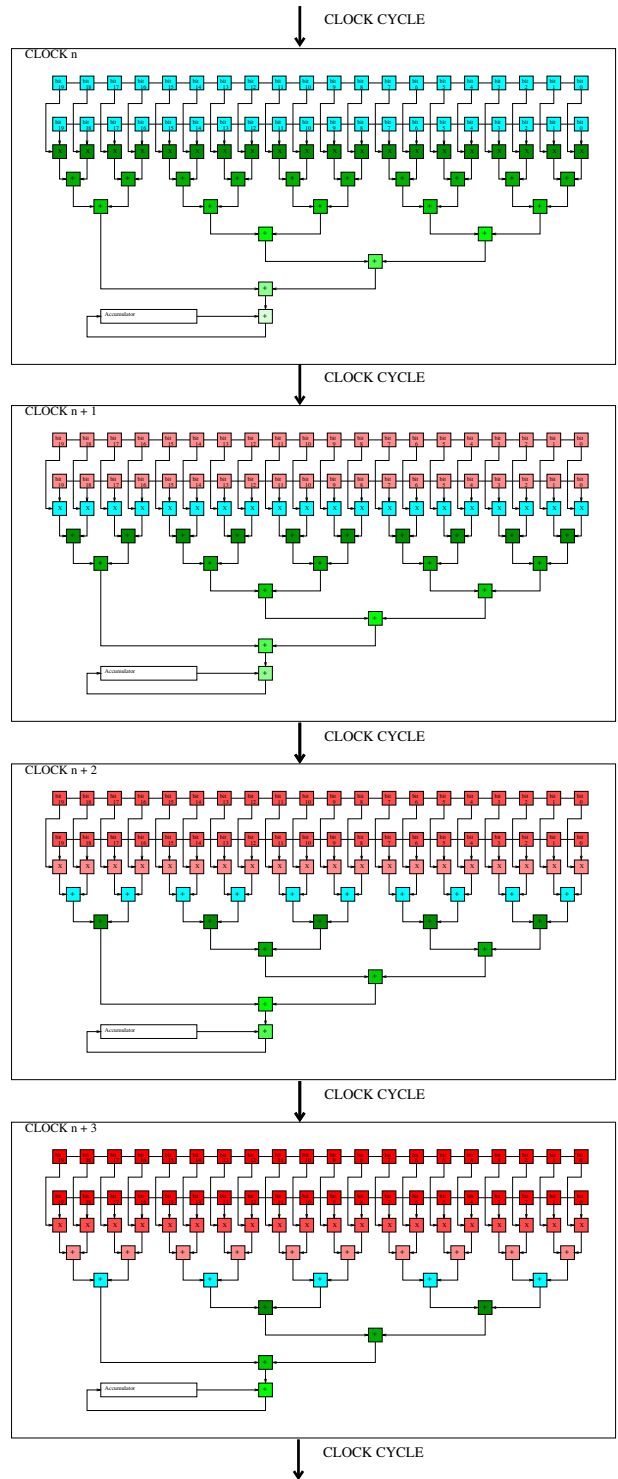


Figure 6.2: Correlator architecture: This plot shows the first four clocks for the cyan data, captured on click n , flowing from the transceiver output through the adder tree and into the accumulator. The Green tones show data captured before clock n and the red show data captured after clock n .

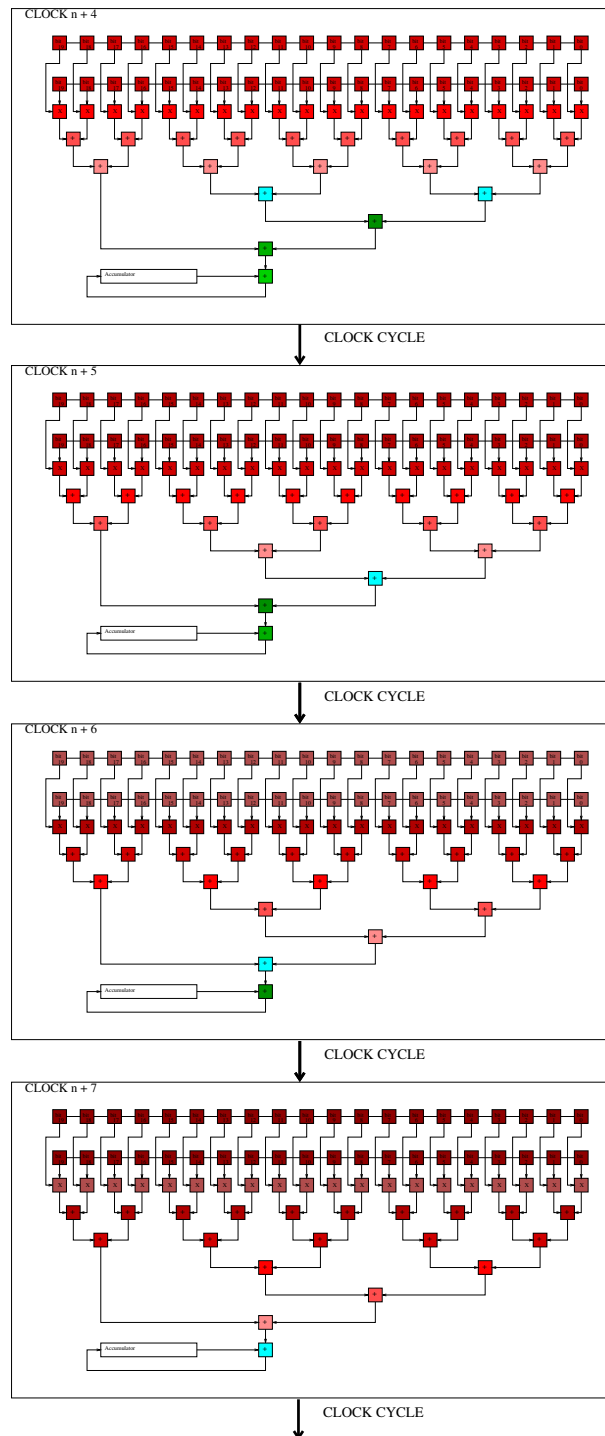


Figure 6.3: Correlator architecture: This plot shows the second four clocks for the cyan data, captured on click n, flowing from the transceiver output through the adder tree and into the accumulator. The Green tones show data captured before clock n and the red show data captured after clock n.

A	B	A \otimes B	A XNOR B
0	0	1	1
0	1	0	0
1	0	0	0
1	1	1	1

Table 6.1: Truth table showing that an XNOR operation can be used to find the correlation between two input 1-bit signals as it is equivalent to the correlation (\otimes) of the two signals.

values then remain in the memory until the data are read and a read flag is asserted. This buffering persists even if a new integration period expires, to ensure that any data currently being read are from the same correlator capture. The integration period maximum was selected so that a perfectly correlated signal would not overflow the data store, which is a 32 bit memory location. If the inputs are perfectly correlated, for every clock cycle 20-bits correlate and therefore 20 is added to the accumulator. This gives a maximum integration time of

$$\frac{2^{32}}{20} \approx 2.15 \times 10^8 \text{ clocks} \approx 1.37 \text{ seconds.} \quad (6.2)$$

The data output rate from the system is such that at this integration time, data can be read out faster than they are produced. Therefore if multiple correlator captures are made, they are from consecutive outputs of the correlator. If these outputs are then summed together, the result is equivalent to an integration being performed with a much longer duration than the single integration period.

When the above correlation approach was adopted for six baselines, 60 lag correlators could be constructed within the single FPGA. This amounts to 360 lags being formed in total.

If the system requirements were altered, such that a frequency coarseness comparable to the analogue case is considered, this would allow 22 16-channel correlators to be constructed, and therefore 22 baselines. The number of baselines relates to the number of antennae by the formula:

$$\text{baselines} = \frac{1}{2}n(n - 1), \quad (6.3)$$

where n is the number of antennae. Twenty-two baselines corresponds to:

$$22 = \frac{1}{2}n(n - 1) \Rightarrow n = 7 \text{ antennae.} \quad (6.4)$$

From this it can be seen that a correlator could be constructed using the above architecture, allowing seven RF inputs and all cross-correlations with a sixteen-point correlation being performed. This corresponds to frequency bins of ≈ 200 MHz:

$$\text{resolution} = \frac{2 \times \text{bandwidth}}{\text{lags}} = \frac{2 \times 1.5625 \text{ GHz}}{16} \approx 200 \text{ MHz}. \quad (6.5)$$

This is without extensive optimization of the internal design. The direct data capture modules that were used in the testing phase to look at raw transceiver output streams are also still present. These use large amounts of the FPGA memory. It was found that the limiting factor for the design was this FPGA memory. It is therefore feasible that if the testing modules were removed, and the design fully optimized, many more baselines could be produced within the device.

6.3 Hardware correlation of single-frequencies

The first test to characterise the response of the six-baseline hardware correlator was to inject a single-frequency signal into two of the inputs and perform the correlation. The RF setup for this experiment can be seen in Figure 6.4. The correlation of the two inputs will show a triangular periodic response whose frequency matches that of the RF input. The zero lag position will also be seen as a peak in the wave, because the two channels were aligned before the correlations were performed. This was done by adding integer-sample delays to the two inputs until the correlation function was centred with respect to the zero lag.

If the other correlation pairs are viewed, the correlation of an RF input with a zero input will show no net correlation. The correlation between the two zero inputs will show a perfect correlation at all lag values. The results for the six baselines can be seen in Figure 6.5 for a 100 MHz signal. Figure 6.6 shows the correlator outputs and recovered spectra for four different RF input frequencies. The plots for the 100 MHz input signal show the expected results of a triangular cross-correlation for the signal-signal pair, and no net correlation for the signal-no signal pairings. A perfect negative correlation is seen in the output of baseline six (the no signal - no signal pair), this was caused by the fact that one of the data streams was inverted within the FPGA. Opposite lines on the differential inputs were being used, therefore when a signal was applied, a 180 degree phase offset was seen. When no signal is applied all the transceivers default to outputting 0's, which are then turned into 1's by the inversion circuit, giving the anti-correlation. The recovered spectra for the signals, shown in Figure 6.6, show a good level of recovery, with the input frequency

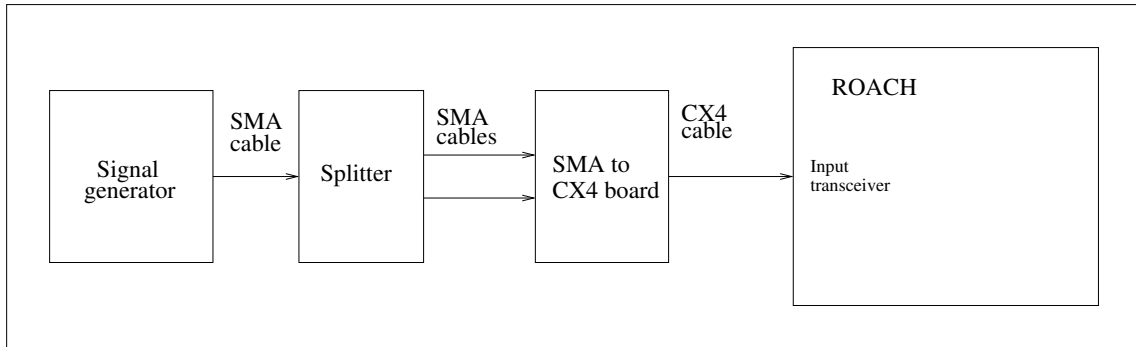


Figure 6.4: RF setup for single-frequency correlation tests.

clearly visible. Low level harmonics can be seen at $\approx 10\%$ the peak value. These are due to the 1-bit nature of the correlation and agree with simulation. The RF signals are being digitised into square waves by the 1-bit sampling. The correlation of two square waves gives a triangular wave correlation function. The Fourier transform of a triangular correlation function gives a set of harmonic peaks of decreasing amplitude, with the primary peak centred on the fundamental frequency of the input signal. The broad peaks seen for some frequencies are caused by the RF frequencies selected not coinciding with the FFT frequencies for the device and therefore the peak is spread over multiple bins.

6.4 Broadband recovery

As with the direct samplers, the next stage of testing after single-frequency use was the injection of a broadband signal. The RF setup for this experiment can be seen in Figure 6.7. This will generate a very different response within the correlator, with a sharp peak around the zero lag and no net correlation elsewhere (assuming a flat broadband signal). In the ideal case of a perfectly flat noise source covering the entire first Nyquist zone, the correlation becomes a delta function at the zero lag. This was not expected from the actual correlation measured due to features in the noise source such as gain slope. These spectral features will add sidelobes to the central peak of the correlation function. The measured correlation and the recovered spectrum can be seen in Figure 6.8. The correlator output shows the expected highly peaked response. The recovered spectrum shows considerable features across the band, however the broadband response is still clearly visible. This shows the systems ability to correlate broadband RF signals.




Figure 6.5: Hardware correlator output for single-frequency 100MHz RF signal into inputs 0 and 1 with the other two inputs not connected. Shown are the six baseline outputs with the cross-correlations for the different pairings corresponding to 1. 0 with 1, 2. 0 with 2, 3. 0 with 3, 4. 1 with 2, 5. 1 with 3 and 6. 2 with 3.

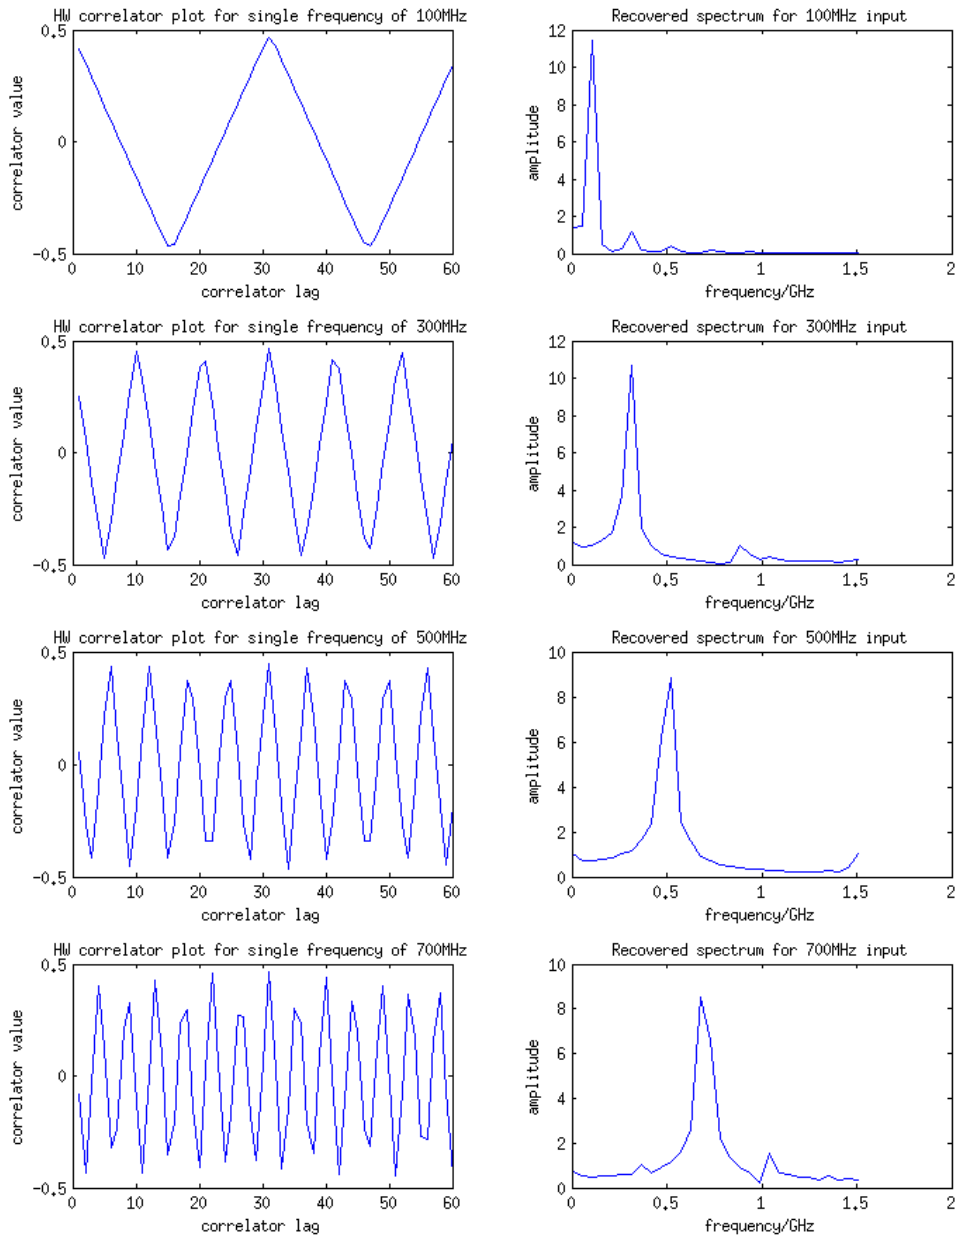


Figure 6.6: (left) Hardware correlator outputs for single-frequency RF signals. (right) The recovered spectra.

