

RADIO ASTRONOMY INSTRUMENTATION FOR REDSHIFTED HYDROGEN LINE SCIENCE



DANNY PRICE
WORCESTER COLLEGE
UNIVERSITY OF OXFORD

A THESIS SUBMITTED FOR THE DEGREE OF
Doctor of Philosophy

TRINITY 2012

You'll wait a long, long time for anything much
To happen in heaven beyond the floats of cloud
And the Northern Lights that run like tingling nerves.
The sun and moon get crossed, but they never touch,
Nor strike out fire from each other nor crash out loud.
The planets seem to interfere in their curves
But nothing ever happens, no harm is done.
We may as well go patiently on with our life,
And look elsewhere than to stars and moon and sun
For the shocks and changes we need to keep us sane.
It is true the longest drought will end in rain,
The longest peace in China will end in strife.
Still it wouldn't reward the watcher to stay awake
In hopes of seeing the calm of heaven break
On his particular time and personal sight.
That calm seems certainly safe to last to-night.

- Robert Frost

Acknowledgements

Throughout the ages, curiosity has driven us: to eat forbidden fruit, to sail to the edge of the world, to boldly go where no man has gone before, and to put cats in boxes then question if they're alive or not. From this perspective, it is not peculiar that scientists strive to build ever more sensitive instruments to peer further into the voids of space.

I have spent the last four years of my life building things to detect and measure the invisible. At times I have lost sight as to why we concern ourselves with such matters, with my motivation drifting down to match those frequencies too low to see. Thank you Sascha for upconverting my motivation back into the visible at these times, and for being a buddy, mentor and all-round bonza bloke. And thank you Mike for your hyphenation-prowess and for funding me these past few years.

I thank the level-3 mariners (Griffin, Charles and Jack), for their friendship, the Keyhaven Incident, and for helping me build the instruments with which to explore the invisible world of radio. I must similarly thank the level-6+ dwellers Ian, Aris, Kris, and Richard for their comradery, wisdom, and excellent company (particularly during Friday tapas). A huge thank you also goes to Adam for his musings on microwave design and the nature of the Universe. And thank you, Steve, for being an inspiration and role model; I will be thinking of you when the first photons fill the SKA with faint whispers of our enigmatic Universe.

I am fortunate to have also had much support outside the walls of the DWB. Thank you Ben and Fred for your sublime OSKAR support and professional coding. Thank you to the HIPSR collaboration — particularly Andrew, Brett, Ettore, Matt, JR and Lister — for your hard work in getting HIPSR off the ground. And thanks to Lister and the two ICRAR tribes for hosting me when back at home, and to Jacinta for being the ICRAR-Oxford counterpoint in the SKA medley.

Finally, I thank those who have stood with me as spacetime shifts beneath our feet. Thank you to Mum & Dad, for nurturing me and shaping who I am, and for the infinite wisdom of Desiderata and the I Ching. Jen: thanks for cooking me two-minute noodles after school and being a wonderful sister. Finally, thank you to my beautiful wife, Lindall, for keeping me sane and making my life a joy. I owe you a puppy.

And so I proclaim the end of 21 years of study, a year for every centimetre of the HI line. Thank you for dealing me blackjack, I love you all.

RADIO ASTRONOMY INSTRUMENTATION FOR REDSHIFTED HYDROGEN LINE SCIENCE

Daniel Charles Price

Worcester College, University of Oxford

A thesis submitted for the degree of Doctor of Philosophy in Trinity Term 2012

ABSTRACT

This thesis presents instrumentation with which to measure the abundance of neutral hydrogen gas in the Universe. Measuring where the Universe's hydrogen is, and tracing how its distribution evolves with time, holds the key to understanding how galaxies evolve, the nature of dark energy, and how the first cosmic structures formed. In particular, this thesis looks at instrumentation for 21-cm intensity mapping telescopes. In 21-cm intensity mapping, the collective emission of many galaxies is measured, without individual detections (Peterson et al., 2009). This technique promises to allow detection of the baryonic acoustic oscillation peaks in the power spectrum of the Universe's matter distribution. Such a detection would increase constraints on cosmological parameters.

There are two main approaches to designing a 21-cm intensity mapping instruments: using a filled aperture instrument such as a single-dish telescope, or using a sparse aperture instrument such as an interferometric array of dipoles. This thesis investigates analogue components for a sparse aperture instrument operating at 1.0-1.5 GHz. As part of this work, a 16-element sparse aperture array was designed and constructed. To test the array's performance, field testing was conducted; the results of which are presented here.

In addition to this, I have designed a new digital spectrometer for redshifted hydrogen line science, named HISPEC. A copy of this spectrometer has been installed on the Parkes 64 m telescope, as a digital signal processor for the 21-cm multibeam receiver. HISPEC has increased instantaneous bandwidth, higher interchannel isolation, and improved quantization efficiency as compared to the existing backend, MBCORR. The HISPEC equipped multibeam receiver is an ideal instrument for 21-cm intensity mapping at redshifts $z < 0.2$.

CONTENTS

1	Introduction	1
1.1	Introduction	1
1.2	A brief history of the Universe	2
1.2.1	Cosmic microwave background	3
1.2.2	Epoch of Reionization	4
1.2.3	Baryonic acoustic oscillations	5
1.3	Measuring BAO with the 21-cm line	6
1.3.1	The 21-cm line	7
1.3.2	21-cm intensity mapping	9
1.3.3	Epoch of reionization instruments	10
1.3.4	Cross-correlated intensity mapping experiments	11
1.3.5	Power spectrum interferometers	11
1.4	Radio astronomy instrumentation	13
1.4.1	Aperture arrays	13
1.4.2	Sparse arrays	15
1.4.3	Interferometric arrays	17
1.4.4	Interferometric power spectrum measurement	18
1.5	Sensitivity and survey speed	19
1.5.1	Point source survey speed	21
1.5.2	Extended source survey speed	22
1.6	HI instruments to measure BAO	23
1.6.1	Single dish case	24
1.6.2	Sparse array case	24
1.6.3	Interferometer case	26
1.6.4	BX case	27
1.6.5	Comparison of instrumentation approaches	28
1.6.6	Proposed intensity mapping instruments	28
1.7	Thesis layout	29

2	Coherency in Radio Astronomy	31
2.1	Introduction	31
2.2	Properties of electromagnetic waves	32
2.2.1	Polarization	33
2.3	Jones and Mueller RIME formulations	34
2.3.1	Hamaker's RIME derivation	35
2.3.2	The 2×2 RIME	37
2.3.3	Applications of the RIME	38
2.3.4	Limitations of the RIME	39
2.4	A limitation of Jones chain decomposition	40
2.4.1	Changes in impedance and the transmission matrix	42
2.5	Coherency in radio astronomy	43
2.5.1	Electromagnetic coherency	43
2.5.2	Voltage and current coherency	45
2.6	Two point coherency relationships	46
2.6.1	A general two point coherency relationship	46
2.6.2	The Radio Interferometer Measurement Equation	48
2.6.3	A 2N-port transmission matrix based RIME	49
2.6.4	The Van-Cittert-Zernicke relationship	50
2.7	Discussion	51
2.7.1	Modelling real analogue components	51
2.7.2	Reflection coefficient	52
2.7.3	Application to calibration	53
2.7.4	Application to end-to-end simulations	54
2.7.5	Conclusions	54
3	D-PAD: An L-band sparse aperture array demonstrator	55
3.1	Introduction	55
3.1.1	Danny's PhD array demonstrator	55
3.2	System overview	57
3.2.1	Signal conditioning requirements	58
3.2.2	Beam pattern	59
3.3	Log periodic dipole array antennas	59
3.3.1	Design and simulation	61
3.3.2	Feeding technique	62
3.3.3	Fabrication	63
3.3.4	Gain and beam pattern measurements	64

3.4	Receiver board	68
3.4.1	Amplifier selection and characterization	70
3.4.2	Circuit board design	73
3.4.3	Performance	73
3.5	Analogue beamformer	74
3.6	Analogue chain	75
3.6.1	Bandpass filter	75
3.6.2	Amplifier	76
3.6.3	Transmission lines	79
3.7	System performance	79
3.7.1	System bandpass	79
3.7.2	Receiver temperature	81
3.7.3	Overall system temperature	82
3.8	Discussion	83
3.8.1	Future work	83
3.9	Conclusions	84
4	Wide-bandwidth digital spectrometers	85
4.1	Introduction	85
4.2	Digital spectrometer implementations	86
4.2.1	Windowing functions	87
4.2.2	Spectrometers for radio astronomy	88
4.3	Polyphase filterbanks	90
4.3.1	Finite impulse response filters	90
4.3.2	Polyphase FIR filters	91
4.3.3	DFT based polyphase filterbanks	92
4.4	Field programmable gate arrays	94
4.5	DSP hardware for radio astronomy	96
4.5.1	Virtex-5 FPGA	96
4.5.2	The CASPER ROACH board	97
4.5.3	iADC digitizer card	99
4.5.4	The Xilinx toolflow	100
4.6	The HISPEC spectrometer	101
4.6.1	Firmware overview	102
4.6.2	Firmware compilation	105
4.7	Results	107
4.8	Discussion	110

4.8.1	Comparison to existing instruments	111
4.8.2	Next generation spectrometers	112
4.9	Conclusions	114
5	D-PAD single station results	115
5.1	Introduction	115
5.2	Deployment	116
5.3	Results	116
5.3.1	GPS satellite measurements	119
5.3.2	Astronomical sources	122
5.4	Discussion	123
5.4.1	Beam pattern	123
5.4.2	System temperature	124
5.4.3	Sensitivity	128
5.4.4	Astronomical sources	130
5.5	Conclusions	131
5.5.1	Future work	131
6	The Parkes HIPSR Project	133
6.1	Introduction	133
6.1.1	The Parkes 64 m telescope	133
6.1.2	Scientific motivation	135
6.2	Initial modes of operation	136
6.2.1	High spectral resolution mode	137
6.2.2	High time resolution mode	137
6.3	Network architecture	138
6.4	Hardware specifications	138
6.5	Monitor, acquisition and control software	141
6.5.1	Data storage	141
6.5.2	User interface	141
6.6	Installation	141
6.6.1	Commissioning tests	142
6.6.2	First light bandpass	144
6.7	Spectral line observations	147
6.7.1	Calibration method	149
6.7.2	Flux scale calibration	151
6.7.3	HIZOA J0836-43	152

6.8	Future Work	156
6.9	Conclusions	157
7	Conclusions	158
7.1	Discussion	159
7.2	Future work	160
7.2.1	Improvements to the HISPEC spectrometer	160
7.2.2	BAO detection with HIPSr	160
7.2.3	A D-PAD based power spectrum interferometer	161
7.2.4	BAO-PAD	162
7.2.5	BAO-PAD sensitivity	163
7.3	Concluding remarks	163
	References	165
A	Schematics	176

CHAPTER 1

INTRODUCTION

1.1 INTRODUCTION

The Square Kilometre Array telescope (SKA, Dewdney et al., 2009) is one of the major science projects of the 21st century . It aims to answer five of the fundamental questions of the Universe:

1. How do galaxies evolve, and what is dark energy?
2. How did the first stars and cosmic structures form?
3. Does General Relativity hold in extreme gravitational potentials?
4. What is the origin of cosmic magnetism and how does it evolve?
5. Is there life on other worlds?

Answering these questions will be a giant leap for our understanding of the Universe. To do so will require the SKA to be orders of magnitude more sensitive than existing telescopes. As the SKA is such an ambitious project, it will take at least a decade from now to complete (completion date is scheduled for 2024, with first science in 2019). In the meantime, there is room for smaller, more targeted experiments to shed light on these fundamental questions.

During my DPhil I have focused on designing instrumentation with which to measure the abundance of neutral hydrogen gas in the Universe. Measuring where the Universe's hydrogen is, and tracing how its distribution evolves with time, holds the key to unlocking the first two of these fundamental questions. To explain why, I will first introduce the history of the Universe as




Figure 1.1: Artist's conception of the evolution of the Universe. As stars are born, the regions immediately around them are ionized. Eventually the intergalactic gas becomes entirely ionized. Image from Loeb (2006).

we currently understand it, before delving more deeply into how we can constrain the cosmology of the Universe with measurements of hydrogen emission.

1.2 A BRIEF HISTORY OF THE UNIVERSE

Cosmology — the study of the Cosmos — aims to explain how the Universe around us formed. Any explanation needs to be well supported by observational data; our current best-fit model of the Universe predicts that it was formed in a 'Big Bang', 13.7 billion years ago (Weinberg, 2008). More precisely, we believe that the Universe began in an extremely compact, hot and dense state, which expanded rapidly and consequently cooled.

After an initial stage of inflation, the Universe remained extremely hot and consisted of a quark-gluon plasma and elementary particles: a 'quark soup'. By some unknown violation of baryon number conservation (baryogenesis), a tiny excess of quarks and leptons was established over their antiparticles. As the Universe continued to expand and cool, quarks and gluons combined to form baryons (protons and neutrons) and antibaryons (antiprotons and antineutrons), which proceeded to annihilate, as did electron and positron pairs. All antimatter was destroyed, leaving the matter dominated Universe that we see today.

All of this happened within the first second of the Universe's life. It was not until 379,000 years later that electrons and nuclei combined to form atoms. This is known as the recombination epoch. A direct consequence of recombination was that the mean-free path for photons became

Table 1.1: The six parameters of the Λ CDM model. Values listed from 7-year WMAP results (Table 8 in Jarosik et al., 2011).

Parameter	Description	Value
$\Omega_b h^2$	Physical baryon density	0.0260 ± 0.00053
$\Omega_c h^2$	Physical dark matter density	0.1123 ± 0.0035
Ω_Λ	Dark energy density	$0.728^{+0.015}_{-0.016}$
Δ_R^2	Curvature fluctuation amplitude	2.441×10^{-9}
n_s	Scalar spectral index	0.963 ± 0.012
τ	Reionization optical depth	0.087 ± 0.014

effectively infinite, causing the decoupling of photons and matter; before this, the Universe was opaque. The photons from this decoupling epoch are thus the first source of radiation that we may probe and measure with telescopes. This signal is known as cosmic background radiation.

1.2.1 COSMIC MICROWAVE BACKGROUND

Since recombination, the Universe has continued to expand. As a result, the cosmic background radiation has been doppler shifted into the microwave regime. Because of this, the signal is more commonly referred to as the cosmic microwave background (CMB). It was first detected in 1965 by Arno Penzias and Robert Wilson (Penzias and Wilson, 1965; Dicke et al., 1965), who received the 1978 Nobel prize for their detection.

The CMB has a nearly perfect black-body radiation spectrum, with a measured temperature of 2.73 K (decreased from an original 3000 K due to the expansion factor of the Universe). It is almost isotropic, but has variations of order a part in 10^5 . These anisotropies encode information about the early Universe. Their existence is well explained by acoustic waves propagating through the baryonic plasma of the early Universe (before recombination). Missions such as the Cosmic Background Explorer (COBE, Smoot et al., 1992), and Wilkinson Map Anisotropy Probe (WMAP, Larson et al., 2011), have taken exquisitely sensitive measurements ($<35 \mu\text{K}$ noise per pixel), of the observed CMB anisotropies. These measurements support Big Bang based cosmological models, and allow us to reject many competing models. The current frontrunner is the Lambda Cold Dark Matter model (Λ CDM), which includes not only the expansion of the Universe, but also the existence of dark energy and dark matter. The Λ CDM model (also known as the standard model of cosmology), is based upon six parameters, from which values such as the Hubble constant and the age of the Universe can be derived. These parameters are shown




Figure 1.2: Angular power spectrum of the cosmic microwave background. The peaks and dips observed in the data are well explained by the acoustic oscillations predicted by the Λ CDM model. Reproduced from Bennett (2006).

in Table 1.1.

Figure 1.2 shows the power spectrum of the CMB, as measured with WMAP and supplementary surveys (points), and as predicted by the Λ CDM cosmological model (red line). The x-axis corresponds to angular scale, and the y-axis is a measure of anisotropy power, in μK . The secondary x-axis, multipole moment, refers to spherical-harmonic multipole moment expansion of the power spectrum. As can be seen, the measured WMAP data are in excellent agreement with the model. The peaks and dips observed correspond to acoustic oscillations in the primordial baryonic plasma, a key prediction of Λ CDM.

1.2.2 EPOCH OF REIONIZATION

Following from recombination is a period known as the Dark Ages. During this epoch, the structure of the Universe was forming through gravitational collapse, but stars had not yet reached the critical mass required for nuclear fusion to begin. The galaxies which we see today formed through gravitational collapse in regions of slight matter over-density, caused by the acoustic oscillations of the early Universe. One by one, the first stars of the Universe switched on, and the Dark Ages transitioned into the ‘Epoch of Reionization’ (EoR). This epoch saw the atoms formed in recombination split (ionized) back into their constituent protons and electrons via radiated

energy from the first stars. An artist's conception of this process is shown in Figure 1.1.

Probing the Dark Ages and Epoch of Reionization is an area of active research within radio astronomy. The LOFAR (de Vos et al., 2009), PAPER (Parsons et al., 2009), and MWA (Lonsdale et al., 2009) telescopes are all searching for the faint signal produced by atomic hydrogen. Detecting this signal over a range of different frequencies (corresponding to different doppler shifts due to expansion) would allow a tomographic map of the atomic hydrogen distribution to be created (Morales and Wyithe, 2010). By analysing this data we will be able to test models and simulations of how the reionization process occurred, and how the first stars and galaxies formed.

1.2.3 BARYONIC ACOUSTIC OSCILLATIONS

The acoustic oscillations of the early Universe are not only encoded into the CMB: they have left their imprint on the large scale matter distribution at all epochs. After reionization, the acoustic waves of the early Universe could no longer propagate. Any wave crests present at this time left behind over-densities of baryonic matter, and similarly rarefactions left behind under-densities. These are referred to as Baryonic Acoustic Oscillations (BAO). That is, the BAO are the peaks and troughs in the angular power spectrum of baryonic matter which correspond to over- and under-densities of matter at specific length scales.

The Sloane Digital Sky Survey (SDSS) detected the BAO through a wide-field optical survey of galaxies (Eisenstein et al., 2005; Percival et al., 2010). This experiment was the first to detect the BAO signature (at a redshift of $z=0.275$). Figure 1.3 show the BAO peak in SDSS data. The bump in the correlation function indicates that there is a higher probability of galaxies being separated at a particular length scale than would be predicted without BAO (pink line). The SDSS data are compared against Λ CDM predictions with a different cosmological parameter, $\Omega_m h^2$; this value corresponds to the physical matter density of the Universe, with Ω_m representing total matter density (dark + baryonic matter + neutrino), and h is the reduced Hubble constant, equal to about 0.7 (km/s)/Mpc.

Combining the anisotropy data from the CMB with BAO data from a later epoch gives constraints on the values of cosmological parameters such as Ω_m . This is as measurements of the BAO at different epochs constitute a 'standard ruler': an object whose size relative to the Uni-




Figure 1.3: Baryon acoustic oscillation peak in SDSS data (Eisenstein et al., 2005). Reproduced from Bennett (2006).

verse remains fixed with the expansion of the Universe (Blake and Glazebrook, 2003; Eisenstein et al., 2005). As of writing, the most recent values for cosmological parameters, as derived from SDSS and WMAP, are given in Jarosik et al. (2011) and Anderson et al. (2012). These values are reproduced in Table 1.1. Honing these values by removing systematic biases and making ever more precise measurements will test the standard model of cosmology and provide us with a deeper understanding of the Universe that engulfs us.

1.3 MEASURING BAO WITH THE 21-CM LINE

An alternative approach to optical galaxy surveys (such as SDSS) is to detect the signal from neutral atomic hydrogen (HI). Hydrogen is the most abundant element in the Universe, and as such accounts for the largest fraction of baryonic mass. HI can be directly detected by its radio emission; by measuring this emission we can infer the kinematics and spatial distribution of HI in the Universe. Loeb and Wyithe (2008) find that a redshifted hydrogen based detection of BAO would constrain the matter power spectrum 1–2 orders of magnitude tighter than optical galaxy surveys such as SDSS–LRG.

There are two approaches which may be employed to extract BAO from redshifted hydrogen measurements. The first is a galaxy redshift survey, in which millions individual galaxies are

detected, and a 3-D map of the Universe is produced. The BAO will then manifest itself as over- and under-densities within this volume. An exciting possibility for detecting BAO through neutral hydrogen is a billion-galaxy spectroscopic redshift survey in HI with the Square Kilometer Array (Rawlings et al., 2004). Such a survey could detect the BAO with very low systematic errors for cosmology.

The second approach is known as 21-cm intensity mapping (Peterson et al., 2009). In this technique, only the the large scale structure is resolved. The advantage of this approach is that HI emission from galaxies below the detection threshold in a redshift survey is not discarded; as a result, less collecting area and sensitivity is required to extract the BAO signal. This technique may allow BAO detection with an instrument costing a fraction of that of the SKA (Chang et al., 2008; Wyithe et al., 2008).

Intensity mapping instrumentation can be divided into filled aperture telescopes, such as single parabolic dishes, and sparse aperture telescopes, such as interferometers. Filled aperture instruments seek to maximize the survey speed for brightness sensitivity, which is achieved by increasing collecting area and decreasing system temperature (see Section 1.5). Survey speed can also be increased by using a multi-pixel feed, so that the instrument has multiple independent fields of view. In contrast, each pair of antennas within an interferometric array is only sensitive to a subset of angular power spectrum modes. As a result, many pairs of antennas (with different inter-antenna separations) are required if one wishes to sample a large fraction of power spectrum modes. A more detailed comparison of the two approaches is given in Section 1.6.

Exploring the viability of intensity mapping instruments for BAO is a key motivation of this thesis. In the section that follows, the underlying physics behind the HI signal is introduced, along formulae with which to estimate the signal strength from clouds of hydrogen gas.

1.3.1 THE 21-CM LINE

The hydrogen line, 21-centimeter line or HI line are all aliases given to the electromagnetic radiation spectral line that is created by a change in the energy state of neutral hydrogen atoms. More specifically, a photon is emitted when an electron's spin flips from parallel with that of the proton, to antiparallel. This emission is known as the hyperfine line and has a central frequency

ν_{10} of

$$\nu_{10} = \frac{8}{3} g_1 \left(\frac{m_e}{m_p} \right) \alpha^2 (R_M c) \approx 1420.41 \text{ MHz}, \quad (1.1)$$

which in free-space is equivalent to a wavelength $\lambda=21$ cm. Here, g_1 is the (dimensionless) magnetic moment for a proton, α is the fine-structure constant, m_e and m_p are the masses of an electron and proton, R_M is the Rydberg constant, and c is the speed of light (Chapter 7, Condon and Ransom, 2010).

To measure the BAO signal using redshifted hydrogen, the 21-cm emission from clouds of HI must be distinguishable from the CMB. The optical depth τ_ν of a cloud of HI is

$$\tau_\nu = \frac{3}{32\pi} \frac{h_P c^3 A_{10}}{k_B T_S \nu^2} \frac{\bar{x}_{HI} n_H}{(1+z)(d\nu_{\parallel}/dr_{\parallel})}, \quad (1.2)$$

where

- h_P is Planck's constant,
- k_B is the Boltmann constant,
- $A_{10} = 2.85 \times 10^{-15} \text{ s}^{-1}$ is the spontaneous emission coefficient of the 21-cm transition,
- T_S is the spin temperature of the HI,
- z is redshift,
- n_H is the density of Hydrogen,
- \bar{x}_{HI} is fraction of non-ionized atomic hydrogen, and
- $d\nu_{\parallel}/dr_{\parallel}$ is the gradient of the proper velocity along the line of sight (Furlanetto et al., 2006; Morales and Wyithe, 2010).

That is, optical depth is proportional to the density of HI and inversely proportional to redshift. The fluctuations in brightness temperature are then largest for clouds with high optical depth:

$$\Delta T = \frac{T_S - T_{CMB}(z)}{(1+z)} (1 - e^{-\tau_\nu}) \approx \frac{T_S - T_{CMB}(z)}{(1+z)} e^{-\tau_\nu}, \quad (1.3)$$

where $T_{CMB}(z)$ is the brightness temperature of the CMB as a function of redshift. As τ_ν is a function of \bar{x}_{HI} , the brightness temperature is expected to change rapidly as atomic hydrogen is

ionized during the EoR, at about $z \sim 6-7$. Arrays such as LOFAR, MWA and PAPER are expected to detect this global step in brightness temperature.

After the epoch of reionization, HI was mainly confined to overdense regions of the intergalactic medium. Galaxy surveys such as HIPASS have demonstrated that most HI is associated with galaxies (Meyer et al., 2004). However, as HIPASS was sensitivity limited, it only reaches out to redshift of $z \approx 0.04$, and as such there are not enough galaxies in the dataset to extract BAO. Large scale galactic redshift surveys, such as a 20,000 square degree survey out to $z \sim 2$ with the SKA, will have the sensitivity and resolution to make the BAO detection through statistical analyses of the spatial distribution of detected galaxies (Abdalla and Rawlings, 2005; Abdalla et al., 2010). Unfortunately, the SKA will not be fully operational until at least 2024. In the meantime, there is room for smaller, targeted HI BAO experiments. This will, however, require a different technique: 21-cm intensity mapping.

1.3.2 21-CM INTENSITY MAPPING

In 21-cm intensity mapping, the collective emission of many galaxies is measured, without individual detections (Peterson et al., 2009). The general idea is that as we are only interested in structure on the scale of BAO, we need just enough resolution to detect the BAO peaks. As resolution decreases, the volume within an instrument's field of view increases, and as such the total intensity of emission increases. This total intensity signal is significantly easier to detect than the signal from individual galaxies.

However, the combined emission from HI is still far less bright than astrophysical foregrounds (primarily synchrotron), which will contaminate the signal. These foreground sources must be subtracted to recover the desired HI emission. Several authors propose that the synchrotron sources may be removed by fitting a power-law spectra (Chang et al., 2008; Peterson et al., 2009). The 21-cm signal should vary rapidly as a function of frequency, so should be separable from the smooth power-law spectra from foregrounds. Foreground removal will require precision calibration and well-understood systematics, which must be taken into consideration when designing an intensity mapping experiment.

So how much resolution does a 21-cm intensity mapping instrument require? Let us first consider a filled aperture approach, with a single-dish telescope. Above the third BAO peak, the

Table 1.2: Relationship between redshifted, observational frequency, and required dish diameter to maintain resolution to resolve third BAO peak.

Redshift z	Frequency MHz	Required dish diameter metres
0	1420	42
1	710	84
2	470	127
4	280	214
6	203	294

structure of the BAO is washed out by non-linear evolution (Chang et al., 2008). If we take this as the smallest spatial scale we need to consider, we require enough resolution to Nyquist sample wavelengths of $35h^{-1}$ Mpc (h is the dimensionless Hubble constant). This corresponds to 20 arcminutes at $z=1.5$, at which point the HI line is redshifted to about 570 MHz. If a parabolic dish was to be used, we would need a diameter of about 100 metres to reach this resolution. This is equal to the largest steerable single-dish telescopes in operation. A table relating redshift, frequency and the required single-dish telescope diameter to maintain Nyquist sampling is shown in Table 1.2.

There is a second requirement of any intensity mapping experiment: that our instrument is stable and sensitive enough to measure brightness fluctuations on BAO levels. The average sky brightness temperature due to 21-cm emission is about $300 \mu\text{K}$, and can be estimated using

$$T_b = 0.3 \left(\frac{\Omega_{\text{HI}}}{10^{-3}} \right) \left(\frac{\Omega_m + a^3 \Omega_\Lambda}{0.29} \right)^{-1/2} \left(\frac{1+z}{2.5} \right) \text{mK}, \quad (1.4)$$

where Ω_{HI} is the cosmic density of neutral hydrogen, and $a = (1+z)^{-1}$ (Chang et al., 2008; Barkana and Loeb, 2007). Reaching this sensitivity in a reasonable time, over large volumes of the sky, requires an instrument with a fast survey speed. That is, an ideal 21-cm intensity mapping instrument would have a wide field of view, while maintaining the required resolution and a high sensitivity.

1.3.3 EPOCH OF REIONIZATION INSTRUMENTS

Although this thesis focuses on post-reionization 21-cm intensity mapping, it should be noted that intensity mapping techniques are applicable at EoR redshifts. Before reionization, a majority

of the Universe's hydrogen was in neutral atomic form. As such, the amplitude of the brightness temperature fluctuations over the sky at a given redshift (Equation 1.3), is of the order of a few milliKelvin. Measuring the EoR power spectrum and detecting the BAO signal is a key motivation of the PAPER, MWA, and LOFAR instruments.

1.3.4 CROSS-CORRELATED INTENSITY MAPPING EXPERIMENTS

The first statistically significant detection of cosmic structure using the 21-cm line was reported in Pen et al. (2009). To achieve this, radio data from the HI Parkes all sky survey (HIPASS; Barnes et al., 2001) was cross-correlated with optical data from the six degree field galaxy redshift survey (6dFGRS; Jones et al., 2004). To compute the cross-correlation between radio and optical data, both the 6dFGRS and HIPASS data were gridded into regular three-dimensional data cubes (redshift data upon the line-of-sight axis), and the correlation between corresponding datum within each data cube was computed. While each galaxy within the 6dFGRS satisfies a 5σ detection threshold, a detection was not applied to the radio data, so the correlation is between the total 21-cm intensity and the distribution of observed galaxies from the 6dFGRS.

A similar strategy was used by Chang et al. (2010), to achieve a statistical detection of HI at a redshift of $z=0.8$ by cross-correlating redshifted HI data from the Green Bank 100 m telescope with optical data from the DEEP2 redshift survey (Davis et al., 2005). Again, the total 21-cm intensity is found to be correlated with the three-dimensional DEEP2 data cube.

The next step for 21-cm intensity mapping experiments is to show auto-correlation of 21-cm data upon a small area of sky. The ATNF observing proposal by Li et al. (2012) — which I am fortunate to be a part of — aims to use the Parkes telescope to do exactly this. This is expected to provide sufficient verification of the technique and calibration procedures to warrant a large volume survey capable of detecting the BAO signal.

1.3.5 POWER SPECTRUM INTERFEROMETERS

An alternative approach to a filled aperture experiment is to build an instrument that is sensitive only to targeted Fourier modes of the angular power spectrum. By definition, the BAO power

spectrum $P(\vec{k})$, is the Fourier transform of the sky brightness $T_b(\vec{r})$. That is,

$$T_b(r_x, r_y, r_z) \rightleftharpoons P(k_x, k_y, k_z), \quad (1.5)$$

where \rightleftharpoons denotes a Fourier transform pair, $\vec{r} = (r_x, r_y, r_z)$ is a spatial location vector, and $\vec{k} = (k_x, k_y, k_z)$ is a vector representing the Fourier mode with wavenumbers k_x , k_y and k_z . As discussed in the next section, an interferometric pair of radio telescopes is sensitive to a particular Fourier mode of the sky related to the spacing between the two. So, instead of building a filled aperture telescope (such as a dish) optimized to measure the sky brightness $T_b(\vec{r})$, we can build a sparse aperture telescope (such as an interferometer) optimized to measure targeted modes of the power spectrum $P(\vec{k})$.

A radio telescope equipped with a spectrometer may measure the redshift z of the HI line, from which the line-of-sight distance D can be inferred by Hubble's law¹

$$H_0 D = cz, \quad (1.6)$$

where H_0 is Hubble's constant. Redshift is related to observed frequency by

$$\nu_{obs} = \frac{\nu_{10}}{1+z}. \quad (1.7)$$

It follows that if we perform a survey of the sky brightness with a spectrometer equipped telescope, we can Fourier transform the data along the frequency axis to recover the line-of-sight modes of the power spectrum, k_{\parallel} . The perpendicular modes, k_{\perp} , are recovered by performing a spatial Fourier transform on the sky brightness distribution. As is discussed in Section 1.4.4, an antenna pair in an interferometer is sensitive to all line-of-sight modes k_{\parallel} , but only a select number of k_{\perp} modes. In contrast, a filled aperture instrument which measures the sky brightness is sensitive all k_{\parallel} and k_{\perp} modes.

The filled aperture and targeted mode approaches require radically different design and instrumentation. In the following section, I discuss the various ways in which the output of antennas may be combined to form a single telescope, and the characteristics of such arrays. I then return to HI BAO instrumentation in Section 1.6.

¹This is the low redshift approximation.




Figure 1.4: Simplified diagram of a four element phased array. The output voltages of each antenna are added together in a process called beamforming.

1.4 RADIO ASTRONOMY INSTRUMENTATION

The fundamental goal of all radio telescopes is to detect electromagnetic radiation from astrophysical sources. Instrument development in radio astronomy is driven by the desire to detect fainter objects with more precision and resolution, at faster surveying speeds. For wide-field intensity mapping, we have a modest requirement on resolution, but require a fast survey speed, which in turn requires high sensitivity.

A common approach to achieving high sensitivity is to use an array of antennas in a single observation to increase the total collecting area. There are two ways in which the voltage outputs of the antennas are combined: weighted summation of all voltages in a operation known as beamforming; or alternatively, antenna pairs are multiplied together to form a set of voltage coherency products, known as correlation. I will refer to instruments in which both beamforming and correlation are performed as ‘BX’ type arrays, and a set of beamformed antennas as an aperture array ‘station’.

1.4.1 APERTURE ARRAYS

Aperture arrays have recently undergone a renaissance in radio astronomy, with several aperture array based telescopes currently under construction and development; for example, the Low Frequency Array in Europe (LOFAR, de Vos et al., 2009), the Long Wavelength Array in America

(LWA, Ellingson et al., 2012), and the Murchison Widefield Array in Australia (MWA, DeBoer et al., 2009). This renaissance has been driven by the attractive properties they offer: they can provide instantaneous coverage of most of the sky, their cost of building scales only linearly with size, and digital signal processing techniques allow calibration to be done electronically. Many of these advantages rely upon high speed digital processing, and as such, have not been realisable in the past.

A phased aperture array is a collection of antennas whose outputs are connected a beamformer, which is a device that forms linear combinations (i.e. weighted sums) of inputs. The result is that the effective radiation power pattern of the antenna array is suppressed in some directions and enhanced in other directions. This allows us to ‘form a beam’ in a particular direction, increasing the array’s sensitivity and angular resolution in that direction. In a conventional beamformer, the beam output $B(t)$ is given by

$$B(t) = \sum_{n=1}^N w_n V_n(t - \tau_n), \quad (1.8)$$

where N denotes the number of antennas (or more generally, sensors), the V_n denote the sensor outputs, the w_n denote weighted gains applied to the antenna outputs, and the τ_n denote the time delays required to point or steer the beam to the specified direction (Mucci, 1984). The process of beamforming can be described mathematically as the inner product of a set of weights with a set of input antenna voltages. If we express these voltages in a column vector

$$\vec{v} = \begin{pmatrix} v_1 & v_2 & \dots & v_N \end{pmatrix}^T, \quad (1.9)$$

then the beamforming operation is the inner product of \vec{v} with a weight vector \vec{w} :

$$B = \vec{w}^H \vec{v} = \sum_{n=1}^N w_n^* v_n. \quad (1.10)$$

As a simple example, let us consider two isotropic antennas on the surface of the earth, separated by a distance d . If a far-field radio source is directly above the antennas (i.e. at zenith), then the wavefront will arrive at the two antennas at the same time and consequently their outputs $V_1(t)$ and $V_2(t)$ will be in phase; however, if the source is located at an angle θ from

zenith, then the wavefront must travel a distance of $d \sin\theta$ further to reach the second antenna. The time delay τ_g for the wavefront to travel this distance is given by

$$\tau_g = \frac{d \sin\theta}{c}, \quad (1.11)$$

where c is the speed of light. This will introduce a phase difference, $\Delta\phi$, of

$$\Delta\phi \equiv 2\pi\nu\tau \text{ mod } 2\pi. \quad (1.12)$$

Consequently, if we wish to add our two antenna outputs in phase we must correct for this phase difference. If we now consider an array of multiple antennas (Figure 1.4), we can calculate the time delays τ_n and apply to each antenna so that their output signals are in phase. This can be accomplished by picking a reference antenna and applying Equation 1.11, or by applying phase delays $\Delta\phi_n$ in the frequency domain.

More advanced beamforming algorithms can be used if we digitize the antenna outputs and then steer the beam electronically with digital delays. A digital beamformer is capable of generating multiple sets of time delays and gains (τ_n, w_n) , so can therefore produce multiple beams on the sky². As a result, a digitally beamformed aperture array is capable of surveying multiple regions of the sky at once greatly increasing survey speed, limited only by computational restrictions.

1.4.2 SPARSE ARRAYS

This thesis explores the use of sparse aperture arrays for radio astronomy; that is, aperture array with inter-element spacings greater than $\lambda/2$. Unlike their dense counterpart ($<\lambda/2$ spacings), sparse aperture arrays are affected by a phenomenon known as grating lobes. Grating lobes are peaks in the radiation power pattern of an array which arise due to the modulo nature of Equation 1.12. That is, a source at position θ_1 may generate a phase difference $\Delta\phi_1$, which is equal to the phase difference $\Delta\phi_2$ arising from a source at position θ_2 . This phenomenon is unique to sparse arrays as the spacing of a dense array ensures that this ambiguity does not

²This can similarly be done with analogue beamformers, however they quickly become complex and may suffer from inter-beam coupling.

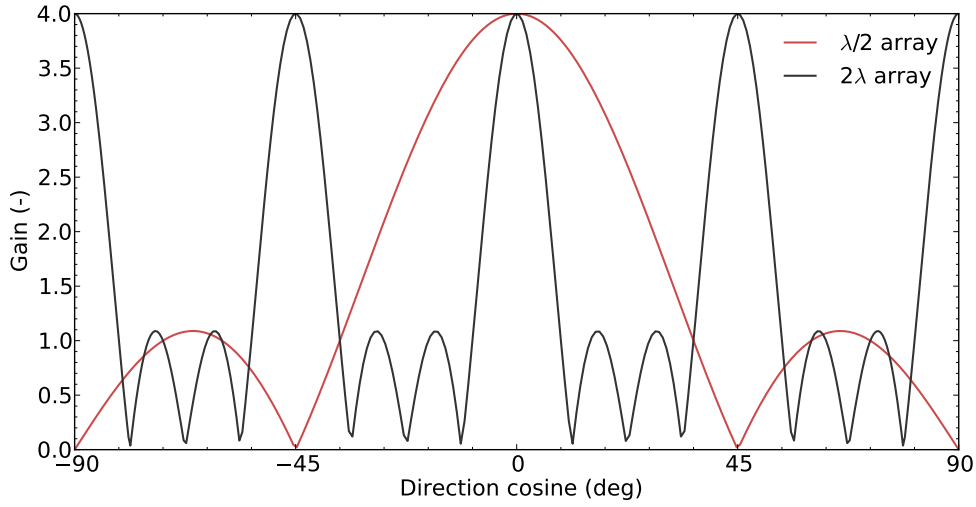


Figure 1.5: Beam pattern for an array of four equally spaced isotropic elements, at $\lambda/2$ and 2λ spacings.

occur. Grating lobes can be shown to occur at angles θ_p such that

$$\sin\theta_p = \sin\theta_0 + \frac{p\lambda}{d_x} \quad (1.13)$$

$$p = \pm(1, 2, \dots), \quad (1.14)$$

for values of p that define an angle with a real sine ($|\sin\theta_p| \leq 1$) (Mailloux, 2005). Here, θ_0 is the direction of the desired scan angle. To ensure the grating lobes do not arise, we have the condition

$$d \leq \frac{\lambda}{1 + \sin\theta_0}. \quad (1.15)$$

From this equation we see that grating lobes cannot appear if $d < \lambda/2$ (i.e. dense), and will always be present if $d > \lambda$ (i.e. sparse).

Figure 1.5 illustrates the affect of grating lobes upon the response of a beamformed arrays. The figure shows the beam pattern for an array of 4 equally spaced isotropic elements, at $\lambda/2$ and 2λ spacings. While the 2λ spaced array has superior frequency resolution (a narrower beam) than the $\lambda/2$ spaced array, grating lobes appear at $\pm 45^\circ$ and $\pm 90^\circ$.

Grating lobes are a major problem for radio astronomy as they cause source confusion, due to ambiguity of the location of the source. However, there are methods to avoid or suppress

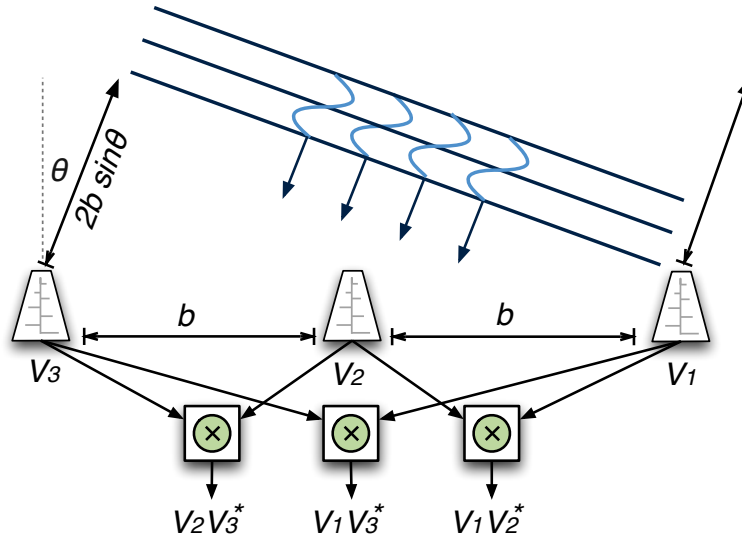


Figure 1.6: Diagram of an interferometer, consisting of three antennas. A plane wave arrives at an angle of incidence θ , which introduces a path difference of $2b \sin\theta$ between antenna pair (1,3), and a path difference $b \sin\theta$ between pairs (1,2) and (2,3). The antenna output voltages $V_1(t)$, $V_2(t)$ and $V_3(t)$ are multiplied together in a device known as a correlator.

the affects of grating lobes. One noteworthy method, in which their affects are negated via inter-array correlation, is discussed in Braun and van Cappellen (2006); this is used in the LOFAR array. In the technique of Braun and van Cappellen (2006), arrays are rotated with respect to each other so that their grating lobes no longer coincide (antenna elements are then counter-rotated so that the polarization coordinate system is the same). This suppresses the grating lobe response in the correlation product of each array pair. Interferometric arrays are discussed in more detail below.

1.4.3 INTERFEROMETRIC ARRAYS

Interferometry and synthesis imaging are fundamental components of modern radio astronomy, and there are many texts devoted to the subject, such as the canonical Thompson et al. (2004) and Taylor et al. (1999). Interferometry relies upon the correlation operation: for a pair of receiver antennas separated by a distance b with voltage outputs $V_1(t)$ and $V_2(t)$, a correlator adds a time delay τ_g analogous to that in Equation 1.11, then computes the time average:

$$\langle V_1(t)V_2^*(t - \tau_g) \rangle_t = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T V_1(t)V_2^*(t - \tau_g) dt. \quad (1.16)$$

A correlator computes this product for each antenna pair in the synthesis array. The distance between them, \vec{b}_n , is known as the baseline length. This is shown in Figure 1.6. In contrast to the inner product computed in beamforming, the cross-correlation operation is formed from the outer product of the voltage vector:

$$X = \vec{v} \otimes \vec{v} = \vec{v} \cdot \vec{v}^H \quad (1.17)$$

$$= \begin{pmatrix} V_1 V_1^* & V_1 V_2^* & \cdots & V_1 V_N^* \\ V_2 V_1^* & V_2 V_2^* & & V_2 V_N^* \\ \vdots & & \ddots & \vdots \\ V_N V_1^* & V_N V_2^* & \cdots & V_N V_N^* \end{pmatrix} \quad (1.18)$$

where X is the cross-correlation matrix, often referred to simply as the correlation matrix. The time average of this matrix, $\langle X \rangle$, is related to the brightness distribution of the sky by a Fourier transform. The Fourier transform relationship between the product $\langle V_p V_q \rangle$ and the sky brightness $I(l, m)$ is shown explicitly by the van Cittert-Zernicke theorem (Zernicke, 1938):

$$\langle V_p V_q \rangle = \int \int I(l, m) e^{-2\pi i(ul+vm)} dl dm, \quad (1.19)$$

where l and m are direction cosines for a sky brightness distribution on the celestial sphere, and u and v are the corresponding Fourier mode coordinates. Each pair of antennas (p, q) in an interferometer is therefore sensitive only to particular Fourier modes; for example, in Figure 1.6, the corresponding (1D) mode is

$$u = \frac{b \cos(\theta)}{\lambda}; \quad (1.20)$$

the more general case (for antennas located on Earth) is given in §4.1 of Thompson et al. (2004).

1.4.4 INTERFEROMETRIC POWER SPECTRUM MEASUREMENT

As was introduced in Section 1.3.5, the BAO power spectrum $P(\vec{k})$, is the Fourier transform of the sky brightness $T_b(\vec{r})$. The (u, v) modes measured by an interferometer correspond directly to the k_\perp modes of the BAO power spectrum. As well as the (u, v) modes of the sky brightness, an interferometer equipped with a spectrometer may measure the redshift of the HI line, from which distance can be inferred. It follows that if we treat the (u, v) modes as a function of

frequency (or redshift), we can Fourier transform along the frequency axis to produce the power spectrum $P(\vec{k})$. In this sense, an interferometer samples all line-of-sight modes k_{\parallel} , but only a select number of modes perpendicular to the line of sight, k_{\perp} . Each antenna pair measures a small subset of k_{\perp} modes, but all line-of-sight k_{\parallel} modes. In contrast, a filled aperture telescope measures all k_{\perp} modes up to its resolution limit, and all k_{\parallel} modes over its redshift (frequency) coverage.

When designing a BAO power spectrum interferometer, one must carefully consider which k_{\perp} modes to measure, and configure antennas accordingly. This task is helped by the rotation of the Earth, which causes the projected baseline length to change as a source is tracked across the sky. As a result, a given antenna pair will measure different k_{\perp} modes as the sky rotates, increasing coverage of k_{\perp} modes at the expense of integration time on each mode.

The total amount of integration time required for a given Fourier mode depends upon the telescope's sensitivity. The sensitivity of a telescope is limited by the amount of noise introduced by the analogue components which amplify, transmit and filter the signal from the antenna; this is discussed in the following section.

1.5 SENSITIVITY AND SURVEY SPEED

The amount of observing time required to reach a target signal-to-noise ratio is dependent upon the amount of noise introduced by the analogue components of a telescope. For an extensive discussion of noise in microwave systems, I direct the reader to Chapter 10 of Pozar (2005), and Chapter 9 of Taylor et al. (1999) for a radio-astronomy focused view. This section gives a brief overview of the factors affecting telescope sensitivity, before introducing figures of merit with which to compare the sensitivity of different instruments.

The signal-to-noise ratio of a telescope is generally defined in terms of equivalent noise temperature. This is not the physical temperature of the receiver, T_p ; it is the equivalent temperature, T_e , which corresponds to a total system noise power of P_N :

$$T_e = \frac{P_N}{k_B \Delta\nu}. \quad (1.21)$$

Here k_B is the Boltzmann constant. The overall equivalent noise of a radio receiver is known as

the receiver temperature, T_{rec} .

The system temperature, T_{sys} , is a measure of the total noise of a telescope, and unlike receiver temperature it includes the noise contribution from the sky and antenna. To mitigate the effect of noise, one can take multiple measurements over bandwidth and time³. The time it takes to reach a target root-mean-square noise temperature is then given by the ideal radiometer equation:

$$\sigma_T = \frac{T_{sys}}{\sqrt{2n_p\Delta\nu\tau}}, \quad (1.22)$$

where n_p is the number of polarizations, and τ is integration time. The smallest detectable signal T_S is given by $T_S = m\Delta T$, where m is a threshold value, generally greater than or equal to 3.

The overall sensitivity of a single-dish radio telescope is given by its flux collecting ability, coupled with its system temperature. A common figure of merit is the system equivalent flux density, SEFD (Chapter 9 of Taylor et al., 1999):

$$SEFD = \frac{2T_{sys}k_B}{A_e}. \quad (1.23)$$

Here, A_e is the effective collecting area of the telescope. The SEFD is measured in Janskys (Jy), where $1 \text{ Jy} = 10^{-26} \text{ Wm}^{-2}\text{Hz}$. It follows from Equation 1.22 and Equation 1.23 that the smallest detectable signal is

$$\Delta S = m \frac{SEFD}{\sqrt{2n_p\Delta\nu\tau}}. \quad (1.24)$$

Equation 1.24 is the sensitivity of a single-dish telescope — or more generally, a telescope with a filled aperture — to a point-like source (i.e. sources that are not resolved). For an interferometer, it is more suitable to look at the sensitivity $\Delta S_{pq}(u, \nu)$ of a single baseline to a single (u, ν) mode:

$$\Delta S_{pq}(u, \nu) = m \frac{SEFD_{pq}}{\sqrt{2n_p\Delta\nu\tau}}, \quad (1.25)$$

where $SEFD_{pq}$ is the system equivalent flux density for the antenna pair (not for the entire array).

³this is true for Gaussian ‘white’ noise, which is both stationary and ergodic.

1.5.1 POINT SOURCE SURVEY SPEED

In designing a telescope for 21-cm intensity mapping, an important metric is the survey speed. This section defines a figure of merit (FoM) which can be used to compare the point source survey speed of different instruments. This FoM is similar to that presented in Cordes (2009); however, it differs in its inclusion of fundamental constants.

As given by Equation 1.24, the time taken to reach a desired flux sensitivity (in Jy) for a given pointing is

$$\tau = \frac{m^2 SEFD^2}{2n_p \Delta\nu} / \Delta S^2. \quad (1.26)$$

To survey a field of solid angle Ω_{survey} requires P pointings, where

$$P = \frac{\Omega_{survey}}{\Omega_{FoV}}. \quad (1.27)$$

Here, Ω_{FoV} is the instantaneous field of view (FoV) of the instrument. The total time taken is then

$$\tau P = \frac{m^2 SEFD^2}{2n_p \Delta\nu \Delta S^2} \frac{\Omega_{survey}}{\Omega_{FoV}}. \quad (1.28)$$

If we rearrange this so that parameters intrinsic to the instrument are equated to survey parameters, and write $\tau P = t$ as the total survey time, we find

$$\frac{\Omega_{survey} m^2}{t \Delta S^2} = \frac{2n_p \Delta\nu \Omega_{FoV}}{SEFD^2} \triangleq PFoM. \quad (1.29)$$

where the right hand side is defined to be the point source survey speed figure of merit (PFoM).

The PFoM may be rewritten as

$$PFoM = \frac{n_p \Delta\nu \Omega_{FoV}}{2k_B^2} \left(\frac{A_e}{T_{sys}} \right)^2. \quad (1.30)$$

As Equation 1.29 relates survey parameters to an instrument's PFoM, one can use the PFoM to find the required survey time t to reach a desired sensitivity ΔS over a patch of sky Ω_{survey} . The PFoM, while only approximate, is a convenient way of comparing telescope survey speed; large PFoMs correspond to instruments with faster point source survey speeds.

1.5.2 EXTENDED SOURCE SURVEY SPEED

Although PFoM is a useful metric for comparing point source sensitivity, intensity mapping experiments must instead be sensitive to large scale brightness temperature fluctuations. For a fully filled aperture telescope, if the source of interest does not fill the beam, then the minimum temperature detectable is

$$\Delta T = \frac{m T_{sys}}{\sqrt{2 n_p \Delta \nu \tau}} \frac{\Omega_B}{\Omega_S}, \quad (1.31)$$

where Ω_S is the solid angle subtended by the source.

For a sparse aperture, the (u, v) plane is not fully sampled. As such, their sensitivity to the sky brightness T_b is smaller by a ‘filling’ factor

$$f = \frac{A_{col}}{A_{phys}}, \quad (1.32)$$

where A_{col} is the total collecting area of the array, and A_{phys} is the physical extent of the array. The brightness temperature sensitivity is then

$$\Delta T = \frac{m T_{sys}}{\sqrt{2 n_p \Delta \nu \tau}} \frac{\Omega_B}{\Omega_S} \frac{1}{f}. \quad (1.33)$$

From Equation 1.33, we may derive a a brightness figure of merit (BFoM) in a similar fashion to the PFoM. Assuming our instrument has sufficient resolution such that $\Omega_B \leq \Omega_S$, we have

$$\tau P = \frac{m^2 T_{sys}^2}{\Delta T^2 2 n_p \Delta \nu} \frac{1}{f}, \quad (1.34)$$

and the BFoM is then defined as

$$\frac{\Omega_{survey} m^2}{t \Delta T^2} = 2 n_p \Delta \nu \Omega_{FoV} \left(\frac{f}{T_{sys}} \right)^2 \triangleq BFoM \quad (1.35)$$

Accordingly, the brightness temperature survey speed of a sparse aperture instrument decreases as f^2 with the filling factor. Comparing Equation 1.30 with Equation 1.35, we note that:

- Increasing the collecting area A_e while keeping f constant will increase the PFoM, but will not increase the BFoM.

- Increasing the filling factor f while keeping the collecting area A_e constant will increase the BFoM, but not the PFoM (so long as the source remains resolved).

This is a key point for intensity mapping instruments. As was introduced in Section 1.3.2, the structure of the BAO is washed out by non-linear evolution above the third BAO peak. So, the ideal intensity mapping instrument has a well filled aperture (over the required diameters given in Table 1.2), a low system temperature, and a large field of view. The next section compares various instrumental approaches which satisfy these requirements.

1.6 HI INSTRUMENTS TO MEASURE BAO

Section 1.4 introduced the various ways in which arrays of telescopes can be combined to increase sensitivity, and Section 1.5 presented some figures of merit for survey speed. We may use Equation 1.35 to find the minimum BFoM to detect the BAO signal within a given observing time, t , by requiring

$$BFoM_{min} \geq \frac{\Omega_{req} m^2}{t \Delta T^2}. \quad (1.36)$$

In Chang et al. (2008), the authors present a figure of $150 \mu\text{K}$ variations over $18h^{-1} \text{Mpc}$ scales for the BAO signal (based on Equation 1.4 above). A large area of sky must be covered ($\Omega_{req} > 1000 \text{ deg}^2$), and for a detection we require $m=3$ or higher. For a year-long observing campaign, this gives

$$BFoM_{min} \geq 3.85 \text{ sr Hz K}^{-2}. \quad (1.37)$$

This number assumes every second within the year is used. Depending on the instrument pointing capabilities and survey strategy, the BFoM will have to be larger than this to account for the loss of integration time over the target field.

It should be stressed that reaching the thermal noise level required is only part of the challenge, and that the BFoM is only an approximation to real survey speed. Detection of the BAO through 21-cm intensity mapping will also rely on a robust calibration and of removal of contaminating foregrounds.

1.6.1 SINGLE DISH CASE

At EoR redshifts ($z \sim 6$), the required dish diameter to resolve the third BAO peak is over 294 m (Table 1.2); building a steerable single dish telescope of this diameter would be prohibitively expensive, due to the weight and physical size of the dish and supporting structures. At redshifts below $z < 0.5$, the required dish diameter is comparable to existing instruments, such as the Parkes 64 m telescope in Australia, and the Robert C Byrd 100 m telescope in Green Bank, West Virginia. This makes intensity mapping with a filled aperture plausible, but the small field of view of these telescopes limits their survey speed.

For example, the Parkes multibeam receiver has 13 dual-polarization feeds, each with a system temperature of ~ 25 K (Staveley-Smith et al., 1996). This receiver operates over $0 \leq z \leq 0.2$. At $z=0.2$, we require a channel bandwidth of $\Delta\nu=6$ MHz to Nyquist sample the third BAO peak of $35h^{-1}\text{Mpc}$. If only the central feed is used, the BFoM is

$$BFoM_{\text{ParkesSB}} = 0.27 \text{ sr Hz K}^{-2},$$

which is less than 1/14 of the required BFoM. Even if the receiver temperature was halved, the BFoM would still not be adequate. If all 13 feeds are used, the FoV increases 13-fold, as does the BFoM:

$$BFoM_{\text{ParkesMB}} = 3.51 \text{ sr Hz K}^{-2} \quad (1.38)$$

This is close to — yet still under — the minimum BFoM. Doubling $\Delta\nu$ (at the expense of line-of-sight resolution), would be sufficient to reach the BFoM requirement; alternatively, one could integrate for a longer period. In conclusion, the Parkes multibeam receiver is a promising instrument for single-dish based 21-cm intensity mapping.

1.6.2 SPARSE ARRAY CASE

To compute the filling factor of a sparse array, one must take into account the effective area of the antenna elements. Once the filling factor is known, the BFoM can then be found using Equation 1.35.

To begin, we note that the product of beam solid angle and effective area is fixed for all

antennas (see for example Kraus and Marhefka, 2002):

$$A_e \Omega_B = \lambda^2, \quad (1.39)$$

and that the directivity of an antenna is defined as

$$D = \frac{4\pi}{\Omega_B} = \frac{4\pi A_e}{\lambda^2}. \quad (1.40)$$

That is, the effective area of an antenna is related to its directivity, with more directive antennas possessing larger effective areas, but smaller solid angles (and hence smaller fields of view).

Now, consider an array consisting of N_A elements, evenly spaced on a grid with inter-element spacing d . For a regularly gridded array of identical antennas, where the collecting areas of each antenna element do not overlap, we may write a filling factor

$$f = \frac{A_{col}}{A_{phys}} = \frac{N_A D \lambda^2}{4\pi d^2}; \quad (1.41)$$

where the collecting areas overlap, we have $f=1$. This is plotted in Figure 1.7 for 3, 6, 9 and 12 dBi directivity antennas. The corresponding solid angles on the sky are 2π (i.e. an entire hemisphere), π , $\pi/2$ and $\pi/4$ steradians, respectively. For an array spanning a 200λ diameter (BAO resolution requirement), the FoV is

$$\Omega_{FoV} = N_{FoV} \Omega_B = \frac{N_{FoV}}{\pi 100^2}, \quad (1.42)$$

where N_{FoV} is the number of beams on the sky. It follows that the BFoM of a sparse array can be approximated as

$$BFoM = \frac{2}{100^2 \pi} n_p N_{FoV} \Delta\nu \left(\frac{N_A D \lambda^2}{4\pi d^2 T_{sys}} \right)^2. \quad (1.43)$$

One may use this equation to explore the relationship between array filling factor and the number of beams required to maintain the minimum BFoM.

For example, let us consider an array operating over 1.0-1.5 GHz, with 9 dBi antenna elements spaced a wavelength apart at 1.0 GHz. From Figure 1.7, we see that the filling factor lies

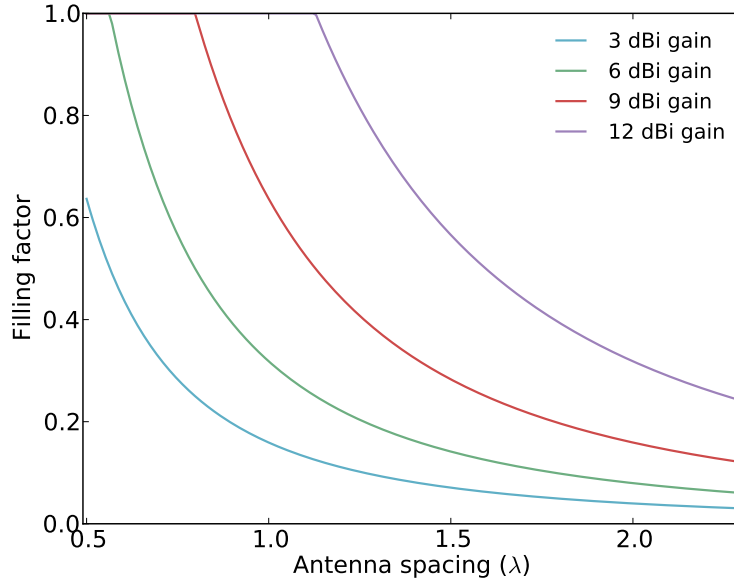


Figure 1.7: Filling factor of an aperture array as a function of antenna spacing. As inter-element spacing in a sparse array increases, the filling factor decreases. The four curves show that the filling factor is dependent upon the gain of the antenna elements.

between 0.7 and 0.3 over this frequency range. If the overall T_{sys} is 50 K, and we again select $\Delta\nu=6$ MHz and $n_p=2$, about 25 beams are required to reach the minimum BFoM at 1.0 GHz, but 140 beams are required at 1.5 GHz.

Such an array would require of order 200^2 antennas. This is fewer than the proposed number of antennas in a single SKA dense aperture array (Faulkner, 2011). While it would be possible to build a 40,000 input beamformer capable of producing 140 beams, a multi-pixel single dish or interferometric array (see below) is likely offer a more affordable alternative.

1.6.3 INTERFEROMETER CASE

The field of view of an interferometer is limited by the FoV of its constituent antennas. For an interferometer comprised of simple dipoles, any patch of sky above the horizon can be imaged using synthesis imaging techniques⁴.

When writing a BFoM for an interferometer, one must pay particular attention to its (u, v) coverage. As discussed in Section 1.4.3, interferometers are only sensitive to particular Fourier modes of the sky, which depend upon the antenna configuration. For Equation 1.35 to be a good

⁴Note that the effective area of the array and (u, v) coverage will change with pointing, due to projection effects.

approximation, we require that each baseline in the interferometer samples a different (u, v) mode; that is, we require that the array is minimally redundant. For highly redundant arrays, certain (u, v) modes will be sampled multiple times, and as such the interferometer will be more sensitive to these modes. Parsons et al. (2011) propose leveraging maximally redundant arrays to measure targeted k_{\perp} modes of the power spectrum in order to detect BAO features. To allow a simple comparison of interferometer, sparse aperture array and single-dish performance, this section assumes a minimally redundant interferometer.

With this caveat, we may use Equation 1.41 to approximate the filling factor of an interferometer, and Equation 1.35 to calculate an approximate BFoM. Assuming antenna elements with 9 dBi directivity, their FoV is $\pi/2$ sr, which is much greater than the required survey field size of $\Omega_{survey}=1000$ deg². Using the same values of T_{sys} , $\Delta\nu$ and n_p as in Section 1.6.2 above, the filling factor required to meet the minimum BFoM is

$$f = \sqrt{\frac{BFoM_{min} T_{sys}^2}{2n_p \Delta\nu \Omega_{survey}}} \approx 0.036,$$

To reach this filling factor requires ~ 1800 antennas of 9 dBi directivity within a 200λ diameter.

Presently, the largest interferometric arrays (MWA and PAPER), are targeting $N_{ant}=128$ antennas. As the required computational power for a correlator is proportional to N_{ant}^2 , such a large correlator is not feasible.

1.6.4 BX CASE

The final case which will be considered is an interferometric array of sparse aperture array stations (i.e. a BX-array). An interferometric array with 128 aperture array stations, each of which consists of 16 antennas (i.e. $B=16$, $X=128$) maintains the required filling factor computed in Section 1.6.3. As beamforming will decrease the FoV of each station, each station would need to have $N_{FoV}=5$ beams on the sky to maintain a 1000 deg² FoV (this assumes an inter-element spacing of λ). Depending on how the antennas are partitioned, the requisite size of the correlator and beamformer changes. For example, with $B=64$, the required correlator size is $X=32$, but N_{FoV} increases to 20. These requirements are orders of magnitude lower than those presented in the extreme cases presented above (i.e. $B=40,000$ and $N_{FoV}=140$ in §1.6.2; and $B=N_{FoV}=1$

and $X=1800$ in §1.6.3).

1.6.5 COMPARISON OF INSTRUMENTATION APPROACHES

To give a brief summary of the four cases considered above:

- Single-dish telescopes with single-pixel feeds do not offer sufficiently high BFoMs for intensity mapping. Multi-pixel feeds are a viable technology for 21-cm intensity mapping over $0 < z < 0.5$. At redshifts above this limit, an alternative approach is required.
- A sparse aperture array with $\sim 40,000$ antenna elements would require up to 140 beams on the sky to meet the BFoM requirement.
- An interferometer imaging a 1000 deg^2 patch of sky would require ~ 1800 antenna elements to meet the minimum BFoM. This would require a correlator with an order of magnitude more inputs than current correlators in operation.
- A BX-type instrument would decrease the size of correlator required by partitioning antennas into aperture array stations. Each aperture array station would need to form multiple beams on the sky to maintain a large field of view. Out of the four cases considered, BX-type arrays are the most attractive candidate for 21-cm intensity mapping at redshifts $z > 0.5$.

A number of other instrumentation approaches have been proposed; an overview of these approaches is given in Section 1.6.6 below.

1.6.6 PROPOSED INTENSITY MAPPING INSTRUMENTS

Over the past few years there have been a considerable number of proposals to build a dedicated 21-cm intensity mapping instrument⁵, see Table 1.3. The proposed instrumentation for these projects differs greatly. For example the CHIME (Dobbs, 2009) and MOURINHO (Zarb-Adami, 2011) projects seek to use arrays of cylindrical reflectors over the redshift range $2.5 \geq z \geq 0.7$. In CHIME, five 100×20 m cylindrical reflectors would be placed alongside each other to form a filled aperture and maximize brightness sensitivity. In the MOURINHO project, the authors

⁵As of writing, while these proposals have been well received, none have received funding.

Table 1.3: Comparison of proposed redshifted hydrogen BAO instruments.

Proposal	Technology summary	Frequency range (MHz)
CHIME	Cylindrical reflectors	400-800
MOURINHO	Cylindrical reflectors	400-800
BINGO	Drift-scan using static parabolic reflector	960-1260
BAOBAB	Interferometric array of dipoles	600-900
SANDGROPER	Interferometric array of sparse apertures	1000-1500
CARPE	Spatial FFT interferometer (MOFF)	350-1500
FFTT	Spatial FFT interferometer	N/A

propose upgrading the existing Medicina Northern Cross radio telescope to build a comparable instrument. Another filled aperture proposal is BINGO (Battye et al., 2012), in which a 40 m parabolic dish is used over 960-1260 MHz. To increase survey speed, they propose that 50 feeds are placed at the focal point. To drive down costs, the dish would be stationary and located at the bottom of a large cliff onto which the feeds are mounted.

A number of sparse aperture instruments have also been proposed. In the BAOBAB proposal (Parsons, 2012), the authors suggest the use of an interferometric array of simple antennas. Much of the digital signal processing hardware would be recycled from the PAPER experiment, which is being upgraded, and BAOBAB would be co-located with PAPER in South Africa. BAOBAB would operate over 600-900 MHz. The CARPE (Morales, 2011), and FFTT (Tegmark and Zaldarriaga, 2009a,b) concepts are compact interferometers in which the correlation operation is replaced with a spatial FFT, which is more computationally efficient (order $n \log_2(n)$ operations), than computing all baselines (order n^2 computations). The spatial FFT approach has recently been verified experimentally by Foster et al. (2012).

Finally, the SANDGROPER proposal (Schediwy, 2010) suggests the use of a BX-type array. The design is similar to the MWA but with a higher frequency of operation of 1000-1500 MHz. The D-PAD sparse array detailed in Chapter 3 of this thesis was precursor work toward Schediwy's SANDGROPER proposal.

1.7 THESIS LAYOUT

The overall success of 21-cm intensity mapping machines depends not only on their sensitivity, but on their calibration; different instrumentation approaches demand different calibration

techniques. It is often unclear what systematics ultimately limit an instrument's performance until after initial prototype and technology demonstrator are constructed and tested.

This thesis tackles some of the instrumentation challenges faced by 21-cm intensity mapping instruments. The introductory chapter elucidates the scientific motivation of the instrumentation presented within this thesis, and gives a general overview of relevant radio astronomy instrumentation. In the next chapter, I discuss the limitations of the Radio Interferometer Measurement Equation in its current form. Being aware of such limitations is important to achieving the level of calibration required for calibration of intensity mapping data.

This chapter is followed by an overview of the D-PAD sparse aperture array project, which I built in conjunction with S. Schediwy during my DPhil. The D-PAD system is then tested and its performance evaluated in Chapter 5. In parallel with this analogue system development, I developed an FPGA based wideband digital spectrometer. This spectrometer has been deployed on the Parkes 64 m telescope. This project, HIPSR, is detailed in Chapter 6. The HIPSR system is an ideal digital signal processor for a 21-cm intensity mapping survey at low redshift, such as that proposed by Pen et. al. (2012).

The final chapter gives a summary of the work presented in this thesis, and discusses future work, before concluding with some closing remarks.

CHAPTER 2

COHERENCY IN RADIO ASTRONOMY

2.1 INTRODUCTION

Having a solid understanding of an instrument's systematics is important for its calibration. As we push for higher sensitivity and imaging dynamic range it is increasingly important to understand and account for instrumental response. For EoR instruments, this is of particular importance, as poor calibration will confound detection of the EoR signature. In this Chapter, I discuss the transmission of electromagnetic signals from astrophysical sources and how we can measure them using arrays of radio telescopes. The motivation is to provide a unified picture of the signal propagation from an astrophysical progenitor right through to the analogue to digital converter in a telescope.

Coherency in electromagnetic fields is a fundamental topic within optics. Its importance in fields such as radio astronomy can not be overstated: interferometry and synthesis imaging techniques rely heavily upon coherency theory (Taylor et al., 1999; Mandel and Wolf, 1995; Thompson et al., 2004). Of particular importance to radio astronomy is the Van-Cittert-Zernicke theorem (vC-Z, Zernicke, 1938) and the radio interferometer Measurement Equation (RIME, Hamaker et al., 1996). The vC-Z relates the brightness of a source to its mutual coherency as measured by an interferometer, and the RIME provides a polarimetric framework to calibrate out corruptions caused along the signal's path.

While the vC-Z theorem dates back to 1938, more recent work such as that of Carozzi and Woan (2009) extends its applicability to polarized measurements over wide fields of view. The

RIME has a much shorter history: it was not formulated until 1996 (Hamaker et al., 1996). Before the RIME, calibration was conducted in an ad-hoc manner, with each polarization essentially treated separately. The framework was expounded in a series of follow-up papers (Sault et al., 1996; Hamaker and Bregman, 1996; Hamaker, 2000, 2006); recent work by Smirnov extends the formalism to the full sky case, and reformulates the RIME using tensor algebra (Smirnov, 2011a,b).

This chapter discusses why the RIME is not suitable for describing the analogue components of a telescope. We present a reformulation of the RIME that considers full electromagnetic coherency statistics (i.e. electric, magnetic and mixed coherencies). We then show that this reformulation does not suffer from the same limitations as the RIME of Hamaker et al. (1996), and that it is similar to the 2N-port transmission matrix from microwave engineering. The relationship between the result reported here and the vC-Z formulation of Carozzi and Woan (2009) is also discussed.