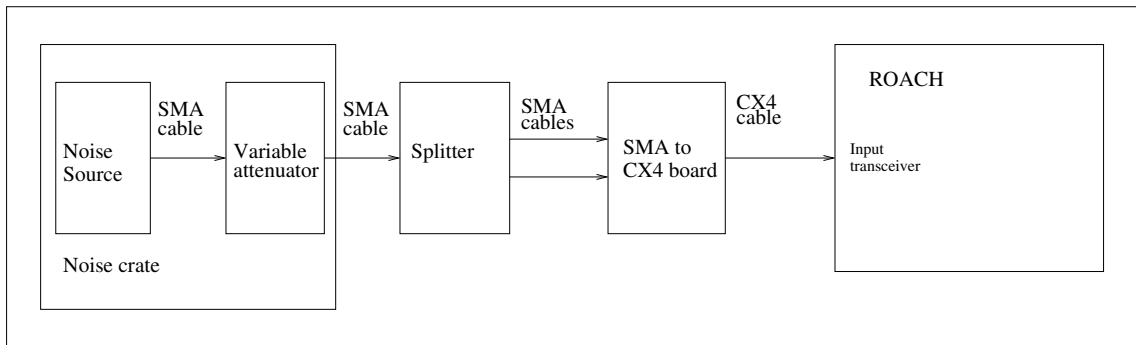


Figure 6.7: RF setup for broadband correlation tests. A noise source with a DC-1.5 GHz low pass filter was used to generate the broadband signal.




Figure 6.8: (Top) Hardware correlator output for broadband RF data with (Bottom) recovered spectrum.

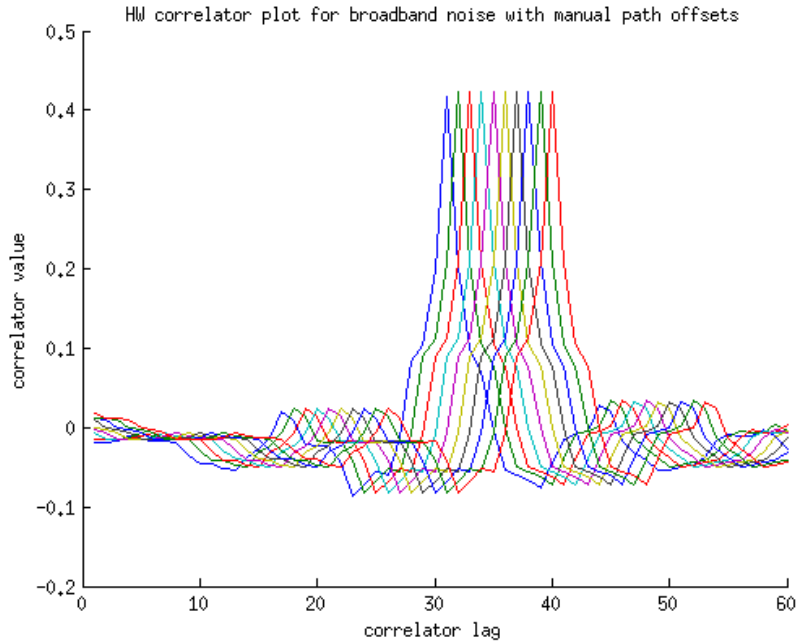


Figure 6.9: Hardware correlator output for broadband RF data with manual path delays added to one of the data streams

6.5 Broadband recovery with manual path delay added

An extension to the above experiment was performed where the integer delay block within the FPGA was used to add differing amounts of delay into one of the input lines. The expected effect of adding path delay to a single line is that the response of the correlator will shift through the lags while maintaining shape. This experiment establishes whether the data streams are being properly analysed, primarily with regards to the interpretation of the 20-bit simultaneous output. If the bit order were being corrupted, then the shape of the correlation function would be observed to change when delay was added. This change can not be seen in the plot of Figure 6.9, which shows the overlaid response for different delay values. This demonstrates that the data flow through the correlator is working as expected.

6.6 Single-frequency with uncorrelated noise

The test of single-frequency injection with uncorrelated noise has already been described in Section 4.3. It was shown that with sufficient isolation between the two channels, a very low-level correlated single-frequency signal could be recovered from

a very noisy environment. This was done to test the distinction between two frequencies within the first and second Nyquist zones of a single sampler, which were all in the first Nyquist zone of an interleaved sampler.

In the context of the characterisation of the correlator it was interesting to take a first Nyquist zone single-frequency signal, and investigate how low a power level could still be detected from within the uncorrelated noise. To do this, the uncorrelated noise was first turned off and the power level required for the single-frequency to trigger the samplers was established. The uncorrelated noise was then switched back on and the power level of the single-frequency signal decreased until it could no longer be detected in the recovered spectrum.

The experimental setup used and the result obtained can be seen in Figures 6.10 and 6.11 respectively.

The integration period for the correlation was selected such that no reduction in the noise floor was seen for longer integrations. This means that the noise floor was entirely being formed by a correlated noise component within the system. It was found that a single-frequency power level of 48 dB below triggering was detectable by the system. This represents the effective output dynamic range of the correlator in this hardware setup.

It is also important to investigate how the correlator handles uncorrelated noise. To do this, an integration period shorter than the above was used, where the correlated noise floor was not reached. In this regime the power level of the uncorrelated noise goes as:

$$\Delta P_{noise} = \frac{k_B T \Delta \nu}{\sqrt{\Delta \nu \tau}} \quad (6.6)$$

where ΔP_{noise} is the noise floor measured in watts, k_B is the Boltzmann constant, T is the effective noise temperature of the system, $\Delta \nu$ is the bandwidth and τ is the integration period. From this it can be seen that the noise floor will integrate down with the square root of the integration period. If the log of the power measured is plotted against the log of the integration period the gradient should be $-\frac{1}{2}$, with a turn off when the correlated noise component starts to dominate:

$$\log_{10}(P_{noise}) = -\frac{1}{2} \log_{10}(\tau) + \log_{10}(k_B T \Delta \nu) - \frac{1}{2} \log_{10}(\Delta \nu) \quad (6.7)$$

This was checked on the system using the previous setup and increasing the integration period from $\approx 5\mu s$ up to $\approx 50ms$. The phase switching period was set

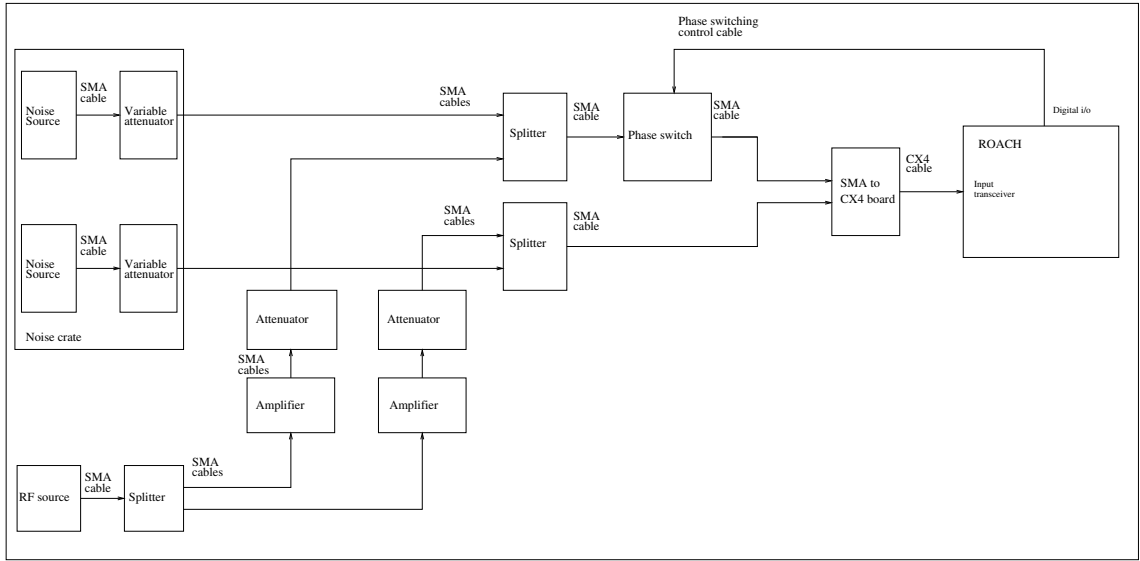


Figure 6.10: RF setup for single-frequency correlation with independent noise fed into the two samplers. The power level of the single-frequency signal was reduced until it could not be detected above the correlated noise level.

at $\approx 5\mu\text{s}$, as integrations must occur over whole numbers of phase switch states to ensure proper cancellation of cross talk.

The result obtained when a 100 MHz single-frequency signal was injected along with the uncorrelated noise can be seen in Figure 6.12. As can be seen the gradient of the line is $-\frac{1}{2}$, as expected, for the middle portion of the graph. At integration periods longer than 3 ms the correlated noise terms dominate and the line hits a limiting value. This correlated noise floor is not an artifact of the sampling environment used but is caused by the RF setup.

At very short integration periods small differences between the integration period and the phase switching period become apparent and start to affect the system by not exactly cancelling the cross talk. For this experiment the phase switching was run much faster than in all other experimentation. This was done to allow the effect of very short integration periods to be investigated. On a real system both the integration period and phase switching periods would be kept longer, meaning matching between their periods would be better achieved. Also, if the phase switching period was increased, the distortion caused at switch-over will only be a very small fraction of the total time, and will therefore only contribute a small amount of cross talk.

The power axis on Figure 6.12 has not been calibrated and therefore an absolute measurement of the noise temperature can not be obtained from this graph.




Figure 6.11: Plot showing how far below triggering level a CW signal can be detected. The power of the CW was reduced from a point where it could trigger the samplers without the uncorrelated noise.




Figure 6.12: The decrease in the uncorrelated noise level as integration period increases with a 100MHz single-frequency correlated signal.

6.7 Broadband signal with uncorrelated noise

The final test that was performed on the correlator system was the injection of a correlated broadband component along with the uncorrelated noise term. This simulates the actual working conditions experienced by radio telescopes where the astronomical signal appears with a low-level broadband spectrum. The setup used can be seen in Figure 6.13, where the setup from Section 6.6 has been modified such that the single-frequency source has been replaced by a broadband noise source with a LPF (Low Pass Filter). This filter was selected at 400 MHz to allow a distinct feature in the broadband noise to be seen when the spectrum was recovered.

The integration period for the correlations was again altered from an initial short value of $\approx 5\mu\text{s}$ up to $\approx 50\text{ms}$. When this was done it was found that the correlated noise floor was no longer being reached. This was due to the removal of two of the amplifiers used to isolate the two RF arms. This should have increased the correlated noise floor as more power can now pass between the two arms of the device, however the reduction in the correlated component shows that the amplifiers must be picking up RF signals, most likely on their power lines, and adding to the over all unwanted correlations.

The maximum integration period was altered to $\approx 5\text{ s}$, by looping a $\approx 50\text{ ms}$ integration to form a longer effective period. The results for this experiment can be seen in the plots of Figure 6.14. The top plot shows the recovered spectrum for the different integration periods. As can be seen, the 400 MHz correlated noise passband is clearly visible in the longer integration period lines. At shorter integration periods the noise level is above the correlated signal level and therefore no 400 MHz cutoff can be seen in the passband.

As with the single-frequency case, in the bottom plot of Figure 6.14, the noise level can clearly be seen to integrate down proportionally to $1/\sqrt{\tau}$. The measured correlated noise starts to dominate the system at an integration period of $\approx 0.1\text{ s}$. The LPF used with the correlated noise source has a large amount of out-of-band attenuation, and therefore is not responsible for the measured correlated noise floor. Imperfections in the sampling will cause power scattering out of the passband used and spread across the entire spectrum. To test whether the correlated noise floor measured is caused by this power scattering, the power level of the correlated noise source was reduced and the experiment repeated. If power scattering was the dominant cause of the noise floor then a decrease in the correlated signal power would result in the noise floor reducing as well. When the experiment was performed this was not

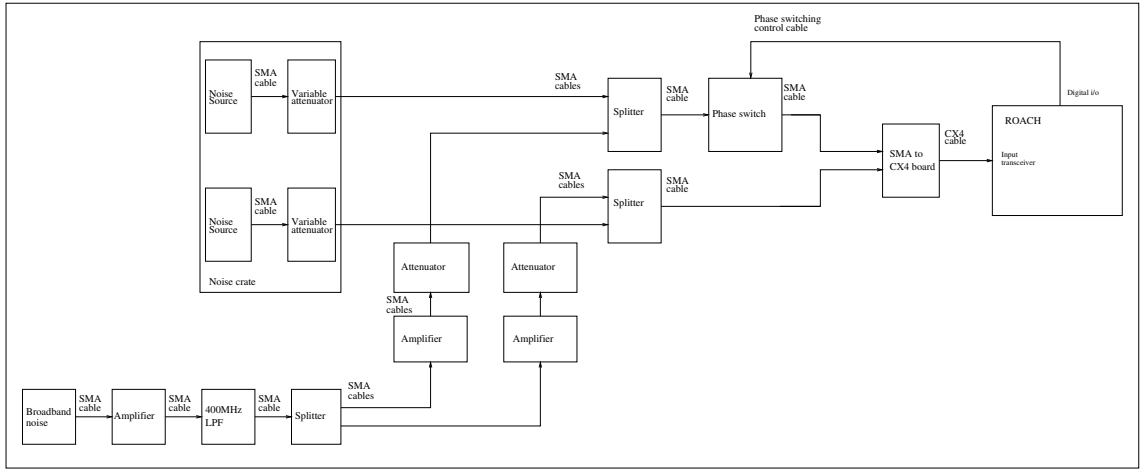


Figure 6.13: Setup for broadband noise correlation with independent uncorrelated noise fed into each sampler. A 400 MHz LPF is used to generate a feature in the band.

seen, showing that the noise floor is being caused by the RF setup used, as in the previous experiment. This agrees with simulations which show the correlated noise floor due to scattering of power should be drastically below the noise floor seen. The exact level of the noise floor was not predicted by the simulation, as a very long integration of noise like signals would be required. This was not feasible to perform, as the running duration of the code quickly became unmanageable as the integration period was increased. The system could also not be simulated with the absence of noise to remove the need to integrate down to find the intrinsic correlated limit. This is because the fact that the samplers are being triggered off independent noise means that quantisation errors caused by the samplers will also not correlate. If the noise sources were removed and the correlated passband allowed to dominate, the quantisation errors would correlate leading to an increased correlated noise floor.