2.2 PROPERTIES OF ELECTROMAGNETIC WAVES

The propagation of electromagnetic radiation across free space is governed by the electromagnetic wave equations:

$$\vec{\nabla}^2 \vec{E} = \mu_0 \epsilon_0 \frac{\partial^2 \vec{E}}{\partial t^2} \quad (2.1)$$

$$\vec{\nabla}^2 \vec{H} = \mu_0 \epsilon_0 \frac{\partial^2 \vec{H}}{\partial t^2}. \quad (2.2)$$

A brief derivation of these equations from Maxwell's equations can be found in texts such as Griffiths (1999). An important characteristic of these equations is that they are linear: there are no cross terms, and all terms are to the first power; because of this linearity, we can decompose solutions into a superposition of monochromatic sinusoids

$$\vec{E}(\vec{r}, t) = E_0 \cos(\omega t - \vec{k}\vec{r} + \phi_E) \quad (2.3)$$

$$\vec{H}(\vec{r}, t) = H_0 \cos(\omega t - \vec{k}\vec{r} + \phi_B). \quad (2.4)$$

In an electromagnetic plane wave, \vec{E} and \vec{H} are always perpendicular to each other and the direction of propagation. The direction of propagation is then given by the Poynting vector \vec{s} , which can be thought of as representing the energy flux (in W/m^2) of an electromagnetic field,

$$\vec{s} = \frac{1}{\mu_0} \vec{E} \times \vec{H}. \quad (2.5)$$

Moreover, their amplitudes are constrained by $H_0 = \frac{1}{c} E_0$. Due to this dependence, often only the electric field is considered.

2.2.1 POLARIZATION

Polarization is a key measurement within astronomy; textbooks such as Tinbergen (2005) are devoted to the subject. Most radio telescopes consist of a pair of orthogonal antennas, each of which is sensitive to a single polarization. Most astrophysical radio emission is noise-like and thus inherently unpolarized. Nevertheless, a number of radio sources, such as pulsars and masers, do emit polarized radiation, and effects such as Faraday rotation by a galaxy's magnetic field can yield polarized signals. Measuring polarization can elucidate information about the source itself, and/or the medium between the source and observer.

THE STOKES PARAMETERS are a set of quantities which fully describe the polarization state of an electromagnetic wave. These four parameters, I , Q , U and V , are related to the amplitudes of perpendicular components of the electric field

$$E_x = e_x(t) \cos(\omega t + \delta_x) \quad (2.6)$$

$$E_y = e_y(t) \cos(\omega t + \delta_y) \quad (2.7)$$

by time averages of these electric field parameters:

$$I = \langle E_x E_x^* + E_y E_y^* \rangle \quad (2.8)$$

$$Q = \langle E_x E_x^* - E_y E_y^* \rangle \quad (2.9)$$

$$U = \langle E_x E_y^* + E_y E_x^* \rangle \quad (2.10)$$

$$V = i \langle E_x E_y^* - E_y E_x^* \rangle \quad (2.11)$$




Figure 2.1: Comparison of linear polarization (left), and circular polarization (right). Image credit: Wikimedia (2010).

The parameter I is a measure of the total power in the wave, Q and U represent the linearly polarized components, and V represents the circularly polarized component (see Figure 2.1). The Stokes parameters have the dimensions of flux density, and they combine additively for independent waves.

2.3 JONES AND MUELLER RIME FORMULATIONS

Before continuing on to derive a more general relationship between the two-point coherency matrix and the voltage-current coherency, I will give a brief overview and derivation of the radio interferometer Measurement Equation of Hamaker et al. (1996). My motivation behind this is to highlight that Hamaker et. al.'s RIME is a special case of a more general (and thus less limited) coherency relationship.

In their seminal Measurement Equation paper, Hamaker et al. (1996) showed that Mueller and Jones calculi provide a good framework for modeling radio interferometers. In optics, Jones and Mueller matrices¹ are used to model the transmission of light through optical elements (Jones, 1941; Mueller, 1948). Mueller matrices are 4×4 matrices which act upon the Stokes vector

$$\vec{S} = (I \quad Q \quad U \quad V)^T, \quad (2.12)$$

whereas Jones matrices are only 2×2 in dimension and act upon the 'Jones vector': the electric

¹Note that matrices are represented with upright block capitals (non-italicized, e.g. V_{pq}), and vectors with italic letters with an arrow, e.g. \vec{e} .

field vector in a coordinate system such that z-axis is aligned with the Poynting vector,

$$\vec{e}(\vec{r}, t) = \begin{pmatrix} E_x(\vec{r}, t) & E_y(\vec{r}, t) \end{pmatrix}^T. \quad (2.13)$$

Jones calculus dictates that along a signal's path, any (linear) transformation can be represented with a Jones matrix, J:

$$\vec{e}_{out}(\vec{r}, t) = J\vec{e}_{in}(\vec{r}, t) \quad (2.14)$$

A useful property of Jones calculus is that multiple effects along a signal's path of propagation correspond to repeated multiplications:

$$\vec{e}_{out}(\vec{r}, t) = J_n \cdots J_2 J_1 \vec{e}_{int}(\vec{r}, t), \quad (2.15)$$

which can be collapsed into a single matrix when required.

The RIME uses Jones matrices to model the various corruptions and effects during a signal's journey from a source right through to the correlator. A block diagram for a (simplified) two element interferometer is shown in Figure 2.2. From left to right, the figure shows the journey of a signal from a source right through to the correlator. The radiation from the source is picked up by two antennas, which I have denoted with subscript p and q . The radiation follows a unique path to both of these antennas; each antenna also has associated with it a unique chain of analogue components that amplify and filter the signal to prepare it for correlation. Each of these subscripted boxes may be represented by a Jones matrix; alternatively an overall Jones matrix can be formed for the p and q branches (the dashed areas).

2.3.1 HAMAKER'S RIME DERIVATION

The derivation of the RIME is remarkably simple and elegant. For a single point source of radiation, the voltages induced at the terminals of a pair of antennas, p and q are

$$\vec{v}_p(t) = J_p \vec{e}_0(t) \quad (2.16)$$

$$\vec{v}_q(t) = J_q \vec{e}_0(t) \quad (2.17)$$

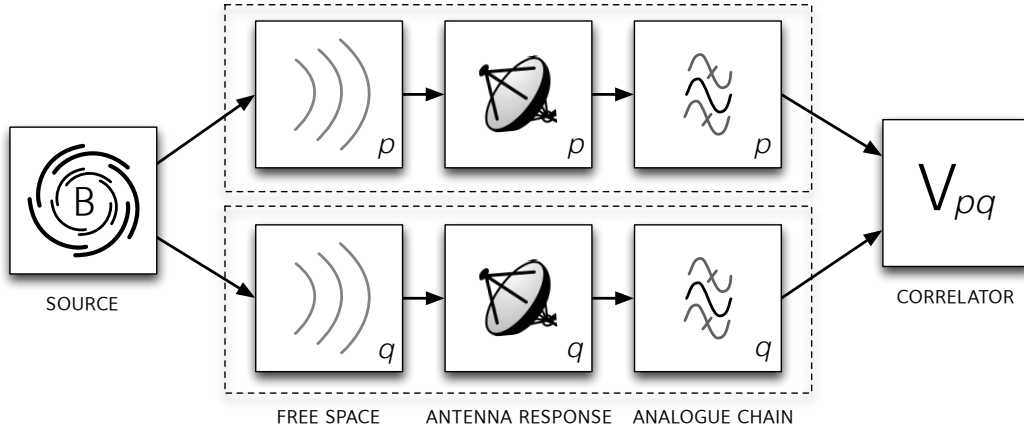


Figure 2.2: Block diagram showing a simple model of an interferometer that can be modelled with the RIME. Each block within the diagram may be represented by a Jones matrix.

In its simplest form, the RIME is formed by taking the outer product of these two relationships. Note that in their original paper, the authors use the Kronecker incarnation of the outer product, which we will denote with \star . We reserve the symbol \otimes for the matrix outer product of two matrices $a \otimes b = ab^H$, where H denotes the Hermitian conjugate transpose². Using the Kronecker outer product, the RIME is given by

$$\langle \vec{v}_p \star \vec{v}_q \rangle = (J_p \star J_q) \langle \vec{e}_0 \star \vec{e}_0 \rangle = (J_p \star J_q) \vec{e}_{00} \quad (2.18)$$

where J_p and J_q are the Jones matrices representing all transformations along the two signal paths, and $(J_p \star J_q)$ is a 4×4 matrix. Here, \vec{e}_{00} is the sky brightness of a single point source of radiation. For a multi-element interferometer, every antenna has its own unique Jones matrix, and a RIME may be written for every pair of antennas.

Due to their choice of outer product, Hamaker et. al. arrive at a coherency vector, as opposed to the coherency matrix of Wolf (1954):

$$\vec{e}_{pq}(r_p, r_q, \tau) = \langle e_p(r_p, t) \star e_q(r_q, t + \tau) \rangle = \begin{pmatrix} \langle E_{px} E_{qx}^* \rangle \\ \langle E_{px} E_{qy}^* \rangle \\ \langle E_{py} E_{qx}^* \rangle \\ \langle E_{py} E_{qy}^* \rangle \end{pmatrix}. \quad (2.19)$$

²For a discussion on the subtleties of outer product definition, see Smirnov (2011b) §A.6.1

The column vector of a point source at \vec{r}_0 is then \vec{e}_{00} ; that is, $p = q$ and $\tau = 0$. The vector \vec{e}_{00} is related to the Stokes vector by the transform

$$\begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 1 & 0 \\ 0 & -i & i & 0 \end{pmatrix} \begin{pmatrix} \langle E_{0x} E_{0x}^* \rangle \\ \langle E_{0x} E_{0y}^* \rangle \\ \langle E_{0y} E_{0x}^* \rangle \\ \langle E_{0y} E_{0y}^* \rangle \end{pmatrix}. \quad (2.20)$$

The quantity $(J_p \star J_q)$ in Equation 2.18 can therefore be viewed as a Mueller matrix. That is, Equation 2.18 can be considered a Mueller calculus based Measurement Equation for a radio interferometer. To summarize, Hamaker et. al. showed that:

- Jones matrices can be used to model the propagation of a signal from a radiation source through to the voltage at the terminal of an antenna.
- A Mueller matrix can be formed from the Jones terms of a pair of antennas, which then relates the measured voltage coherency of that pair to a source's brightness.

Showing that these calculi were applicable and indeed useful for modeling and calibrating radio interferometers was a huge step forward in radio polarimetry.

2.3.2 THE 2×2 RIME

In a later paper, Hamaker (2000) presents a modified formulation of the RIME, where instead of forming the coherency vector from the outer product (\star), the coherency matrix is formed from the outer product (\otimes):

$$E_{pq} = \langle \mathbf{e}_p \otimes \mathbf{e}_q \rangle = \begin{pmatrix} \langle E_{px} E_{qx}^* \rangle & \langle E_{px} E_{qy}^* \rangle \\ \langle E_{py} E_{qx}^* \rangle & \langle E_{py} E_{qy}^* \rangle \end{pmatrix} \quad (2.21)$$

The resulting coherency matrix is then shown to be related to the Stokes parameters by $E_{00} = B$, where

$$B = \begin{pmatrix} I + Q & U - iV \\ U + iV & I - Q \end{pmatrix}. \quad (2.22)$$

The equivalent to the RIME of Equation 2.18 is

$$\langle \mathbf{v}_p \otimes \mathbf{v}_q \rangle = \langle (J_p \mathbf{e}_0) \otimes (J_q \mathbf{e}_0) \rangle = J_p \langle \mathbf{e}_0 \otimes \mathbf{e}_0 \rangle J_q^H \quad (2.23)$$

or more simply,

$$V_{pq} = J_p B J_q^H. \quad (2.24)$$

This approach avoids the need to use 4×4 Mueller matrices, so is both simpler and computationally advantageous. This form is also cleaner in appearance — fewer indices are required — so is preferred by authors such as Smirnov (and myself).

Smirnov (2011a) takes the 2×2 version of the RIME as a starting point and extends the RIME to a full sky case. By treating the sky as a brightness distribution $B(\sigma)$, where σ is a direction vector, each antenna has a Jones term $J_p(\sigma)$ describing the signal path for a given direction. The visibility matrix V_{pq} is then found by integrating over the entire sky:

$$V_{pq} = \int_{4\pi} J_p(\sigma) B(\sigma) J_q^H(\sigma) d\Omega. \quad (2.25)$$

This is a more general form of Zernicke's propagation law, (§4.4-32 in Mandel and Wolf, 1995). Smirnov goes on to derive the Van-Cittert Zernicke theorem from this result; I return to vC-Z later in this chapter.

2.3.3 APPLICATIONS OF THE RIME

The RIME is used extensively in calibration of radio interferometric data. Both the MeqTrees (Noordam and Smirnov, 2010) and CASA (NRAO, 2011) software packages use the Jones formalism of the RIME to model the response of an interferometer. Calibration is achieved by inverting these matrices to extract the underlying source brightness.

The RIME is also a useful tool for simulating the response of an interferometric array. The OSKAR2 software package (Mort et al., 2012), uses the RIME to model the response of large interferometric arrays. OSKAR2 simulates observations by taking a sky model and evaluating the RIME over the sky to form visibilities. These visibilities can then be fed back into data reduction packages to test the quality of the calibration and imaging routines.

2.3.4 LIMITATIONS OF THE RIME

The Jones formalism of the RIME has certain limitations. For one, it does not account for the effect of noise. In reality, the analogue components of a telescope will add noise to the signal, which limits the sensitivity of the interferometer. One must integrate over time to decrease the effect of noise and increase sensitivity. The stability of the interferometer and source confusion dictate the maximum integration time achievable for a given instrument.

Another complication arises due to the Jones formalism's assumption that the Jones matrices are constant over time and frequency. For a correlation of length Δt over a bandwidth $\Delta \nu$, the visibility matrix is more precisely defined as

$$\langle V_{pq} \rangle = \frac{1}{\Delta t \Delta \nu} \int_{t_0}^{t_1} \int_{\nu_0}^{\nu_1} J_p(t, \nu) B J_q^H(t, \nu) d\nu dt, \quad (2.26)$$

as shown in Equation 22 of Smirnov (2011a). The effect of long integrations and wide bandwidths is to decrease the amplitude of the measured visibilities; it is commonly known as decorrelation. To achieve long integrations, one must evaluate the visibilities multiple times over the desired integration length. The same applies for bandwidth: most correlators channelize the bandwidth into evenly spaced bins, so that each bin can be treated separately.

In the final installment of his four paper series, Smirnov introduces a general tensor formalism of the RIME (Smirnov, 2011b). This formalism is better capable of describing mutual coupling between antennas, beamforming, and wide field polarimetry. These capabilities come at a cost: the tensor formalism is significantly more complex and far less intuitive than the Jones formalism.

Another limitation of the RIME is that it does not consider the state of the magnetic field. All formulations of the RIME to date — including the tensor formulation — suffer this limitation. In free space, the magnetic field coherency can be easily derived from the electric field coherency. However, at the boundary between two media, the magnetic field must also be considered. In the following section, I show that one must account for the magnetic field when considering the analogue components of a telescope.

2.4 A LIMITATION OF JONES CHAIN DECOMPOSITION

I will now show decomposing a telescope's overall system Jones matrix into subcomponents can yield invalid results. To illustrate this, I contrast Jones calculus with transmission matrix analysis, a well established method used in microwave engineering. At radio frequencies, a voltage is only well defined over a clearly specified terminal pair. In microwave network theory, such a pair of conductors is called a 'port'; microwave components are then modeled as networks of ports. A coaxial cable is a simple example of a 2-port network, and a power divider is an example of 3-port network. I refer the reader to the Chapter 4 of Pozar (2005) for an introduction to microwave networks.

The voltages, currents and impedances on the ports of an N-port network are linearly related to each other. For example, the 2-port transmission, or *ABCD* matrix, relates the voltages and currents of a 2-port network:

$$\begin{pmatrix} V_1 \\ I_1 \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} V_2 \\ I_2 \end{pmatrix}, \quad (2.27)$$

this is shown in Figure 2.3. If two 2-port networks are connected in cascade (see Figure 2.3), then the output is equal to the product of the transmission matrices representing the individual components:

$$\begin{pmatrix} V_1 \\ I_1 \end{pmatrix} = \begin{pmatrix} A_1 & B_1 \\ C_1 & D_1 \end{pmatrix} \begin{pmatrix} A_2 & B_2 \\ C_2 & D_2 \end{pmatrix} \begin{pmatrix} V_3 \\ I_3 \end{pmatrix}, \quad (2.28)$$

Like the Jones matrix, the *ABCD* matrix allows a signal's path to be modelled through multiplication of matrices representing discrete components. While the Jones matrix acts upon a pair of orthogonal electric field components, the *ABCD* matrix acts upon a voltage-current pair at a single port.

Jones matrices are not suitable for describing cascaded analogue systems. Consider the transmission of a voltage $V(t)$ through two analogue components connected in cascade. If one models each of these components, a and b , with a Jones matrix, J_1 and J_2 respectively, then the

voltage after component a will be

$$\begin{pmatrix} V_2 \\ 0 \end{pmatrix} = \begin{pmatrix} J_1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} V_1 \\ 0 \end{pmatrix}. \quad (2.29)$$

Or more simply, $V_2 = J_1 V_1$. Similarly, after component b the relationship is

$$V_1 = V_3 / J_2 J_1. \quad (2.30)$$

In contrast, the $ABCD$ matrix formulation predicts

$$V_2 = B_1 I_1 + A_1 V_1 \quad (2.31)$$

after the first component. If the impedance at the input port is Z , we can use the impedance relation³ $Z = V/I$ and conclude that

$$J_1 = 1 / (A_1 + \frac{B_1}{Z_1}). \quad (2.32)$$

After the second component the $ABCD$ prescription gives

$$V_3 = A_2 (\frac{B_1}{Z_1} + A_1) V_1 + B_2 (\frac{D_1}{Z_1} + C_1) V_1, \quad (2.33)$$

so for Equation 2.30 and Equation 2.33 to be equal, we require that

$$J_2 = \frac{J_1^{-1}}{A_2 J_1^{-1} + B_2 (C_1 + \frac{D_1}{Z_1})}. \quad (2.34)$$

This is only satisfied if:

- We allow J_2 to be dependent upon J_1 . However, if we allow this, we have not successfully modelled our components as discrete devices.
- The first component has $C_1 = -D_1 / Z_1$.

³This relation assumes all port-to-port impedances are identical

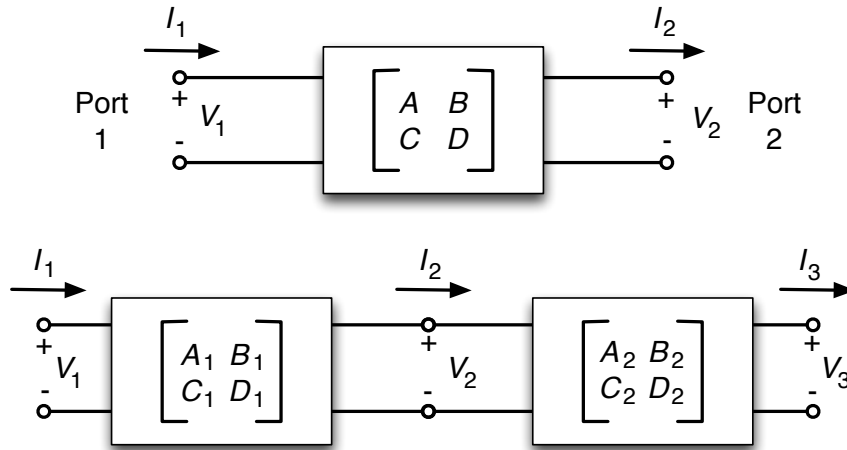


Figure 2.3: Top: The ABCD matrix for a 2-port network. Bottom: Connecting two components in cascade. Diagram adapted from Pozar (2005).

- The first component has $C_1 = D_1 = 0$.
- The second component has $B_2 = 0$.

Jones calculus is therefore only valid for describing 2-port networks if at least one of the last three cases are true, or if the off diagonal terms of the transmission matrices are small enough that we can neglect their contribution. Otherwise, the Jones matrices of each element are dependent upon every preceding Jones matrix, which is hardly the decomposition we set out to achieve.

2.4.1 CHANGES IN IMPEDANCE AND THE TRANSMISSION MATRIX

The impedance relation $Z = V/I$ used in the preceding section is a simplification that is, in general, not valid for N-port networks. The general relationships between the voltages and currents on a component's ports are encoded in the impedance matrix Z :

$$\begin{pmatrix} V_1 \\ V_2 \\ \vdots \\ V_N \end{pmatrix} = \begin{pmatrix} Z_{11} & Z_{12} & \cdots & Z_{1N} \\ Z_{21} & \ddots & & \vdots \\ \vdots & & & \vdots \\ Z_{N1} & \cdots & \cdots & Z_{NN} \end{pmatrix} \begin{pmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{pmatrix}, \quad (2.35)$$

where Z_{ab} is the port-to-port impedance from port a to port b. The elements in the 2-port transmission matrix are related to these port-to-port impedances by

$$A = Z_{11}/Z_{21} \quad (2.36)$$

$$B = \frac{Z_{11}Z_{22} - Z_{12}Z_{21}}{Z_{21}} \quad (2.37)$$

$$C = 1/Z_{21} \quad (2.38)$$

$$D = Z_{22}/Z_{21}. \quad (2.39)$$

Comparing this to the conditions placed on Equation 2.34 above, we note:

- $B = 0$ for all components that have a perfect impedance match between ports.
- C approaches zero as Z_{21} approaches infinity; this represents an open circuit.
- For $C = -D/Z_{21}$ to hold, we would require a negative port-to-port impedance. This can be satisfied by active components but is very rarely encountered in practice.

With these conditions in mind, we may conclude that Jones calculus is valid for describing components with perfect or close to perfect impedance matches. In the sections that follow, I probe a little deeper into how impedance mismatches are related to electromagnetic waves and coherency statistics.

2.5 COHERENCY IN RADIO ASTRONOMY

In this section, I introduce the coherency matrix \mathbb{W}_{pq} and the voltage-current coherency matrix \mathbb{V}_{pq} . In Section 2.6, I formulate relationships between source brightness and measured coherency based upon these matrices.

2.5.1 ELECTROMAGNETIC COHERENCY

To begin, I introduce the coherence matrices of Wolf (1954), that fully describe the coherence statistics of an electromagnetic field. We may start by introducing $\vec{e}(\vec{r}, t)$ and $\vec{h}(\vec{r}, t)$ as the complex analytic representations of the electric and magnetic field vectors at a spacetime point

(\vec{r}, t) :

$$\vec{e}(\vec{r}, t) = \begin{pmatrix} E_x(\vec{r}, t) & E_y(\vec{r}, t) & E_z(\vec{r}, t) \end{pmatrix}^T \quad (2.40)$$

$$\vec{h}(\vec{r}, t) = \begin{pmatrix} H_x(\vec{r}, t) & H_y(\vec{r}, t) & H_z(\vec{r}, t) \end{pmatrix}^T \quad (2.41)$$

The coherence matrices are then defined by the formulae

$$E_{pq}(\vec{r}_p, \vec{r}_q, \tau) = \left\langle \left\langle E_j^*(\vec{r}_p, t) E_k^*(\vec{r}_q, t + \tau) \right\rangle \right\rangle \quad (2.42)$$

$$H_{pq}(\vec{r}_p, \vec{r}_q, \tau) = \left\langle \left\langle H_j^*(\vec{r}_p, t) H_k^*(\vec{r}_q, t + \tau) \right\rangle \right\rangle \quad (2.43)$$

$$M_{pq}(\vec{r}_p, \vec{r}_q, \tau) = \left\langle \left\langle E_j^*(\vec{r}_p, t) H_k^*(\vec{r}_q, t + \tau) \right\rangle \right\rangle \quad (2.44)$$

$$N_{pq}(\vec{r}_p, \vec{r}_q, \tau) = \left\langle \left\langle H_j^*(\vec{r}_p, t) E_k^*(\vec{r}_q, t + \tau) \right\rangle \right\rangle. \quad (2.45)$$

Here, j and k are indices representing the x, y, z subscripts from Cartesian coordinates. E_{pq} and H_{pq} are called the electric and the magnetic coherence matrices, and M_{pq} and N_{pq} are called the mixed coherence matrices. The subscripts p and q correspond to the spacetime points (\vec{r}_p, t) and $(\vec{r}_q, t + \tau)$, respectively. We may arrange these into a single 6×6 matrix \mathbb{W}_{pq} that is equivalent to the time averaged outer product of the electromagnetic field column vectors at spacetime points (\vec{r}_p, t) and $(\vec{r}_q, t + \tau)$:

$$\mathbb{W}_{pq} = \left\langle \left(\begin{pmatrix} \vec{e}_p \\ \vec{h}_p \end{pmatrix} \otimes \begin{pmatrix} \vec{e}_q \\ \vec{h}_q \end{pmatrix} \right) \right\rangle = \begin{pmatrix} E_{pq} & M_{pq} \\ N_{pq} & H_{pq} \end{pmatrix} \quad (2.46)$$

This matrix fully describes the coherence properties of an electromagnetic field at two points in spacetime. We will refer to this matrix as the two point coherency matrix. It is worth noting that:

- When $\vec{r}_p = \vec{r}_q$ and $\tau = 0$ we retrieve the ‘EM sixtor matrix’ of Bergman and Carozzi (2008). Bergman and Carozzi (2008) show this sixtor matrix is related to what they refer to as ‘canonical electromagnetic observables’: a unique set of Stokes-like parameters that are irreducible under Lorentz transformations. These are used in the the analysis of electromagnetic field data from spacecraft.
- When $\vec{r}_p = \vec{r}_q$ and $\tau = 0$, and we choose a coordinate system with z in the direction of

propagation (i.e. along the Poynting vector), E becomes what Smirnov (2011a) refers to as the brightness matrix, B . To highlight the relationship between \mathbb{W}_{00} and the brightness matrix, we use the notation $\mathbb{B} \triangleq \mathbb{W}_{00}$, and will refer to \mathbb{B} as the brightness coherency.

2.5.2 VOLTAGE AND CURRENT COHERENCY

A radio telescope converts a free space electromagnetic field into a time varying voltage, which we then measure after signal conditioning (e.g. amplification and filtering). As such, radio interferometers measure coherency statistics between time varying voltages.

One may model the analogue components of a telescope as a 6-port network, with three inputs ports and three output ports. I propose this so that there is an input-output pair of ports for each of the orthogonal components of the electromagnetic field. We can then define a set of voltages $v(t)$ and currents $i(t)$:

$$\vec{v}(t) = \begin{pmatrix} V_x(t) & V_y(t) & V_z(t) \end{pmatrix}^T \quad (2.47)$$

$$\vec{i}(t) = \begin{pmatrix} I_x(t) & I_y(t) & I_z(t) \end{pmatrix}^T. \quad (2.48)$$

In practice, most telescopes are single or dual polarization, so only the x and y components are sampled. Nonetheless, it is possible to sample all three components with three orthogonal antenna elements (Bergman et al., 2005). The voltage response of an antenna is linearly related to the electromagnetic field strength (Hamaker et al., 1996), and the current is linearly related to voltage by Ohm's law (§7.1.1 in Griffiths, 1999), so we may write a general linear relationship

$$\begin{pmatrix} \vec{v}(t) \\ \vec{i}(t) \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} \vec{e}(\vec{r}, t) \\ \vec{h}(\vec{r}, t) \end{pmatrix}, \quad (2.49)$$

where A , B , C and D are block matrices forming an overall transmission matrix \mathbb{T}' . We can now define a matrix of voltage-current coherency statistics that consists of the block matrices

$$V_{pq}(\tau) = \left(\left\langle \left\langle V_j^*(t) V_k^*(t + \tau) \right\rangle \right\rangle \right) \quad (2.50)$$

$$I_{pq}(\tau) = \left(\left\langle \left\langle I_j^*(t) I_k^*(t + \tau) \right\rangle \right\rangle \right) \quad (2.51)$$

$$K_{pq}(\tau) = \left(\left\langle \left\langle V_j^*(t) I_k^*(t + \tau) \right\rangle \right\rangle \right) \quad (2.52)$$

$$L_{pq}(\tau) = \left(\left\langle \left\langle I_j^*(t) V_k^*(t + \tau) \right\rangle \right\rangle \right), \quad (2.53)$$

these are analogous to (and related to) the electromagnetic coherency matrices above⁴. In a similar manner to the two-point coherency matrix, we define \mathbb{V}_{pq}

$$\mathbb{V}_{pq} = \left\langle \left(\begin{array}{c} \vec{v}_p \\ \vec{i}_p \end{array} \right) \otimes \left(\begin{array}{c} \vec{v}_q \\ \vec{i}_q \end{array} \right) \right\rangle = \begin{pmatrix} V_{pq} & K_{pq} \\ L_{pq} & I_{pq} \end{pmatrix}, \quad (2.54)$$

which we will refer to as the voltage-current coherency matrix. This is analogous to the ‘visibility matrix’, V_{pq} , of Smirnov (2011a).

2.6 TWO POINT COHERENCY RELATIONSHIPS

Now we have introduced the two-point coherency matrix \mathbb{W}_{pq} and the voltage-current coherency matrix \mathbb{V}_{pq} , we can formulate relationships between the two. In this section, I first formulate a general coherency relationship describing propagation from a source of electromagnetic radiation to two spacetime coordinates. I then show that this relationship underlies both the RIME and the vC-Z relationship.

2.6.1 A GENERAL TWO POINT COHERENCY RELATIONSHIP

Suppose we have two sensors, located at points \vec{r}_p and \vec{r}_q , which fully measure all components of the electromagnetic field. Assuming linearity, the propagation of an electromagnetic field $\vec{f}_0 = \left(\vec{e}(\vec{r}_0, t) \quad \vec{h}(\vec{r}_0, t) \right)^T$ from a point \vec{r}_0 to \vec{r}_p and \vec{r}_q can be encoded into a 6×6 matrices, \mathbb{T}_p

⁴Note spatial location r is no longer relevant as the voltage propagates through analogue components with clearly defined inputs and outputs.

and \mathbb{T}_q :

$$\vec{f}_p = \mathbb{T}_p \vec{f}_0 \quad (2.55)$$

$$\vec{f}_q = \mathbb{T}_q \vec{f}_0 \quad (2.56)$$

The coherence between the two signals \vec{f}_p and \vec{f}_q is then given by the matrix \mathbb{W}_{pq} :

$$\mathbb{W}_{pq} = \langle \vec{f}_p \otimes \vec{f}_q \rangle \quad (2.57)$$

$$= \langle (\mathbb{T}_p \vec{f}_0) \otimes (\mathbb{T}_q \vec{f}_0) \rangle \quad (2.58)$$

$$= \langle \mathbb{T}_p (\vec{f}_0 \otimes \vec{f}_0) \mathbb{T}_q^H \rangle \quad (2.59)$$

$$= \mathbb{T}_p \mathbb{B} \mathbb{T}_q^H \quad (2.60)$$

we can write this in terms of block matrices

$$\begin{pmatrix} E_{pq} & M_{pq} \\ N_{pq} & H_{pq} \end{pmatrix} = \begin{pmatrix} A_p & B_p \\ C_p & D_p \end{pmatrix} \begin{pmatrix} E_{00} & M_{00} \\ N_{00} & H_{00} \end{pmatrix} \begin{pmatrix} A_q & B_q \\ C_q & D_q \end{pmatrix}^H \quad (2.61)$$

This is the most general form that relates the coherence at two points \vec{r}_p and \vec{r}_q , to the electromagnetic energy density at point \vec{r}_0 .

In radio astronomy, antennas are used as sensors to measure the electromagnetic field. Following from Equation 2.61, we may write an equation relating voltage and current coherency:

$$\mathbb{V}_{pq} = \mathbb{T}'_p (\mathbb{T}_p \mathbb{B} \mathbb{T}_q^H) \mathbb{T}'_q{}^H. \quad (2.62)$$

As the \mathbb{T}' and \mathbb{T} matrices are both 6×6 , we can collapse these matrices into one overall matrix. Equation 2.62 then becomes

$$\mathbb{V}_{pq} = \mathbb{T}_p \mathbb{B} \mathbb{T}_q^H, \quad (2.63)$$

which is the general form that relates the voltage-coherency matrix \mathbb{V}_{pq} to the brightness coherency \mathbb{B} .

Equation 2.63 is the central result of this thesis. It is a general case which relates the EM field at a given point in space-time to the voltage and current coherencies in between pairs of telescopes. In the sections that follow, I show that generalized versions of the Van-Cittert-Zernicke theorem and Radio Interferometer Measurement Equation may be formulated based upon this coherency relationship, and that the common formulations can be derived from these general results.

2.6.2 THE RADIO INTERFEROMETER MEASUREMENT EQUATION

By comparing Equation 2.24 with Equation 2.61, it is apparent that the Jones formulation of the RIME is retrieved by setting all but the top left block matrices to zero, such that we have

$$V_{pq} = A_p E_{00} A_q^H. \quad (2.64)$$

But under what assumptions may we ignore the other entries of Equation 2.61? To answer this, we note that for monochromatic plane waves in free space E and H are in phase and mutually perpendicular:

$$\vec{e}(\vec{r}, t) = \begin{pmatrix} e_x(\vec{r}, t) & e_y(\vec{r}, t) \end{pmatrix}^T \quad (2.65)$$

$$\vec{h}(\vec{r}, t) = \frac{1}{c_0} \begin{pmatrix} -e_y(\vec{r}, t) & e_x(\vec{r}, t) \end{pmatrix}^T, \quad (2.66)$$

where c_0 is the magnitude of the speed of light. In such a case, all coherence statistics can be derived from the 2×2 brightness matrix B. Carozzi and Woan (2009) show that the field coherencies can be written

$$\mathbb{B} = \begin{pmatrix} E_{pq} & M_{pq} \\ N_{pq} & H_{pq} \end{pmatrix} = \begin{pmatrix} B & BF^T \\ FB & FBF^T \end{pmatrix} \quad (2.67)$$

where F is the matrix

$$F = \frac{1}{c_0} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \quad (2.68)$$

Under these conditions, the rank of \mathbb{B} is 2, so the relationship in Equation 2.61 is overconstrained. It follows that the 2×2 RIME is perfectly acceptable — and indeed preferable to Equation 2.61

— for describing coherency of plane waves that propagate through free space.

There are numerous situations in which we cannot assume that we have monochromatic plane waves. This includes near field sources where the wavefront is not well approximated by a plane wave; propagation through ionized gas; and situations where we choose not to treat our field as a superposition of quasi-monochromatic components. Most importantly, the assumptions that underlie the 2×2 RIME do not hold within the analogue components of a telescope, where the signal does not enjoy free space impedance.

2.6.3 A 2N-PORT TRANSMISSION MATRIX BASED RIME

For a dual polarization telescope, a 4-port description (2-in 2-out) of our analogue system is more appropriate than the general 6-port description. Using the result 2.67 above and Equation 2.63, we can write a relationship

$$V_{pq} = \begin{pmatrix} A_p & B_p \\ C_p & D_p \end{pmatrix} \begin{pmatrix} B & BF^T \\ FB & FBF^T \end{pmatrix} \begin{pmatrix} A_q & B_q \\ C_q & D_q \end{pmatrix}^H, \quad (2.69)$$

Here, all block matrices have been reduced in dimensions from 3×3 to 2×2 . This version of the RIME retains the ability to model analogue components, but uses the approximations of the vC-Z to express \mathbb{B} in terms of the regular 2×2 brightness matrix B . The transmission matrices here are similar to the 2N-port transmission matrices as defined by Faria (2002).

The main strength of Equation 2.69 over the 2×2 RIME is that the transmission matrices may be broken down into a chain of cascaded components, whereas this is not supported by Jones calculus. That is, we may write

$$\begin{pmatrix} A_p & B_p \\ C_p & D_p \end{pmatrix} = \begin{pmatrix} A_{np} & B_{np} \\ C_{np} & D_{np} \end{pmatrix} \cdots \begin{pmatrix} A_{1p} & B_{1p} \\ C_{1p} & D_{1p} \end{pmatrix} \quad (2.70)$$

for any number of components n in the analogue chain. In effect, this 4-port relationship allows us to ‘open up’ our telescope and describe analogue components discretely; this cannot be done with Jones calculus. We discuss how real analogue components could be modelled in §2.7.1 below.

2.6.4 THE VAN-CITTERT-ZERNICKE RELATIONSHIP

The Van-Cittert-Zernicke theorem (vC-Z), is the cornerstone of radio interferometry. It states that the measurements we make with an interferometer are related to a (far-field) source brightness by a Fourier Transform, under certain circumstances.

Suppose we have N sources of radiation. Their contributions will add up linearly. but the signal transmission matrix is different for each source and sensor pair:

$$\mathbb{W}_{pq} = \sum_{n=1}^N \mathbb{T}_{np} \mathbb{B}_n \mathbb{T}_{nq}^H \quad (2.71)$$

In radio astronomy, we make the simplifying assumption that the sources are all projected upon a 'celestial sphere'. Then we have

$$\mathbb{W}_{pq} = \int_{4\pi} \mathbb{T}_{np}(\sigma) \mathbb{B}(\sigma) \mathbb{T}_{nq}^H(\sigma) d\Omega \quad (2.72)$$

This is a more general form of Zernicke's propagation law, (eqn 4.4-32 in Mandel and Wolf, 1995). We may derive the Van-Cittert-Zernicke theorem from this general result, but for radio astronomy applications, the full 6×6 generalization is unnecessary. In deriving the standard Fourier Transform vC-Z relationship of Zernicke (1938), the following assumptions allow simplifying approximations to be applied:

- That our source is in the far-field. In the case of radio astronomy they are many orders of magnitude farther than the spacing between our interferometer elements. As a result, the electromagnetic field across the interferometer is approximately a plane wave.
- That our source is quasi-monochromatic. As a consequence of the linearity of Maxwell's equations, radiation from any source can be treated as a superposition of quasi-monochromatic components;
- That our source is propagating in free space.

These are precisely the same assumptions as those governing the 2×2 RIME. So, our general form of Zernicke's propagation law does not lead to a new vC-Z relationship. This reinforces the declaration by Carozzi and Woan (2009) that their wide-field, polarimetric vC-Z relations

provides a complete description of the second-order coherence of the electromagnetic field.

Nonetheless, the assumptions made in vC-Z relationships do not apply within the analogue components of telescopes. It follows that frameworks such as the 2×2 RIME, which make use of such assumptions, are not well suited for describing chains of cascaded analogue components.

2.7 DISCUSSION

We are fortunate to be able to use results from both microwave network theory and optics to explain a signal's journey from an astrophysical source right through to the end of a telescope's analogue chain. Microwave network theory is better suited to describe the analogue components of a radio telescope, and Jones matrices are more suitable for describing free space propagation of plane waves.

There are two methods of combining these two fields. The first is to simply use Jones matrices to describe free space propagation, and microwave network methods to describe discrete analogue components. An overall 'system Jones' matrix may be derived to describe instrumental effects, but this matrix should not be decomposed into a Jones chain. The second method is to use the $2N$ -port RIME from Section 2.6.3 above.

2.7.1 MODELLING REAL ANALOGUE COMPONENTS

Voltage, current and impedance are somewhat abstract concepts at microwave frequencies, so engineers often use scattering parameters to quantify a device's characteristics. Scattering parameters related the incident and reflected voltage waves on the ports of a microwave network. The scattering matrix, S , is given by

$$\begin{pmatrix} V_1^- \\ V_2^- \\ \vdots \\ V_n^- \end{pmatrix} = \begin{pmatrix} S_{11} & S_{12} & \cdots & S_{1n} \\ S_{21} & & & \vdots \\ \vdots & & \ddots & \vdots \\ S_{n1} & \cdots & \cdots & S_{nn} \end{pmatrix} \begin{pmatrix} V_1^+ \\ V_2^+ \\ \vdots \\ V_n^+ \end{pmatrix}, \quad (2.73)$$

where V_n^+ is the amplitude of the voltage wave incident on port n , and V_n^- is the amplitude of the voltage wave reflected from port n . More complete coverage is given in Chapter 4 of Pozar

(2005).

For a dual polarization system (we will label the polarizations x and y) with negligible crosstalk, we can model the analogue chain for each polarization as a discrete 2-port network. Assuming that the analogue chains have the same number of components (but not that the components are identical), the transmission matrix for each pair of components is

$$T = \begin{pmatrix} \tilde{A}_x & 0 & \tilde{B}_x & 0 \\ 0 & \tilde{A}_y & 0 & \tilde{B}_y \\ \tilde{C}_x & 0 & \tilde{D}_x & 0 \\ 0 & \tilde{C}_y & 0 & \tilde{D}_y \end{pmatrix}, \quad (2.74)$$

where the elements are from the ABCD matrices of the x and y polarizations, and are given by the relations

$$\tilde{A} = \frac{1 + S_{12}S_{21} + S_{22} - S_{11}(1 + S_{22})}{2S_{12}} \quad (2.75)$$

$$\tilde{B} = -Z_0 \frac{(1 + S_{11} - S_{12}S_{21} + S_{22} + S_{11}S_{22})}{2S_{12}} \quad (2.76)$$

$$\tilde{C} = \frac{1}{Z_0} \frac{-1 + S_{11} + S_{12}S_{21} + S_{22} - S_{11}S_{22}}{2S_{12}} \quad (2.77)$$

$$\tilde{D} = \frac{1 + S_{11} + S_{12}S_{21} - S_{22} - S_{11}S_{22}}{2S_{12}} \quad (2.78)$$

Here, Z_0 is the characteristic impedance of the analogue chain, which in most telescopes is set to 50 or 75 ohms. We have added tildes to the ABCD parameters, as we are using the inverse definition to that in Pozar (2005).

If the system has significant crosstalk between polarizations, the analogue chain is more accurately modelled as a 4-port network. In this case the relationships are not so simple, and the off-diagonal entries of the block matrices of T will no longer be zero.

2.7.2 REFLECTION COEFFICIENT

When a voltage wave travelling upon a transmission line is subject to a change in impedance or discontinuity, some of its energy is reflected. The total voltage on a transmission line can be

written as a sum of incident V_o^+ and reflected V_o^- waves:

$$V(z) = V_o^+ e^{-i\beta z} + V_o^- e^{i\beta z}. \quad (2.79)$$

The ratio of the amplitude of the reflected and incident waves defines the voltage reflection coefficient, Γ :

$$\Gamma = \frac{V_o^-}{V_o^+} = \frac{Z_L - Z_0}{Z_L + Z_0}, \quad (2.80)$$

where Z_0 is the impedance of the transmission line before the impedance change, and Z_L is the impedance after. For a 2-port network, the reflection coefficient $\Gamma \equiv S_{11}$ at the input port, and $\Gamma \equiv S_{22}$ at the output port. The transmission matrix in Section 2.7.1 above is therefore dependent upon the reflection coefficient at the ports.

In contrast, Jones calculus does not consider reflected voltages, nor reverse transmission. This means that effects seen in microwave networks — such as voltage standing waves, which are caused by superposition of the reflected and incident waves on a transmission line — cannot be modelled using Jones calculus.

2.7.3 APPLICATION TO CALIBRATION

A central conclusion of this chapter is that Jones calculus is not suitable for describing analogue components. This does not, however, mean that Jones calculus cannot be used to calibrate a system. As has been discussed, Jones calculus provides an appropriate framework for describing free space propagation of electromagnetic waves. Once the free space wave is converted into a voltage, microwave network methods are more appropriate for describing the propagation through discrete analogue components.

One can avoid encountering issues by simply forming a ‘system Jones’ matrix to describe the overall instrumental response. This matrix should fully describe the entire analogue chain of the telescope. It should be stressed that one cannot decompose this matrix into a Jones chain, describing each component with its own Jones matrix. Nevertheless, one could decompose this matrix into a time-varying and a static matrix for the purposes of describing a time-varying phenomena within the analogue system. This approach treats the analogue systems of a telescope as a ‘black box’: the overall response cannot be attributed to any component in particular. For

many applications, such an approach is perfectly valid.

On the other hand, a black-box approach does not give any insight into how the system operates. There are cases in which the effects of a single component in the chain is of interest. To give an example, one may be interested in how upgrading a specific component in a telescope will affect observed visibilities. Another example would be if one wished to model how thermal expansion of transmission cables affects the overall system response. In such cases, the 2N-port RIME from Section 2.6.3 above is more suitable than the Jones formulation.

2.7.4 APPLICATION TO END-TO-END SIMULATIONS

Software packages such as OSKAR2 (Mort, 2010; Mort et al., 2012) use the 2×2 RIME to simulate the response of aperture arrays and interferometric arrays. OSKAR2 has been developed as a tool for simulating the performance of aperture arrays for the SKA.

By bootstrapping such software to use the 2N-port port RIME as its measurement framework, we could better simulate how single components affect the performance of the telescope. One could then conduct a full end-to-end simulations to investigate how different amplifier and filter implementations affect visibilities. In this way, a 2N-port RIME based version of OSKAR2 could be used as a tool to generate datasets which more accurately reflect the performance of aperture arrays. Having such a capability would be advantageous for designing and specifying analogue components for the SKA.

2.7.5 CONCLUSIONS

The radio interferometer Measurement Equation provides a powerful framework for describing a signal's journey from an astrophysical source to the receiver of a radio telescope. Nevertheless, the Jones calculus it employs is in general not sufficient to describe the components within the analogue chain of a radio telescope. Within the telescope, methods from microwave network theory, such as transmission matrices and scattering parameters, are more appropriate.

The 2N-port RIME we have presented here is formulated from a general electromagnetic basis, which gives us better insight into the fundamental physics involved. The 2N-port RIME prescription is capable of modelling real components connected in cascade; this may be of use in calibration of existing telescopes and simulation of future telescopes.

CHAPTER 3

D-PAD: AN L-BAND SPARSE APERTURE ARRAY DEMONSTRATOR

3.1 INTRODUCTION

We are well on our way toward detecting the Epoch of Reionization (EoR), with the LOFAR, MWA, and PAPER projects all publishing projected timescales for EoR statistical detection (Beardsley et al., 2012; Parsons et al., 2011; Zaroubi et al., 2012). All three of these projects are arrays of simple antennas, whose signals are combined digitally to act as a single telescope. These antennas have no moving parts; instead, digital signal processing techniques are used to form beams on (and images of) the sky, achieving unprecedented fields of view and survey speeds.

The acoustic oscillations of the early universe have left their imprint on the matter distribution at all epochs, not just during the EoR. As discussed in Chapter 1, this signature is known as Baryonic Acoustic Oscillations (BAO). In this chapter, I discuss the development of a sparse aperture array demonstrator that operates over 1.0-1.5 GHz, which retains the core characteristics that have made sparse arrays an attractive technology for EoR experiments. This work paves the way for a future sparse array which can be used for BAO detection.

3.1.1 DANNY'S PHD ARRAY DEMONSTRATOR

D-PAD (Danny's PhD Array Demonstrator), is a sparse aperture array that operates over L-band (1.0-1.5 GHz), which I built in conjunction with Sascha Schediwy during my DPhil. The D-PAD

project aims to investigate the viability of sparse aperture arrays for astronomy over the 1.0–1.5 GHz band; in particular, the viability of sparse arrays for BAO detection experiments.

There are some notable advantages to designing at 1.0–1.5 GHz over lower frequencies, such as the SKA-low range of 70–450 MHz. Firstly, the fractional bandwidth is much lower. Designing an antenna with a good impedance match to its LNA becomes significantly harder as fractional bandwidth increases. Secondly, the physical size of resonant electromagnetic structures such as antennas scales with wavelength. This means that the antennas will be easier to store, easier to deploy, and less costly to fabricate. Thirdly, and perhaps of most significance, there are many commercially available (and affordable) low noise amplifiers and other analogue components at this frequency range. A working system can be assembled without every component being of bespoke design.

Another consideration that must be made is the radio frequency interference (RFI) environment. In the United Kingdom, the radio spectrum is heavily occupied by terrestrial sources such as mobile communications, television and radio broadcasting. In particular, digital television broadcasting occupies much of the 400–800 MHz band, and GSM-900 mobile phone communication band spans 890–960 MHz (OFCOM, 2010). This makes radio astronomy exceptionally troublesome at these frequencies. The exceptional signal strength of FM radio band (87.5–108 MHz) was a factor which motivated the LOFAR to build two aperture arrays: one which operates over 10–80 MHz, and another which operates from 120–240 MHz.

The 1.0–1.5 GHz band, while far from being RFI-free, has significantly less RFI occupancy than lower frequencies. This band also corresponds to the third Nyquist zone of a digitizer running at a sampling rate 1.0 GHz, which means that it can be sampled directly (i.e. without being mixed down to baseband, which requires a local oscillator and mixer). Another advantage is that the 500 MHz bandwidth completely covers the SKA-low range of 70–450 MHz. As such, all the digital systems of D-PAD can be repurposed for testing of SKA-low antennas and analogue hardware. The culmination of all these factors was the motivation for designing D-PAD to operate over 1.0–1.5 GHz.

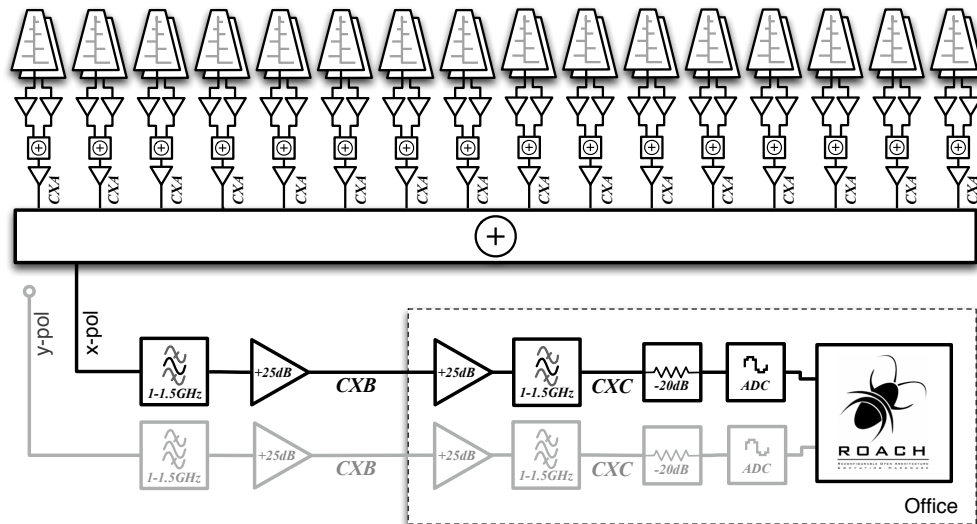


Figure 3.1: Block diagram of the D-PAD single station. Sixteen dual polarizations antennas are amplified, beamformed, and filtered before digitization. Only one polarization is shown in the diagram.

3.2 SYSTEM OVERVIEW

The D-PAD hardware was specifically designed to test the viability of sparse aperture arrays over 1.0–1.5 GHz. A system consisting of 16 dual polarization antennas was produced, each of which has an associated chain of analogue components. A 16 antenna system is large enough to allow different antenna configurations to be tested, while small enough to be affordable and easily reconfigured. I will refer to this system as the D-PAD ‘single station’.

Each of the 16 D-PAD single station antennas were amplified using the receiver board (Section 3.4), before being beamformed (phased to zenith) by an analogue power combiner (Section 3.5). After power combining, the signal was amplified (Section 3.6.2), then filtered using the bandpass filter of Section 3.6.1. The signal was then transported from the telescope site to a close-by processing zone. For commissioning tests, the telescope was located on the rooftop of astrophysics department, and all electronic hardware was located in a nearby office. Inside the office, a second stage of signal conditioning was applied, in which the signal is further amplified filtered, and attenuated so that it was at a level suitable for digitization. After this, the signal was digitized, processed, and then analysed.

A CASPER iADC was used to digitize the signals, and a ROACH board was used for digital signal processing. These are described in further detail in the following chapter, namely Sections 4.5.3

and 4.5.2, respectively. I will refer to the collection of analogue components from and including the power combiner, up to but not including the digitizer, as the D-PAD ‘analogue chain’. As beamforming and signal conditioning is applied to both polarizations discretely, two essentially identical analogue chains are required for one station.

3.2.1 SIGNAL CONDITIONING REQUIREMENTS

Before digitization, the voltage signals at the terminals of the antennas must be amplified to a suitable power level and filtered to prevent the aliasing of out-of-band signals. It is important that the signal-to-noise ratio (SNR) of the received signal is not degraded by the amplification, transport and filtering of the signal. As discussed in Section 3.7.2, the order of analogue components affects the system temperature. D-PAD aimed to achieve a balance between system temperature and complexity, using commercially available components where available.

For digitization, the signal from the antenna array must be amplified and filtered so that it is at a suitable power level. The digitizer used in D-PAD is 8 bit¹, and has a full scale analogue input voltage of ± 500 mVpp (Atmel, 2009). This means the minimum distinguishable voltage is

$$\frac{500 \text{ mV}_{pp}}{2^8 \sqrt{2}} = 1.38 \text{ mV}_{RMS}, \quad (3.1)$$

or -44.1 dBm; similarly, the maximum detectable signal is 4.0 dBm. If we allow 5 bits of overhead for strong RFI sources (and astrophysical signals, which are much weaker), we arrive at a nominal input power level of about -27 dBm.

To work out the minimum amplification required, we first need to find the equivalent noise power of the system. The noise power (in dBm units), is given by

$$N_{sys} = 10 \log_{10} \left(\frac{k T_{sys} \Delta \nu}{1 \times 10^{-3} \text{ W}} \right). \quad (3.2)$$

For a system temperature of 100 K, with $\Delta \nu = 500$ MHz, the noise power is -91.6 dB; so, we require of order 65 dB of gain to reach the iADC’s nominal power level.

¹The effective number of bits (ENOB) of this ADC is ~ 7.2 ; this is discussed further in §4.5.3

3.2.2 BEAM PATTERN

The D-PAD antennas were laid out in a planar 4×4 grid with 2λ spacing at 1.5 GHz ($\lambda=20$ cm).

The beam pattern of a planar array can be expressed analytically as

$$E(r) = f(\theta, \phi)F_x(\theta, \phi)F_y(\theta, \phi), \quad (3.3)$$

where $f(\theta, \phi)$ is the beam pattern of a single antenna element, and $F_x(\theta, \phi)$ and $F_y(\theta, \phi)$ are two linear array factors in the x- and y-directions:

$$F_x(\theta, \phi) = \sum_{k=1}^K e^{jk_0(k-1)d_x \sin\theta \cos\phi} \quad (3.4)$$

$$F_y(\theta, \phi) = \sum_{k=1}^K e^{jk_0(k-1)d_y \cos\theta \sin\phi}, \quad (3.5)$$

where $k_0 = 2\pi/\lambda$. For our 4×4 array at 1.5 GHz, $d_x = d_y = 2\lambda$, so the beam pattern may be written as

$$E(r) = f(\theta, \phi) \sum_{n=1}^4 e^{j4\pi(n-1)\sin\theta \cos\phi} \sum_{m=1}^4 e^{j4\pi(m-1)\sin\theta \sin\phi}, \quad (3.6)$$

Figure 3.2 shows simulations of the D-PAD beam pattern at 1.0 and 1.5 GHz. The FWHM of the beam (through either $\phi = 0$ or $\theta = 0$) is $\sim 20.8^\circ$ at 1.0 GHz, and $\sim 13.9^\circ$ at 1.5 GHz. The first nulls appear at $\pm\lambda/D$; for our 4×4 grid $D=1.2$ m this corresponds to $\pm 14.3^\circ$ at 1.0 GHz, decreasing to $\pm 9.6^\circ$ at 1.5 GHz. Measurements of the beam pattern are detailed in Section 5.4.1.

Now that I have described the design specifications of the D-PAD project, I will introduce the analogue components of the system, beginning with the D-PAD antenna.

3.3 LOG PERIODIC DIPOLE ARRAY ANTENNAS

In sparse aperture arrays, antennas are spaced more than a half wavelength apart. This means their dimensions can be physically larger than that possible within a dense array, allowing the use of higher gain antenna structures. The D-PAD antenna is a dual polarization log-periodic dipole array (LPDA) high gain antenna, which has been designed specifically for the D-PAD system. Each antenna element consists of four separate LPDA ‘blades’, which are arranged into a

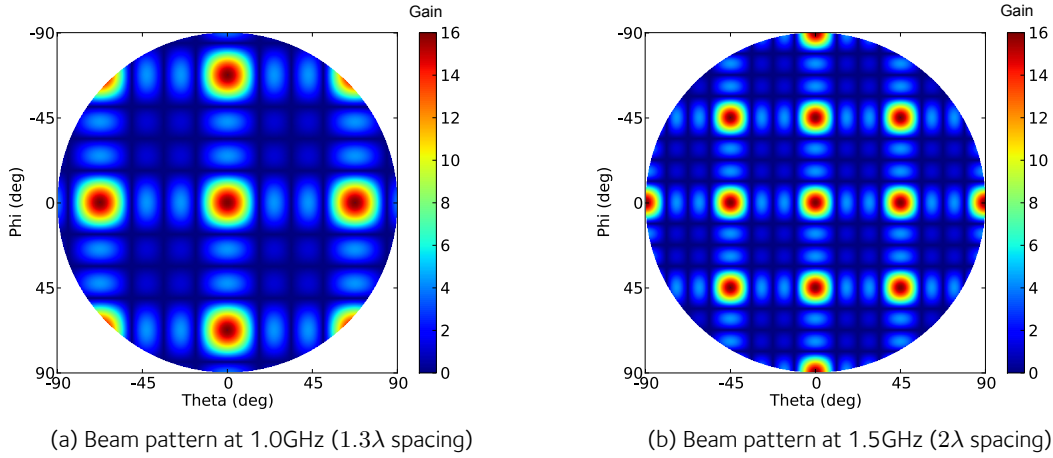


Figure 3.2: OSKAR simulations of D-PAD single station beam pattern, at 1.0 and 1.5 GHz. Isotropic antennas have been used in this simulation, so that the grating lobe structure is clear, and so the beam pattern is equivalent to the array factor.

pyramidal structure (to improve beam-pattern symmetry), as shown in Figure 3.3; each blade is a fully functional single polarization antenna in its own right. Arranging the blades into a pyramidal structure gives two perpendicular pairs of antenna blades, one for each polarization. The output of each blade pair is amplified then power combined, resulting in one output signal per polarization. Each pair of blades has a more symmetric beam pattern than that of a single blade.

As the name suggests, an LPDA is an array of dipoles, whose lengths, L_n and spacings D_n , are determined by a geometric series. The dipole elements of each LPDA blade are constrained by the ratio

$$\tau = \frac{D_{n+1}}{D_n} = \frac{L_{n+1}}{L_n} \quad (3.7)$$

In general, a value of $0.8 < \tau < 0.98$ is chosen for performance and size reasons (Straw, 2007). LPDA antennas perform well over a wide frequency range, making them desirable for wideband applications. LPDAs also provide directivity maximized in the direction of the array. This occurs as the parasitic dipole elements of the array are inductive at wavelengths longer than their resonant lengths, causing them to act as a reflector. Conversely, when capacitive (wavelength shorter than resonant length), they act as a director. Arranging the dipole elements in size order ensures that the radiation pattern is maximized in the desired direction (perpendicular to dipole elements, also known as end-fire). The dipole elements are connected together by a parallel plate waveguide, to which a 50Ω coaxial cable was soldered to feed the antenna.




Figure 3.3: CADFEKO models of the D-PAD antenna. a) The pyramidal structure is comprised of four separate antenna blades. b) Simulated far-field radiation pattern for a single antenna element. c) Simulated far-field radiation pattern for an angled pair – note the increase in beam symmetry and decrease in beam width.

3.3.1 DESIGN AND SIMULATION

The four-blade design of the D-PAD antenna is based upon work by Paul Grimes on dual polarization LPDAs covering a frequency range of 300–700 MHz (P. Grimes, personal communication). Each angled pair of blades is power combined to improve the symmetry of the radiation pattern, and the two pairs of blades are orthogonal to each other so the antenna is sensitive to polarization. The Allen Telescope Array feed employs a similar dual polarization design (Engargiola, 2002).

The D-PAD antenna blade is based on a design in which the antenna structure is printed onto a dielectric substrate (Pantoja, 1987, 1986). I generated the antenna blade geometry using the Antenna Magus software package², which implements the design procedures presented in Peixeiro (1988) and Straw (2007). The Antenna Magus model was then exported into FEKO³, where its performance was simulated and optimized. A target impedance of 50Ω was used. The simulated beam pattern for a single LPDA blade is shown in Figure 3.3b, and the impedance over the band is shown in Figure 3.4.

After optimizing the performance for a single LPDA blade, I simulated an angled pair of blades (see Figure 3.3c). This improves the directivity and symmetry of the beam pattern. Note that the second blade is a mirrored version of the first. If this is not done, a phase-inverting power com-

²<http://www.antennamagus.com/>

³<http://www.feko.info/>

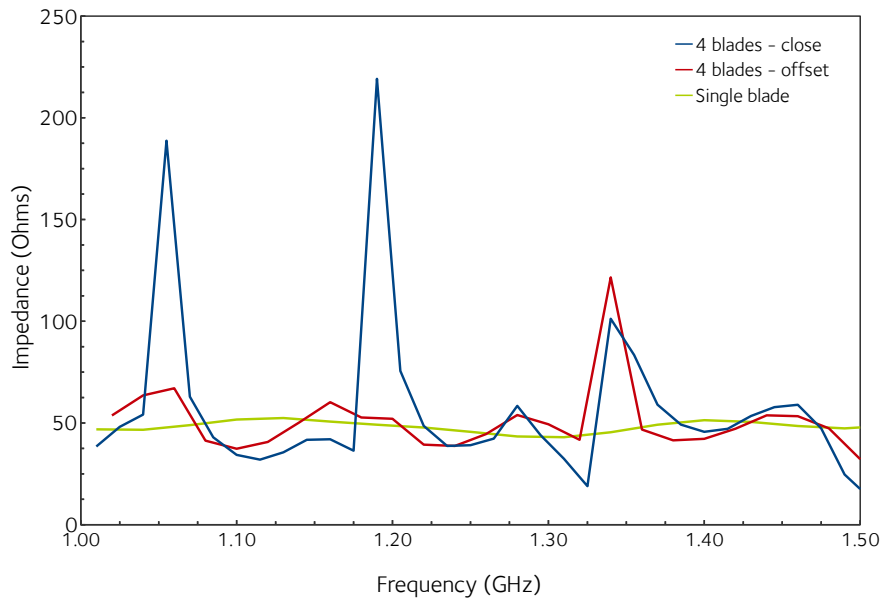


Figure 3.4: Simulated data showing the effect of fringe field coupling between orthogonal D-PAD blades. This coupling affects the impedance as seen at a blade's output terminal. At particular resonant frequencies, the impedance deviates significantly from the desired 50Ω . This decreases the gain of the antenna.

biner must be used to combine the two signals constructively. Such combiners are significantly more expensive than in-phase combiners.

To create a dual polarization antenna, I arranged four LPDAs into a pyramidal structure within FEKO (see Figure 3.3a). These simulations showed that fringe-fields of adjacent blades couple together, as is presented in Figure 3.4. This decreases the directivity of the antenna, which is undesirable. This issue was also found in the Allen Telescope Array feed. One solution to this problem is to simply move the blades further apart, so the fringe fields do not couple as strongly. However, doing so affects the radiation pattern. To find a compromise between the ideal pattern of Figure 3.3c and the level of fringe-field coupling, I altered both the angle of divergence from the pyramid's apex, and the separation between the top dipole elements. The optimal parameters were found experimentally, after fabrication of the antenna blades.

3.3.2 FEEDING TECHNIQUE

An attractive property of LPDA designs is that they do not require a balun. A balun (*balance-unbalance*) connects a balanced two-conductor line to an unbalanced line. Without a balun, cur-




Figure 3.5: Each of the four blades comprising the D-PAD LPDA is fed with a coaxial cable based infinite balun. In this figure, the antenna is connected to the receiver board (this is introduced in §3.4).

rent may flow down the outside conductor of the feed line, making it part of the radiating system.

A formal balun is not necessary for the LPDA design, as an ‘infinite’ balun may be used. In the infinite balun approach, the copper shielding of the coaxial cable is exposed and soldered directly to the feedline of the antenna. The inner conductor is then soldered onto the feedline of the other side of the antenna. For a TEM mode on a coaxial cable there is an inherent current balance: the net current on the inside conductor is equal to the net current on the outside conductor.

To feed the D-PAD LPDA, a coaxial cable (RG316) is soldered directly to the parallel plate waveguide feedline. The coaxial cable is terminated with a male SMA connector. Figure 3.5 shows the four coaxial cables of a D-PAD antenna mated to the receiver board of of Section 3.4.

3.3.3 FABRICATION

To fabricate the antenna, the geometry was exported from FEKO to file and a photolithography mask was produced. This mask was then sent to a circuit board manufacturer (ALR Services), who then printed the antenna onto 2.4mm thick FR-4 substrate.

In the second stage of fabrication, two pairs of antenna blades are glued together using epoxy resin. To do this, the antennas were arranged in the required pyramid structure on a specially designed rig, glued, then left to dry. A fully fabricated D-PAD antenna element is shown in Figure 3.6. In total, 20 antennas were made using this technique, of which 16 were selected for the




Figure 3.6: Fabricated D-PAD antenna element. The antenna is printed onto FR-4 substrate using photolithography.

D-PAD single station deployment detailed in Chapter 5.

3.3.4 GAIN AND BEAM PATTERN MEASUREMENTS

The radiation properties of the D-PAD antennas were measured in an anechoic chamber at the Oxford Engineering department. This chamber is a Faraday shielded room where every surface is lined with Eccosorb® radiation absorbent foam. The chamber simulates a reflectionless free space (less than -40 dB reflectivity at 1 GHz), suitable for testing antennas and electromagnetic devices without interference from reflections or external sources. The chamber's dimensions allow a pair of test antennas to be separated by up to a maximum of 4.0 m. The far-field distance (Fraunhofer zone) of an antenna is given by

$$R = \frac{2L^2}{\lambda}, \quad (3.8)$$

where L is the maximum dimension of the antenna. For the D-PAD antenna at 1.0 GHz, $R = 0.6$ m, so the chamber is suitable for far-field beam pattern measurements.

EXPERIMENTAL SETUP Figure 3.7 shows the experimental setup I used to measure the beam pattern and frequency response of the D-PAD antenna. One D-PAD antenna was placed on a mechanical turntable on one side of the anechoic chamber and another on a stationary platform on the other side. The two antennas were separated by a distance $d = 3.5$ m (tip to tip) and

aligned at boresight. The antenna on the turntable was then used as a transmitter (Tx), and the other antenna as a receiver (Rx). The Tx antenna transmitted a signal of known input power and frequency from port 1 of a VNA; this was picked up by the Rx antenna which was connected to Port 2 of the VNA. The measured S_{21} parameter represents the full signal transmission from port 1 to port 2.

BORESIGHT GAIN To find the frequency response of the antenna, one must first remove the contribution of the connective cables and free space path loss. The free space path loss (FSPL) is given by

$$FSPL = \left(\frac{4\pi d}{\lambda} \right)^2, \quad (3.9)$$

and the ratio of the received power P_{Rx} to transmitted power is then given by

$$\frac{P_{Rx}}{P_{Tx}} = \frac{G_{Tx} G_{Rx}}{FSPL}, \quad (3.10)$$

where G_{Tx} is the gain of the Tx antenna, and G_{Rx} is the gain of the Rx antenna (Kraus and Marhefka, 2002). In the setup detailed above, the two antennas are the same. By the antenna reciprocity theorem, $G_{Rx} = G_{Tx}$, so the gain of the D-PAD antenna is therefore given by

$$G_{ant} = \frac{4\pi d}{\lambda} \sqrt{\frac{P_{Rx}}{P_{Tx}}}. \quad (3.11)$$

Figure 3.8 shows the measured frequency response of the D-PAD antenna; that is, the gain as a function of frequency. The figure shows both the response of a single blade, and the response of the D-PAD antenna. To arrive at this result, the contribution from the cables and power splitters were measured using the VNA and then subtracted from the overall gain (in dB). The average (linear) gain over 1.0–1.5 GHz is 5.1 ± 0.5 , with a maximum of 6.2 occurring at 1.35 GHz, and a minimum of 2.5 occurring at 1.11 GHz⁴. The largest source of uncertainty in this experiment is due to the distance d , which may be measured from the tip, centre, or ends of the antenna, such that we have $\Delta d \approx 0.3$ m. This gives rise to a $\sim 17\%$ uncertainty for the antenna gain. Ideally, a larger anechoic chamber would be used, so that $d \gg \Delta d$. One could also compare against an antenna of known gain (this is known as the three-antenna method, see Kraus and Marhefka,

⁴here, the uncertainty (\pm) is standard deviation over the 1.0–1.5 GHz band.

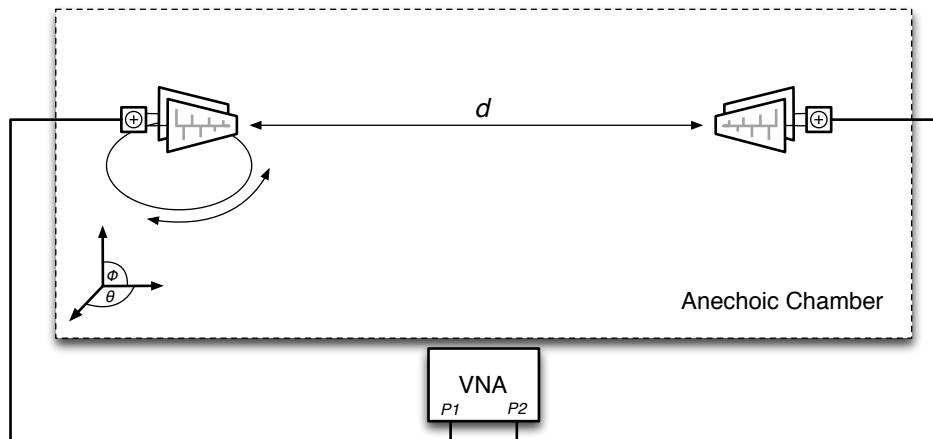


Figure 3.7: Experiment setup for D-PAD beam pattern measurements in anechoic chamber. A pair of D-PAD antennas were placed in the anechoic chamber, separated by a distance $d=3.5$ m. One is used as a transmitter (Tx), the other as a receiver (Rx). The Tx antenna was placed on a turntable which can be scanned over 360 degrees.

2002).

In comparison to the simulated response (dashed line), the antenna exhibits a number of dips in frequency performance. This is likely due to impedance mismatch between the antenna's spine, and the soldered coaxial cable which acts as an infinite balun. The dips are exacerbated when the blades are arranged into the final pyramidal structure, as was seen as fringe field coupling in simulation.