6.8 Correlating interleaved signals

The above sections discuss the use of the correlator with non-interleaved RF signals. During the course of testing the direct sampling method of Chapter 4, correlations were required for interleaved signals. The intuitive solution to this problem would have been to interleave the sampling streams within the FPGA before the correlation was performed. The correlation could then be performed on the interleaved inputs to calculate the desired response. This would have required a complete redesign of the correlator architecture and would also of required the FPGA code to be reprogrammed

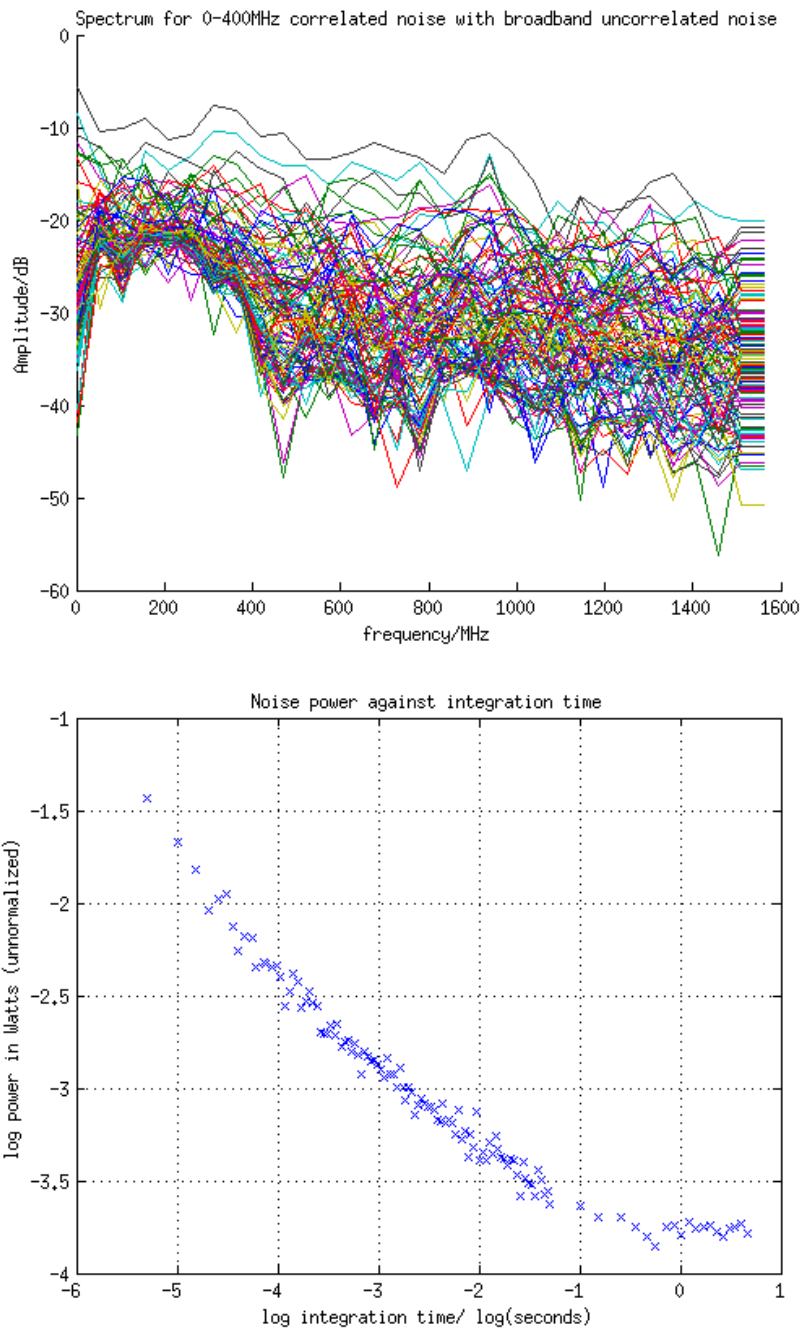


Figure 6.14: (Top) Spectrum of 0-400 MHz broadband correlated noise added to uncorrelated noise. The integration time is increased to integrate down the uncorrelated component with the different integration periods plotted as different coloured lines. The upper lines correspond to an integration time of $5 \mu\text{s}$, while the lowest lines correspond to an integration time of $\approx 5 \text{ s}$, in which the output spectrum is detected with a signal-to-noise ratio of $\approx 15 \text{ dB}$. (Bottom) The decrease in the uncorrelated noise level as integration period increases.

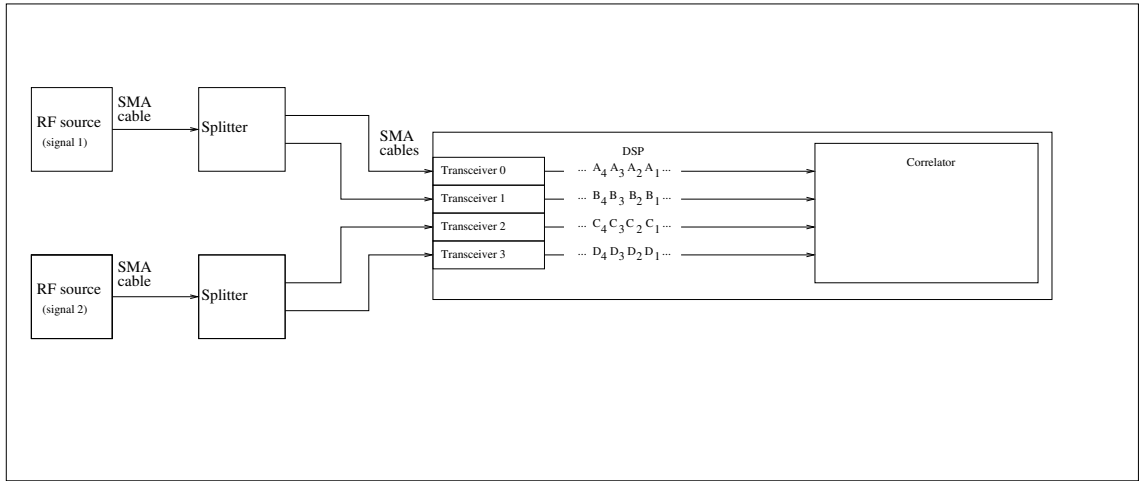


Figure 6.15: Example setup for how an interleaved correlator can be formed from the cross-correlations of the individual streams.

when switching between a non-interleaved and interleaved setup. It was therefore decided to adopt another method for calculating the interleaved correlations.

It was realised that all the information required to calculate the correlation between two interleaved signals was contained within the four cross-correlations of the individual streams. These four cross-correlations are already calculated within the six baselines of the current correlator hardware.

Assuming the setup of Figure 6.15, with transceivers 0 and 1 acting as an interleaved pair to sample signal 1 and with transceivers 2 and 3 acting as an interleaved pair to sample signal 2. The four output streams (A, B, C and D) can be interleaved to increase the sampling rate of the original signals, such that,

$$\text{signal1} = (A_1 B_1 A_2 B_2 A_3 B_3 \dots), \quad (6.8)$$

$$\text{signal2} = (C_1 D_1 C_2 D_2 C_3 D_3 \dots). \quad (6.9)$$

From this it can be seen that the zero lag of the correlation function is formed by

$$\text{lag}_0 = (A_1 \times C_1) + (B_1 \times D_1) + (A_2 \times C_2) + (B_2 \times D_2) + \dots = \sum_i (A_i \times C_i) + \sum_i (B_i \times D_i). \quad (6.10)$$

This can be seen to be the addition of the zero lag values for the correlation between A and C and the correlation between B and D.

The first lag is formed by

$$\text{lag}_1 = (B_1 \times C_1) + (A_2 \times D_1) + (B_2 \times C_2) + (A_3 \times D_2) + \dots = \sum_i (B_i \times C_i) + \sum_i (A_{i+1} + D_i), \quad (6.11)$$

which can be seen to be the addition of the zero lag for the correlation between B and C with the first lag of the correlation between A and D. In this manner all the correlation values between the interleaved pairs can be formed from additions of the four cross-correlations of the individual streams.

Figure 6.16 shows the sampled data for a 100 MHz RF signal fed into the four transceivers with uncorrelated noise also added. The four cross-correlations are plotted, showing the expected responses with shifts in the positions cause by the offsetting of the sampling points. The bottom plot shows the result when the individual correlations are combined to form a single correlation output for the interleaved pairs. This shows a smooth sine response with no discernible distortion from the interleaving technique.

6.9 Conclusion

This chapter has demonstrated the ability of the correlator constructed to perform the six cross-correlations on the four input streams in parallel with 60 lags per baseline. The frequency resolution obtained is ≈ 50 MHz and the bandwidth for the system is 1.5625 GHz. With sixteen lags per baseline a seven-input correlator could be implemented on a single Virtex 5 chip, and probably even more inputs would be possible with further optimisation of the FPGA memory usage.

The system also shows the expected response to the injection of uncorrelated noise, with the level integrating down as $1/\sqrt{\tau}$.

The technique required for performing correlations on interleaved signals is also shown. This allows the bandwidth to be increased to 3.125 GHz while maintaining ≈ 50 MHz frequency resolution.




Figure 6.16: (Top) four cross correlations between transceiver outputs for a 100 MHz RF input. (Bottom) The combined correlator output giving the interleaved correlation response.

Chapter 7

RF and IF systems for broadband interferometry

In the previous chapters two distinct correlators have been discussed. These are a 2-20 GHz single-baseline analogue correlator and a 1.5625 GHz six-baseline digital correlator. The bandwidths of these correlators are such that specialist RF and IF devices are required to feed in the astronomical signals. It was decided to develop the required devices in-house as the majority of the devices required were not commercially available.

The following chapter outlines the development and characterization of the required RF components. These have been used throughout the experiments described in this thesis. Some of these devices had been developed prior to the research of this thesis and some have been designed and constructed as part of the research.

Throughout the investigations, the flatness of the passband for the devices was considered as this, along with the bandwidth, constrains the usability of the devices. The flatness of the passband is important as the correlators being considered have a certain power at which their optimum sensitivity is achieved. If the power level at one part of the band is lower than this, a reduction in the signal-to-noise will be seen.

The devices considered in the following sections operate at a bandwidth of ~ 20 GHz. This was selected as it corresponds to the bandwidth of the analogue correlator system. The development of RF devices at this frequency has become considerably easier in recent years due to the availability of broadband packaged devices. These can be soldered directly to RF circuits using standard surface-mount techniques to produce the required systems.

7.1 RF and IF components

Key components required for an ultra-broadband system are:

1. Amplifiers
2. Filters
3. Path compensation
4. Variable attenuation
5. Phase switches
6. Slope compensation

The following sections outline the development and characterisation of these devices.

7.1.1 2-20 GHz amplifier

Prior to the research of this thesis, a 2-20 GHz broadband amplifier was developed by the group for the C-BASS project[24]. It is based on two Hittite HMC462LP5 chips¹ in series with a common power connection to generate a gain of 26 dB over the desired frequency range. The board layout for this device can be seen in Figure 7.1, with Figure 7.2 showing the S-parameter graphs for forward and backwards transmission as well as reflection.

As can be seen, the band has a slope of approximately 3 dB over the 20 GHz with a fall off in gain after 22 GHz.

7.1.1.1 Noise Figure

An important characteristic of an amplifier is its noise figure. This signifies how much noise the device adds to the system when used.

To measure the noise generated by the amplifier, the output power is measured with an input signal of known power. The signal used needs to have a comparable noise level to that expected from the system. If the noise level used is too high then a large uncertainty will exist in the noise figure measured. Noise figure meters exist that can perform the entire noise calculation in a single device. These contain their own noise generator with a variable output power. The noise power generated

¹<http://www.hittite.com/>




Figure 7.1: Circuit layout for the 2-20 GHz amplifier designed by the experimental cosmology group




Figure 7.2: S-parameters for the 2-20 GHz amplifier




Figure 7.3: Amplifier noise measurement theoretical setup.

is quite large and therefore only suitable for noisy devices. This was not used for the amplifier being tested as its expected noise figure is very low and would not be accurately measured by the noise figure meter. Instead, the signal used was a 50Ω load connected to the input. The temperature of this load was controlled by placing it either at room temperature or submerging it in liquid nitrogen. The noise power output from a load at a temperature T_{load} is given by:

$$P_{load} = k_B T_{load} \Delta\nu, \quad (7.1)$$

where k_B is the Boltzmann constant and $\Delta\nu$ is the bandwidth of the system. The load is connected to the amplifier input by a semi-rigid coaxial cable. As this cable presents loss in the system it will also contribute slightly to the noise level by:

$$P_{cable} = k_B T_{cable} \times (1 - G_c) \times \Delta\nu, \quad (7.2)$$

where G_c is the gain of the cable. When measuring the power at the output of the system, the gain and loss terms also need to be taken into account. If the setup shown in Figure 7.3 is created then the power measured is given by:

$$X = P_{load} G_1 G_2 G_c + P_{cable} G_1 G_2 + P_{amp} G_1 G_2, \quad (7.3)$$

where X is the power measured and G_1 and G_2 are the gains of the two amplification stages. These equation assumes that the noise generated by the other cables and amplifiers in the setup can be ignored, since their contribution to the relative noise power is attenuated by a factor of G_1 . If the power measurement is taken for two load temperatures, room temperature and liquid nitrogen temperature, two power readings are generated. If the ratio of the two readings is taken the following is found:

$$Y = \frac{X_{Hot}}{X_{Cold}} = \frac{P_{loadHot} G_1 G_2 G_c + P_{cable} G_1 G_2 + P_{amp} G_1 G_2}{P_{loadCold} G_1 G_2 G_c + P_{cable} G_1 G_2 + P_{amp} G_1 G_2}, \quad (7.4)$$