At higher frequencies, the measured response is lower than the simulated. This is likely due to the simulation neglecting the dielectric loss of the FR-4 material onto which the antenna structure is printed.

BEAM PATTERN Armed with this boresight gain value, the Tx antenna was scanned using the mechanical turntable and off-axis power measurements were taken at regular 5 degrees intervals of over a full 360 degrees. By doing so, one can find the antenna gain as a function of pointing direction, $G(\theta, 0)$. By rotating both the Rx and Tx antenna 90 degrees about boresight and repeating the process, the beam pattern can be sampled over $G(0, \phi)$. In a similar fashion, by changing the pointing direction on one axis and scanning about the other, one can build up a full 3D beam pattern of the antenna. I have chosen a coordinate system where θ and ϕ are in the E- and H- planes of the antenna's x-polarization, respectively. The directivity can be computed

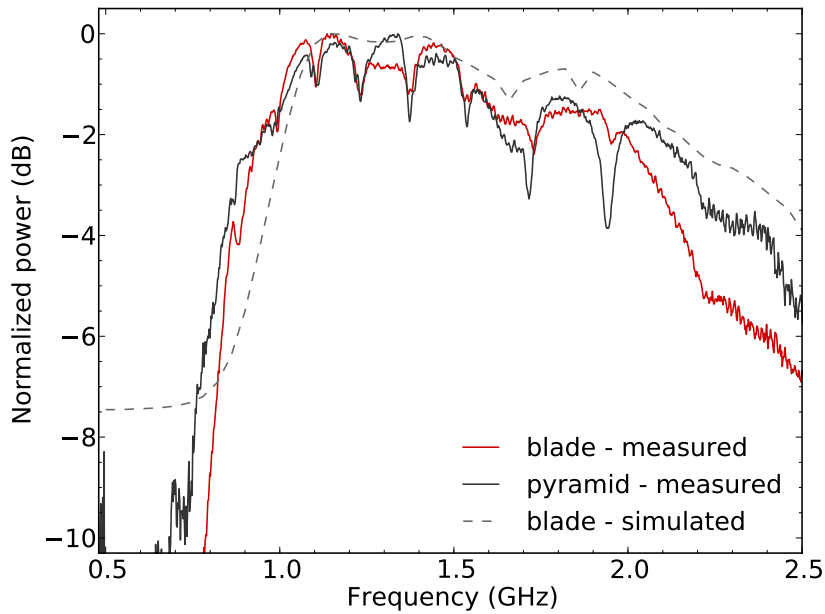


Figure 3.8: Frequency performance of the D-PAD LPDA (black), compared to that of a single blade (red), and simulated performance (dotted grey).

from the ratio of maximum to average power of the normalized pattern $P_n(\theta, \phi)$:

$$D = \frac{P_n(\theta, \phi)_{max}}{P_n(\theta, \phi)_{avg}}. \quad (3.12)$$

Note that this must be evaluated over all (θ, ϕ) . To do so, I have computed the average of the E- and H- planes and have assumed the pattern is circularly symmetric. Figure 3.10 shows the directive gain for the D-PAD antenna at 1.0, 1.25 and 1.5 GHz (in decibels above isotropic, dBi). The average directivity over the band is 6.14 ± 0.6 .

CROSS POLARIZATION Figure 3.9 shows the cross polarization performance of the D-PAD antenna. Here, we have used the Ludwig-3 definition (Ludwig, 1973). This was measured by rotating the Tx antenna around boresight, and measuring power output at the Rx antenna. The cross polarization is given by the ratio of Rx power at 90 degree rotation to the Rx power without rotation. Minimum values of cross polarization correspond to maximal fringe coupling, and local minima of the gain in Figure 3.8.




Figure 3.9: Cross polarization performance of the D-PAD antenna, as measured in the anechoic chamber.

3.4 RECEIVER BOARD

To maintain a low system temperature, the signal from the antenna must be amplified at the earliest stage possible (before noise is introduced). Ideally, a differential amplifier would be connected at the feed point of the D-PAD antenna. Unfortunately, there is no commercially available differential amplifier over the D-PAD operational range with a suitably low noise figure. As such, the first component in the signal chain is a low noise amplifier board, which connects directly to the infinite balun of the D-PAD antenna. This is known as the D-PAD receiver board.

The receiver board was designed specifically for the D-PAD antenna. To connect to the antenna blades, the board has four SMA inputs. The input for each blade is amplified, and then the polarization pairs of blades are power combined. After this, the signals are amplified further and then routed to SMA outputs. The board is mounted directly to the bottom of the D-PAD antenna on a steel plate, see Figure 3.11. The design and implementation is discussed in more detail below.

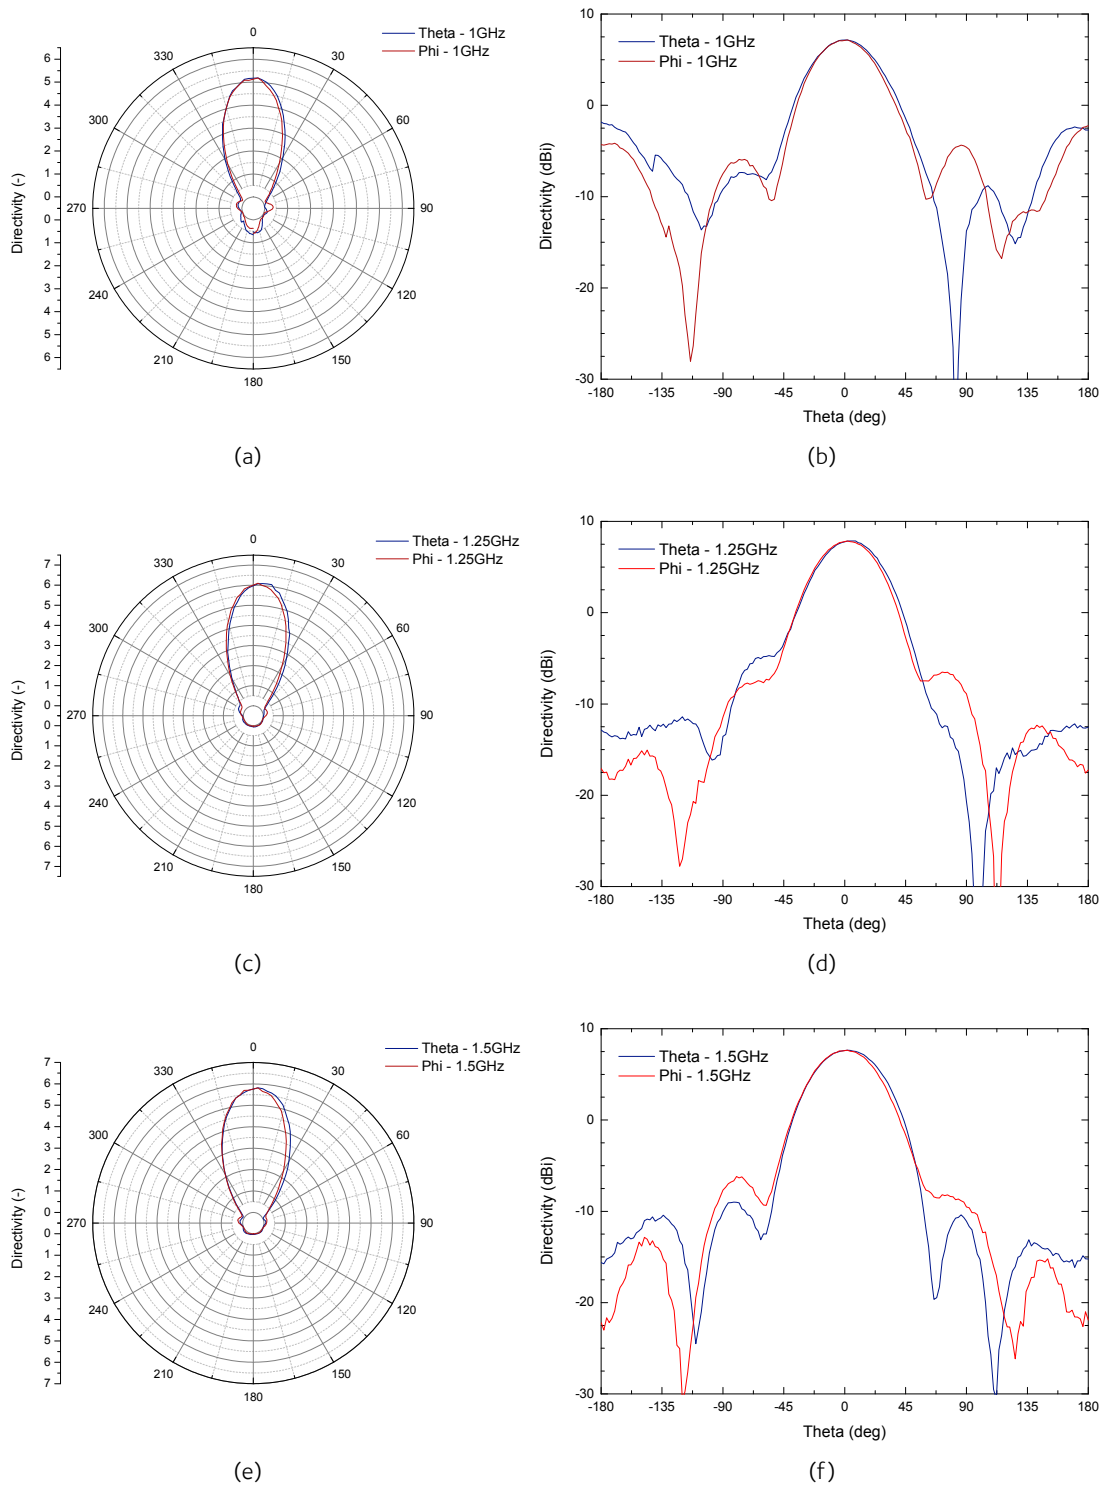


Figure 3.10: Directivity of the D-PAD LPDA, shown on the left in polar-linear, and regular-dB on the right. a) and b): 1GHz; c) and d): 1.25GHz; and e) and f): 1.5GHz.




Figure 3.11: D-PAD receiver board. Input pairs (X0, X1) and (Y0, Y1) are amplified, power combined, then further amplified, before being output on X2, and Y2 respectively.


3.4.1 AMPLIFIER SELECTION AND CHARACTERIZATION

For the first stage of amplification, the Avago MGA-633P8 low noise amplifier (LNA) was selected (Avago, 2010b). The MGA-633P8 is a Gallium Arsenide (GaAs), pseudomorphic high electron mobility transistor (pHEMT). At 900 MHz, it is specified of having 18 dB of gain and a 0.37 dB noise figure, which corresponds to a noise temperature of 25 K. A second amplification stage is implemented using the Avago MGA-30889 LNA. This is a similar GaAs pHEMT amplifier which supplies 15.5 dB gain at 900 MHz but has a higher noise figure of 1.9 dB (Avago, 2010a). While the noise figure is higher, the chip has a far flatter gain slope, making it more suitable for the second amplification stage where noise figure is less critical.

GAIN CHARACTERISTICS To characterize the MGA-633P8 LNA, an evaluation board consisting of the LNA, bias circuitry and SMA input/output connectors was purchased. Its frequency response was then measured using an Anritsu 37369C Vector Network Analyser (VNA). I have reported here the S_{21} (forward gain from input to output port), and S_{11} (reflected power) parameters. The response of the LNA evaluation board is shown in Figure 3.12a. As can be seen, the amplifier supplies around 18 dB at the lower limit of the D-PAD band (1.0 GHz), which decreases smoothly to around 13.5 dB at the upper limit of 1.5 GHz.



(a) S_{11} (reflection) and S_{21} (forward gain) of the Avago MGA-633P8 LNA development board.



(b) S_{11} (reflection) and S_{21} (forward gain) of the D-PAD receiver board. The S_{11} coefficient is higher (worse) than that measured on the LNA development board.

Figure 3.12: Transmission and reflection characteristics of the LNA development board and D-PAD receiver.

NOISE TEMPERATURE MEASUREMENT To ensure that the manufacturer's noise figure claim was realistic, I measured the noise temperature of the Avago LNA using a technique known as the Y-factor method. In this technique, two matched loads at temperatures T_H (hot load) and T_C (cold load) are connected to the device under test (DUT), and the output power is measured for each case. The output power consists of noise power generated by both the DUT and the source resistor, so at two different temperatures the two power measurements are

$$N_H = GkT_H\Delta\nu + GkT_{DUT}\Delta\nu, \quad (3.13)$$

$$N_C = GkT_C\Delta\nu + GkT_{DUT}\Delta\nu. \quad (3.14)$$

In these two equations, T_{DUT} is the noise temperature of the DUT and $G\Delta\nu$ is the the gain-bandwidth product of the amplifier. The ratio of these two equations is known as the Y-factor:

$$Y = \frac{N_H}{N_C} = \frac{T_H + T_{DUT}}{T_C + T_{DUT}} > 1 \quad (3.15)$$

The noise temperature of the DUT is then given by

$$T_{DUT} = \frac{T_H - YT_C}{Y - 1}. \quad (3.16)$$

To apply this method, I connected a 50Ω terminated load on a semi-rigid steel cable to the input of the LNA evaluation board. In this case, the DUT is the semi-rigid cable together with the LNA board. The noise contribution from the amplifier is then extracted by applying the noise temperature cascade rule; that is:

$$T_{DUT} = T_{cable} + \frac{T_{LNA}}{G_{cable}} \quad (3.17)$$

To measure the noise power, the output of the board was connected to a TTi PSA2701T spectrum analyser. The hot load power measurement N_H was taken at room temperature, $T_H = 295 \pm 1\text{K}$. A cold load measurement was taken by placing the terminator in liquid nitrogen, $T_C = 77\text{K}$. Applying the Y-factor method and solving for T_{LNA} , we found a noise temperature of $T_{LNA} = 35 \pm 20\text{K}$. This is consistent with the manufacturer's claim of 0.37 dB, or 25 K.

3.4.2 CIRCUIT BOARD DESIGN

The circuit board layout was designed by Hamid Dabiri-Khah Oxford Physics Central Electronics. It is a dual-substrate board, 50×75 mm in size, consisting of a low loss substrate (Rogers 4350B), affixed to 1 mm thick FR-4. The low loss substrate is required as any noise introduced before the first stage of amplification is detrimental to the system temperature; the FR-4 adds thickness and rigidity. A component listing and schematic diagrams can be found in Appendix A. To connect to the D-PAD antenna, the receiver board has four SMA inputs. A 50Ω impedance coplanar waveguide connects each input to an Avago MGA-633P8 LNA, which applies the first stage of amplification. The path length between the SMA input and the LNA has been made as short as possible to minimize ohmic losses which increase noise temperature. After amplification, the signals from matching antenna blades are power combined using a Mini-Circuits GP2S+ surface mount chip, before a secondary amplification stage using the MGA-30889 LNA. The amplified signal is then output via two SMA connectors, one for each polarization.

After the design was finalized, 20 copies of the printed circuit board were fabricated. Figure 3.11 shows a photograph of a receiver board mounted to the steel plate of the D-PAD antenna. The soldered pins of the four input SMAs are visible at the top of the photo; these SMA connectors stick through mounting holes drilled into the steel plating. The two SMA outputs are visible at the bottom edge of the board, along with the DC power supply feedthrough.

3.4.3 PERFORMANCE

The receiver can be considered a single 6-port network; or, if one assumes the two polarizations are perfectly decoupled, two 3-port networks. As vector network analysers are only 2-port, measurements were made by terminating all but two ports. Figure 3.12b shows the typical S_{11} and S_{21} as measured for an input-output pair of a single polarization. The forward gain (S_{21}) of the system is seen to be between about 30 dB at 1.0 GHz, smoothly decreasing to ~ 26 dB at 1.5 GHz.

For the evaluation board, the S_{11} is consistently under -15 dB, meaning no more than 2% of an incident voltage wave is reflected. In comparison, the receiver board has a higher S_{11} over the operational band, peaking at about -12 dB, which corresponds to about 6% reflection. Correspondingly, the mismatch loss of receiver board is higher than that of the evaluation board.




Figure 3.13: S_{11} (reflection) and S_{21} (forward gain) of all 16 channels of the MiniCircuits ZC16PD-252-S+ power combiner. Each channel was measured with all SMA inputs terminated.

Mismatch loss is defined as


$$ML_{dB} = 10 \text{Log}_{10}(1 - \Gamma^2), \quad (3.18)$$

where $\Gamma = S_{11}$ is the reflection coefficient. If $S_{11} = -15$ dB, the mismatch loss is -0.14 dB; for $S_{11} = -12$ dB this increases to -0.28 dB.


It follows that the noise temperature of the receiver board will be higher than that of the evaluation board. As the board is a 3-port network, one cannot apply standard (i.e. 2-port) techniques to measure the noise temperature — there is no IEEE definition of noise temperature in multi-port systems. Instead, we may treat this mismatch loss as a passive noise source that is present in the receiver board, but not in the evaluation board. In Section 3.7.2 below, overall receiver temperature is computed, treating this mismatch loss as a passive source of noise.

3.5 ANALOGUE BEAMFORMER

Aperture arrays gain much of their power and flexibility through their beamformer. The ability to form multiple beams on the sky can improve survey speeds by orders of magnitude over single beam implementations. Of course, this requires enough signal processing backend to process



(a) Bandpass filter PCB, without SMA connections and folded copper casing



(b) Bandpass filter with copper casing.

Figure 3.14: Photos of D-PAD bandpass filter, with (b) and without (a) copper casing.

each beam. For initial testing, a power combiner and a set of equal length cables was chosen to phase the D-PAD station to zenith. This is arguably the simplest beamformer possible, so is ideal for our preliminary studies.

To combine the signals, a Mini-Circuits ZC16PD-252-S+ 16-way power splitter/combiner is used. This power combiner is quoted as having 0.3 dB typical amplitude unbalance, and 5 degree typical phase unbalance. VNA measurements show that this is indeed the case, however there is a significant gain ripple across the operational band; see Figure 3.13. For his astrophysics summer project, Andrew Hickling designed an alternative 16-input, 3-stage Wilkinson power divider with better phase and amplitude performance (Hickling, 2010). This board would be an excellent replacement for the Mini-Circuits combiner.

3.6 ANALOGUE CHAIN

The D-PAD analogue chain provides filtering, further amplification, and signal transport from the field to the bunker. In this section I introduce the remaining components which make up the analogue chain.

3.6.1 BANDPASS FILTER

In the D-PAD single station, the digitizer card is run at a sampling frequency of 1 Gsample/s, which gives a bandwidth of 0.5 GHz. The 1.0-1.5 GHz band is directly sampled; this falls in the third Nyquist zone of the digitizer. The D-PAD bandpass filter prevents frequencies outside this


range from aliasing in.

While the topic of microwave filter design is extensive, there are several software packages which expedite the process. A good example of this is Ansoft Designer⁵ 3.5, which contains a microwave filter design tool. The theory and design equations implemented by Ansoft Designer are discussed in texts such as Hong and Lancaster (2001) and Matthaei et al. (1980). In a microstrip filter, all required capacitance and inductance are generated through changes in microstrip thickness and structural discontinuities. The edge coupled filter consists of a series of half wavelength resonators which are positioned with adjacent resonators parallel to each other along half of their length. Due to their proximity to one another, adjacent resonators are capacitively coupled to one another.


I used the Ansoft Designer filter design tool to design and simulate a 7-pole edge coupled microstrip 1.0-1.5 GHz bandpass filter. A photomask was then printed and photolithographic techniques were used to etch the filter geometry onto 1.6 mm thick FR-4 substrate. This photolithography was done by Paul Pattison in the Oxford Central Electronics Workshop. FR-4 was chosen as it is significantly cheaper than most other substrates. While the loss tangent of FR-4 is higher than other substrates, the signal is amplified significantly before filtering, so the contribution to the receiver temperature is negligible. After etching, SMA ends were soldered on and the substrate was enclosed in a folded copper case. Photos of the filter with and without its case are shown in Figure 3.14. The filter's bandpass, as measured on the VNA, is shown in Figure 3.15a.

3.6.2 AMPLIFIER

The receiver board does not supply enough amplification to reach the levels required by the digitizer card, so further amplification is required. This amplification is supplied with MiniCircuits ZRL-2150+ amplifiers. The amplifier supplies ~ 25 dB of gain over the D-PAD operational band, and has a quoted noise figure of 1.4 dB. The measured response of this amplifier as measured with the VNA is shown in Figure 3.15b.




(a) S11 (reflection) and S21 (forward gain) of a typical D-PAD bandpass filter.




(b) S11 (reflection) and S21 (forward gain) of a typical MiniCircuits ZRL-2150+ amplifier.

Figure 3.15: VNA measurements of D-PAD analogue hardware performance.



(a) S11 (reflection) and S21 (forward gain) of a typical D-PAD CXA type cable (RG223, SMA terminated).



(b) S11 (reflection) and S21 (forward gain) of a typical D-PAD CXB type cable (RG316, TNC terminated). Note TNC to SMA adaptors are connected in this measurement.

Figure 3.16: VNA measurements of D-PAD coaxial cable characteristics.

3.6.3 TRANSMISSION LINES

To connect the D-PAD receiver boards to the beamformer, a set of matching SMA cables were ordered from Gigatronix. These cables are 50Ω , SMA terminated RG 223, specified to 1500 (+2.5/-0) mm in length. To transport the signal from the rooftop to the office, a set of two 50Ω , TNC terminated 2000 mm (20 metre) RG213 cables were also used. SMA plug to TNC adaptors are attached to either end of this cable. The VNA measured frequency responses of these cables are shown in Figure 3.16.

3.7 SYSTEM PERFORMANCE

Although the performance of each component was tested discretely, the overall performance of the system is of most concern. The order in which the components are connected affects the receiver temperature, and as shown in Chapter 2, interactions between components must also be considered. In this section, I present the overall system performance for the D-PAD system.


Every component in the system was labelled with a unique identifier of three letters followed by three digits (e.g. REC001). The letters form an abbreviated name, for example CXA for coaxial cable type A. In the sections that follow, I use the following three letter abbreviations for subscripts: ANT for antenna; REC for receiver; CXA for coaxial cable type A; GMA for the first gain amplifier; BPA for the first bandpass filter; GMB for the second gain amplifier; CXB for coaxial cable type B; BPB for the second bandpass filter; and, CXC for coaxial cable type C. The order in which they are connected is given in Table 3.1.

3.7.1 SYSTEM BANDPASS

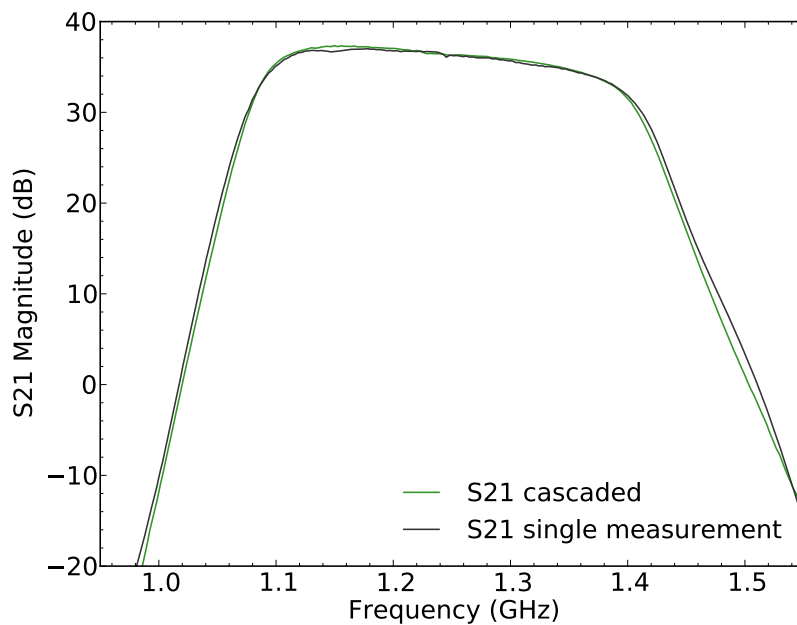
Figure 3.17a shows the bandpass of the analogue chain for a single polarization (from BPA to CXC). This was measured by cascade of each component, and also as one DUT (see Chapter 2). The maximum gain over the band is 37 dB. The 3 dB bandwidth is 1.09–1.37 GHz, while the 10 dB bandwidth is 1.07–1.43 GHz.

To stop spurious signals aliasing in from other Nyquist zones, the signal must be attenuated outside the 1.0–1.5 GHz band. As can be seen in Figure 3.17b, the bandpass response at 1.0 and

⁵<http://www.ansys.com>



(a) D-PAD bandpass, over 0.5-3.0 GHz.



(b) Zoomed-in version of D-PAD bandpass, from 1.0-1.5 GHz.

Figure 3.17: Bandpass of the D-PAD analogue chain, from BPA to CXC. This does not include the gain contribution from the receiver board or beamformer.

Table 3.1: Cascaded receiver temperature and gain totals, using a mismatch loss of -0.14 dB.

		Discrete		Cascaded / Total	
Component		Noise Figure (dB)	Gain (dB)	Gain (dB)	Noise Temp (K)
Rec board	Mismatch loss	0.14	-0.14	-0.14	9.50
	LNA A	0.37	18	17.86	36.14
	Combiner	0.9	(2.1)	19.96	37.23
	LNA B	1.9	15.5	35.46	38.83
Rooftop	Coax A	0.85	-0.85	34.61	38.85
	Beamformer	3.2	(8.8)	43.41	38.96
	Filter A	2.2	-2.2	41.21	38.97
	Gain Block A	1.4	25	66.21	38.98
	Coax B	5.3	-5.3	60.91	38.98
Office	Gain Block B	1.4	25	85.91	38.98
	Filter B	2.2	-2.2	83.71	38.98
	Coax C	0.85	-0.85	82.86	38.98
	Attenuator	20	-20	62.86	38.98

1.5 GHz is -48 and -35 dB, respectively, which is sufficient to suppress out-of-band sources. Over the GSM-900 mobile phone band (0.89-0.96 GHz), which is a particularly strong source of RFI, the signal is suppressed by over 80 dB.

3.7.2 RECEIVER TEMPERATURE

The receiver temperature is a measure of the overall noise performance of a receiving system. For the D-PAD single station, the receiving system consists not only of the receiver board, but of every component up until the digitizer. Each of these components adds a certain amount of noise, increasing the receiver temperature.

If one knows the noise figure (or noise temperature) of the individual stages, the noise figure of a cascade connection of stages can be determined. The noise characteristics of a cascaded system are dominated by the characteristics of the first stage, since the effect of the second stage is reduced by the gain of the first. We can generalize this result to an arbitrary number of stages, such that

$$T_{cas} = T_{e1} + \frac{T_{e2}}{G_1} + \frac{T_{e3}}{G_1 G_2} + \dots, \quad (3.19)$$

this is known as the Friss formula (Chapter 10 of Pozar, 2005). The contribution of each component — as given by Equation 3.19 — is shown in Table 3.1. The table shows that the noise temperature is dominated by the receiver board's mismatch loss and the noise temperature of

first LNA of the receiver board.

The table gives only a single representative value for the receiver temperature; in reality this will vary across the frequency band. The receiver board has an S_{11} under -15 dB over 1.0–1.3 GHz, but this then rises to -12 dB at 1.5 GHz. This doubles the mismatch loss from -0.14 to -0.28 dB, which in turn increases the receiver temperature from 39 K to 50 K. Taking this into account, the overall receiver temperature is thus 44 ± 6 K.

3.7.3 OVERALL SYSTEM TEMPERATURE

The overall system temperature T_{sys} depends on the sky, the antenna, and the receiver. The system temperature may be calculated by

$$T_{sys} = T_{ant} + T_{AP} \left(\frac{1}{\varepsilon_1} - 1 \right) + T_{LP} \left(\frac{1}{\varepsilon_2} - 1 \right) + \frac{1}{\varepsilon_2} T_{rec}. \quad (3.20)$$

where ε_1 is the antenna's thermal efficiency, which is close to unity; T_{AP} is the antenna's physical temperature; T_{LP} is the physical temperature of the transmission line connecting the antenna to the receiver; ε_2 is the line efficiency; and, T_{rec} is the receiver temperature (§12-2 of Kraus and Marhefka, 2002). The antenna temperature, T_{ant} , is formally defined as

$$T_{ant} = \frac{1}{4\pi} \int_0^{2\pi} \int_0^\pi P_n(\theta, \phi) T_b(\theta, \phi) \sin\theta d\theta d\phi, \quad (3.21)$$

that is, T_{ant} is the average sky brightness as weighted by the antenna's normalized power pattern. If $T_b(\theta, \phi) = T_{sky}$ is constant over the entire sky, the integral reduces to $T_{ant} = T_{sky}$. The average sky temperature over the 1.0–1.5 GHz frequency band is between 3–6 K (§12-1 of Kraus and Marhefka, 2002).

For D-PAD, the transmission line is the coaxial cable soldered onto the antenna. VNA measurements showed a 50cm length of this cable with female SMA connectors has less than 0.3 dB of loss; that is, no more than 6.7% of power is lost through ohmic losses. Anechoic chamber measurements show that over the operational band, $6.1 \pm 3.2\%$ of incident power on the transmission line's input port is reflected (i.e. the antenna's S_{11} is under -10 dB). The line efficiency is therefore $\varepsilon_2 = 0.87 \pm 0.03$. Plugging in all these measured values and using a temperature $T_{AP} = T_{LP} = 290$ K, the system temperature evaluates to $T_{sys} = 102 \pm 12$ K.

3.8 DISCUSSION

For aperture arrays to be a competitive technology for BAO detection the cost per antenna must be extremely low, as many antennas would be required to reach a suitable sensitivity. The amplifiers, filters and transmission lines associated with each antenna element must also be low cost, as must the beamformer. D-PAD is a first step to show that this is possible. Neglecting engineering time, the total cost of the analogue components (including platform and radome) of the D-PAD single station is approximately £4000, which equates to under £250 per antenna. The cost would be significantly lower if components were ordered in large quantities and mass produced.

Sparse arrays gain their flexibility and survey speed from their beamformer. Digital beamformers are particularly flexible and well suited to forming multiple beams (Armstrong, 2012), but for very large arrays they are currently prohibitively expensive. In D-PAD, the cost of the digitizer is about £1000, and the digital signal processing hardware is about £2000 (see Chapter 4). If this hardware was used to digitize each antenna, 16 digitizers and 8 signal processing boards would be required, at a combined cost of £32,000, or £2000 per antenna. As such, one would expect that the cost per unit of the digital systems is likely to dominate hardware costs in a sparse array. Driving the per antenna digitization cost down will be an important task toward building the SKA-low and SKA-dense aperture arrays.

Many of the proposed BAO experiments cite a target system temperature of 50 K (Section 1.6.6). Although D-PAD has a system temperature of 102 ± 12 K, a T_{sys} of 50 K is not an unattainable target. Some possible optimizations to reduce the T_{sys} are discussed in further detail below.

3.8.1 FUTURE WORK

There are a few areas where the performance of the D-PAD system could be further optimized. The two most notable issues are the coupling between fringe fields of the antenna, and the losses due to the infinite balun. Mitigating these issues would boost system performance by lowering the system temperature, and increasing the aperture efficiency (which in turn increases effective area). Should a differential LNA become available, it may be possible to integrate the LNA directly onto the antenna blade, removing balun and connector losses.

A potential approach to reducing fringe field coupling is to interleave the LPDA blades, as is done by de Lera Acedo (2011) in their SKA-LA antenna design. However, this would preclude the use of LPDA designs printed onto substrate, which would likely increase cost. One could also investigate alternative dual polarization antenna structures.

Another area of research is the design of a more flexible analogue beamformer, as an alternative to the Mini-Circuits combiner. MSc student Andrew Hickling has designed a 16-way Wilkinson power divider based beamformer, which has a flatter gain slope and less phase imbalance between inputs. An extension to this would be an analogue beamformer capable of forming multiple beams on the sky. For example, an Butler matrix beamformer may be designed (§1.3.2 of Mailloux, 2005). Alternatively, beamforming could be performed in the digital domain, although this is likely to be orders of magnitude more expensive.

3.9 CONCLUSIONS

In this chapter, I have reported on the design of a sparse aperture array demonstrator called D-PAD which operates over 1.0–1.5 GHz. The D-PAD hardware is inexpensive and shows promising performance, which could be improved iteratively if a sparse array for BAO detection was to be constructed. Implementing an antenna with improved fringe-field coupling and a low noise differential LNA would improve performance by decreasing system temperature and increasing effective collecting area. I suggest that these areas warrant further research if a 1.0–1.5 GHz sparse array is to be built for radio astronomy.

In Chapter 5, I report on the deployment of the D-PAD system, in which the characteristics of the array as a whole were measured. In particular, the beam pattern, system temperature, and bandpass shape were characterized. A discussion of the viability of a sparse array based instrument for BAO detection is then given in Chapter 7. In the next chapter, I present a wide-bandwidth digital spectrometer suitable for use in redshifted hydrogen spectroscopy.

CHAPTER 4

WIDE-BANDWIDTH DIGITAL SPECTROMETERS

4.1 INTRODUCTION

The radio spectrum spans over six orders of magnitude: from wavelengths of tens of metres (\sim MHz) to sub-millimetre wavelengths (\sim THz). The amount of this bandwidth that can be measured by a telescope is limited by instrumental factors. The first limit arises from the frequency response of the feed and receiver, whose responses are only optimal over a portion of the spectrum. Many parabolic dish based telescopes combat this by equipping a suite of feeds and receivers, each optimized for a certain frequency range. For example, the Parkes 64 metre telescope has over 14 receivers spanning 0.44–24 GHz. Similarly, the Square Kilometre Array will consist of three different telescope technologies, all of which are specially suited to different parts of the radio spectrum.

The second limit is imposed by the hardware used to process the signal. In most modern radio telescopes, signals are digitized and processed in the digital domain. The bandwidth of the analogue to digital converter, data transport bandwidths, and the processing power of the hardware, all limit the instantaneous bandwidth achievable.

In this chapter, I discuss the design of a wide-bandwidth spectrometer specifically suited for HI detection. This is implemented on the Parkes 64 metre single dish telescope (Chapter 6), and on the D-PAD sparse aperture array (Chapter 5). HI-line observations require fine frequency resolution so that spectral features can be discerned, and enough bandwidth so one can observe the HI line over a large range of different redshifts. The spectrometer design presented here can

process up to 512 MHz of instantaneous bandwidth. This bandwidth allows the spectrometer to cover a large redshift range (e.g. $0 \leq z \leq 0.5$), as will be required for BAO detection through HI observations.

4.2 DIGITAL SPECTROMETER IMPLEMENTATIONS

In radio astronomy, a spectrometer is a device that computes the power spectral density of a signal. The power spectral density of a waveform is related to its autocorrelation function by the Wiener-Khinchin theorem. The theorem states that the relationship between a stationary ergodic signal $x(t)$, its power spectrum $S_{xx}(\nu)$, and its autocorrelation $r_{xx}(\tau)$ is given by

$$S_{xx}(\nu) = \int_{-\infty}^{\infty} r_{xx}(\tau) e^{-2\pi i \nu \tau} d\tau. \quad (4.1)$$

That is, the autocorrelation

$$r_{xx}(\tau) = \langle x(t)x(t - \tau) \rangle \quad (4.2)$$

is related to power spectral density by a Fourier transform. In the discrete case, the relationship becomes

$$S_{xx}(k) = \sum_{k=-\infty}^{\infty} \langle x(n)x(n - k) \rangle e^{-2\pi i k \tau}, \quad (4.3)$$

which is a discrete convolution. It follows from the convolution theorem that

$$S_{xx}(k) = \langle |X(k)|^2 \rangle, \quad (4.4)$$

where $X(k)$ denotes the Discrete Fourier Transform (DFT) of $x(t)$:

$$X(k) = \sum_{n=0}^N x(n) e^{-2\pi i n k / N} \quad (4.5)$$

There are two distinct classes of digital spectrometers: ones that approximate $S_{xx}(k)$ through numerical approximation to Equation 4.3, and those that evaluate Equation 4.4. The first class are called lag-autocorrelators, or auto-correlation spectrometers (ACS). The second class are known as a discrete Fourier transform filterbanks (FTF).

ACS and FTF architectures differ in both their spectral response and the number of computations required to implement them. For regularly sampled data, the Fast Fourier Transform (FFT) algorithm may be used to compute the DFT (see, for example, Brigham, 1988), which reduces the number of computations required from $O(N^2)$ for to $O(N\log(N))$. In Chapter 4 of Taylor et al. (1999), it is shown that the ratio of multiplies for the two architectures is

$$R_{\frac{ACS}{FTF}} = \frac{n_t}{2\log_2(n_t)}, \quad (4.6)$$

where n_t is the number of samples per FFT or lag correlation. So in general, FTF architectures require fewer computations than their equivalent ACS counterpart. Romney also shows that the spectral response of FTF and ACS architectures differ, with ACS architectures having a *sinc* response, and FTF architectures possessing a *sinc*² response. The result is that interchannel isolation is better in FTF architectures (first sidelobe at ~ -13.6 dB), than in ACS architectures (~ -6.8 dB).

4.2.1 WINDOWING FUNCTIONS

The spectral response of an FTF spectrometer can be improved by windowing the data before performing the DFT; in ACS architectures a windowing function may be applied after the auto-correlation step. An in-depth discussion of windowing functions is given in Gade and Herlufsen (1987).

For filterbanks, the Hamming and Hann windows are commonly applied windowing functions. The Hamming window has coefficients given by

$$w(n) = 0.54 - 0.46\cos\left(\frac{2\pi n}{N-1}\right), \quad (4.7)$$

which are optimized to minimize the level of the first sidelobe. These coefficients are similar to those of the Hann window¹:

$$w(n) = 0.5\left(1 - \cos\left(\frac{2\pi n}{N-1}\right)\right), \quad (4.8)$$

¹Named after Julius von Hann, but often erroneously referred to as 'Hanning'.

which offers faster sidelobe rolloff, but a higher first-sidelobe level. These should be compared to the Dolph–Chebyshev window

$$w(n) = \frac{\cos\{N\cos^{-1}[\beta\cos(\frac{\pi k}{N})]\}}{\cosh[N\cosh^{-1}(\beta)]}, \quad (4.9)$$

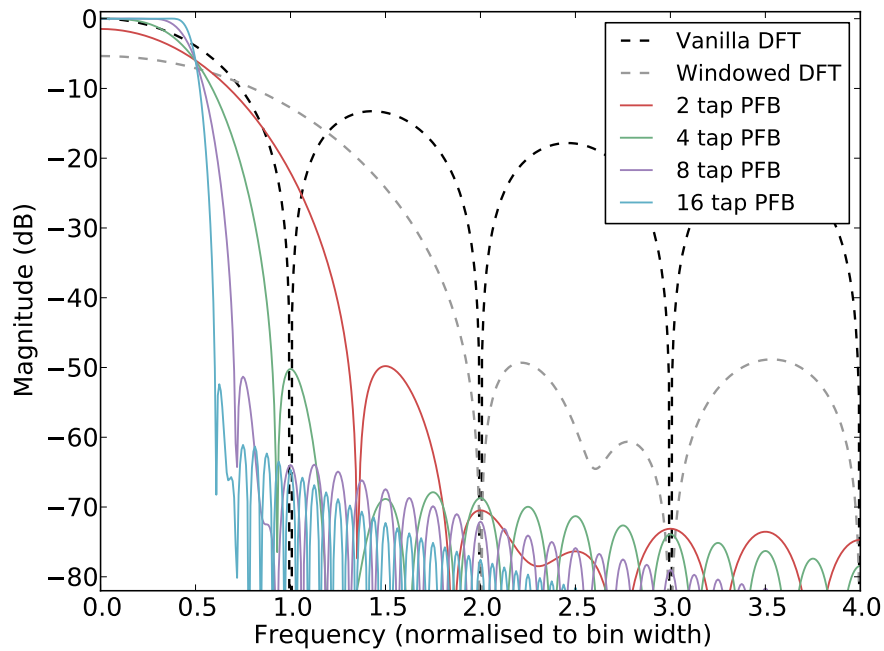
where $\beta = \cosh(N^{-1}\cosh^{-1}(10^\alpha))$, which sets the average sidelobe level to be -20α dB, but offers no sidelobe rolloff.

A more dramatic improvement may be achieved by using an lowpass filter frontend before the DFT (Bellanger et al., 1976), to form what is known as a polyphase filterbank (PFB). This approach is detailed further in the following section. Figure 4.1a compares the spectral response of a standard DFT ('vanilla'), Hamming windowed DFT, and polyphase filterbank, for a single channel. An example showing a 16-channel, 4-tap Hamming window-based polyphase filterbank is shown in Figure 4.1b.


4.2.2 SPECTROMETERS FOR RADIO ASTRONOMY

The first digital spectrometer used for radio astronomy was developed by Weinreb (1963). This 1-bit ACS was used to observe the 18-cm wavelength hydroxyl (OH) absorption line in the spectrum of Cassiopeia A, providing the first evidence of OH in the interstellar medium (Weinreb et al., 1963). The first reference to FTF spectrometers for radio astronomy can be found in Chikada et al. (1987); however FTF spectrometers did not enjoy widespread adoption until much later. This is due to implementation issues, such as increased data output rates as compared to ACS implementations; see Bunton (2000), for a discussion of these issues.

In comparison to both ACS and FTF architectures, polyphase filterbanks offer vastly improved interchannel isolation. For radio astronomy purposes, high interchannel isolation is important so that spectral features are not smeared out, and so the spectrometer is more resilient to the high levels of radio frequency interference (RFI) emitted from terrestrial sources. Polyphase filterbanks are therefore the best candidate for radio astronomy spectrometers, if computationally affordable.



(a) Affect of number of taps on the frequency response of a single polyphase filter bin. Hamming window coefficients are used in these data.



(b) A 4-tap, 16 channel polyphase filterbank, spanning a 1 MHz band. The middle eight channels are highlighted to indicate the overlap between channels.

Figure 4.1: Polyphase filter bank (PFB) spectrometer characteristics.




Figure 4.2: N-tap FIR filter block diagram. An FIR filter applies a weighted sum to the input sequence $x(t)$, outputting the filtered signal $y(t)$.

4.3 POLYPHASE FILTERBANKS

A filterbank is simply a collection of filters, a simple example being a highpass and lowpass filter pair. I discuss here filterbanks where each filter has identical passband characteristics and evenly spaced central frequencies, such as the 16 channel filterbank shown in Figure 4.1b. This equally spaced form is the most commonly implemented for radio astronomy applications.

This section begins with a short overview of digital filtering techniques; I refer the reader unfamiliar with digital signal processing to the excellent Lyons (2004) and Smith (2002). A comprehensive overview of polyphase filters is given by Vaidyanathan (1990), and in Chapter 2 of Harris (2004); I will give a brief introduction here. In the diagrams that follow, the symbol \otimes denotes multiplication of time samples, and \oplus denotes addition. The symbol z^{-n} is used to denote a time delay of n units, and the symbol $\downarrow D$ is used for downsampling by a factor D and $\uparrow U$ for upsampling by a factor U .

4.3.1 FINITE IMPULSE RESPONSE FILTERS

A finite impulse response (FIR) filter is the windowed moving average of an input sequence $x(t)$. An FIR filter computes the sum

$$y(t) = \sum_{k=0}^{K-1} h(k)x(t-k), \quad (4.10)$$

where $y(n)$ is the output sequence, and $h(k)$ is a set of K coefficients used for weighting. The upper summation bound, K , is called the number of taps. A streaming implementation of an FIR filter is shown in Figure 4.2.

If we downsample after an FIR by $\downarrow D$, we only keep the outputs $t = mD$, $m \in \mathbb{Z}^+$. A $\downarrow D$ downsampled filter will alias spectra centred at any multiple of the output sample rate to baseband. In such cases it is more efficient to only compute the terms we wish to keep:

$$y(mD) = \sum_{k=0}^{K-1} h(k)x(mD - k). \quad (4.11)$$

One way we can accomplish this is to use a polyphase decimating filter, which is discussed below.

4.3.2 POLYPHASE FIR FILTERS

A common technique in DSP is to decompose an input sequence $x(t)$ into a set

$$\mathbb{P} = \left\{ x_k(t) \quad k \in (0, P - 1) \right\} \quad (4.12)$$

of P sub-sequences, $x_k(t)$, each of which is given by

$$x_k(t) = (\downarrow P)(z^{-k})x(t). \quad (4.13)$$

This is known as polyphase decomposition. Even and odd decomposition of the signal $x(t)$ is achieved when $P = 2$:

$$x_0(t) = \{x(0), x(2), x(4), \dots\} \quad (4.14)$$

$$x_1(t) = \{x(1), x(3), x(5), \dots\}. \quad (4.15)$$

More generally, an input stream may be decomposed into P different ‘phases’.

Polyphase filter structures are often more efficient than standard finite impulse response filters when used in sample rate conversion. A $\downarrow P$ decimating FIR filter of length $K = MP$ can

be constructed from P discrete FIR filter ‘branches’, each acting upon a different phase; that is

$$y(t) = \sum_{p=0}^{P-1} \sum_{m=0}^{M-1} h_p(m)x_p(t-m), \quad (4.16)$$

This is known as a decimating polyphase filter. For example, a $P = 4$ branch polyphase filter with $M = 7$ taps per sub-filter would compute the sum

$$\begin{aligned} y(t) = & \sum_{m=0}^7 h_0(m)x_0(t-m) + \sum_{m=0}^7 h_1(m)x_1(t-m) \\ & + \sum_{m=0}^7 h_2(m)x_2(t-m) + \sum_{m=0}^7 h_3(m)x_3(t-m), \end{aligned} \quad (4.17)$$


and has an output equivalent to a $4 \times 7 = 28$ tap standard FIR filter with 4:1 downsampling.

Decimating polyphase filter structures are more efficient than standard FIR filter based downsampling techniques. If $\downarrow D$ downsampling occurs after the moving average of Equation 4.10, we compute D sums, but keep only 1 in D of these. This is inefficient. In comparison, Equation 4.16 only computes the output values that are of interest. A comparison of two decimating filters is shown in Figure 4.3. Figure 4.3b is a polyphase type structure, but decimation occurs after summation. Conversely, Figure 4.3a shows a more efficient implementation where decimation occurs before summation. By the noble identities (see Vaidyanathan, 1990), the output of this second filter is identical to that in Figure 4.3b.

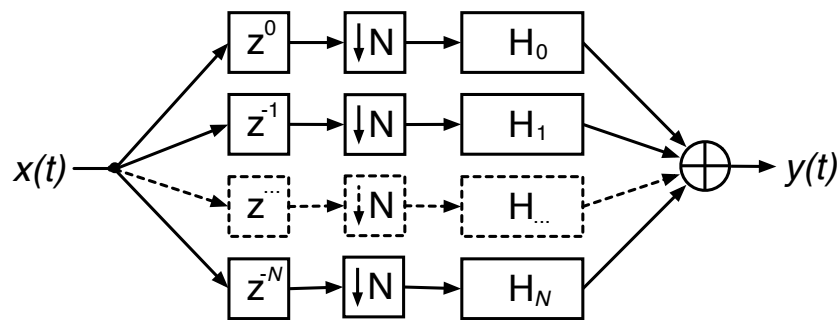
4.3.3 DFT BASED POLYPHASE FILTERBANKS

A computationally efficient filterbank with high interchannel isolation can be constructed from an FFT preceded by a prototype polyphase FIR filter frontend. Such an implementation exploits the fact that a lowpass filter with coefficients $h(k)$, can be converted into a quadrature (complex) bandpass filter with central frequency ω_k by multiplying the coefficients by $e^{i\omega_k t}$. That is,

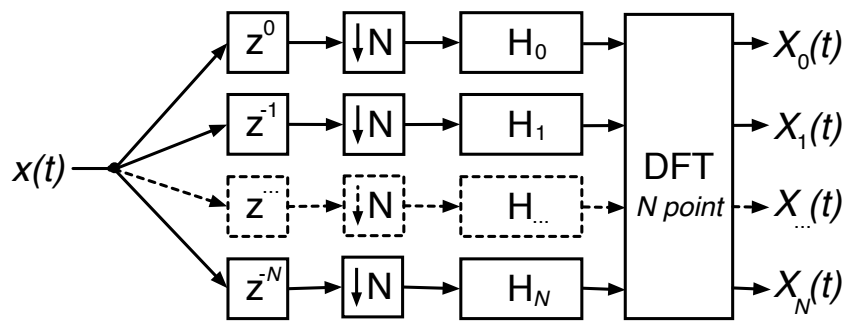
$$h_{bpf}(k) = h(k)e^{i\omega_k t}. \quad (4.18)$$



(a) An inefficient implementation of a decimating filter.



(b) N:1 decimating polyphase filter.



(c) N-channel polyphase filter bank.

Figure 4.3: Development of a polyphase filterbank from a prototype filter (a). In (b), the down-sampling step is moved using the noble identities for filters. In (c), the adder is replaced with an N- point DFT, creating a filterbank.

Now, suppose we have a decimating lowpass polyphase filter, such as that in Figure 4.3b. The output of each branch is

$$y_p(t) = \sum_{m=0}^{M-1} h_p(m)x_p(t - m), \quad (4.19)$$

where $h_p(m)$ are coefficients from our prototype lowpass filter. If we follow this by a DFT, as in Figure 4.3c, we then have

$$Y(k) = \sum_{p=0}^{P-1} y_p(t)e^{-2\pi ikp/P} \quad (4.20)$$

$$= \sum_{p=0}^{P-1} \sum_{m=0}^{M-1} h_p(m)e^{-2\pi ikp/P} x_p(t - m). \quad (4.21)$$

Comparing this form to Equation 4.11 and Equation 4.18, we recognise that the output of this structure — Figure 4.3c — is equivalent to a set of $\downarrow P$ downsampling polyphase filters:

$$\mathbb{F} = \left\{ h_k(m), \quad k \in (0, P - 1) \right\} \quad (4.22)$$

where the central frequency of each filter is shifted by an amount $2\pi ik/P$. For data sampled at the Nyquist rate f_s , this filterbank consists of N filters spanning $-f_s/2$ to $f_s/2$, with each filter separated by $f_s/2N$. For real sampled data the negative frequencies contain no extra information, so they need not be computed. This leaves a bank of $N/2$ filters, evenly spaced over a bandwidth $f_s/2$. This is the structure implemented in the HISPEC spectrometer detailed in this chapter.

This polyphase filterbank (PFB) structure was first proposed by Bellanger et al. (1976), and has been previously used in radio astronomy applications, see for example Zimmerman and Gulkis (1991) and Parsons et al. (2006). The rest of this chapter discusses the implementation of a 8192 channel polyphase filterbank for radio astronomy, on Field Programmable Gate Array based digital hardware.

4.4 FIELD PROGRAMMABLE GATE ARRAYS

A Field Programmable Gate Array (FPGA), is a reconfigurable integrated circuit that is used for digital signal processing. FPGAs consist of programmable logic components and memory ele-




Figure 4.4: Example FPGA prototyping boards. The Xilinx Virtex-5 FPGA is shown on the left, and an Altera Cyclone series FPGA is shown on the right.

ments with reconfigurable interconnections. Complex signal processing structures, such as digital filters and FFTs, can be made by interconnecting these multiple logic components. Some example FPGA boards are shown in Figure 4.4.

FPGAs have several advantages over CPUs and GPUs in certain applications. Firstly, they are more energy efficient, in some cases orders of magnitude so. For example, Kestur et al. (2010) report that FPGAs achieve between 2.7 to 293 times better energy efficiency for certain basic linear algebra subroutines (BLAS) than CPU and GPU based BLAS implementations. In addition, FPGAs are better suited to high bandwidth applications, as the I/O (input/output) bandwidth on CPUs and GPUs are limited by the speed of their host's PCIe bus, whereas FPGA bandwidth is limited by the number and speed of I/O pins upon the chip. FPGAs are not without their disadvantages: they are significantly harder to program, as algorithms must be mapped into a physical circuit within the FPGA's fabric. Each FPGA has a limited amount of logic and memory available, so algorithms which require large amounts of data to be stored in memory are hard to implement. The cost of the FPGA chip and FPGA design software is also a factor which may limit their applicability.

For cases where energy efficiency is the prime concern, FPGAs are outclassed by application specific integrated circuits (ASICs). ASICs are chips designed for one particular use; they trade off the flexibility of reconfigurability for smaller footprints and energy consumption. Kuon and Rose (2006) report that an FPGA is on average 40 times larger and 3.2 times slower than an ASIC implementation of logic-only circuits, and that FPGAs consume on average 12 times more dynamic power than equivalent ASICs. The disadvantages of ASICs are their lack of reconfigurability; their significant fabrication cost (of order US\$1M); and, longer lead times due to their

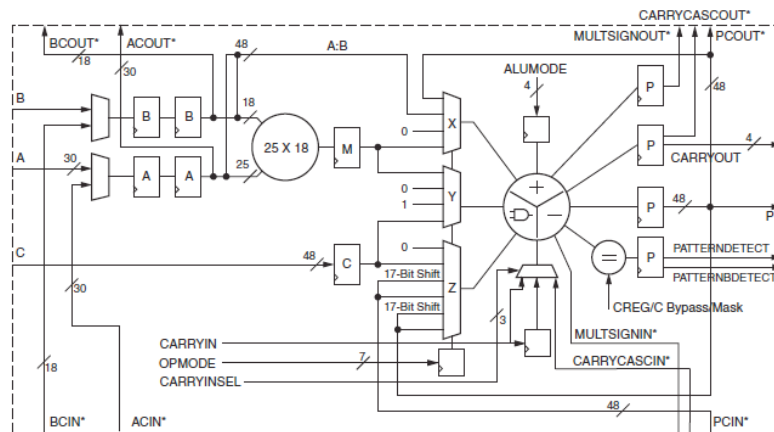


Figure 4.5: Schematic of the Xilinx DSP48E slice, from DPS48E usage guide (Xilinx, 2012).

bespoke nature.

If energy efficiency and bandwidth are not limiting factors, and a desired algorithm can be performed on a CPU or GPU, then it should be implemented on one. This is as the development time is invariably lower, the code is easier to read, and debugging is far easier. The choice between FPGAs, CPUs, GPUs, and ASICs is therefore not a simple one: it is highly dependent on the application, the timescale, and how much money one has to spend.

4.5 DSP HARDWARE FOR RADIO ASTRONOMY

For high bandwidth radio astronomy applications, FPGAs are an attractive technology, and have received much attention in recent years. See for example the Altera-powered Uniboard project (Szomoru, 2010), and the Xilinx-powered CASPER iBOB, BEE2 and ROACH FPGA boards (Parsons et al., 2006). A recent addition to the CASPER family is the ROACH-2, a Virtex-6 based board. The spectrometer design detailed in this chapter is designed to run upon the CASPER ROACH board, which is introduced below. Central to the ROACH is the Virtex-5 FPGA.

4.5.1 VIRTEX-5 FPGA

The Virtex-5 FPGA is a high performance FPGA designed by Xilinx Incorporated². It is constructed using a 65 nm CMOS manufacturing process. Since my DPhil. began in 2009, Xilinx have released

²<http://www.xilinx.com/support/documentation/virtex-5.htm>

the Virtex-6 and Virtex-7 FPGAs, which are based respectively on 40 nm and 28 nm manufacturing processes.

The basic logic elements for Xilinx FPGAs are Configurable Logic Blocks (CLBs). These provide combinatorial and synchronous logic as well as distributed memory and shift registers. In addition to CLBs, each Virtex-5 contains block RAMs (BRAMs); these RAM modules are arranged into columns on the chip fabric and are 36 kb each.

Much of the DSP versatility of the Virtex-5 comes from elements called DSP48E slices. These slices consist of a 25×18 bit multiplier followed by data routing multiplexers. This multiplex is followed by an element which can be used as an adder or two-input logic unit. A circuit schematic of the DSP48E is shown in Figure 4.5.

Peak performance of the Virtex-5 is achieved when every slice is being used constructively every clock cycle. In reality, this level of utilization is never achieved; nevertheless, if one managed to utilize all 640 DSP48E slices of a SX95T Virtex-5 at the maximum clock rate of 550 MHz, 352 billion multiply accumulate operations each second (MACs) could be performed. Hickish et al. (2012) shows that each DSP48Es can be used to compute several 4-bit multiplications per clock cycle, so performance up to 1TMAC is theoretically achievable on one Virtex-5 FPGA.

4.5.2 THE CASPER ROACH BOARD

The FPGA boards produced by the Collaboration for Astronomical Signal Processing and Electronic Research (CASPER) are based around the Xilinx Virtex series of FPGA. The ROACH, or Reconfigurable Open Architecture Computing Hardware board, is a SX95T Virtex-5 FPGA based signal processing board designed for radio astronomy applications. A simplified block diagram of the board is shown in Figure 4.6.

INPUT/OUTPUT INTERFACES Four CX4 connectors allow up to 40Gb/s of bandwidth for data transfer over 10GbE or XAUI protocols. Daughter boards, such as analogue to digital converters, may be connected via the two Z-DOK interfaces, each with 40 differential pairs which connect to the FPGA I/O pins. High speed data transfer may also be done through the Q-Strip® connector (QSH). General Purpose Input/Output (GPIO) pins are also provided for control of miscellaneous sensors and external devices. For temporary storage of large vectors, the FPGA is connected to

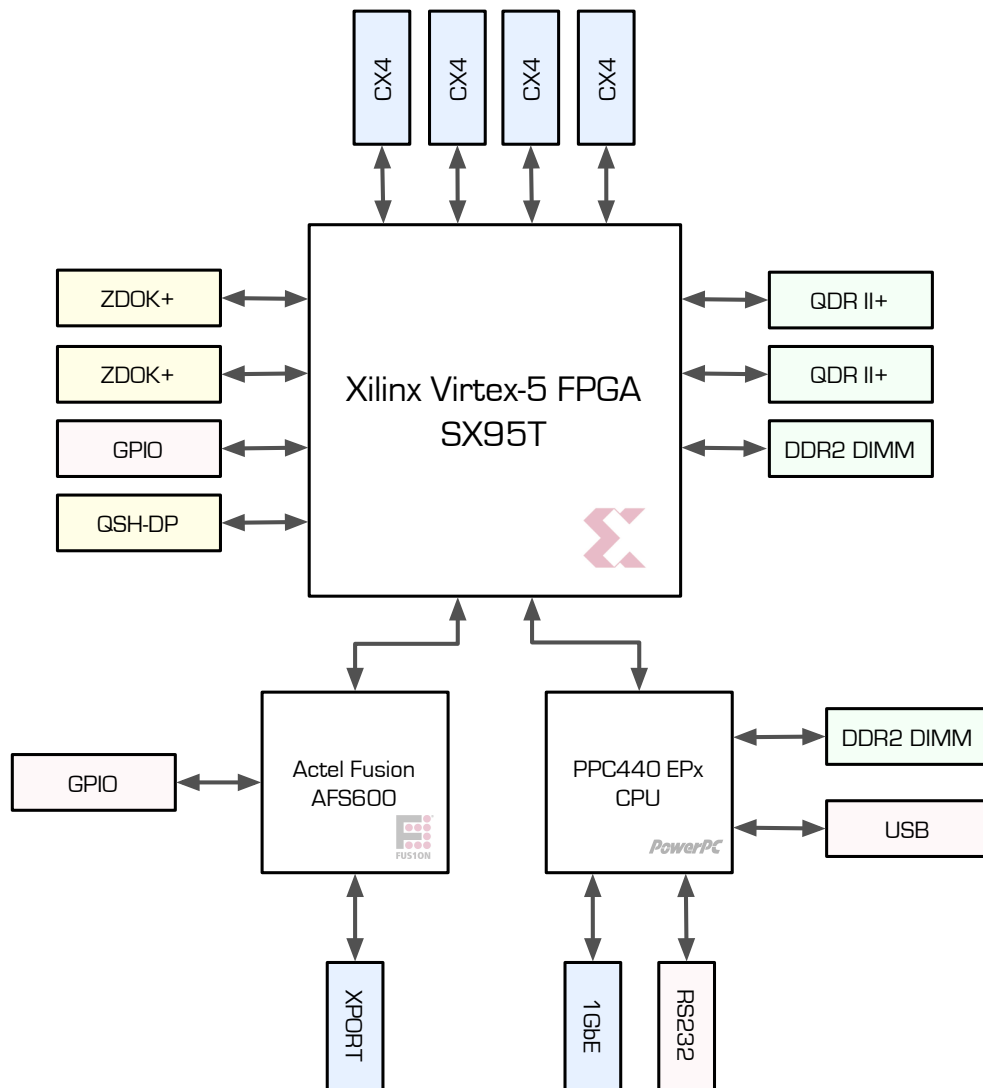


Figure 4.6: Block diagram of the CASPER ROACH board, showing input/output interfaces and peripheral memories.




Figure 4.7: The CASPER ROACH board, revision 1.00. Image credit: CASPER collaboration

quad data rate random access memory (QDRII+ RAM) and double data rate memory (DDR2).

POWERPC The Virtex-5 FPGA is connected to a PowerPC chip by an On-Chip Peripheral Bus (OPB) controller. The PowerPC runs a customized Linux kernel and provides command and control for the board. It has its own dedicated DDR2 and Flash memory, 1Gb Ethernet port, USB port, and SD memory card socket.

ACTEL FUSION An Actel Fusion FPGA controls the ATX power supply's soft power toggling and monitors the voltage and current draw of key power supplies. It also monitors fan speeds and component temperatures. Remote access to the Actel FPGA is enabled by an XPORT Ethernet adaptor, which encapsulates serial I/O with the Actel into an Ethernet format.

4.5.3 IADC DIGITIZER CARD

The iADC, or ADC2x1000-8, is an 8-bit, 2-input analogue to digital converter (ADC). The iADC is a Eurocard-sized daughter board which connects to the ROACH via the Z-DOK connector. Its centrepiece is an Atmel/e2V AT84AD001B digitizer chip, which runs at speeds of up to 1Gsample/sec. Inputs are connected to the iADC's SMA jacks; a MiniCircuits ADTL2-18 balun (30 MHz-1.8 GHz bandwidth) converts the single-ended transmission line to differential, as is required by the digitizer chip. When in operation, an external synthesized clock signal must be presented as a clock source. A pulse per second (1PPS) may also be connected; this signal is

Table 4.1: Atmel AT84AD001B performance characteristics (Atmel, 2009).

Sampling rate	1Gsample/sec, 2GS/s interleaved
Number of bits	8
Effective number of bits (ENOB)	6.8
Differential analogue input voltage (full scale)	500mVpp
Spurious free dynamic range	-54 dBc
Differential non-linearity	0.25 LSB
Integral non-linearity	0.5 LSB
Bit error rate	10^{-13} at 1GSample/sec

passed through to the FPGA so multiple ROACH boards may be synchronized to within a clock cycle. Using this ADC, two signals with a bandwidth of 512 MHz may be digitized. Alternatively, the AT84AD001B offers an interleave mode which sacrifices one input but is then able to digitize 1024 MHz from a single input.

Two important characteristics for spectrometers are crosstalk levels between signals, and the linearity of the digitizer's response. Crosstalk of the iADC has been measured by Parsons (2009) to be under 28 dB, and was found that ~ 60 dB of isolation can be achieved by subtracting the stable crosstalk component. Linearity was tested by Lipsey (2012, personal communication), and found to be in accordance with the manufacturer's quoted specification (differential 0.25 LSB, integral 0.5 LSB). Further specifics are listed in Table 4.1.

4.5.4 THE XILINX TOOLFLOW

All the FPGA firmware detailed in this thesis was designed using the Xilinx ISE toolflow³. The toolflow has been integrated into MATLAB Simulink⁴, which is a high level graphical programming environment.

The CASPER collaboration provide Simulink based component libraries designed specifically for radio astronomy instrumentation. These libraries include implementations of DSP structures such as FFTs, FIR filters, vector accumulators, digital mixers and complex multipliers. These components are parameterized, so that their specifications can be changed by the user.

The CASPER libraries are built upon the Xilinx System Generator blockset for Simulink. Each Xilinx block has a set of configurable parameters, input ports and output ports. Simulink allows systems to be developed by connecting blocks together, and provides a simulation environment

³<http://www.xilinx.com/products/design-tools/ise-design-suite/>

⁴<http://www.mathworks.com.au/products/simulink/>

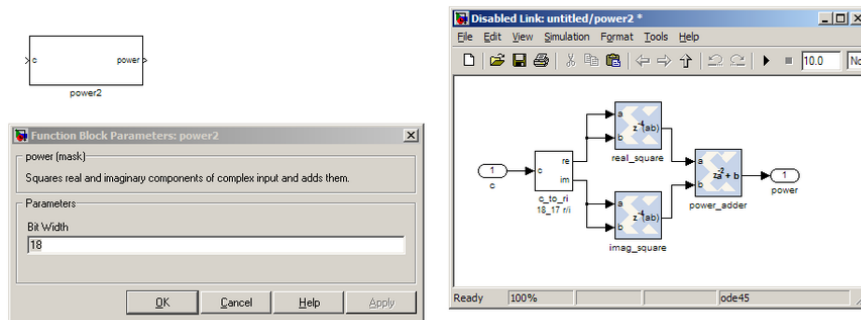


Figure 4.8: Example of a CASPER simulink block. This particular block computes the power of a complex valued input signal.

for analysing signal properties. Once a model has been constructed, an HDL (hardware description language) netlist can be compiled through Xilinx System Generator. From this point, the Xilinx toolflow can be used to synthesize and implement FPGA specific firmware.

The Xilinx toolflow maps Simulink designs into physical circuits upon the FPGA's fabric, using a heuristic place-and-route algorithm. This process can take several hours for designs with high resource utilization (using a computer with a 2.66 GHz Intel Core-2 Duo CPU, 8 GB RAM). After the design is placed and routed, a bitstream is then generated with which to program the FPGA. Loading a pre-compiled bitstream onto an FPGA takes less than a second, hence FPGAs are considered 'field programmable'.

4.6 THE HISPEC SPECTROMETER

In this section, I introduce an FPGA based wide-bandwidth spectrometer called HISPEC, designed for redshifted HI line spectroscopy. This spectrometer is used in the HIPSR project detailed in Chapter 6, and the D-PAD single station testing detailed in Chapter 5. HISPEC is based on an 8192-channel, 4-tap polyphase filterbank. It digitizes both inputs from a dual-polarization antenna and applies a PFB to each input before computing all four Stokes parameters and accumulating the signal. The spectrometer can digitize and process up to 512 MHz of bandwidth instantaneously.

The HISPEC spectrometer runs on an SX95T Virtex-5 FPGA equipped ROACH board. I designed the FPGA gateway using the Xilinx ISE toolflow, and the MATLAB Simulink block libraries provided by the CASPER collaboration (Parsons et al., 2006); in addition, I developed new blocks

whenever the CASPER provided blocks were unsuitable or simply did not exist.

4.6.1 FIRMWARE OVERVIEW

A block diagram of the HISPEC firmware is shown in Figure 4.9. Only the top-level blocks are shown; each of these blocks is created from Xilinx primitives. The data can be visualized as flowing from down the arrows and through the processing blocks, from left to right. In the diagram, all Input/Output interfaces are shown in light yellow. In the firmware description that follows, I will use the notation a_b to refer to a bit numbers with a decimal point preceding bit b . For example, 8_7 is an 8-bit signal, with 7 fractional bits. All arithmetic in the FPGA is fixed point, and is represented in 2's complement form when signed. From left to right, Figure 4.9 shows:

- Many of the CASPER library blocks use a vector warning system to initialize, synchronize and reset blocks. The warning signal is a single boolean 'pulse' of one cycle. Once generated, this propagates through the design, and care is taken to ensure that the first valid datum arrives at each block one clock cycle after the warning. This is commonly referred to as a 'sync pulse'. In the HISPEC design, this pulse is generated by the *sync* block.
- The first block, *adc*, links the iADC digitizer to the Simulink model as an external HDL defined peripheral. The *adc* block presents four time samples from the iADC in parallel per FPGA clock cycle. These samples are fixed point, 8_7 bit signed 2's complement numbers in the range $[-1,1)$. For diagnostics, a digital noise source *noise* of identical bitwidth can be used as a test vector for downstream blocks. The selection between *adc* and *noise* is achieved using the multiplexer block *mux*, which is controlled by the register *mux_sel*.
- The signals pass through the 4-tap, 16384-branch, Hamming-windowed *pfb_fir* polyphase filterbank frontend. Each filter coefficient 18_17 bits, and the number of bits representing the signal grows from 8_7 to 18_17 bits through this filter. The *pfb_fir* is provided by the CASPER toolflow.
- After this, a 16384 point Radix-2 biphase pipelined FFT is applied, *fft0* and *fft1*, to each signal. 18-bit fixed point arithmetic is used throughout, with the data cast back down to 18-bit after each multiply. The output is 18-bit real, 18-bit imaginary, but is represented

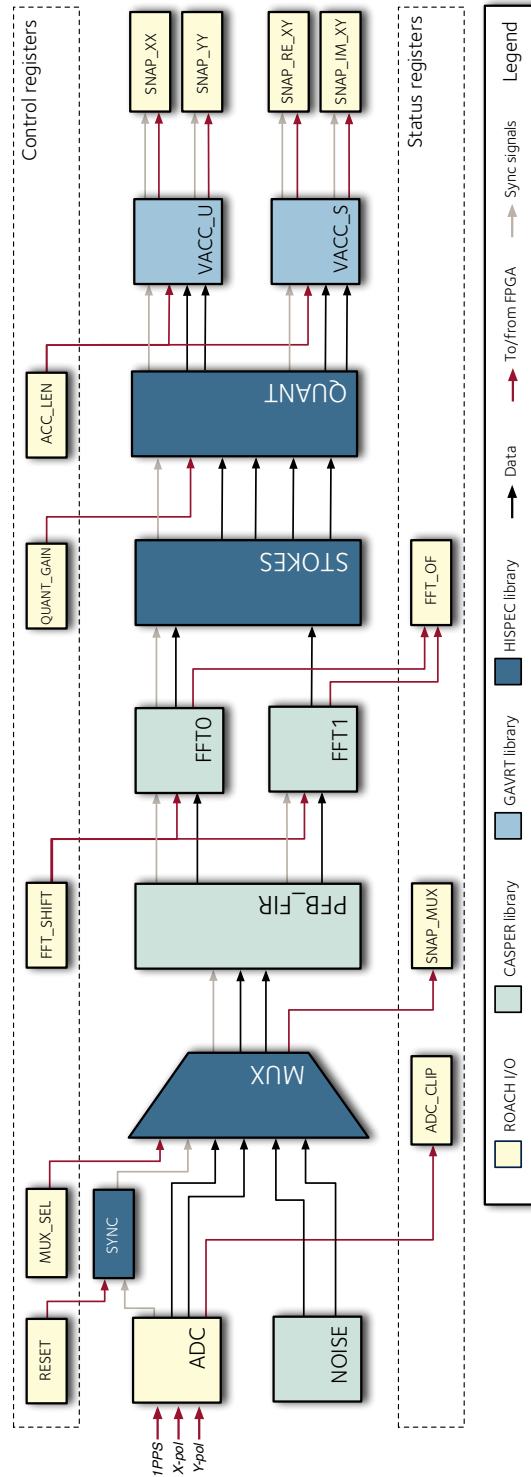


Figure 4.9: Simulink diagram of the HISPEC spectrometer. Only the top-level blocks are shown. The control and status registers are software accessible for real-time configuration via the PowerPC. Spectral data is stored in the shared BRAMs snap_xx, snap_yy, snap_re_xy and snap_im_xy, which is also software accessible.

as a 36_0 quantity. This block is a core component of the CASPER library and has been optimized by many members of the collaboration.