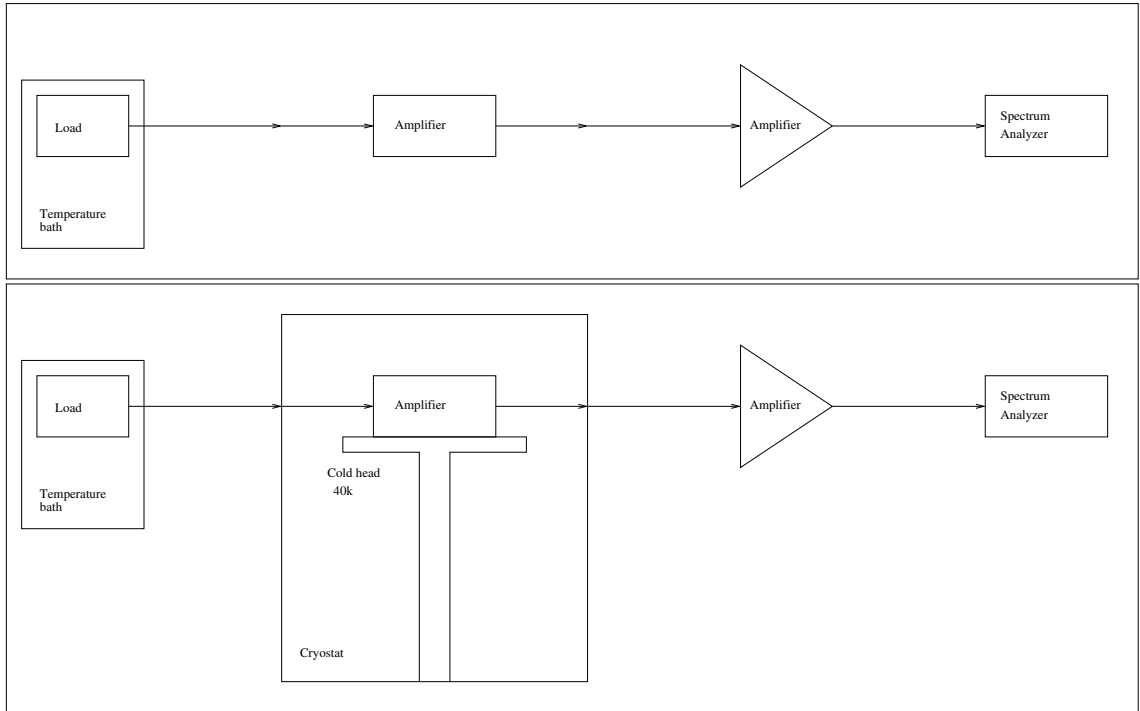


Figure 7.4: Amplifier noise measurement: (Top) bench setup and (Bottom) cryostat setup.

This can be rearranged to give the noise power from the amplifier:

$$P_{amp} = \frac{P_{loadHot}G_c + P_{cable} - Y(P_{loadCold}G_c + P_{cable})}{Y - 1}. \quad (7.5)$$

As can be seen, the gain terms for the amplifiers have cancelled out of the equation. The final stage is to rewrite the above equation in terms of effective noise temperatures, giving:

$$T_{amp} = \frac{T_{loadHot}G_c + T_{cable} - Y(T_{loadCold}G_c + T_{cable})}{Y - 1} \quad (7.6)$$

The above describes the theoretical model for finding the noise temperature for the amplifier. The noise figure will vary based on the temperature of the amplifier. To test the full range of capabilities the noise measurement was made for the amplifier at both room temperature and also cryogenically cooled to 40 K. The amplifier was not designed to operate down to cryogenic temperatures, however, it was decided to investigate its potential use. The two setups required for this are shown in Figure 7.4.




Figure 7.5: Setup for measuring loss in cables. The thermal tether, in this case a copper strap, links the cables to the cold head at their centre.

In order to measure the noise power, it is essential to measure the loss in the cables. For the room temperature experiment this is straightforward as they can simply be connected to the VNA and the loss measured just prior to the noise measurement. For the cryogenic setup the situation is more complicated. The loss in the cables shows a high dependence on their temperature. This means that the loss must be measured with the cables at the correct temperature. To do this, the setup of Figure 7.5 was created, where if the cables are assumed to be identical, the loss measured between the input and output of the cryostat is twice that of the cable in the original setup.

Consideration also has to be placed on the value of T_{cable} , as the temperature is not constant over the entire cable. To calculate this, the cable is subdivided into small elements, with each element contributing to the overall loss and noise power for the cable. The temperature profiles could not be accurately measured, instead an

exponential profile based on heat transfer theory, shown in Figure 7.6, was adopted. The top plot is also used for the bench setup as a temperature gradient exists between the liquid nitrogen cooled load and the room temperature amplifier.

The results for the room-temperature and cryogenically-cooled amplifier are shown in Figure 7.7. As can be seen, the noise temperature is substantially lower for the cryogenically cooled case. The noise temperature is also seen to increase at higher frequencies, this is expected as larger losses will be present both in the cables and in the amplifier.

The specification for the Hittite amplifier used gives a noise figure of ≈ 2 dB to ≈ 5 dB between 2-20 GHz. This can be converted to a noise temperature by:

$$NF = 10 \times \log\left(1 + \frac{T_n}{T_0}\right) \quad (7.7)$$

where NF is the noise figure in dB, T_n is the noise temperature and T_0 is room temperature. The noise temperatures generated from the Hittite data sheet can be seen in the green line on the top plot of Figure 7.7. These can be seen to be slightly lower than the measured values. The differences will be caused by loss in the SMA connectors and coplaner waveguide leading into the chip, as well as the solder connection between the chip and the board.




Figure 7.6: Simulated temperature profile for the cable (Top) outside and (Bottom) inside the cryostat.

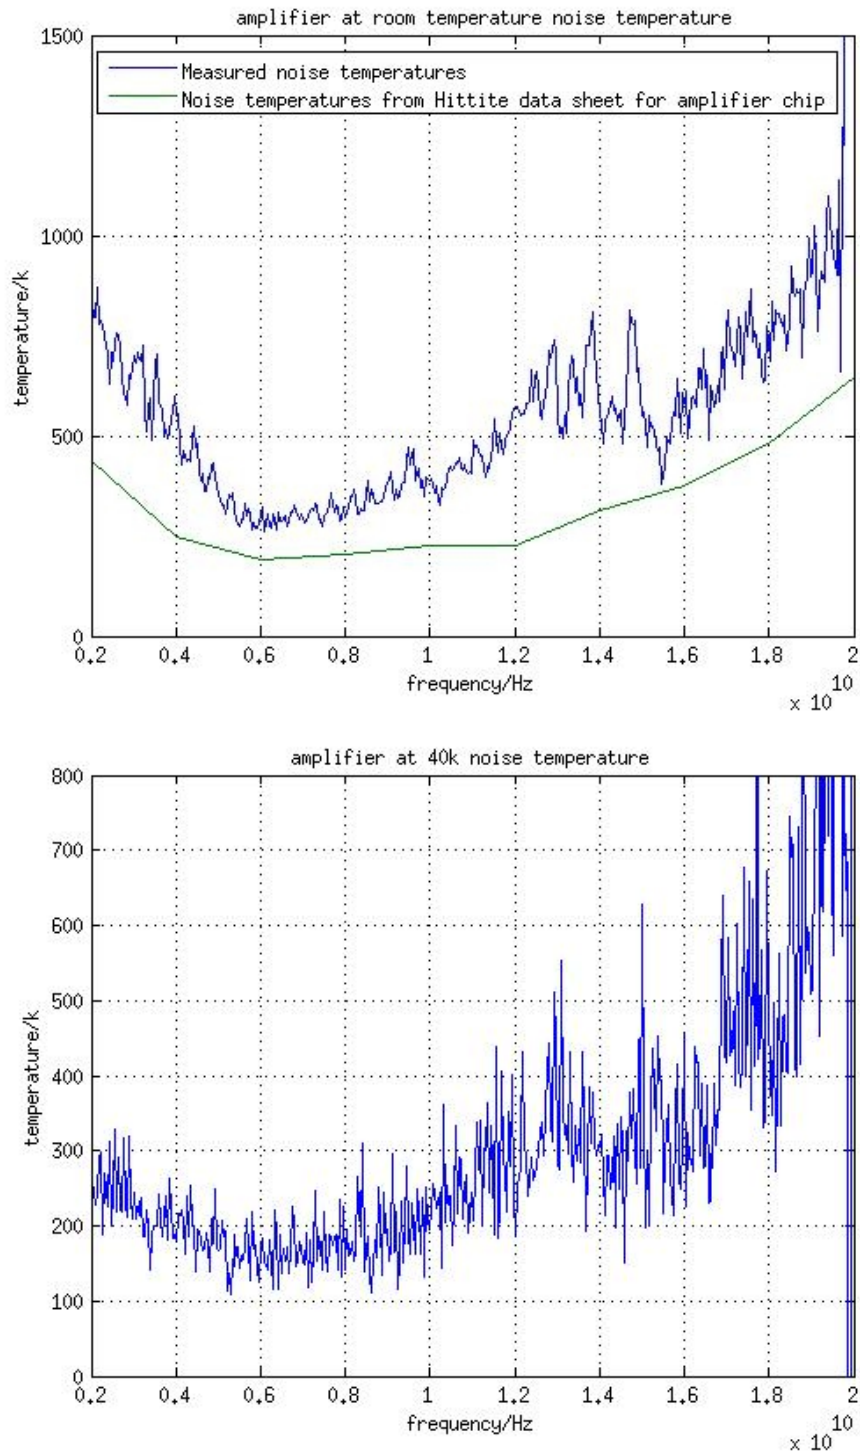


Figure 7.7: (Top) Measured noise temperature for a 2-20 GHz amplifier at room temperature in blue with the noise response from the Hittite data sheet for the amplifier chip in green. (Bottom) Measured noise profile for a 2-20 GHz amplifier at 40 K

7.1.2 Low pass filters

An important aspect of any IF system is its ability to isolate the desired frequency band. The above amplifier operates up to a maximum frequency of 20 GHz. It was therefore necessary to investigate a device capable of isolating this band.

Prior to the work of this thesis, an 18 GHz LPF (Low Pass Filter) was developed. This was produced using RF design techniques by another group member[22] and can be seen in Figure 7.8. The band produced operates up to 18GHz instead of 20 GHz as a slightly different application was being considered. Figure 7.9 shows the S-parameter graphs for the device where a steep roll off is seen after 18GHz. This is not ideal for the 20 GHz band being considered, but was sufficient for initial testing. A slightly scaled design of this filter could easily be produced with a cut-off frequency of 20 GHz.

7.1.3 Path compensation

When performing correlations on broadband signals it is important to ensure that the propagation delays between the two signal paths into the correlator are equal. As sources track across the sky the geometric orientation of the dishes causes the relative path delays within the system to change. This means that any alignment procedure must contain a dynamic element and therefore can not be compensated for using a fixed cable length. Instead, a configurable path delay element is required for each dish within the IF, capable of adding the required propagation delay to align the signals in the correlator. Ideally an infinite number of relative path options would be allowed by the system. This would enable the correlation function to be exactly centred with the zero lag. No broadband system, however, exists that can perform this. Instead a discrete step was chosen to allow the path compensation to be achieved by switching in different stripline lengths. This system was developed by a previous group member[22] prior to this thesis research. The system was developed for use with the analogue correlator described in Chapter 2. The smallest step was chosen to be of smaller size than a single lag difference in this correlator. Path steps increasing in a binary sequence were then chosen to allow any multiple of the minimum path length to be constructed by a unique combination of all the path lengths. The lag spacing within the analogue correlator is 7.5 mm, therefore the path compensation minimum step was chosen as 2.5 mm to allow the correlation function to be properly aligned within one lag. The path compensator delay lengths used can be seen in Table 7.1.




Figure 7.8: Photo of the 18 GHz low pass filter




Figure 7.9: S-parameters for the 18 GHz LPF

Theoretical compensator value/mm	Path compensator 1/mm	Path compensator 2/mm
2.5	4.8	4.3
5	7.9	6.9
10	10.2	10.3
20	20.3	21.3
40	41.5	41.9
80	83.5	83.3
160	166.5	162.9

Table 7.1: Measured path delays in path compensators. Two of each nominal path value were constructed.

The RF construction of these devices uses two broadband single-pole, dual-throw (SPDT) RF switches to select one of two RF paths. The two RF paths differ in path length by the required delay value. An 80 mm path compensator is shown in Figure 7.10. In this photo the switches can be seen as the two surface mount IC on the left of the board. These switch the RF down one of the two copper microstrip tracks. Also included on the board is a small digital logic circuit used to control the switches. This links to the switches using the two yellow wires seen jumping from bottom right to middle left. The forward transmission for the nominal 2.5 mm path compensator is shown in Figure 7.11. The plots shows an approximately flat phase for the first path, and then a slope in phase vs frequency for the second path, which represents a change in path of ≈ 4.8 mm. The other S-parameters are shown in Figure 7.12. As can be seen the device is symmetric in its response, as the switches can work in either direction. As the path difference is constructed using physical stripline length differences, the accuracy of this system is limited by the photo-etching construction technique used. It was necessary to measure the actual path lengths achieved. The VNA was used as it can measure the path delay between its two inputs. This was done in the two states of the path compensator and the difference taken. The actual compensation delay of the devices was then found and can be seen in Table 7.1.

The values being generated by the actual path compensators vary considerably from those desired with the fractional error being greater for the smaller path lengths. This was initially perceived to be an issue with the system, however it can be used to dramatically improve the resolution of the path compensation, as it removes a degeneracy inherent in the original design. Each arm of the IF contains a discrete path compensation unit that allows the full range of path delay to be added. The path compensation is achieved by the difference between the delay added to each arm. With the ideal path compensation values only net delays of 2.5 mm multiples are




Figure 7.10: 80 mm Path compensator using two RF switches to direct the signal down one of two tracks. The track options differ in length and therefore provide a path delay. Multiple of these units is required in any final system. Where each path compensator would possess a different delay value.

achievable, with different configurations generating the same net delay. This means that in theory, with seven path units per arm, 255 possible values are achievable. When the real path values are considered this degeneracy is lifted with ≈ 16000 possible values achievable.

The previous discussion has not specified the upper limit to the series of path compensators. The smallest step size was defined by the correlator architecture, however the upper limit is set by the geometric properties of the telescope system. This is determined by the separation of the dishes and the observation angles being used. If the dishes move from horizon to horizon, enough path compensation is required to correct for the delays experienced. For example, with the GUBBINS system described in Section 1.5 the dish separation is 0.5 m. This means that for horizon-to-horizon observations ± 0.5 m of path compensation would be required. This would mean eight path compensators in each arm of the IF.

As can be seen from the S-parameter plots of Figure 7.12, one of the defining features of the path compensator devices is the band slope introduced. This is ≈ 5 dB over the 20 GHz of bandwidth for a single device, and will naively sum to ≈ 40 dB if the chain of eight per arm is used. This represents a considerable quantity of slope in the system and requires correction to stop a reduction in sensitivity being seen across




Figure 7.11: Forward transmission for the two signal paths through the nominal 2.5 mm path compensator. (Top) magnitude and (Bottom) phase difference are both shown. The actual delay corresponds to ≈ 4.8 mm.