- The *stokes* block then computes auto- and cross-correlations for the two signals, from which Stokes parameters can be determined. Denoting the two input signals as X and Y , the required correlations are the real-valued XX^* and YY^* , and the complex-valued XY^* . As the FFT uses 18-bit complex fixed point arithmetic, the dynamic range is $10\log_{10}(2^{2 \times 18}) = 108.37$ dB. The heart of the *stokes* block is a DSP48E slice based complex multiplier. The multiplication yields 36_34 bit numbers. The autocorrelations are computed using unsigned arithmetic, and the cross correlations with signed.
- The *quant* block converts the signal down to 18_0 bits in preparation for vector accumulation. Bit selection is achieved by multiplication with a 16_0 bit gain term, growing the signal to 52_0 bits, from which a 18_0 bit slice is taken. As the data output from the *stokes* block is above 25 bit, this multiply is conducted using distributed logic, not DSP48E slices.
- The *vacc_s* and *vacc_u* blocks (signed for cross correlations, unsigned for autocorrelations), are vector accumulators with 32_0 bit output which uses BRAMs for vector storage. Data are accumulated by feeding the output of the BRAM back into a preceding adder with an appropriate time delay. Up to 2^{16} accumulations cycles can be run before there is risk of signal overflow. These blocks are from the GAVRT blockset, which is a sub-library supplied by CASPER, originally designed for the Goldstone Apple Valley radio telescope.
- Finally, the four *snap* blocks on the right hand side of the diagram store the signal in BRAM, which is shared between the FPGA and Power PC.

Along with these signal processing blocks are a set of control registers, which allow the user to set the accumulation length, FFT shift schedule (which determines at which stages the signal is requantized), and to select between the digital noise source and the iADC. These are read-writable from the PowerPC, and read-only from the FPGA. A number of status registers are also included; these are read-only by the PowerPC and read-writable by the FPGA. The status registers are used for monitoring the data quality: boolean flags are raised for ADC clipping and

Table 4.2: HISPEC resource utilization, as a fraction of total resources available on the Virtex 5 FPGA.

BRAMs	202/244	82%
DSP48Es	222/640	34%
Slice logic	11203/14720	76%

signal overflows, and a count of successful accumulations is kept. A number of diagnostic shared BRAMs are interspersed throughout the design for debugging and signal monitoring.

Several of the pins of the FPGA are connected to a bank of GPIO (general purpose input/output) headers. Eight of these headers are connected to a panel of diagnostic LEDs. In the HISPEC design, boolean flag is raised whenever ADC clipping and signal overflow occurs. This flag is passed to the GPIO pins which creates a voltage and causes the LEDs to flash. The 1PPS signal and the accumulation count also trigger an LED to flash, akin to a ‘heartbeat’ expressing board health.

4.6.2 FIRMWARE COMPILATION

One of the challenges of FPGA firmware design is ensuring an implementation uses less circuit logic than the amount available on the FPGA fabric. More specifically, the number of nets in the HDL netlist may not exceed the number of available components on the FPGA chip. To fit dual 16384-point FFTs onto the Virtex-5 required balancing BRAM and distributed memory usage. The CASPER FFT block is parameterized, so that coefficients and delays can be stored in distributed memory if they are under a given size; this decreases the BRAM usage at the expense of slices and routing complexity. To find a suitable balance, I compiled the FFT multiple times with different memory storage parameters until I found an appropriate implementation.

Along with such physical constraints are timing constraints: the output of each logic element must reach the input of the next element within a clock cycle. If the Xilinx place-and-route tool cannot find a solution in which all timing constraints are satisfied, the implementation is said to ‘not meet timing’. As the FPGA clock speed increases, the period of time available to route a signal decreases, and as such the likelihood of meeting timing also decreases.

To digitize 512 MHz of bandwidth, the FPGA must be clocked at 256 MHz. While the HISPEC design compiles at clock speeds up to 200 MHz, attempting to compile the design at 256 MHz fails. To find which nets were the most troublesome to route, I analysed the HISPEC design using

the Xilinx PlanAhead software package⁵. Among the features of PlanAhead is a floor-planning tool, which highlights which nets are not meeting constraints. Once identified, issues can then be removed by modifying the design, or by asserting placement constraints.

The majority of timing issues for HISPEC arose from the *pfb_fir* frontend. In the *pfb_fir* block, filter coefficients are stored in BRAM. Routing the coefficients from the BRAM to the filter taps (which use DSP48E slices as multipliers), required routing over path distances that were unachievable at the high clock speed. The solution I implemented was to force the coefficients to be routed through an intermediate register, which required modification of the *pfb_fir*. These modifications increase the overall latency of the block without altering the overall implementation.

An alternative strategy to meeting timing at high clock speeds is to place constraints on the position of nets. While I experimented with this approach, solving the underlying timing issues within the *pfb_fir* by adding latency proved adequate and constrained placement was ultimately not required. Figure 4.10 shows the final mapping of the HISPEC spectrometer into the FPGA's logic blocks.

Table 4.2 shows the resource utilization of the final design. The BRAM usage of both the filter frontend, FFT and vector accumulator grow in proportion to the number of channels; as such, a design similar in implementation but with a larger number of channels would not fit into the FPGA fabric, as too many BRAMs would be required. Adding additional taps to the filter frontend also increases BRAM usage.

In order to build larger full Stokes spectrometers on a single FPGA, one would have to optimize their designs to use fewer BRAMs, or use a different FPGA with a larger number BRAMs. This design could be ported to the ROACH-2, which has a Xilinx Virtex-6 XC6VSX475T FPGA as its centrepiece. This chip has 1064 BRAMs⁶, so one would expect that the number of channels could be quadrupled to 2^{16} without major changes to the design. The recently released Virtex-7 series sports up to 1880 BRAMs and 3360 DSP48Es⁷, and this upward trend will continue in future FPGAs.

⁵<http://www.xilinx.com/tools/planahead.htm>

⁶<http://www.xilinx.com/support/documentation/virtex-6.htm>

⁷<http://www.xilinx.com/products/silicon-devices/fpga/virtex-7/index.htm>




Figure 4.10: Xilinx PlanAhead circuit diagram of HIPSR spectrometer compiled at a clock speed of 256MHz, before floor planning (left), and after (right). Components associated with the PFB FIR frontend are highlighted in gold, and the two FFTs in red and green.

4.7 RESULTS

To measure the response of the spectrometer to a reference signal, I connected a TTI TGR2050 synthesized RF signal generator to the spectrometer's inputs. This signal generator operates over 150 kHz to 2000 MHz, with a setting resolution of 10 Hz. The TG2050 has an output power of -127 to $+7$ dBm, with ± 2 dBm accuracy. Harmonically related signals are specified to <-25 dBc (decibels relative to the carrier). Further specifications can be found in TTI (2006).

In the paragraphs that follow, the response of the spectrometer to a single tone, the spectral isolation between channels, and the noise floor of the spectrometer are discussed. Field testing, in which the spectrometer is connected to a radio antenna, is detailed in subsequent chapters.

SINGLE TONE RESPONSE The response of the digitizer to a 1 MHz sine wave is shown in Figure 4.11. These 8-bit data were extracted from a shared BRAM on the FPGA fabric, which captures 4096 FPGA clock cycles worth of data. This is useful for debugging and ensuring the power level of the input signal is suitable. The filterbank's response to a -30 dBm, 10 MHz signal is shown in Figure 4.12. The first harmonic is ~ 12 bits below the carrier, or -36.1 dBc — well below the manufacturer quoted <-25 dBc.

SPECTROMETER NOISE FLOOR Figure 4.13 shows the response of the spectrometer with no input signal; i.e. with terminators on the signal inputs. In the figure, the lowest 18 bits of the 36-bit




Figure 4.11: iADC response for a 1 MHz sine wave. The dBm power value neglects transmission loss of ~ 3 dB over the coaxial cable's length




Figure 4.12: Spectrometer response to a 10 MHz tone. The harmonics of the signal can be seen at a lower level; the first harmonic is -36.1 dBc. Power is expressed on a \log_2 based scale (equivalent to number of bits) on the left axis, and as power in decibels on the right axis ($10\log_{10}$ values).




Figure 4.13: Noise floor of the spectrometer (integration time of 2 s). The reference clocks of the PCI bus, FPGA, and DDR2 memory can be seen in the data. The autocorrelation XX^* is shown in black, and YY^* is in red.

autocorrelations were selected, and an accumulation length of 2^{14} clock cycles was used. There is a clear structure apparent in the data, particularly around multiples of 25 MHz. This structure arises to the imperfect nature of the digitizer: the sample and hold process introduces noise, distortion and aperture jitter.

There are also some self-induced interference sources apparent in the spectrum of Figure 4.13. At 133 MHz the PCI bus clock is visible, and at 4096 (200 MHz), the derived FPGA clock can be seen. The DDR2 clock can also be seen, at 333 MHz. The non-linear response of this noise floor can be mitigated by setting the gain coefficients of the *quant* block such that the least significant bit is not selected. A requisite for this is that the RMS power of the digitized signal is much greater than the 1-bit toggle level.

SPECTRAL ISOLATION In order to probe the response of a single polyphase filterbank channel, I measured the response of a single channel as the input frequency was scanned. So that high frequency resolution could be achieved, data capture was automated with a Python script that controls the signal generator via GPIB and collects data from the spectrometer via Ethernet. An




Figure 4.14: Measured response of a single polyphase filterbank channel of HISPEC, compared to the theoretical response of a 32768-tap Hamming windowed FIR filter with 32-bit floating point coefficients.

input signal power of -15 dBm was used to measure the sidelobes, and decreased to -30 dBm to measure the main lobe. The measured filter response is shown in Figure 4.14, along with the response of an equivalent 32768-tap Hamming windowed filter. As can be seen, the responses match well over the main lobe and first sidelobe, but deviate after this. This is due to rounding errors within the *pf_b_fir* and FFT, which use 18-bit coefficients. Nevertheless, these sidelobes lie below the 54 dB of dynamic range available after the *quant* block, so cannot be seen unless the signal is clipping elsewhere, which is unsatisfactory.

4.8 DISCUSSION

The HISPEC spectrometer has been designed for wide-bandwidth spectroscopy at L-band wavelengths. Consequently, it is illuminating to compare its performance and specifications with existing instruments in use at major science facilities.

4.8.1 COMPARISON TO EXISTING INSTRUMENTS

Table 4.3 presents the specifications of three L-band spectrometers currently in use at major facilities, and contrasts them against HISPEC. MBCORR is the main spectrometer in use at the Parkes 64 m telescope (Staveley-Smith et al., 1996); GBTS is the lag correlator in use at the Robert C. Byrd 100 m dish at Green Bank, West Virginia (GBT, 2011); and FFTS is the L-band spectrometer in use at the Effelsberg 100 m telescope (Klein et al., 2006).

QUANTIZATION EFFICIENCY Signal quantization affects the signal-to-noise ratio (SNR) of a signal. The fractional loss in SNR is known as the efficiency factor, n_Q ; Thompson et al. (2007) gives convenient formulae to calculate n_Q . The GBT Spectrometer and MBCORR use 3- and 4-level quantization, and corresponding efficiency factors of 0.8812 and 0.8098, respectively. In comparison, HISPEC and FFTS use 256-level quantization (8-bit), which yields efficiency factors close to unity. As shown by Thompson et al. (2007), the efficiency factor depends upon the root mean square (RMS) value of the input signal. If one sets the input power to be 6 dB below the full-scale level, then an efficiency factor of 0.9999 is achieved. However, for L-band applications it is desirable to have some headroom so that intermittent RFI sources do not cause clipping. A more suitable choice is to set the RMS power of the signal to 30 dB below the full-scale level; this yields an efficiency factor $n_Q=0.9984$ while allowing 24 dB of headroom for RFI sources.

MAXIMUM CHANNELS AT FULL BANDWIDTH These three spectrometers are all flexible: bandwidth may be traded off for increased spectral resolution. HISPEC is also bandwidth agile as the reference clock signal can be set as low as 200 MHz, resulting in 100 MHz of bandwidth. However, this is not an efficient use of the FPGA's resources because the number of operations performed per second drops accordingly. For this reason, the table compares maximum channels at maximum instantaneous bandwidth, as opposed to maximal spectral-resolution modes.

INTERCHANNEL ISOLATION An ideal filter channel would have a width equal to the frequency spacing, with unity response over its bandwidth $\Delta\nu$ and zero response elsewhere. As a perfect FIR filter would require an infinite number of coefficients, all filterbank implementations suffer from

imperfect responses: neighbouring channels overlap, and there is always some level of spectral leakage between channels.

The interchannel isolation of a spectrometer is dependent upon the windowing function used. FFTS uses a 3-term Blackman-Harris window, while a Hann window is routinely applied to data from MBCORR and GBTS. For a Hann-windowed ACS, the first sidelobe level is -16 dB. In contrast, the Blackman-Harris window of FFTS gives -60 dB sidelobe suppression. HISPEC uses a Hamming window, a version of the Hann window modified to minimize the first sidelobe level, which gives a sidelobe level of -50 dB. Note that the first sidelobe of a Hamming-windowed FFT lies 6 dB higher, at -44 dB.

Applying a windowing function increases the width of the main lobe. The first null in the response of a non-windowed FFT occurs at ± 1 channel (see Figure 4.1a), giving a first null bandwidth (FNBW) of $2\Delta\nu$. For a Hann window, the main lobe width is doubled; for a Blackman-Harris window, it is quadrupled to $8\Delta\nu$. In comparison, a 4-tap Hamming-windowed PFB has a FNBW of $1.88\Delta\nu$. This means that neighbouring channels overlap less, so spectral resolution is higher. At the location of the first sidelobe of the Blackman-Harris window, the 4-tap Hamming window is -80 dB, a full 20 dB lower.

POLARIMETRY As MBCORR, GBTS, and FFTS do not compute the cross correlations of two input signals, only the I and Q Stokes parameters may be formed. HISPEC does compute the cross correlations, so the U and V Stokes parameters may also be formed. As such, HISPEC is suitable for polarimetric observations, where FFTS, MBCORR and GBTS are not.

4.8.2 NEXT GENERATION SPECTROMETERS

HISPEC is one of many planned modes of the HIPSr project detailed in Chapter 6, which will likely replace the aging MBCORR. There is similar upgrade work taking place at Green Bank and Effelsberg.

XFFTS A successor to FFTS, named XFFTS, has recently been installed at Effelsberg (Klein et al., 2012). XFFTS is a polyphase filterbank implemented on a Virtex-6 based board with an integrated e2v 4×1.25 GS/s, 10-bit digitizer. The Virtex-6 chip in XFFTS is more powerful than the

Table 4.3: Comparison of HISPEC with current L-band spectrometers. Note that all of these spectrometers have multiple modes of operation; I have listed here maximum channels at maximum instantaneous bandwidth.

	MBCORR Parkes 64 m	GBTS Green Bank 100 m	FFTS Effelsberg 100 m	HISPEC This work
Sampling	4 level ($n_Q=0.8812$)	3 level ($n_Q=0.8098$)	256 level ($n_Q=0.9948$)	256 level ($n_Q=0.9948$)
Max bandwidth	64 MHz	800 MHz	500 MHz	512 MHz
Max chans (full BW)	1024	8192	16384	8192
Implementation	ACS (1024 lag)	ACS (8192 lag)	FTF (32768 point)	PFB (4 tap, 16384 point)
Windowing function	Hann	Hann	Blackman-Harris	Hamming
First sidelobe level	-16 dB	-16 dB	-60 dB	-50 dB
Main lobe FNBW	$4\Delta\nu$	$4\Delta\nu$	$8\Delta\nu$	$1.88\Delta\nu$
Polarimetry	I, Q	I, Q	I, Q	I, Q, U, V

ROACH Virtex-5 and as such each XFFTS board is capable of digitizing and more bandwidth and processing more channels than a HISPEC programmed ROACH.

While both XFFTS and ROACH are Xilinx FPGA based, the design of the two are markedly different. The ROACH does not have an on-board digitizer, but instead has a ZDOK connector with which to connect daughter boards. ROACH boards are designed to be interconnected by up to four 10GbE/XAUI connections, whereas XFFTS board only offers 1GbE (they are not designed to be interconnected). The QDR, DDR and PowerPC of the ROACH are also absent from the XFFTS. The absence of such features allows the XFFTS board to fit onto a eurocard — the same size as a single ROACH daughter board.

VEGAS The CASPER collaboration have also designed a Virtex-6 based instrument, ROACH-2. The ROACH-2 board is similar in design to ROACH, but has a more powerful FPGA and increased I/O throughput and memory bandwidth. A number of ROACH-2 boards have been installed at Green Bank for the VEGAS project (Versatile GBT Astronomical Spectrometer), which will replace the GBTS spectrometer and add new functionality (Werthimer, 2012). The ROACH-2 boards are paired with a National ADC08300 based 3Gsample/s digitizer card. In total, 8 ROACH-2 boards will be installed, and will be connected by 10Gb Ethernet to a CPU/GPU based high performance processing cluster.

This heterogeneous architecture is expected to greatly increase the flexibility of the system, by allowing algorithms not suitable for the FPGA to be run on the CPU/GPU nodes, and

vice-versa. The HIPSR system introduced in Chapter 6 also uses a heterogeneous architecture; however, the HISPEC firmware does not require any CPU/GPU based signal processing as all necessary tasks are conducted on the FPGA.

4.9 CONCLUSIONS

Polyphase filterbanks have attractive characteristics for radio astronomy applications. They offer excellent interchannel isolation and can be implemented using efficient FFT based structures. Field Programmable Gate Arrays are a viable platform on which to implement wide bandwidth polyphase filterbank spectrometers. I have implemented such a spectrometer using the Xilinx toolflow and CASPER block libraries. In this chapter I detailed the implementation and presented results on its performance.

The HISPEC spectrometer is only one of many planned modes of operation for the HIPSR system. An overview of HIPSR, its installation, and first-light results with the HISPEC firmware are detailed in Chapter 6. In the next chapter, I present results from the D-PAD single station testing, in which the HISPEC design was used as the digital backend.

CHAPTER 5

D-PAD SINGLE STATION RESULTS

5.1 INTRODUCTION

Aperture arrays are an attractive technology for low frequency radio astronomy, as they offer superb sky coverage. Unlike a traditional single-pixel feed-based telescope, a single aperture array can have multiple beams pointed on the sky — so long as this is supported by the beamformer. If omnidirectional antennas are used in tandem with a sufficiently powerful digital beamformer, a fan of beams which cover the entire sky can be formed. This functionality comes at a price: a discrete analogue chain is required for each antenna. Also, the number of antennas required to reach a target collecting area grows as frequency increases. Duplicating an antenna and its analogue chain costs money, so at high frequency aperture arrays become prohibitively expensive. The cost of the beamformer is also a factor that must be considered.

As was discussed in Chapter 1, intensity mapping experiments require a large field of view so they can quickly survey large areas of the sky. As such, aperture arrays should be considered as a potential technology for 21-cm intensity mapping instruments. At metre wavelengths, sparse aperture arrays are indeed a viable technology, as is evidenced by the LOFAR, MWA, and LWA telescopes. The D-PAD project aims to show that sparse aperture arrays are a viable technology for redshifted hydrogen surveys over 1.0-1.5 GHz (L-band), which corresponds to wavelengths of 20-30 cm.

In this Chapter, I report on the deployment of the first D-PAD station. While each individual component may be tested in the lab, the performance of the array as a whole can only be de-

terminated by field testing. The aim of the deployment was to measure the characteristics of the station, such as the system bandpass, beam pattern, and system temperature.

This chapter begins with an overview of when, where and how the station was deployed. This is followed by results, which are then analysed to extract performance metrics. I conclude with a summary and some discussion on possible future work.

5.2 DEPLOYMENT

The D-PAD single station was tested during February 2012, on the rooftop garden of the Oxford Astrophysics department, located in the Denys Wilkinson Building (DWB). The DWB is located at latitude and longitude (51.76° , -1.26°). While the rooftop is a convenient location for testing, altitudes below 30 degrees are occluded by surrounding buildings. The station was arranged on a bespoke platform, see Figure 5.1a. The antennas were placed on an evenly spaced grid with 40cm between antennas, measured from the antenna centre. This corresponds to 2λ at 1.5 GHz, and 1.33λ at 1 GHz. The response of a 4×4 grid can be modelled both analytically and by simulation, so is a good configuration for initial system tests.

To protect the antennas and electronics from the elements, a radome was fashioned out of polyethylene sheeting and PVC piping. These plastics are transparent at radio wavelengths. As the radome has a large surface area and low weight, it was held taut by nylon ropes fixed to the ground, to prevent it from being blown over by wind gusts. Results from the testing campaign, which lasted eight days, are presented below.

5.3 RESULTS

A typical bandpass for the single station is shown in Figure 5.2. The bandpass was measured with the HISPEC spectrometer from Chapter 4, which was then compared against a TTI PSA2071T spectrum analyser. Data from HISPEC is shown as a function of time in Figure 5.3, over a two day period February 27th-28th. It is immediately apparent there are strong narrow-band radio sources occupying much of the spectrum. By consultation with the UK frequency allocation table, it is seen that many of these sources correspond to commercial broadcasting frequencies (OFCOM, 2010). Such sources are referred to as radio frequency interference (RFI), as they



(a) D-PAD single station in 4×4 square configuration



(b) D-PAD single station radome

Figure 5.1: Photos from D-PAD single station testing on Oxford Astrophysics rooftop, April 2012.




Figure 5.2: Typical single station bandpass, as measured with HISPEC and a TTI spectrum analyser. Spikes in the bandpass correspond to terrestrial (or satellite) radio interference sources.

impede our ability to detect astronomical sources. There are a few particularly prominent RFI sources that warrant special attention:

- The peaks at and around 1090 MHz correspond to the Automatic Dependent Surveillance-Broadcast (ADS-B) aircraft tracking system. All commercial aircraft within the European Union are required to transmit navigational information to air-traffic control over ADS-B; we are picking this signal up as the aircraft passes through the beam. Less activity is seen at night, presumably due to decreased air traffic.
- The intermittent signal centred at 1176.45 MHz which can be seen in Figure 5.3 is produced by NAVSTAR Global Positioning System satellites (GPS). As of writing, there are only two GPS satellites which operate at this frequency, and only one is above the horizon at any one time. These satellites were used to measure the station's beam pattern and are discussed below in Section 5.3.1.
- The intermittent signals centred at 1227.6 MHz also correspond to GPS satellites. There are 31 satellites in total that operate at this frequency.
- The large peak at 1394 ± 5 MHz corresponds to a license exempt band which may be

used for audio/video transmission. Given the overwhelming signal strength, it is highly likely that this is produced locally by the Oxford University video surveillance system. In Figure 5.3, this has been blanked.

From the point of view of a radio astronomer, the high level of RFI occupancy is disappointing. In particular, the large peak at 1394 MHz would make the Oxford physics rooftop a poor site choice for a redshifted HI survey. Nonetheless, this affords the opportunity to confirm the interchannel isolation of the HISPEC polyphase filterbank spectrometer.

5.3.1 GPS SATELLITE MEASUREMENTS

GPS is an omnipresent satellite based navigational system operated by the United States of America Department of Defense (gps.gov, 2012). It operates over five frequency ranges, which are referred to as L1, L2, L3, L4, and L5. These are centred at 1575.42, 1227.60, 1381.05, 1379.913, and 1176.45 MHz, respectively. L1 and L2 are the main bands available for general (civilian) usage. L3 is reserved for nuclear detonation detection, and L4 is currently only used in ionospheric studies. The final band, L5, is proposed to carry a civilian ‘safety of life’ signal (Erker et al., 2009). As of writing, only two of the fleet of newer GPS satellites that support this band have been launched.

GPS satellites complete an orbit twice each sidereal day. As the Earth completes a full rotation in one sidereal day, GPS satellites repeat their transit as viewed from Earth once every sidereal day. Ephemeris information for the GPS constellation is available online from the Celestrak website¹. It is supplied in Two-Line Element set format (TLE), which can be read by ephemeris packages such as pyEphem², and satellite tracking programmes such as GPredict³.

Figure 5.4a shows the collected power over the L2 GPS satellite band. To identify which satellites contribute to the measured power level, the trajectory of each satellite as seen from Oxford was computed using pyEphem. Figure 5.4b shows the altitude of GPS satellites over one sidereal day. As can be seen, three satellites pass nearly directly overhead: BIIA-24, BIIR-4, and BIIRM-3. However, there is consistently at least three satellites within the D-PAD single station beam at any time, and it is not possible to differentiate between them (without decoding the

¹<http://www.celestrak.com/NORAD/elements/>

²<http://rhodesmill.org/pyephem/>

³<http://gpredict.oz9aec.net/>






Figure 5.3: D-PAD station bandpass, over a 2 day period. Color scale is power, in dB.



(a) Collected power over the GPS L2 band, over a 24 hour period.



(b) L2 GPS satellite altitudes as observed from Oxford. These data are computed from published ephemeris information.

Figure 5.4: D-PAD observations of GPS L2 satellites. Three L2 GPS satellites come within two degrees of zenith: BIIA-24, BIIR-4, and BIIRM-3.




Figure 5.5: Collected power over the GPS L5 band, over a 24 hour period. BIIF-1 and GPS BIIF-2 are the only two L5 GPS satellites currently operational.

transmissions).

In comparison, there are currently only two operational L5 satellites: BIIF-1 and BIIF-2, and only one of these is above the horizon at any given time (see Figure 5.5). The observed power from these two satellites is also shown in Figure 5.5. The peaks in collected power around 12:00 LST may be caused by an additional undocumented satellite — both BIIF satellites are below the horizon. In the following section, we use these data to probe the D-PAD single station beam pattern.

5.3.2 ASTRONOMICAL SOURCES

Figure 5.6 shows the collected power over the RFI free band 1410-1415 MHz as a function of sidereal time. These data are from a single rotation of the station, with a 10 second integration time. No astronomical sources can be discerned from these data. This is discussed further in Section 5.4.4 below.




Figure 5.6: Collected power over 1410-1415MHz, over a 12 hour period. No astronomical sources can be discerned from the data.

5.4 DISCUSSION

While satellites are in general an unwanted source of RFI for radio telescope, their signals can be used to measure the performance of the array. For example, Ries (2007) used the transmissions from Orbcomm satellites to measure the beam pattern of the PAPER antenna. Similarly, L2 GPS satellites were used by Olofsson et al. (2009), to measure the beam pattern of the EMBRACE aperture array. I have employed similar methods here to measure the beam pattern of the D-PAD array, using the L5 GPS carrier.

5.4.1 BEAM PATTERN

Here, I have used the L5 signal from BIIF-1 and BIIF-2 to make some initial beam pattern measurements of the D-PAD single station. The trajectory of the L5 satellites through the D-PAD beam is shown in Figure 5.7a. The North celestial pole is aligned with 0 degrees in azimuth. While neither satellite passes directly through the zenith, they do pass within the expected main beam, which should have nulls at $\pm\lambda/D = 0.15$ radians, where $D = 160.0$ cm and $\lambda = 25.4$ cm at L5. In Figure 5.7, the observed power (Figure 5.7) has been mapped onto these trajectories.

So that different slices of the beam could be measured, the D-PAD platform was rotated by

0, 45, 90, and 135 degrees ($\pm 5^\circ$) clockwise⁴, yielding 8 slices of the beam in total. While more slices would have been advantageous, this was not possible due to time constraints (the room in which the backend hardware was located was only available for a week).

These data have been combined in Figure 5.8a. To do so, the data were gridded in 8 degree steps, and nearest neighbour interpolation was applied. Only the middle 45 degrees of the beam is shown; outside this field reflections from buildings and metallic structures corrupt the data. This is particularly unfortunate as the first grating lobes are expected at about $\pm 60^\circ$ at 1.17 GHz (see Figure 5.8a). The expected FWHM of the D-PAD array is 18.4° (see §3.2.2); the measurements give a FWHM of $(20.0 \pm 5)^\circ$ with a possible beam offset from zenith of a few degrees in the $-\theta$ direction. Some sidelobes can be seen on the $\theta = 0$ and $\phi = 0$ planes, however much of the expected sidelobe structure of Figure 5.8a is not apparent. This is mainly due to a lack of data: only 8 satellite crossings were measured, so only features that lie upon these beam slices are present. Given a longer observation campaign or a steerable platform, many more data points could be taken to expose the fine grain features of the beam.

5.4.2 SYSTEM TEMPERATURE

The SEFD of a telescope can be estimated from a drift scan of a source of known brightness. For example, in Ellingson et al. (2012), the SEFD of the LWA1 station is computed from drift scans of Cassiopeia A. In Olofsson et al. (2009), a similar technique was used to calculate the system temperature of the EMBRACE array, using the L2 carrier signal of GPS BIIR-02.


In the method of Ellingson et al. (2012), the SEFD is computed from the ratio of power when the source is overhead, P_1 , to the power when the source is absent, P_0 . For a single polarization, this ratio is shown to be

$$\frac{P_1}{P_0} = \frac{\frac{1}{2}kT_{sys} + \frac{1}{2}SA_e}{\frac{1}{2}kT_{sys}} = \frac{S}{SEFD} + 1. \quad (5.1)$$


Rearranging this and solving for SEFD, one finds

$$SEFD = S\left(\frac{P_1}{P_0} - 1\right)^{-1}. \quad (5.2)$$

⁴0 degrees corresponds to alignment of a selected platform's edge with the celestial pole.




(a) Trajectory of L5 GPS satellites as observed from Oxford, UK.




(b) Beam slice at 1176MHz, showing measured power.

Figure 5.7: GPS transit observation with D-PAD single station.



(a) Simulated D-PAD single station array factor at 1.17GHz



(b) D-PAD single station beam pattern, as measured with GPS satellites (nearest neighbour interpolation).

Figure 5.8: Comparison of D-PAD beam pattern with simulated array factor, at 1.17 GHz.

To apply this to the GPS data, we must first compute the flux density of the L5 carrier signal. Assuming isotropic radiation, the relationship between a source's intrinsic luminosity L and its flux density S is

$$S = \frac{L}{4\pi d^2}. \quad (5.3)$$

In this case, d is the distance between the GPS satellite and the D-PAD station. The luminosity (W/Hz) can be approximated from the total radiated power (W) by $L = P_{rad}/\Delta\nu$, where $\Delta\nu$ is the bandwidth of the carrier signal. Combining this with Equation 5.2 yields

$$SEFD = \frac{P_{rad}}{4\pi d^2 \Delta\nu} \left(\frac{P_1}{P_0} - 1 \right)^{-1}. \quad (5.4)$$

Navtech (2012) (among numerous other websites) lists a satellite altitude of $d = 20,183$ km, a bandwidth of ± 12 MHz, and a circularly polarized equivalent isotropically radiated power (EIRP) of 23.0 dBW. While Navtech states that these numbers are based upon the official interface specification (Dunn, 2011), the EIRP is not listed in this document (it may in fact be classified). As such, the actual EIRP may differ significantly. Plugging these values into Equation 5.4 (assuming EIRP is accurate to ± 3 dBW⁵), one arrives at an SEFD of $(1.88 \pm 1.12) \times 10^6$ Jy. Referring back to Equation 1.23, the system temperature is

$$T_{sys} = \frac{SEFD}{2k_B} A_{eff} \quad (5.5)$$

In this case, the factor of 2 in the denominator corresponds to the loss of power incurred by measuring the circularly polarized signal with a single linearly polarized antenna. The effective area for a single antenna is given by

$$A_{eff} = \frac{G\lambda^2}{4\pi}, \quad (5.6)$$

so for the array at 1174 MHz $A_{eff} = 16 \times 0.026\text{m}^2 = 0.42\text{m}^2$. It follows that the measured system temperature lies between 89 ± 53 K, which is in agreement with the 102 ± 12 K value calculated from the measurements taken in Chapter 3.




Figure 5.9: Power spectrum of the gain fluctuation ($\Delta G/G$) of an RFI free channel. Both polarizations exhibit a knee frequency of ~ 11 s.

5.4.3 SENSITIVITY

The sensitivity of the D-PAD single station is limited by the amount of RFI-free channels and the gain stability of the system. For an ideal radiometer the minimum detectable signal is constrained by the rms noise

$$\sigma^2 = \frac{T_{sys}}{\sqrt{2n_p\Delta\nu\tau}}; \quad (5.7)$$

a signal must be larger than this value (typically at least 5 times) for it to be detectable. In practice, systematic errors dictate a maximum integration time over which the noise does not decrease (and may in fact increase). A small change in gain ΔG will cause a change in power

$$\Delta P = \Delta G k T_{sys}, \quad (5.8)$$

which is equivalent to a change in system temperature

$$\Delta T_G = T_{sys} \left(\frac{\Delta G}{G} \right). \quad (5.9)$$

⁵It is also possible that the EIRP is not accurate at all; if so the SEFD and T_{sys} presented here are also inaccurate.

As independent errors add in quadrature, the overall noise level (for a single channel) is given by

$$\sigma_T^2 = \sigma_{noise}^2 + \sigma_G^2 \quad (5.10)$$

$$= \left(\frac{T_{sys}}{\sqrt{\Delta\nu_{chan}\tau}} \right)^2 + \left(\frac{\Delta GT_{sys}}{G} \right)^2. \quad (5.11)$$

The gain stability of a receiver is then specified as the time at which $\sigma_{noise}^2 = \sigma_G^2$. System stability is discussed in further detail in Condon and Ransom (2010).

Figure 5.9 shows the power spectrum of the measured gain fluctuations for a interference free channel ($\Delta\nu_{chan} = 48.8\text{kHz}$). The knee frequency of this plot corresponds to a maximum integration time $\tau = 11.0$ seconds before the variance from gain instability begins to dominate. Taking gain instability into account, the maximum sensitivity of D-PAD is given by

$$\Delta S \approx \frac{SEFD}{\sqrt{2n_p\Delta\nu\tau}} \Big|_{\tau=11.0s}. \quad (5.12)$$

As the noise in each channel is independent only if the channel is not polluted with RFI, here n_{chans} is the number of RFI free channels. Using an SEFD of 1.88×10^6 Jy, with $n_{chans} = 800$ RFI free channels selected from the bandpass (50 MHz), the resultant sensitivity is 40 Jy, and the corresponding minimum detectable signal is then 200 Jy. There are only three astronomical sources with flux above this threshold that are viewable from Oxford: Cassiopeia A, Cygnus A, and the Sun.