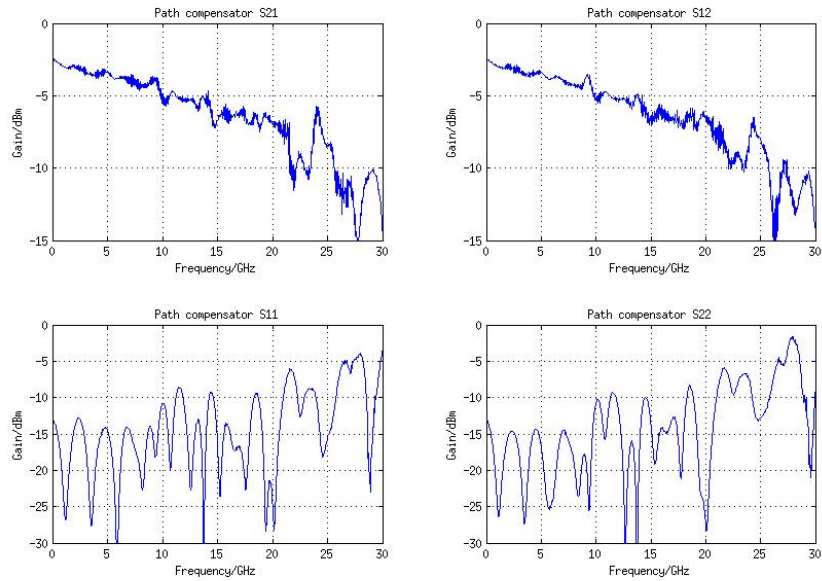


Figure 7.12: S-parameters for the 2.5 mm path compensator.

the band. This will be discussed in Section 7.1.6.

7.1.4 Variable attenuation

To provide broadband variable attenuation, a VVA (Voltage-Variable Attenuator) was used. This is an RF component that uses a FET circuit to generate a broadband attenuation proportional to a voltage input. The device used can be seen in Figure 7.13 along with its attenuation response for a selection of control input voltages. The device was developed prior to the research of this thesis[25] and is based on a Hittite HMC346LC3B². The response can be seen to be flat up to a frequency of 20 GHz with a ≈ 3 dB slope across the band. This slope increases with attenuation. At attenuation levels greater than 20 dB the response can be seen to have considerable distortion, this sets the limit on the attenuation obtainable from the device.

7.1.5 Phase switching

As in the digital correlator experiments of Section 6.6, phase switching can be used in an IF system to deal with cross-talk effects. It can also be used, as in the case of the analogue correlator, to ensure that the correlated signal is modulated and therefore

²<http://www.hittite.com/>





Figure 7.13: (Left) Photo of the variable voltage attenuator with (right) its response for different input signals. A broadband flat response can be seen down to attenuations greater than 20 dB.

not blocked by DC coupling. These functions require a 180-degree broadband phase shift to be switched on and off by a digital control.

The system developed was produced by Christian Holler[22] and uses a broadband SPDT RF switch to direct the signal between two options, one containing a broadband 180-degree phase shift and the other containing no phase shift. This can be seen in the RF circuit layout of Figure 7.14.

The 180-degree phase shift is produced by transitioning a microstrip into a balanced stripline and then back into a microstrip. The output microstrip is upside-down in relation to the input microstrip, which means that the signal and ground lines have been exchanged, resulting in a 180 degree phase shift. As the phase shift is performed using a geometric transform, the bandwidth of the device is only limited by the bandwidth of the transition between microstrip and stripline. Figure 7.15 shows photos of the top and bottom of the phase shift device. The transitions between microstrip and stripline are clearly visible. This bandwidth exceeds the 30 GHz measurement made on the VNA. Although the phase shift has a bandwidth greater than 30 GHz, the splitter and RF switch used only operate up to ≈ 20 GHz and therefore limit the bandwidth of the phase switching system. Figure 7.16 shows the different signal paths within the device. As can be seen a 180 degree difference is generated across the entire 18 GHz band. An error of ± 6 -degrees can be seen, this will not affect the cancellation of cross talk, as a perfect 180 degree phase-shift will still be implemented within the DSP. The phase error will, however, cause a loss in signal correlation, decreasing the signal-to-noise ratio.

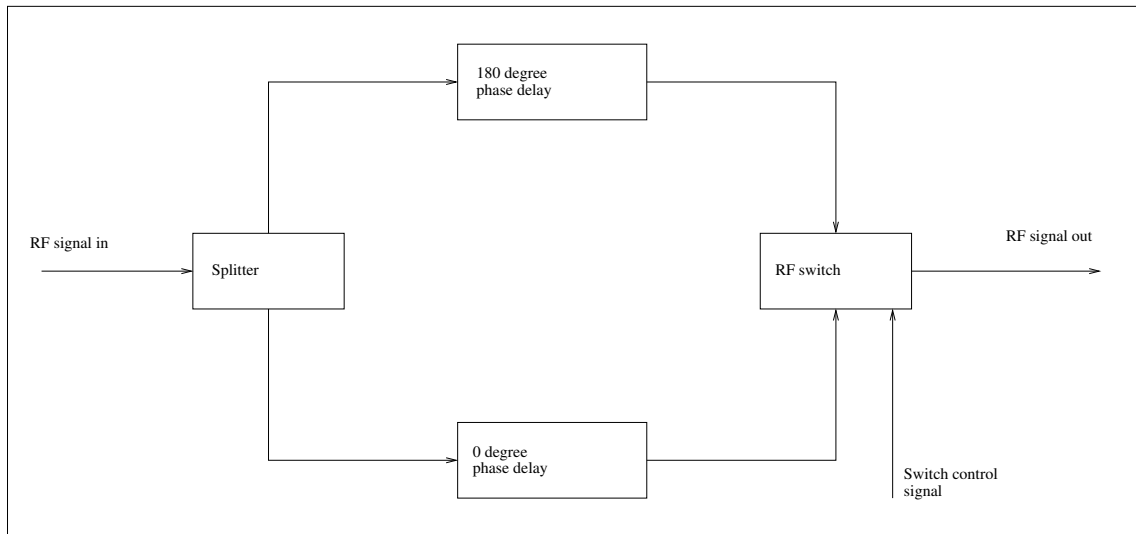


Figure 7.14: RF setup for the phase switching circuit using a 180 degree phase delay block and an RF switch.

7.1.6 Slope compensation

From the characterisation of the above devices, it was seen that a negative band slope is present in most of the responses. This is most pronounced in the path-compensator units which each exhibit a ≈ 5 dB slope across the 20 GHz band. The path compensators are designed to be run in sets to provide the necessary correction and will therefore contribute a large amount of slope to any system produced. To counteract this band slope, either the path compensators (the primary source of band slope) needed to be altered to give a flatter response, or a device was needed to compensate for the slope seen. No easy modification could be made to the path compensators. Also, band slope is seen in all the devices to a small extent and therefore having a component available to counteract this effect would prove helpful. A standalone slope compensator device was therefore designed.

The requirement for the slope compensator was to provide a smooth broadband response with a positive gain slope. After initial investigation it was decided that no single device could be constructed that could possess enough positive slope. It was therefore decided to produce multiple slope compensators that, when summed together, provided the necessary band response. This also had the effect of allowing the exact amount of slope to be tuned by adding and removing compensators. Previous to the work of this thesis, another project member had developed a simple slope compensator for the frequency range required. This was constructed by cutting the track linking the two chips in a 2-20 GHz amplifier and adding a series capaci-




Figure 7.15: Photos of the top and bottom of the phase shift device. The transitions between microstrip and stripline can clearly be seen.

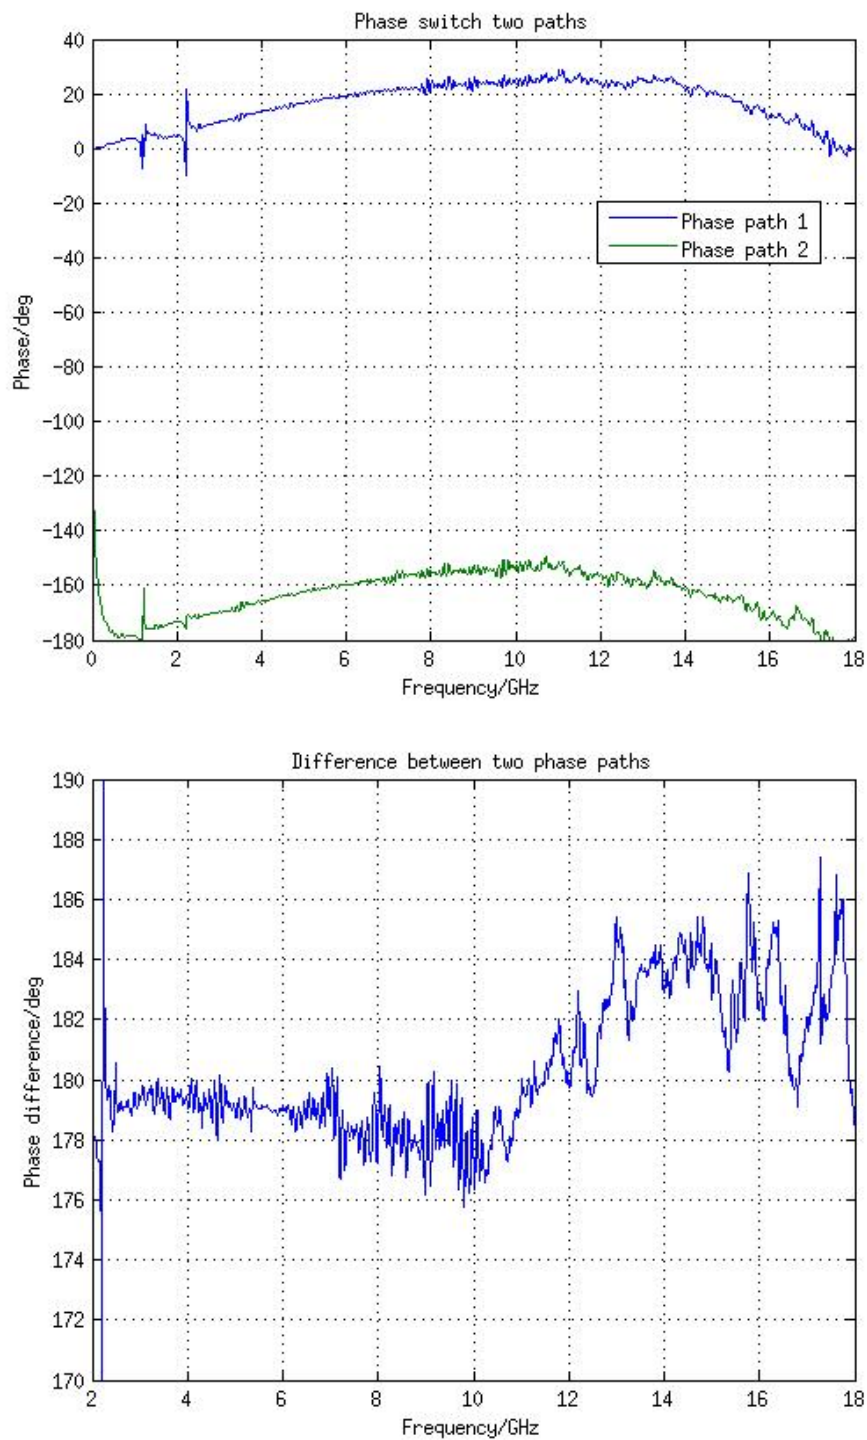


Figure 7.16: (Top) Forward transmission phase for the two paths through the 180 degree phase module with (Bottom) the difference shown.




Figure 7.17: Initial prototype slope compensator. This uses a series capacitor between the two amplification stages of a 2-20 GHz amplifier.

tor. This produced the forward transmission seen in Figure 7.17. As can be seen, a good positive slope is produced from 2-14 GHz, however the passband quickly flattens and then drops off at higher frequencies. It was decided that alternate methods for producing the required slope should be investigated, with an aim to cover the entire 2-20 GHz band.

On investigating possible solutions to the problem of broadband slope compensation the main design idea found was a resistive Tee-attenuator[36]. This is a passive device, shown in Figure 7.18, which allows impedance matched attenuation. It can, however, be modified with the addition of frequency dependent components (capacitors and inductors) to provide a varying attenuation across the band[36]. The modified circuit diagram can be seen in Figure 7.19, along with a theoretical band response. At low frequencies the capacitors are in a high impedance state and the inductor is in a low impedance state. This means that the circuit appears as an unmodified matched attenuator. At high frequencies, the capacitors enter a low impedance state with the inductor switching to a high impedance. This means that the series resistors are shorted and the shunt line is open. The high frequency signal therefore sees a straight transmission line to the output, offering no attenuation. Figure 7.19 shows

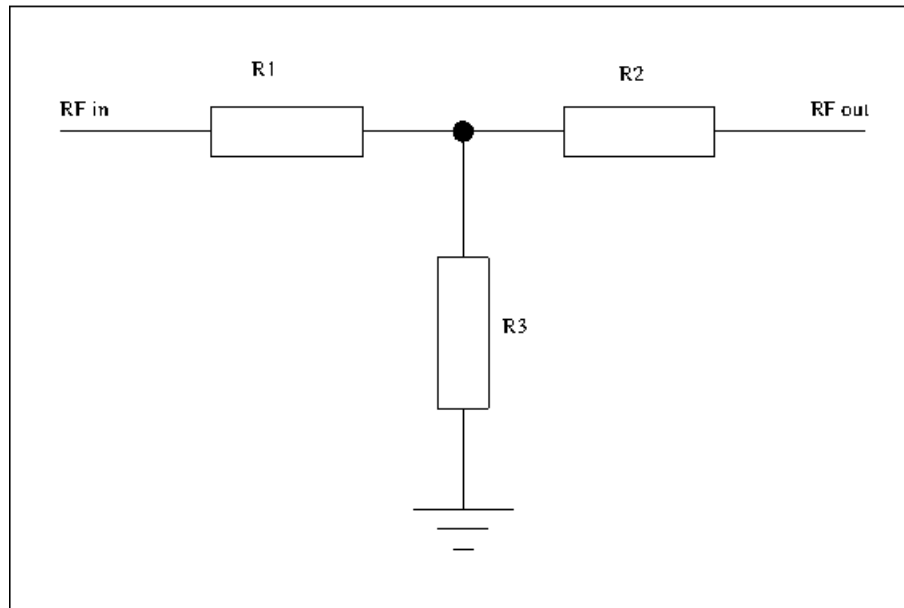


Figure 7.18: Resistor Tee-attenuator. This device allows a matched attenuation circuit to be constructed.

that a smooth positive band slope of magnitude ≈ 7 dB can be achieved over the 20 GHz bandwidth. Given that each path compensator introduces ≈ 5 dB of slope, this would mean that one slope compensator circuit would be required for every one-and-a-half path compensators used. As the device is a passive attenuator circuit, it does result in a lower signal level and therefore an increase to the overall gain in the IF chain is required to compensate. The simulated response shown in Figure 7.19 is for one configuration of the device. Different slope amounts and impedance matching can be produced by varying the component values used.

With the above in mind, two different implementations were considered. The first of these was to simply produce the device as a standalone passive component and add extra amplification using the standard amplifiers. This adds modular simplicity but does require the addition of many new RF components in the chain. There was also a concern that as the device is sensitive to the specific impedances used in the circuit, then the chaining together of multiple passive devices could produce a non-linear response where the total passband is not the sum of the individual contributions. The method chosen to counter all these effects was to add an active isolation circuit into the device. The simplest example of this was to sandwich the passive slope circuit between the two amplifier chips used in the 2-20 GHz amplifier. The broadband amplifier chips show a high degree of isolation which results in the complete device showing a positive overall gain. For the purposes of experimentation,


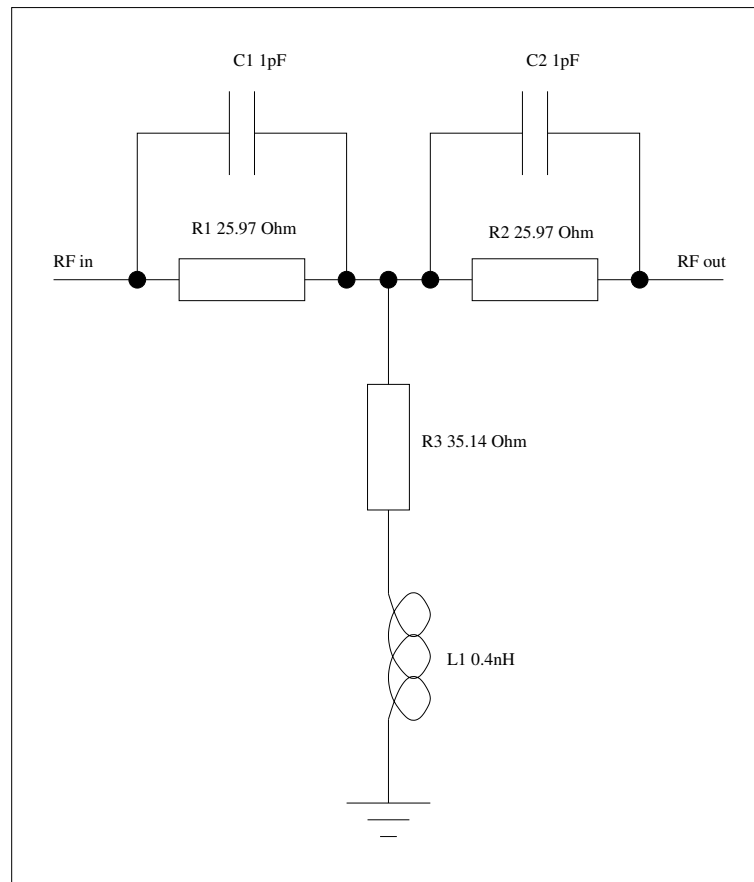


Figure 7.19: Resistor Tee-attenuator with frequency dependence (Top) circuit and (Bottom) response. This device allows a matched frequency dependent attenuation circuit to be constructed.

both the passive and active devices were developed. This was done as the passive design had to first be developed in isolation before it could be amalgamated with the broadband amplifier circuit. This approach also meant that the effects of the passive slope compensator could be viewed separately from those of the amplifier.

The circuit was required to work over a very broad bandwidth of between 2 and 20 GHz. This complicated the construction of the device as most discrete components will not operate up to these frequencies without parasitic effects dominating. Discrete devices do exist that have been specifically designed to operate up to 20 GHz, however, these devices still have sufficient parasitic reactances that full circuit modelling is required to achieve a final design. To ensure that a simple and practical design could be produced careful attention was paid to the method of implementing each of the circuit elements (resistors, capacitors and inductors). In each case the relative merits of discrete components vs RF circuit techniques were investigated.

If the resistors are first considered. It is possible to generate resistive elements using thin-film resistive layers[35]. However, this requires the application and etching of a resistive compound to the board and ultimately was decided to be overly complicated for the desired device. Instead it was decided to use high-frequency resistors in the standard range of values for which full S-parameter models were available. Having settled on discrete devices for the resistors, the possibility for capacitors and inductors using RF track gaps and track width changes was investigated. This was done by constructing a full RF model of the system, using the S-parameters for the resistors. The track widths and positions were then altered to observe the effects on the passband shape. It was found that the necessary capacitance values were not easily obtainable, however the correct inductive values could be generated. Detailed circuit designs were then carried out using discrete resistors, discrete capacitors and microstrip inductors.

For the final product two different resistor value pairs were selected that could generate different levels of slope. These values are shown in Table 7.2, along with the theoretical slope. When these two devices were modelled the graphs of Figure 7.20 were obtained for their forward transmission and reflection properties. These plots show a worse impedance matching than the original theoretical model, so it was necessary in the final product to attempt to add secondary impedance matching devices. This was done using RF stubs that help improve the reflection response of the circuit. The response was also further altered by running optimisation routines on the inductive element to maximise the slope.





Figure 7.20: Simulated response for Tee-attenuator slope compensator using (Top) 30Ω and 18Ω resistors and (Bottom) 50Ω and 18Ω resistors.

	Design 1	Design 2
R1/Ohms	18	18
R2/Ohms	18	18
R3/Ohms	50	30
C1/pF	0.1	0.1
C2/pF	0.1	0.1
L1/nH	0.4	0.4
Estimated slope/dB	18.64	14.59

Table 7.2: Two designs for the slope compensator with estimated slope over the 20 GHz bandwidth.

To provide the desired capacitance values, high frequency RF capacitors were sourced with the required capacitance values. When these were combined with the RF resistors and inductive track element previously discussed the circuit diagram of Figure 7.21 was produced. This shows the standard Tee-attenuator design with the modifications of added discrete capacitors and a series inductive track change in the ground circuit. A stub has also been added to the input transmission line to help match the impedance of the incoming circuit and minimise any reflections. Both the size and shape of the stub, as well as the dimensions of the inductive element were run through the optimisation program within Ansoft Designer³ to try and minimise the reflection while providing a smooth band slope. This design was constructed using both the 50/18 and 30/18 resistor pairings. After testing, it was found that the performance of the 50 Ω based slope design provided a larger and smoother slope than the 30 Ω version, so the 50 Ω design was adopted. The S-parameter results obtained from this circuit can be seen in the graphs of Figure 7.22. The response obtained shows a high degree of smoothness over the 20 GHz pass band as well as 15 dB of slope. The S11 response is also consistently below the 10 dB mark showing an adequate, if not perfect, reflection property.

The passive slope circuit described above was then added to the amplifier design resulting in the circuit seen in Figure 7.23. The response is shown in Figure 7.24 with the passive slope passband being modified by the amplification profile already discussed. The reverse transmission and reflection properties for this device are also much improved on the passive device which should help limit any resonant effects in the IF chain.

³<http://www.ansys.com/>

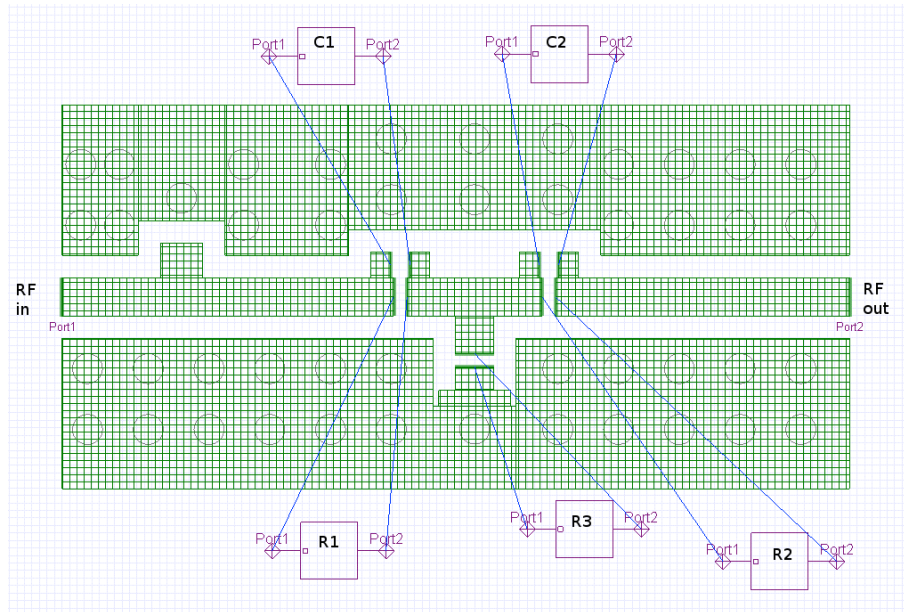


Figure 7.21: RF layout for slope compensator with discrete resistors and capacitors and an inductive track element. The resistors and capacitors are labelled as in the Tee-attenuator circuit. A matching stub can also be seen just after the RF input. This is used to minimise the RF reflections from the device.




Figure 7.22: Passive slope compensator S-parameters.

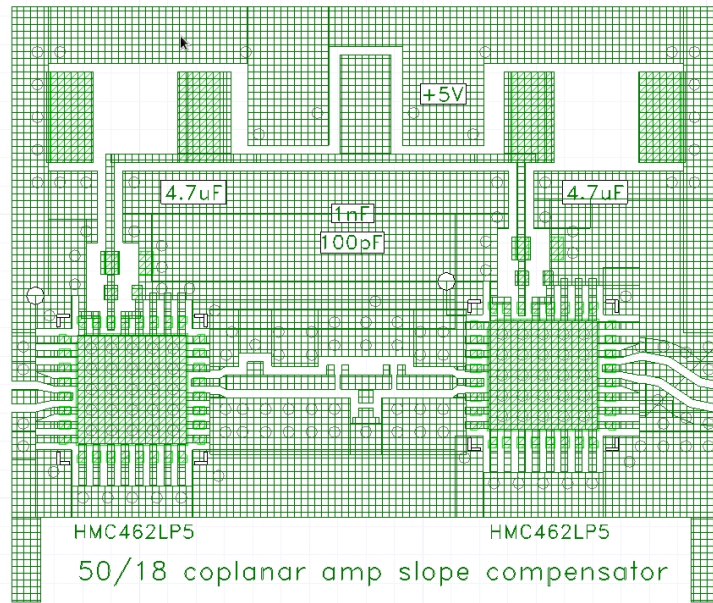


Figure 7.23: RF layout for the passive slope compensator circuit placed between the two amplification stages of the 2-20 GHz amplifier.




Figure 7.24: Active slope compensator S-parameters.

7.1.7 Reduced-bandwidth amplifier

As well as the ultra-broadband devices described above, the digital sampling system required the development of several RF devices in the frequency range of DC-1.5 GHz. This does not fall in the range of some of the devices mentioned above. The first device that was developed was a reduced bandwidth amplifier.

The reduced bandwidth amplifiers that were produced were developed in conjunction with another group member for both this thesis project and another similar application. The amplifier design was based around a BeRex BG14A chip⁴ which shows low band slope up to 3 GHz. As the device requires DC biasing, AC coupling capacitors were required on both input and output, as well as an inductor to stop RF contamination of the power lines. The device was chosen to minimise the slope across the band while still providing relatively high gain. The amplifier chip is rated down to a frequency of 5 MHz. This is sufficiently low for the systems being considered as DC response is not actually required. The capacitors and inductor used were selected to provide the best response across the entire band up to 3 GHz. The circuit layout and final product can be seen in Figures 7.25 and 7.26.

For the application of the digital correlator, the low-end response below 100 MHz is not important as this exists as a single frequency bin, and can be ignored without dramatically affecting the spectral performance of the system. This is not the case for the other group member's application which had a lot higher sensitivity to the bottom of the band. The capacitors and inductor selected for the device were very broadband high-precision devices, which were expensive to buy, and it was therefore decided to investigate whether they could be replaced with lower-cost equivalents while still maintaining the required system performance. This has the effect of reducing the band performance, specifically at the low frequency end. The inductor used has a unique footprint and would therefore be hard to replace without a complete board redesign. The capacitors used, however, were a standard 0402 package and could therefore be swapped out to a more standard part. The performance reduction seen was below the lower frequency limit of the lab VNA at 40 MHz, it was therefore deemed acceptable for the system being constructed. This result can be seen in the S21 parameter graphs for the two different models (with the high frequency capacitors and standard capacitors) shown in Figure 7.27 with the full S-parameter matrix for the device with standard capacitors shown in Figure 7.28.

⁴<http://www.berex.com/>

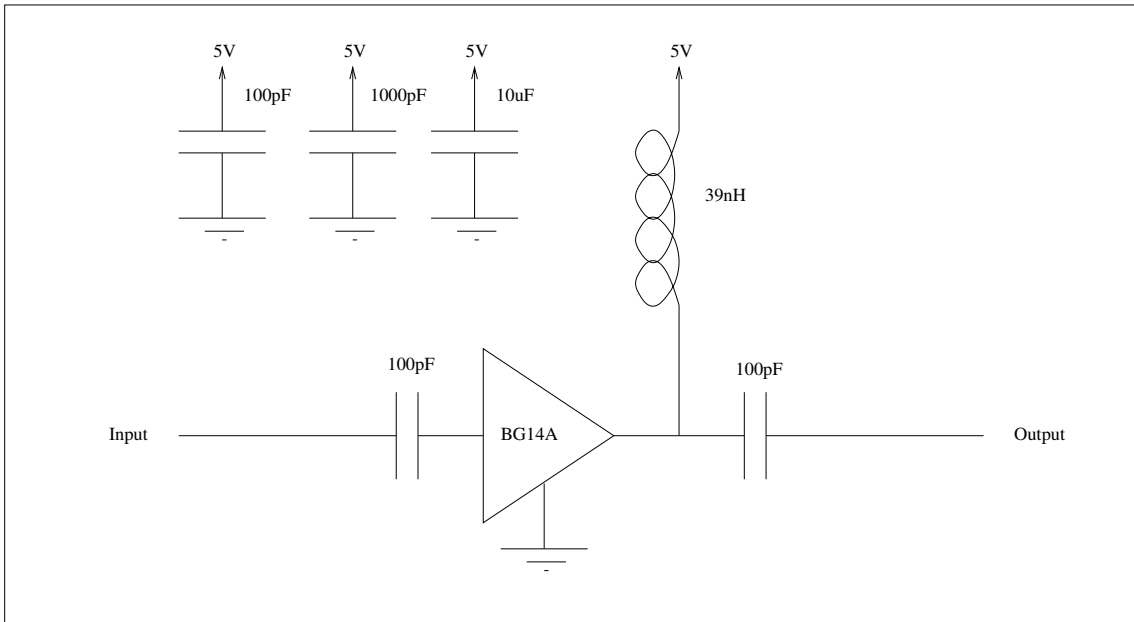


Figure 7.25: Reduced bandwidth amplifier design.




Figure 7.26: Photo of a reduced bandwidth amplifier.




Figure 7.27: Comparison between forward transmission for different capacitors used in reduced bandwidth amplifier.




Figure 7.28: S-parameters for the reduced bandwidth amplifier.



Figure 7.29: Photo of the 1.5 GHz low pass filter.

7.1.8 Low-pass filter

As well as the amplifiers described above, anti-aliasing filters were also required for the digital system. These were designed to isolate the band up to 1.5 GHz to ensure sufficient roll-off before the end of the first Nyquist zone.

Unlike the 18 GHz low pass filter which was made using RF circuit design, the 1.5 GHz filter uses discrete capacitors and inductors to form the necessary passive filter circuit. This was developed in-house by the electronics department and packaged into a thin copper tube to aid RF shielding and provide good grounding. The finished device along with S-parameter graphs for its response can be seen in Figures 7.29 and 7.30.

7.1.9 Digitally-controlled attenuator

It was decided to develop an attenuator device specifically for the digital system. This was done since digitally controlled devices were found that could operate over the frequency range and allow simple control. The device used was a Hittite digital variable attenuator chip HMC939LP4⁵. This uses a 5-bit binary number to specify the attenuation value of a broadband signal in 1 dB increments up to a maximum of 31 dB. An Arduino microcontroller was used to control the digital inputs of the device, with the values being set from a computer interface. The entire device, complete with control hardware and a single reduced bandwidth amplifier was mounted on an aluminium board to form part of a test apparatus for the digital sampling setup. The device can be seen in Figures 7.31 and 7.32 with its band response for forward

⁵<http://www.hittite.com/>




Figure 7.30: S-parameters for the 1.5 GHz LPF.

transmission in Figure 7.33. This plot shows a smooth response with a small slope across the band and many achievable power levels. A consistent feature is seen at ≈ 1.75 GHz, which appears to be inherent in the attenuator. It exists outside of the 1.5 GHz band of the digital correlator and is only of small extent and was therefore ignored.




Figure 7.31: Photo of the AGC board.




Figure 7.32: Close up of the variable attenuator.




Figure 7.33: Forward transmission for the AGC in different power settings. A small feature exists at 1.75 GHz, outside the band of the digital correlator, and will therefore be ignored.

7.2 IF chain

In the above section the RF devices required for the testing of the analogue and digital correlators have been discussed. It is also of interest to consider how these devices might be combined to form the IF chains for a full correlator system.

7.2.1 Basic IF chain construction

The two correlators operate over very different frequency ranges, therefore two different IF chains were designed. The key components for the IF chains are amplification, anti-aliasing filtering and, for the analogue correlator, path compensation. Schematic representations of the two systems can be seen in Figure 7.34. These show the RF components required by each antenna. This setup would be duplicated for each antenna.

A key parameter of the IF was the level of amplification required. This was calculated based on the power level required by the correlators and the expected noise power in the system. Both correlators operate at a signal power of approximately 0 dBm. A variable gain (AGC) unit was added to the IF to allow fine tuning of the signal level. The power level of the signals received by the correlators will be dominated by the noise temperature of the first amplifier. If this noise temperature is assumed to be ≈ 500 K, the equivalent power in the band can be found using,

$$P = k_B T \Delta\nu, \quad (7.8)$$

where P is the power, k_B is Boltzmann constant, T is temperature and $\Delta\nu$ is the bandwidth of the system. For the analogue system this leads to a power level of:

$$P = k_B T \Delta\nu \approx -69 \text{ dBm}, \quad (7.9)$$

and for the digital system gives:

$$P = k_B T \Delta\nu \approx -80 \text{ dBm} \quad (7.10)$$

This means that net IF gains of around 70 dB and 80 dB are required for the analogue and digital systems respectively.

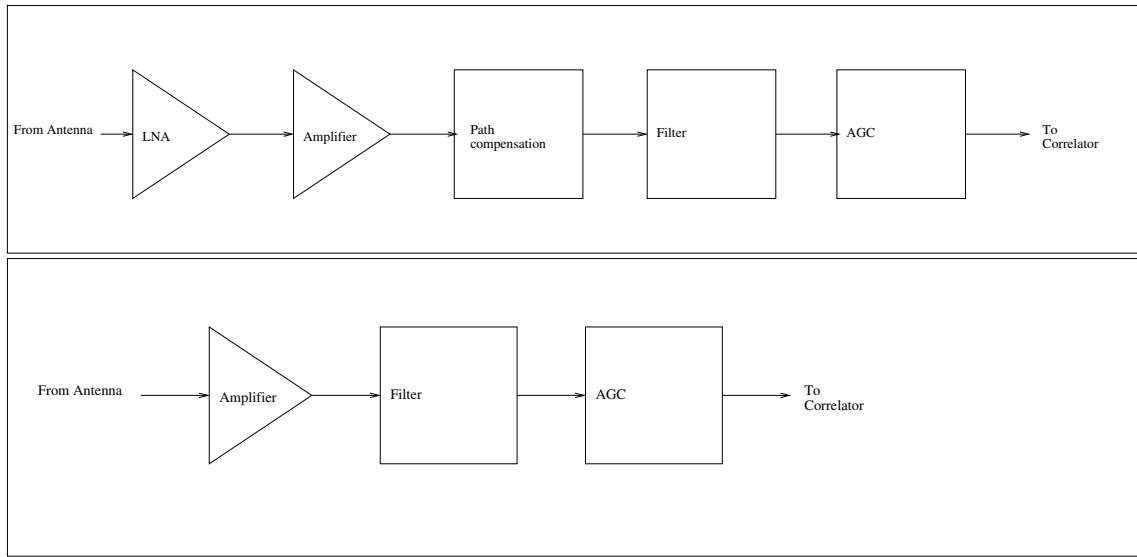


Figure 7.34: Basic setup for (Top) analogue and (Bottom) digital correlators IF.

7.2.2 IF chain complete with S-parameters

It is possible to expand on the IF designs of Figure 7.34 using the actual device properties found in the characterisation section earlier in this chapter.

This was done by assuming that all devices behave linearly and therefore the total passband of the system can be found by adding together the S21 spectra of the individual components. This represents an ideal case where no reflections are accounted for. In the real system the interactions between devices will cause a degraded response with standing waves existing between the components.

When the final system is constructed consideration will have to be placed on the ordering of the components used. Chains containing large amounts of gain are prone to oscillate and therefore are ideally avoided. Badly matched devices (those showing poor S11 and S22 spectra) should also not be put together as a large standing wave will be caused between them.

Despite the above limitations, a first approximation of the response of the system can be generated. It also predicts how many of the actual RF components are required to achieve the necessary gain, band isolation and path compensation. Band slope was not accounted for at this stage, as a greater understanding of the full extent of the slope was first required. The new plots of Figures 7.35 and 7.36 show the full analogue and digital system IF chains. The layouts show a slight excess in gain which can be balanced with the AGC and fixed value attenuators. The analogue layout also shows a total of 2557.5 mm of path compensation from ten devices, which would be sufficient

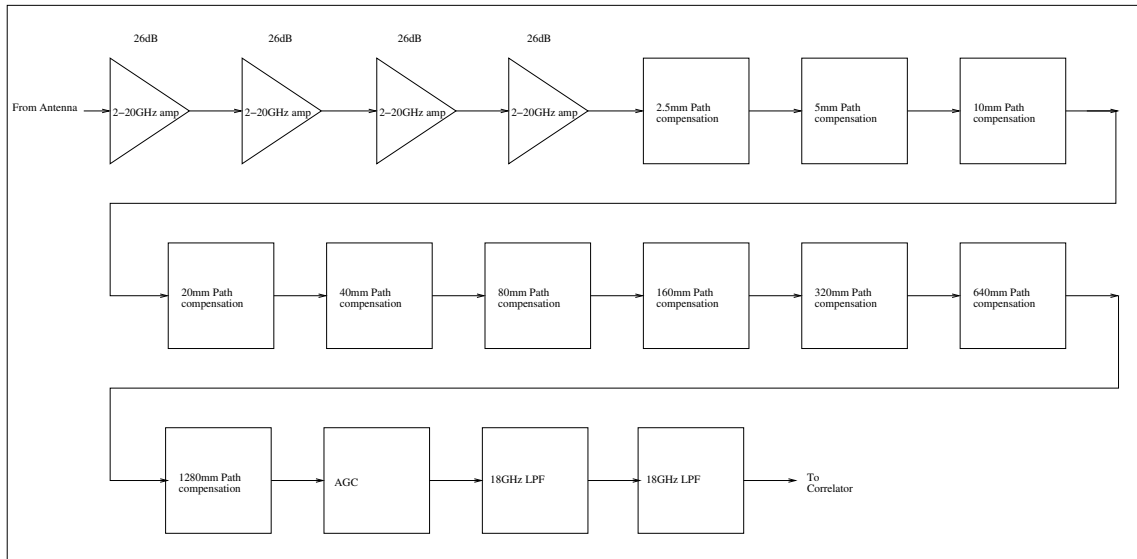


Figure 7.35: Full IF layout for the analogue systems. The total gain is ≈ 104 dB. The system also incorporates ≈ 2557.5 mm of path compensation.

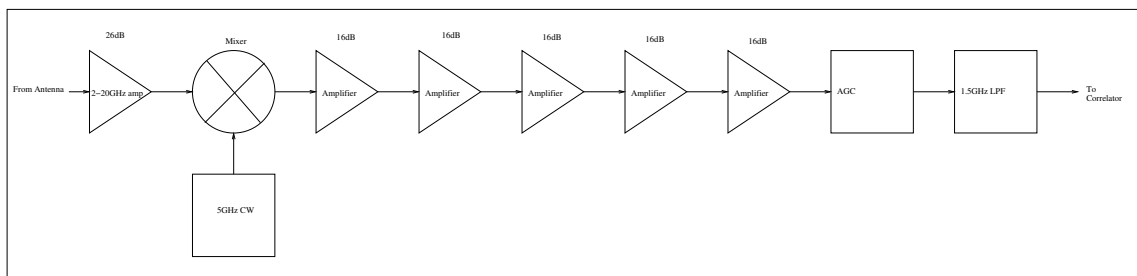


Figure 7.36: Full IF layout for the digital systems. The total gains is ≈ 106 dB. Note the amplifier blocks in the digital system are made from the reduced bandwidth amplifiers previously discussed.

for a short baseline device, such as GUBBINS. Four LPFs are used in the analogue system to generate a sufficient level of band isolation.

Simulated forward transmission characteristics for these IF chains are seen in Figure 7.37. The results show the required gain and passband isolation. The required band slope compensation for the analogue system was then identified from this plot as ≈ 60 dB over the 20 GHz bandwidth.

It was previously found that each slope compensator device generated ≈ 15 dB of positive slope across the 20 GHz band. This means that four slope compensators were required to compensate for the band slope seen.

The final analogue IF chain, incorporating the slope compensator devices, can be seen in Figure 7.38 with the simulated response generated shown in Figure 7.39. This plot shows how the overall band slope has dramatically been reduced. However,




Figure 7.37: Simulation of the IF passband for the (Top) analogue and (Bottom) digital system. These are constructed using the individual passband responses of the components

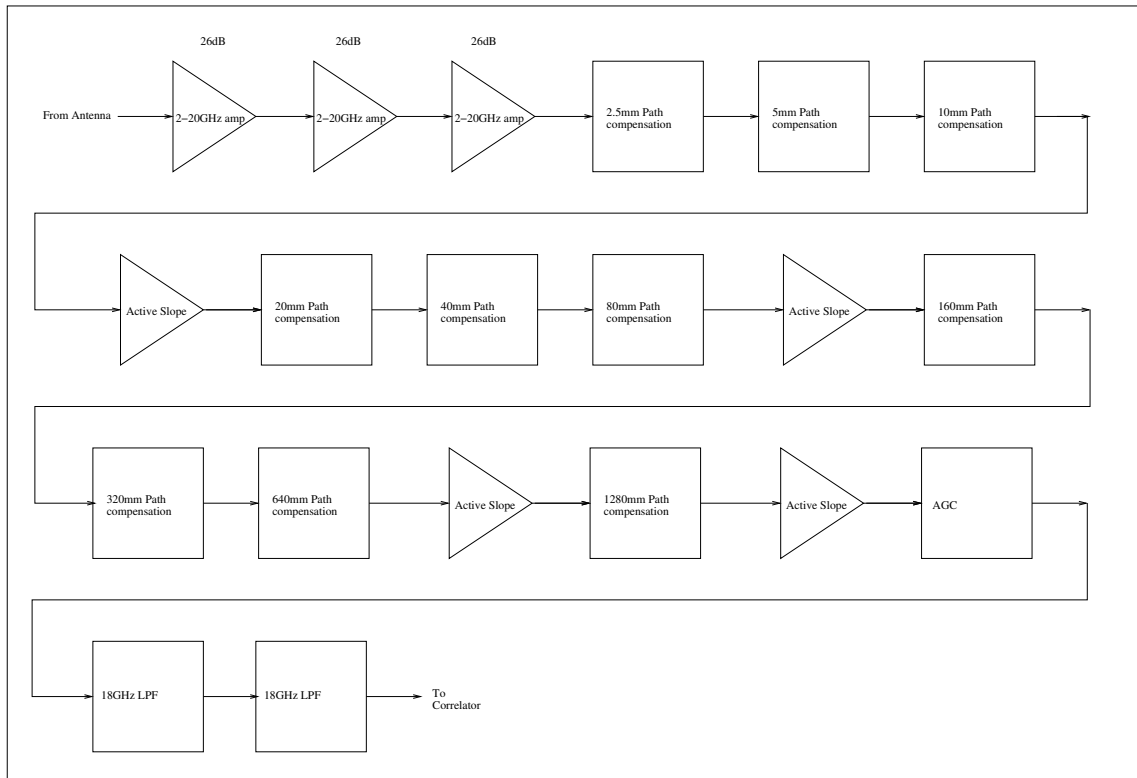


Figure 7.38: Full IF layout for the analogue system with slope compensation.

two problems can be seen. The first of these is a slight remaining negative slope. More slope compensation can be added, however due to its discrete nature another compensator results in a pronounced positive slope. The second problem seen is the presence of large features in the passband. These are of order 10 dB and can be seen across the entire band. As this plot is only generated from simulated data this effect will also be re-investigated when the final IF is constructed as the features seen will likely differ.




Figure 7.39: Simulation of the IF passband for the analogue system with the addition of slope compensation. This is constructed using the individual passband responses of the components. Also shown is the old response without the slope compensation present.

7.3 Conclusion

This chapter has shown the testing and development of the associated ultra-broadband RF and IF devices required in radio astronomy to support a correlator over the 2-20 GHz frequency range. The bandwidth achieved and passband flatness are shown to be acceptable with the addition of a broadband slope compensator.

The design of two IF systems, one for the analogue correlator and one for the digital, was also shown. The simulated overall passband for the systems shows the required bandwidths and also highlights how the slope compensators can be used to counter band slope.

Chapter 8

Conclusion

8.1 Overview

The research of this thesis can be divided into the following three main areas.

8.1.1 Analogue correlator

In Chapter 2, the analogue correlator previously constructed by Christian Holler was tested, including the simulation and implementation of a basic calibration. This showed the operation of the correlator over a reduced 10 GHz bandwidth. The feasibility of calibrating the correlator was also shown using a basic matrix method.

8.1.2 Digital sampling techniques and digital correlator

In Chapter 3, a new sampling technique was described and implemented using the high-speed serial transceivers on a Xilinx Virtex 5 FPGA to act as high-speed direct samplers. The response of the system was seen to be severely limited by the RF connection method used. Despite this connection method, ENOB measurements for the individual samplers were seen to be close to the theoretical maximum across the band (≈ 0.8 bits at 500 MHz). This was compared to an external high-speed sampling technique that produced a slightly improved signal-to-noise ratio.

In Chapter 4, the direct sampling method was extended with interleaving techniques. Two samplers were combined into an interleaved system capable of distinguishing between signals in the first and second Nyquist zones of a single sampler.

In Chapter 5, two samplers were combined in a reference-based multi-level sampler which achieved a maximum ENOB of ≈ 1.2 bits. Four samplers were then combined to make a 6.25 GS/s 1.5-bit sampler. This achieves a slightly lower ENOB performance than expected due to the poor RF connectivity used on the ROACH board.

In Chapter 6, a full six-baseline, 60-lag correlator was constructed from the direct sampling technique and basic tests performed. The correlator was seen to integrate noise down correctly, with a correlated noise floor for the system at -48 dBm. The entire correlator and sampling setup was incorporated onto a single FPGA device, showing the feasibility of incorporating a large amount of the telescope backend electronics into a small environment. This has the effect of reducing the complexity of the system as well as the cost and power usage.

8.1.3 RF and IF device development

In Chapter 7, a selection of ultra-broadband devices were documented that would be necessary for any radio telescope constructed. Those RF devices which had not yet been developed, such as the slope compensator, were also designed and constructed. The slope compensators showed a smooth 15 dB positive gain slope over the 20 GHz band.

A full IF system was also simulated for both the analogue and digital correlators. The simulated system response was seen to have sufficient bandwidth and amplification, with passband slope issues addressed with the use of the slope compensator devices.

8.2 Comparing analogue to digital approach

This thesis describes two distinct correlator designs, a 2-20 GHz analogue correlator and a DC-1.5625 GHz digital correlator. The main difference between the performance of the two systems is the bandwidth achieved. At this stage of development it is not feasible to produce a digital system with a comparable bandwidth to the analogue. This limitation is mostly due to the maximum sampling rate achievable, but also by the amount and speed of the DSP being used.

The analogue correlator also has better signal-to-noise performance as there are no digitisation losses. This disparity would be reduced by moving to higher bit-depth systems with the digital sampling technique used.

The RF signal design of the analogue correlator is simpler than the digital system, as standard RF circuit design techniques can be used to produce a correlator board. No allowance has to be made for the large amount of digital circuitry required for the digital system. Digital boards are now, however, being commercially produced with clean high-frequency inputs and universal connectors. This means only the external interface-board is unique to the sampling technique, which can easily be developed.

Although the design of the analogue correlator circuit is simple, the construction techniques required are more complicated. In the case of the GUBBINS correlator cards, many issues were experienced in their fabrication, for example, the boards were prone to cracking during production. The multiplier chips were also individually bonded to the boards, with it proving difficult to produce a single board with all sixteen multiplier chips successfully bonded. These difficulties are not seen in the digital system as production methods already exist to make the required FPGA boards.

As well as the increased ease of production the digital system also has three main advantages over the analogue. The first of these is the frequency resolution achievable. It is very straightforward, using the digital system, to construct correlators with a fine spectral resolution. This is not practical with the analogue device as a very large correlator board would be required to accommodate all the different lags.

The second advantage of the digital technique is that path compensation and some of the IF filtering can be performed digitally within the FPGA. This simplifies the IF chain used, reducing complexity, loss and importantly, band slope.

Lastly, the digital approach is reconfigurable, whereas the analogue approach is not. When an analogue correlator is constructed, all characteristics are rigidly set by the design. In a digital system certain aspects, such as the frequency resolution, can be changed depending on the requirements of the observation.

8.3 Future research

The research documented in this thesis shows considerable development in multiple fields of study. The following section highlights the key expansions to this research that would be performed if more time was available.

8.3.1 Analogue correlator

The research carried out with the analogue correlator has established basic signal correlation and introduced the concept of calibration in both a theoretical and practical sense. Further research is required in both of these areas to fully establish the capabilities of the system produced.

The main experiment that would be performed is the extension of the reduced bandwidth calibration, using single-frequencies, to a calibration covering the entire frequency band and using all four correlator cards. This would allow the full correlator to be tested in a laboratory setup.

The broadband calibration could also be implemented in a laboratory setup, using broadband noise sources. The equivalence of the two calibration methods could therefore be checked and any practical differences identified.

Finally, on-sky testing of the calibration method could be performed using astronomical signals. This is discussed more in Section 8.3.3.

8.3.2 Digital sampling and correlation

The digital sampling system has received a large amount of development and shows a strong result. However a constant issue was experienced due to the imperfect nature of RF signal injection into the FPGA. To perform further and more conclusive testing of the system, a cleaner RF interface into the high-speed transceivers would be required. One method for providing this would be to move to a ROACH 2 FPGA board¹.

8.3.2.1 ROACH 2

A ROACH 2 FPGA board is very similar to the ROACH board used in the thesis research, but has a newer Virtex 6 FPGA. It also has a mezzanine board interface to the high-speed transceivers. This allows different physical interfaces to be used, by constructing a converter card from the desired interface to the mezzanine connector. On the ROACH 2 board, the signal routing between the mezzanine connector and the FPGA is designed to offer very high signal quality. In this manner it is possible to implement any desired connector standard in a modular fashion, with high data quality achievable. This means that an SMA based mezzanine board could be produced, which would provide far superior signal quality directly into the FPGA than was seen with the ROACH interface used in this thesis.

The ROACH 2 also offers faster serial transceivers, with 36×6.6 GS/s GTX transceivers present in the device [46]. The increase in system resources would also allow a ≈ 15 antenna, 16-lag, correlator to be constructed, instead of the seven antenna current limit. The new system would have a bandwidth of 3.3 GHz and a resolution of ≈ 412 MHz. The VHDL code environment currently used would have to be redeveloped to work with the new hardware, however many of the components required would be common to both systems, simplifying the upgrade process.

¹<https://casper.berkeley.edu/>

8.3.2.2 Multi-bit interleaved correlator

A simple four-input, six-baseline, 1-bit, 1.5625 GHz-bandwidth lag correlator was developed using the current ROACH platform. This system was constructed to allow basic testing of the samplers. If a new ROACH 2 setup was constructed a much larger and more complex correlator could be developed. For example, the 36×6.6 GS/s transceivers could theoretically be linked together to form a single baseline system with anything between a ≈ 60 GHz 1-bit system to a 3.3 GHz 4-bit system.

A plan has always existed to provide a completely reconfigurable system, where the interleaving and bit-depth can be dynamically selected during operation. This could be achieved as the interleaving alignment is provided on-chip and the bias settings could easily be controlled using fast DACs. This would allow different interleaving and bit-depth systems to be selected during operation, with different configurations quickly being tested and compared.

8.3.3 On-sky testing

Finally, during Chapter 5 a full IF system was designed for the two correlators. As all the RF components have been developed it would be possible to construct a simple on-sky test setup. This could either be connected to a current radio telescope, or a custom antenna/mount system could be produced. This would then allow astronomical sources to be observed to fully check the correlators and IF's capabilities for astronomical applications.

Bibliography

- [1] L Ball, R Braun, et al. ATNF Science Priorities Science in 2010 - 2015 . Technical report, Australia Telescope National Facility, 2008. Version 2 (http://www.atnf.csiro.au/observers/planning/more_information.html).
- [2] A. Baudry. The ALMA Correlators: Technical details, Performance and Status of the Main Array Correlator. *ALMA Newsletter*, 7:18–31, January 2011.
- [3] A. J. Beasley and S. N. Vogel. CARMA: specifications and status. In T. G. Phillips and J. Zmuidzinas, editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4855 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 254–264, February 2003.
- [4] R. L. Brown, W. Wild, and C. Cunningham. ALMA - the Atacama large millimeter array. *Advances in Space Research*, 34:555–559, January 2004.
- [5] M. A. Clark, P. C. La Plante, and L. J. Greenhill. Accelerating Radio Astronomy Cross-Correlation with Graphics Processing Units. *ArXiv e-prints*, July 2011.
- [6] R. P. Escoffier, G. Comoretto, J. C. Webber, A. Baudry, C. M. Broadwell, J. H. Greenberg, R. R. Treacy, P. Cais, B. Quertier, P. Camino, A. Bos, and A. W. Gunst. The ALMA correlator. *Astronomy and Astrophysics*, 462:801–810, February 2007.
- [7] AMI Consortium: Zwart et al. The Arcminute Microkelvin Imager. *Monthly Notices of the Royal Astronomical Society*, 391:1545–1558, December 2008.
- [8] Australia Telescope National Facility. Wideband Correlator. Technical report, Australia Telescope National Facility, Dec 1999. Preliminary (<http://www.atnf.csiro.au/technology/electronics/docs/sest/sest.html>).

- [9] S. Garrington, editor. *e-MERLIN & EVN, ALMA Commissioning Skills Workshop*, 2012. AdvancedRadioAstronomyMaterial (<http://www.alma.ac.uk/events/past-events/advancedradioastronomymaterial/>).
- [10] B. Gilbert. A precise four-quadrant multiplier with subnanosecond response. *IEEE J. Solid-State Circuits*, 3:365-373, 1968.
- [11] T. Girard. Understanding effective bits. Technical report, Signatec. Application Note: AN95091.
- [12] P. K. Grimes et al. GUBBINS: A Novel Millimeter-Wave Heterodyne Interferometer. In E. Bryerton, A. Kerr, and A. Lichtenberger, editors, *Twentieth International Symposium on Space Terahertz Technology*, page 82, April 2009.
- [13] A. I. Harris. Gilbert cell multiplier measurements from 2-18.5 GHz. Technical report, University of Maryland.
- [14] A. I. Harris. Gilbert cell multiplier measurements III: Response from 2-30 GHz. Technical report, University of Maryland.
- [15] F. J. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66-1:51-83, January 1978.
- [16] David Hawkins. Private communication.
- [17] Hittite Microwave. *HMCAD5831LP9BE: 26 GS/s 3-bit analog-to-digital converter w/ overrange, inhibit, and 1:2 DEMUX*. Datasheet.
- [18] C. Holler. *Correlator and antenna design for the Arcminute Microkelvin Imager (AMI)*. PhD thesis, Astrophysics Group, Cavendish Laboratory, 2003.
- [19] C. M. Holler et al. A 6-12 GHz analogue lag-correlator for radio interferometry. *Astronomy and Astrophysics*, 464:795-806, March 2007.
- [20] C. M. Holler et al. A 2-20 GHz Analog Lag-Correlator for Radio Interferometry. *IEEE transactions on instrumentation and measurement*, VOL. 61, NO. 8, August 2012.
- [21] C. M. Holler, M. E. Jones, A. C. Taylor, A. I. Harris, and S. A. Maas. A 2-20 GHz Analog Lag-Correlator for Radio Interferometry. Technical report, University of Oxford.

- [22] Christian Holler. Private communication.
- [23] P Horowitz and W Hill. *The Art of Electronics*. Cambridge University Press, 1989.
- [24] O. G. King et al. The C-Band All-Sky Survey (C-BASS): Design and implementation of the northern receiver. *ArXiv e-prints*, October 2013.
- [25] Oliver King. Private communication.
- [26] C.-T. Li et al. AMiBA Wideband Analog Correlator. *The Astrophysical Journal*, 716:746–757, June 2010.
- [27] K.-Y. Lin et al. AMiBA: System Performance. *The Astrophysical Journal*, 694:1629–1636, April 2009.
- [28] K. Y. Lo et al. AMiBA: Array for microwave background anisotropy. In J. C. Wheeler and H. Martel, editors, *20th Texas Symposium on relativistic astrophysics*, volume 586 of *American Institute of Physics Conference Series*, pages 172–177, October 2001.
- [29] E Momjian and R Perley. Comparison of the performance of the 3-bit and 8-bit samplers at C(4-8GHz), X(8-12GHz) and Ku(12-18GHz) bands. Technical report, NRAO, 2013. EVLA memo 166 (<http://www.aoc.nrao.edu/evla/memolist.shtml>).
- [30] S Padin et al. A Wideband Analog Correlator for Microwave Background Observations. *IEEE transactions on instrumentation and measurement*, 50, October 2001.
- [31] A. Parsons, D. Backer, A. Siemion, H. Chen, D. Werthimer, P. Droz, T. Filiba, J. Manley, P. McMahon, A. Parsa, D. MacMahon, and M. Wright. A Scalable Correlator Architecture Based on Modular FPGA Hardware, Reuseable Gateway, and Data Packetization. *Publications of the Astronomical Society of the Pacific*, 120:1207–1221, November 2008.
- [32] R Perley, editor. *The Jansky Very Large Array: An Exemplar Array*, ALMA Commissioning Skills Workshop, 2012. AdvancedRadioAstronomyMaterial (<http://www.alma.ac.uk/events/past-events/advancedradioastronomymaterial/>).

- [33] R. A. Perley, C. J. Chandler, B. J. Butler, and J. M. Wrobel. The Expanded Very Large Array: A New Telescope for New Science. *apjl*, 739:L1, September 2011.
- [34] R. A. Perley, C. J. Chandler, B. J. Butler, and J. M. Wrobel. The Expanded Very Large Array: A New Telescope for New Science. *The Astrophysical Journal Letters*, 739:L1, September 2011.
- [35] A Pham et al. Development of integral passive components for multilayer organic MCMs at millimeter wave frequencies. *IEEE Transactions on Advanced Packaging*, 25:98–101, February 2002.
- [36] D Pozar. *Microwave Engineering*. Wiley, 2005.
- [37] K Rauch. Third-generation CARMA correlator FPGA configuration. Technical report, CARMA, 2013. CARMA Memo 58 (<http://www.astro.umd.edu/kzhang/carmaweb/memos/>).
- [38] M Revnell. The EVLA 8-BIT digitizer performance. Technical report, NRAO, 2005. EVLA memo 98 (<http://www.aoc.nrao.edu/evla/memolist.shtml>).
- [39] P Roberts. SEST sampler. Technical report, Australia Telescope National Facility, 1998. Preliminary (<http://www.atnf.csiro.au/people/Paul.Roberts/sestsam.pdf>).
- [40] G Taylor, C Charilli, and R Perley. *Synthesis Imaging In Radio Astronomy II*. Sheridan books, 2008.
- [41] A Thompson, J Moran, and G Swenson jr. *Interferometry and Synthesis in Radio Astronomy*. Wiley-Interscience, 2001.
- [42] P. S. Udomprasert, B. S. Mason, and A. C. S. Readhead. The Sunyaev-Zel’dovich Effect with the Cosmic Background Imager. In *Constructing the Universe with Clusters of Galaxies*, 2000.
- [43] W. E. Wilson et al. The Australia Telescope Compact Array Broad-band Backend: description and first results. *Monthly Notices of the Royal Astronomical Society*, 416:832–856, September 2011.
- [44] Xilinx. Virtex-5 Family Overview. Technical report, Xilinx, 2009.

- [45] Xilinx. Virtex-5 FPGA RocketIO GTP Transceiver: User Guide. Technical report, Xilinx, 2009.
- [46] Xilinx. Virtex-6 Family Overview. Technical report, Xilinx, 2012.