The issue of gain instability is not unique to D-PAD, it affects all single station radio telescopes. For this reason, many single dish telescopes employ receiver architectures that makes differential measurements. For example, the pseudo-correlation radiometer compares the sky signal (T_{sky}) to a well calibrated reference load (T_{load}). In this case, the differential measurement between the two introduces a change in system temperature

$$\Delta T_G = (T_{sky} - T_{load}) \left(\frac{\Delta G}{G} \right). \quad (5.13)$$

For well matched sky and load temperatures this becomes negligible. The main drawback of this technique is that it introduces a factor of $\sqrt{2}$ increase in receiver temperature; this is a small price to pay for the insulation against gain variations. A thorough comparison of receiver architectures

is given in §3.3 of King (2009).

A similar improvement can be made by taking interferometric measurements between two or more stations: the gain fluctuations of each station are independent and thus uncorrelated. While this chapter details single station performance, D-PAD is intended as a precursor to an interferometric array. As such, the sensitivity limit presented here may be surpassed if more than one station is used.

5.4.4 ASTRONOMICAL SOURCES

As shown in Figure 5.6, no astrophysical sources were detected during the observation campaign. This is in spite of the fact that D-PAD has the required sensitivity to do so. I address this here.

The three brightest astronomical sources in the sky at 1.4GHz are Cassiopeia A (Cas A), Cygnus A (Cyg A), and the Sun. Cas A is a supernova remnant at a distance of about 11000 light years. It is located in the constellation of Cassiopeia at right ascension, declination ($23^h 23^m 26^s$, $58^\circ 48'$). Cyg A is a radio galaxy located in the Cygnus constellation at right ascension, declination ($19^h 59^m 28^s$, $40^\circ 44' 02''$). The flux densities of Cas A and Cyg A can be modelled by

$$\log_{10}(S_{CasA}) = 5.745 - 0.770 \log_{10}(\nu) \quad (5.14)$$

$$\log_{10}(S_{CygA}) = 4.695 + 0.085 \log_{10}(\nu) - 0.178 \log_{10}(\nu)^2, \quad (5.15)$$

where S is flux in Jansky and ν is frequency in MHz, as shown in the canonical Baars et al. (1977). At $\nu=1.4$ GHz, these formula give $S_{CasA}=2.08$ kJy and $S_{CygA}=1.6$ kJy. As the Sun emits black body radiation, its flux density is given by

$$S = \frac{2kT_{sun}}{\lambda^2} \left(\theta_{sun} \times \frac{\pi}{180} \right)^2, \quad (5.16)$$

where $\theta_{sun} = 0.52^\circ$ is its angular size and $T_{sun} = 5778$ K is its temperature. This evaluates to 29.5 kJy at 1.4 GHz.

During the observation period, the sun rose at 7:00am and set at 5:30pm, reaching a maximum altitude of about 30 degrees. This means that the Sun never reached a high enough altitude to directly traverse through the station's beam. Even so, a portion of its power will be reflected off surrounding buildings and structures, which will cause the collected power to fluctuate in a

complicated way. This confounds observations of Cyg A and Cas A, which reach their maximum altitude at about 9:45am and 1:30pm, respectively. As the station cannot be pointed, one must rely on a drift scan approach to detect the sources. Compared to GPS measurements, only a modest power increase will be seen as the source enters the beam — not enough to stand out over the complex gain variations displayed in Figure 5.6.

The situation could be remedied by moving the station to a site not overshadowed by buildings and metallic structures. Additionally, if the station was positioned on a steerable platform, the station itself could be scanned across the source on shorter timescales. The observations could also be done at a different time of year, when the Sun is below the horizon during the transit of Cas A and Cyg A.

5.5 CONCLUSIONS

Radio frequency interference is an unfortunate reality of radio astronomy. However, in certain cases, we may use these artificial sources of radiation as test signals with which to measure the characteristics of our instruments. In this chapter, the GPS L5 carrier signal was used to both measure the beam pattern of the D-PAD array, and also determine its system temperature. These measurements would not have been possible without the spectral isolation afforded by the polyphase filterbank based HISPEC spectrometer detailed in Chapter 4.

5.5.1 FUTURE WORK

The Oxford physics department rooftop is not an ideal place from which to conduct radio astronomy. Reflections from surrounding buildings limit the field of view, and there are strong RFI sources which occupy the spectrum. If future work on the array was to be done, it would be wise to find a more suitable site. Within Oxfordshire, one might consider Begbroke Science Park, or Wytham woods. Both of these sites are owned by the University of Oxford and are in less populated areas. However, Begbroke is located very close to Oxford airport, so the ADS-B band around 1090MHz is likely to be considerably polluted. The other site, Wytham woods, is used primarily for Zoological research; there is very little support infrastructure. Further afield, the Chilbolton and Jodrell Bank radio observatories would be fitting locations.

The beam pattern measurements could be improved by taking more slices through the beam with more accurate rotation angles. This could be achieved by positioning the array on a steerable platform. To remove the uncertainty of the GPS carrier signal strength, the power of the GPS signal could be measured with a reference antenna with a known SEFD.

CHAPTER 6

THE PARKES HIPSR PROJECT

6.1 INTRODUCTION

The work on the HISPEC spectrometer, detailed in Chapter 4, was motivated by the need for a new wide-bandwidth spectrometer for the Parkes telescope. In this chapter, I discuss the HIPSR (HI-Pulsar) signal processing system which I designed and installed in collaboration with Swinburne University, the International Centre of Radio Astronomy Research (ICRAR), and the Australia Telescope National Facility (ATNF).

HIPSR is a new digital signal processor for the Parkes 21 cm multi-beam receiver, which will provide larger instantaneous bandwidth than the existing backend, as well as providing increased processing power for pulsar observations. Moreover, this new signal processor is able to provide finer spectral resolution, and increased dynamic range through 8-bit sampling.

This chapter begins with an overview of the Parkes 64 m telescope and multibeam receiver, then details the scientific motivation behind the HIPSR project. After this, particulars of the hardware and installation process are presented. This is followed by first-light results; in particular, an observation of the HI-massive galaxy HIZOA J0836-43. The chapter closes with discussion of future work and some concluding remarks.

6.1.1 THE PARKES 64 M TELESCOPE

The CSIRO Parkes 64 m radio telescope is located at the Parkes Observatory, approximately 365 km West of Sydney in New South Wales, Australia. Opened in October 1961, it has been




Figure 6.1: Parkes 64 m telescope (L), and 21-cm multibeam receiver (R). The multibeam receiver is located at the prime focus of the 64 m telescope.

in operation for over 50 years, and remains the third highest cited single-dish telescope (CSIRO, 2011). The Parkes Observatory is located 414.80 m above sea level at latitude $-32^{\circ}59'54.3''$ South, longitude $148^{\circ}15'48.6''$ East. The Observatory is part of The ATNF, which is run by Astronomy and Space Science, a division of the Commonwealth Scientific and Industrial Research Organisation (CSIRO).

The 64 m telescope is equipped with an L-band multibeam receiver, which covers a frequency range of 1230–1530 MHz. The multibeam receiver, located at the prime focus of the Parkes 64 m telescope, consists of 13 cryogenically cooled dual polarization feeds in a hexagonal array, as shown in Figure 6.1 (Staveley-Smith et al., 1996). Each of the 13 beams has two linear polarizations, referred to as ‘pol A’ and ‘pol B’. The average receiver temperature is ~ 21 K, and each beam has a beam width of $14'.4$ (FWHM).

The receiver was designed primarily for redshifted 21-cm HI line observations over $0 \leq z < 0.2$ and was installed in 1996, along with the MBCORR digital signal processing (DSP) system. Although now over 15 years old, MBCORR remains the primary signal processing system for HI spectroscopy with the multibeam receiver.

An additional DSP backend was installed at Parkes in 2008, designed for high time resolution spectroscopy with the multibeam receiver (McMahon, 2008; Keith et al., 2010). This backend was known as the Berkeley–Swinburne Pulsar Recorder (BPSR), and consisted of 26 field programmable gate array based spectrometers connected by 10Gb Ethernet to a cluster of 13

servers, which wrote the spectra to disk. These spectra were then processed offline. As these spectrometers processed the entire usable bandwidth of the receiver at high time resolution, the system proved remarkably capable for pulsar and fast transient work. However, as the BPSR system only produced coarsely chanelized spectra ($\Delta\nu \sim 390$ kHz), it could not be used for HI spectroscopy.

In this chapter, I report on the installation of HI-Pulsar (HIPSR), an upgrade to BPSR which provides new high spectral resolution modes for HI observations. The HIPSR system digitizes the entire usable bandwidth of the multibeam receiver, has increased resilience to radio frequency interference, and provides a significant increase in overall computational power in comparison to MBCORR and BPSR. It is intended that HIPSR will replace MBCORR as an ATNF national facility instrument, once all required modes of operation are reproduced.

6.1.2 SCIENTIFIC MOTIVATION

INTENSITY MAPPING As was introduced in Chapter 1, the baryonic acoustic oscillations (BAO) of the early Universe have left an imprint in the matter distribution of the Universe. Measuring the angular power spectrum of the sky $P(\vec{k})$ through redshifted hydrogen measurements would yield constraints on the value of cosmological parameters. While optical surveys such as SDSS have detected hundreds of thousands of galaxies — enough to detect BAO — HI surveys such as ALFALFA (Giovanelli, 2005) and HIPASS (Meyer et al., 2004) have detected an order of magnitude fewer. Intensity mapping techniques sidestep the need for individual detections, greatly increasing survey speed. The HIPSR backend is advantageous for intensity mapping as it allows the full bandwidth of the 21-cm multibeam receiver to be processed instantaneously. This increases the instantaneous redshift range, and consequently more k_{\parallel} modes of the BAO power spectrum $P(\vec{k})$ can be observed at once. While the prime motivation for this thesis is BAO detection through redshifted hydrogen, an upgrade to the MBCORR backend would be beneficial for all future redshifted hydrogen observations with Parkes.

HI SPECTRAL STACKING A noteworthy new technique which is pushing the redshift boundary for 21-cm HI detection is HI spectral stacking. HI spectral stacking uses a-priori knowledge from optical surveys, such as SDSS, to recover statistical spectral information from an ensemble of




Figure 6.2: A stacked HI spectrum by Delhaize et. al., from data collected with the Parkes multi-beam receiver and 2dFGRS (Delhaize and Staveley-Smith, 2012). A strong statistical detection (black) is achieved by stacking ~ 4000 sources; a reference spectrum (green) and profile width (red dashed) are also shown. Image courtesy of J. Delhaize.

low signal-to-noise radio spectra. In Delhaize and Staveley-Smith (2012), 21-cm emission data from HIPASS is used, along with photometric redshift data from the 2 degree field galaxy redshift survey (2dFGRS Colless et al., 2001), to achieve an 18σ HI detection of galaxies out to a redshift of $z=0.14$. To do so, each radio spectrum is shifted to a common reference using 2dF redshift data, and then the spectra are averaged to form a statistical ‘stacked’ spectrum. An example spectrum is shown in Figure 6.2. This stacked spectrum encodes the average HI mass of galaxies at redshifts up to $z=0.14$, which helps constrain Ω_{HI} , the cosmic HI mass density of the universe. How the HI content of galaxies changes with cosmic epoch is not currently well constrained by observation or simulation.

6.2 INITIAL MODES OF OPERATION

While numerous observation modes are planned, for commissioning we focused on the two first-light modes of operation: a spectrometer for redshifted HI surveys, and a port of the BPSR pulsar recorder to ROACH.

Table 6.1: Comparison of HIPSR with MBCORR 64 MHz mode.

Specification	HIPSR-400	MBCORR
Instantaneous bandwidth	400 MHz	64 MHz
Sampling	8 bits	2 bits
Number of channels	8192	1024
Spectral resolution (Velocity resolution)	48.8 kHz (10.3km/s)	62.5 kHz (13.4km/s)
Implementation	polyphase filterbank (4-tap Hamming)	lag correlator (1024-lag, 3 level)
Inter-channel isolation	~50 dB	~6.8 dB
Polarimetry	Full stokes I, Q, U, V	Autocorrelation only I, Q
Integration length	0.5-10s	2-5s

6.2.1 HIGH SPECTRAL RESOLUTION MODE

The existing multibeam correlator, MBCORR, was commissioned in 1996. MBCORR is a 2-bit, lag autocorrelator. In comparison to MBCORR, the spectrometer firmware detailed in Section 4.6 offers a factor of six increase in bandwidth while still achieving higher spectral resolution, vastly improved inter-channel isolation, and increased dynamic range through 8-bit sampling. To use this spectrometer with the multibeam system, 13 copies of this spectrometer are run in parallel, with each spectrometer digitizing and processing a single beam of the receiver. A direct comparison to MBCORR is shown in the right hand column of Table 6.1. Note that as the spectrometer only needs to process 400 MHz of bandwidth to digitize the entire IF, a sample rate of 800 Msample/s has been used in this comparison.

6.2.2 HIGH TIME RESOLUTION MODE

The High Time Resolution Universe survey (HTRU) is an ongoing pulsar and fast transient survey using the Parkes multibeam receiver on the 64 m telescope. It has already discovered a slew of pulsars, dynamic radio transients and most recently, provided compelling evidence for a planet made out of diamond (Keith et al., 2010; Bates et al., 2011; Bailes et al., 2011; Keith et al., 2012; Burke-Spolaor et al., 2011).

Signal digitization, channelization and packetization for the HTRU is presently performed on first generation CASPER hardware (McMahon, 2008). I have reproduced this functionality on ROACH by porting the Virtex-II FPGA firmware written for BPSR to Virtex-5. While some mod-

est improvements have been made to the FPGA firmware (more polyphase filterbank taps and improved bit precision), the main benefit to the HTRU survey — and to pulsar science with Parkes — is from the increased processing power provided by a new CPU/GPU beowulf cluster which will replace the BPSR servers.

The development of all pulsar processing modes is being lead by Swinburne University, and as such is not discussed in further detail in this thesis.

6.3 NETWORK ARCHITECTURE

The HIPSR system may be described as a multi-node DSP ‘frontend’ connected to a multi-node DSP ‘backend’ by a 10Gb Ethernet (10GbE) switch. The FPGA based DSP frontend conducts signal processing tasks such as signal filtering and channelization; a general purpose GPU (GPGPU) based DSP backend performs higher-level, moderate to low-bandwidth DSP such as running pulsar processing pipelines. A diagram of the system architecture is shown in Figure 6.3.

Packetized DSP systems are becoming increasingly common within radio astronomy, see for example the PAPER correlator (Parsons, 2009), and the VEGAS spectrometer for the Green Bank telescope (Werthimer, 2012). The inclusion of a switch greatly increases the flexibility of the system as data can be dynamically routed between nodes. It does, however, require all data to be packetized along with metadata that describes packet’s data payload. Nevertheless, using a switched 10GbE based packetized architecture instead of defining and commissioning a bespoke backplane drastically decreased development time: all interconnect is off-the-shelf, industry standard, commodity hardware.

6.4 HARDWARE SPECIFICATIONS

The DSP frontend consists of 13 Reconfigurable Open Architecture Computing Hardware boards (ROACH), each of which has a Xilinx Virtex-5 FPGA as its centrepiece. Each ROACH has a 2-input, 1Gsample/sec digitizer mezzanine card installed, which allows for a total of 26 inputs to be digitized (i.e. every IF of the multibeam receiver). The DSP backend consists of 8 server nodes, each of which is equipped with two Nvidia C2070 GPGPUs. This hardware is discussed in further detail below.

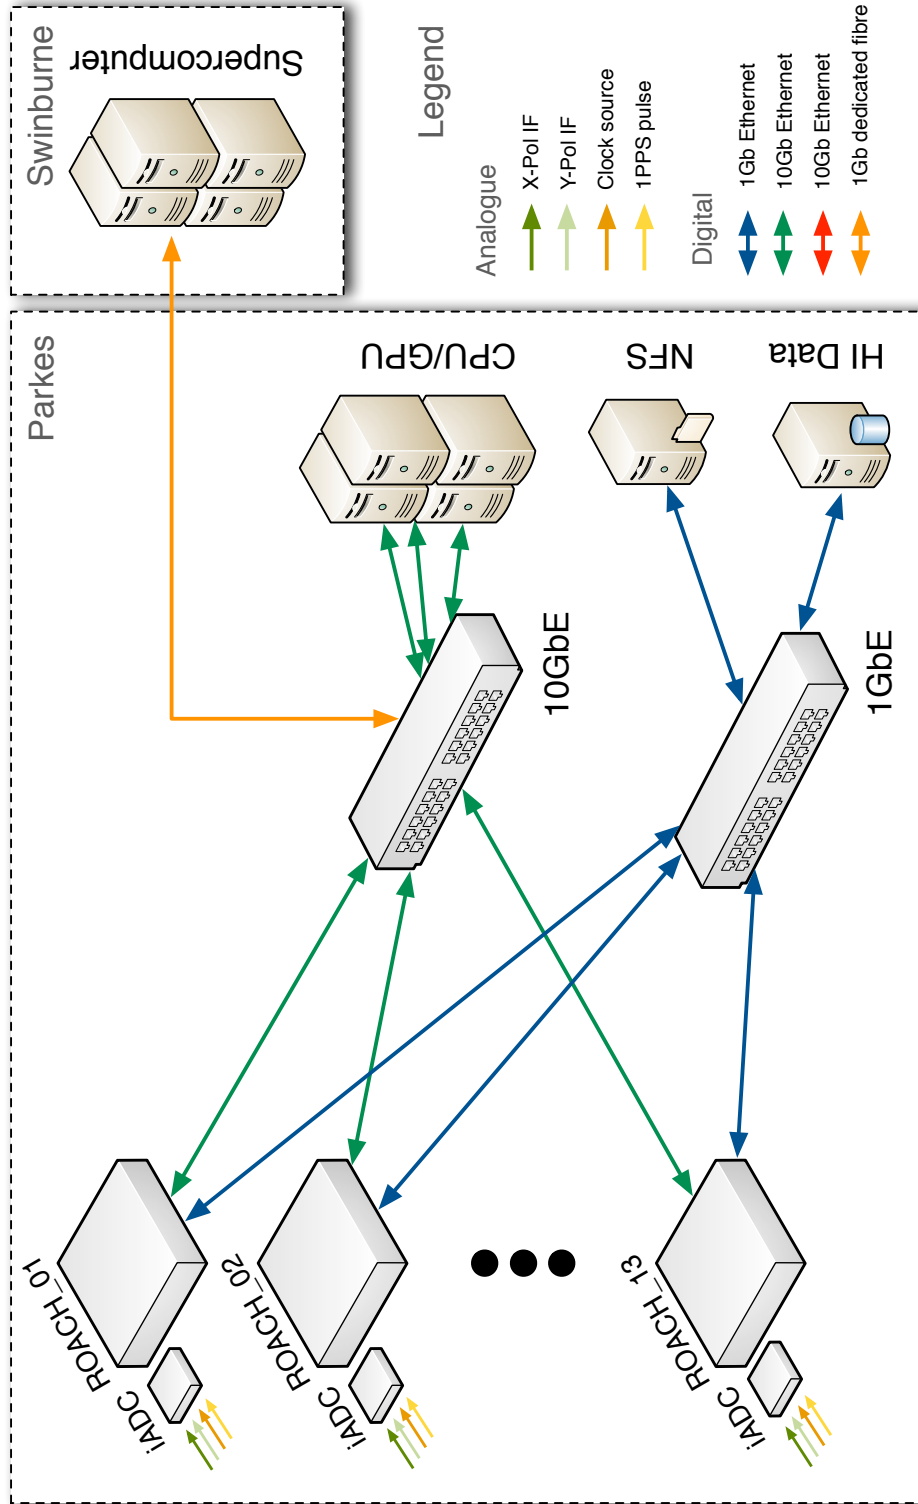


Figure 6.3: HIPSR network diagram. All processing nodes are connected via 10Gb Ethernet. The HIPSR system is connected to a supercomputer cluster at Swinburne University via a dedicated 1 Gb Ethernet fibre-optic connection.

DIGITAL FRONTEND To digitize the 26 IFs of the multibeam receiver, we selected the CASPER iADC digitizer card. Each CASPER iADC can digitize 2 signals, so a total of 13 iADCs are required to digitize the 26 IFs from the receiver. These 13 iADCs are connected to the ZDOK connector of 13 ROACH boards. The ROACH boards process these signals, and then output the processed signals over Ethernet.

NODE INTERCONNECT As is shown in Figure 6.3, the DSP frontend and backend nodes are connected via a 10Gb Ethernet switch. This Cisco 4900M switch¹ is configured with 24 CX4-type ports, 13 of which are connected to the DSP frontend ROACH boards, and 8 of which are connected to the server nodes. The switch allows bi-directional data flow between nodes. The 10GbE switch also has a dedicated 1GbE fibre link to the Green II supercomputer² at Swinburne University in Melbourne. All nodes are also connected by 1GbE, through a Cisco 3750 switch³. For low bandwidth applications where all of the necessary DSP can be conducted on the FPGA, all data may be read off the 1GbE ports of the ROACH boards.

GPGPU SERVER NODES The GPGPU server nodes are bespoke systems, built to specification by Silicon Graphics Pty. Ltd. Each of the 8 server nodes are comprised of dual 2.66 GHz Intel Xeon six-core CPUs, 48 GB DDR3 memory, and dual Nvidia Tesla C2070 GPGPUs. While it is intended that a majority of DSP will be conducted on the GPGPUs, the Intel Xeon CPUs may also be used, if required. As each server has a single CX4-type 10GbE network interface card (NIC), the input data rate for each server is limited to 10 Gb/s. Given that there are 26 IFs which must be processed by 8 server nodes, the DSP frontend must decrease the data rate⁴ from each IF from 6.4Gb/s to about 3.0Gb/s.

DATA SERVER In addition to the GPGPU server nodes, there is a server which provides monitor and control for the DSP frontend, and provides data storage. This server has a 2.4 GHz Intel Xeon four-core CPU, and 24 GB DDR3 memory. As it is not used for DSP, it does not have a GPGPU. This server runs a DHCP daemon and provides a network file system (NFS) to the ROACH boards.

¹<http://www.cisco.com/en/US/products/ps9310/index.html>

²<http://astronomy.swin.edu.au/supercomputing/green2/>

³<http://www.cisco.com/en/US/products/ps7077/index.html>

⁴digitized data rate is $2 \times 400\text{MHz} \times 8\text{bit} = 6.4\text{Gb/s}$

6.5 MONITOR, ACQUISITION AND CONTROL SOFTWARE

During observations, data from all 13 ROACH boards must be collected, along with metadata from the TCS. The KATCP protocol (Cross et al., 2009) is used to communicate with ROACH boards, while TCS uses a TCP/IP ASCII keyword based system. A multithreaded python script receives these data, collates them and then writes them to file. In addition, the python server script forwards processed data on to a graphical user interface for real-time monitoring.

6.5.1 DATA STORAGE

Data are stored in Hierarchical Data Format⁵ 5, which offers numerous advantages over other file formats, such as FITS. Of particular regard to radio astronomy data, HDF5 is designed to handle very large (>1TB) datasets, offers on-the-fly compression using SZIP, supports parallel I/O, and contains optimizations for random access patterns. For these reasons, HDF5 is gaining traction in the radio astronomy community: it is used by LOFAR and the Medicina BEST-2 telescope for data storage (Anderson et al., 2010; Foster et al., 2012), and will likely be used in MeerKAT (J. Manely, personal communication).

6.5.2 USER INTERFACE

Data can be monitored in real-time using a graphical user interface (GUI). The GUI is built using Qt⁶, Python⁷ and PySide⁸. The Matplotlib python library⁹ is used for plotting data. The GUI has three main displays: an overview showing the bandpass of all 13 beams, an plot of overall power versus time, and a window which displays the bandpass of a single selected beam. This is shown in Figure 6.4; additional data visualizations will be added in future versions of the GUI.

6.6 INSTALLATION

The HIPSR hardware was installed over the period 11–22 March 2012. The HIPSR hardware was installed into RFI shielded racks R and S (see Figure 6.5), located on the first floor of the telescope

⁵<http://www.hdfgroup.org/HDF5>

⁶<http://qt.nokia.com/>

⁷<http://www.python.org/>

⁸<http://www.pyside.org/>

⁹<http://matplotlib.sourceforge.net/>

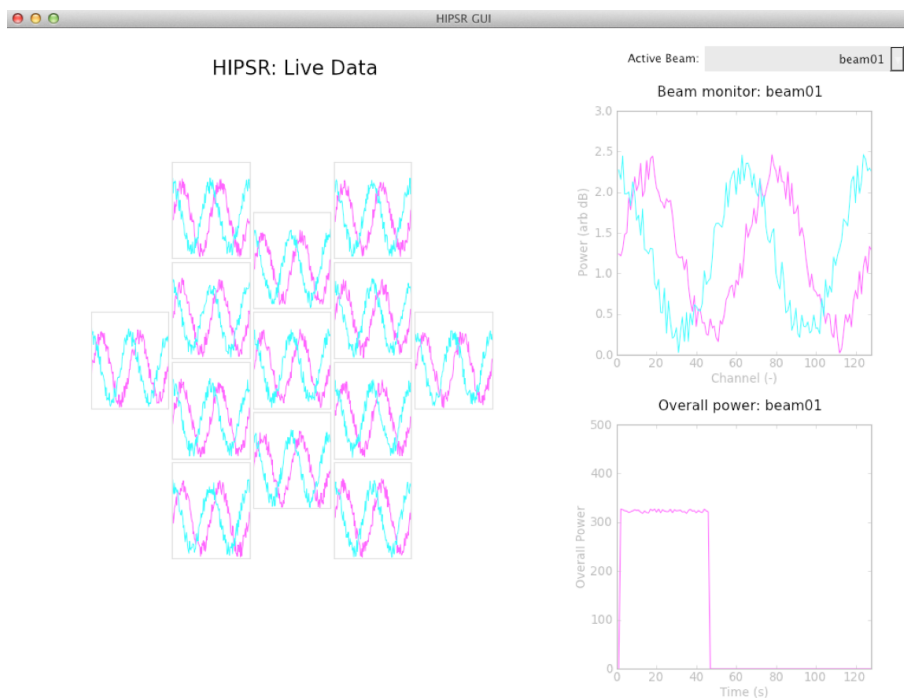



Figure 6.4: HIPSR GUI running on Mac OS X. The bandpass for all 13 beams is displayed on the left; the right hand side displays the overall power and bandpass of a selected beam.

building (this room sits directly underneath the telescope dish). Each of the 13 beams has two linear polarizations, giving a total of 26 IFs. A set of BNC cables were laid under the floor from the multibeam IF distribution panel to BNC feedthroughs installed on the racks, see Figure 6.6. A further set of cables connect the BNC feedthrough to the SMA input of the iADC cards. A third set of cables connects the ‘sync’ input of the iADC to a 1PPS distribution unit, and a clock signal is distributed to each board by a fourth set of SMA cables.

Each server node and ROACH board is connected to a 1GbE switch by an Ethernet cable (CAT-5E STP). Similarly, each server node and ROACH is connected to a 10GbE switch via CX4 type cables.

6.6.1 COMMISSIONING TESTS

After hardware installation and configuration, a series of tests were conducted to verify the ROACH boards were fully operational. To test the FPGAs were programming correctly, a simple diagnostic bitstream that flashes the front panel LEDs at varying rates was loaded. This firmware uses the internal FPGA clock for timing. Following this, a similar test was run, but with a bitstream using the external clock and the 1PPS signal. This verified that the clock and 1PPS were func-



(a) HIPSR hardware, rack R

(b) HIPSR hardware, rack S

Figure 6.5: HIPSR hardware installed into Parkes 64 m telescope RFI shielded racks.



Figure 6.6: BNC cabling from multibeam receiver to HIPSR cabinets. Cables are colour coded for easier identification.

tioning correctly, and that the boards could be synchronized. After these commissioning tests, we conducted our first-light observations.

6.6.2 FIRST LIGHT BANDPASS

During first light observations, the telescope was pointed at Zenith (i.e. stowed); there were no continuum sources of note in the field of view during this time, so the sky may be considered a cold load with $T_{sky} \sim 5$ K. The spectrometer firmware detailed in Section 4.6 was used to digitize the full 400 MHz of bandwidth available from the multibeam receiver. The bandpasses for all 13 beams are shown in Figure 6.7. In this observation, a 2 second integration time, which corresponds to 400×10^6 FPGA clock cycles. A mixing frequency of 1195 MHz was used to convert the 1195-1595 MHz band down to 0-400 MHz. This mixing frequency was chosen so that HIPSR could be tested without affecting MBCORR observations. The first 50MHz of the downmixed band is suppressed by a lowpass filter which was installed to mitigate strong interference from the Thuraya satellite.

As is seen in Figure 6.7, there are considerable differences in total power between IFs; for example, polarization B is ~ 5 dB lower than polarization A in the central beam. This will be remedied for future observations by installing attenuators to achieve a better match (< 1 dB between polarizations per beam, < 3 dB between beams). This aside, the total power levels are still appropriate: the 8 bit ADC gives ~ 48 dB of dynamic range, and all signals are well above the noise

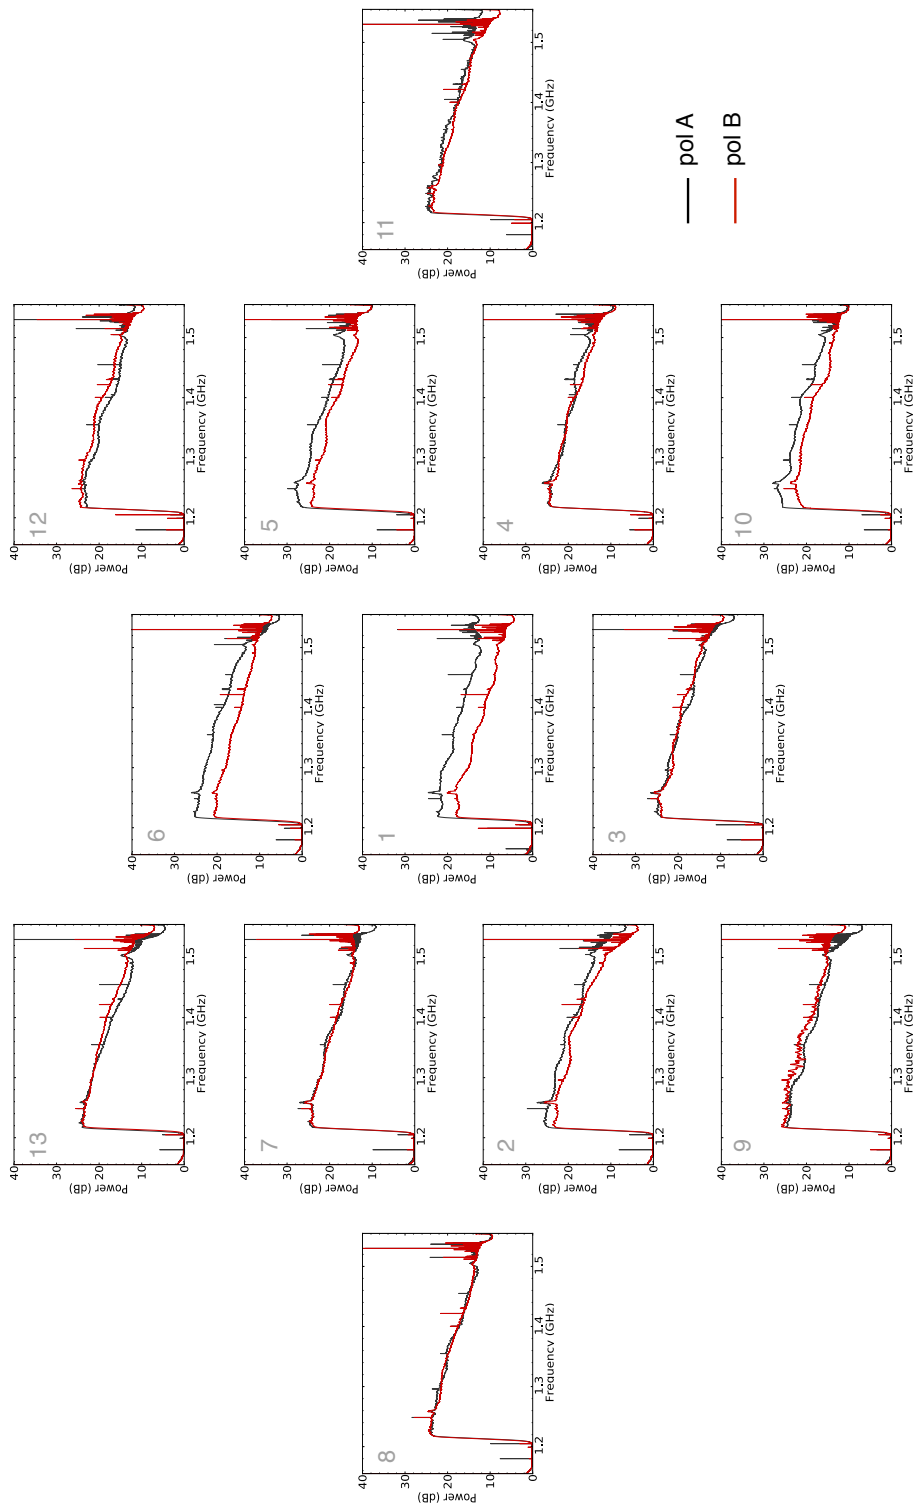


Figure 6.7: HIPSr bandpass for all 13 beams. An integration time of 2 s was used. All data was taken simultaneously from 13 parallel copies of HISPEC, running on the HIPSr ROACH boards.




Figure 6.8: Example of HIPSR and MBCORR response to narrowband RFI. The (unsmoothed) data output from MBCORR is visibly affected by Gibbs ringing.

floor and well below levels that would cause clipping.

The signal from galactic HI can be clearly seen at 1420.4 MHz; other narrowband spectral features, such as those at ~ 1200 MHz, are from man-made sources. A zoomed plot of a narrowband RFI source in the central beam is shown in Figure 6.8 and is juxtaposed against a parallel observation made with MBCORR. The comparison highlights the superior performance of HIPSR for spectral line work, and its resilience against strong signals.

Figure 6.9 shows the uncalibrated Stokes I, Q, U and V parameters for an observation of Hydra A (see below). As the feeds are linearly polarized, the Stokes parameters are given by

$$I = AA^* + BB^* \quad (6.1)$$

$$Q = AA^* - BB^* \quad (6.2)$$

$$U = 2\text{Re}\{AB^*\} \quad (6.3)$$

$$V = 2\text{Im}\{AB^*\}. \quad (6.4)$$

Where $*$ denotes the complex conjugate. The power mismatch between polarizations A and B is reflected in the non-zero power of Stokes Q, and the non-zero values seen in the U and V correspond to the real and imaginary components of the cross correlation AB^* . Polarimetric observations and calibration studies will be conducted at a later date, after the input power of each polarization is better matched with a set of attenuators.

6.7 SPECTRAL LINE OBSERVATIONS

To show the viability of the HIPSR system for HI surveys, this section presents an observation of a faint (~ 50 mJy) HI massive galaxy HIZOA J0836-43, and of a known calibrator source, Hydra A. The observations of Hydra A are used to deduce a set of calibration coefficients, which are then applied to the HI profile of J0836-43. This is then compared to the original detection of J0836-43 from archival data, and shown to be in agreement.




Figure 6.9: Stokes parameters as a function of frequency for observation of Hydra A.

6.7.1 CALIBRATION METHOD

Position switching is a common calibration technique employed in single dish radio telescopes, in which an on-source spectrum is taken then differenced with an off-source spectrum. The difference in power is then proportional to the increase in system temperature. While the technique is discussed in introductory texts such as O’Neil (2002), I have derived a more general result here that explicitly shows dependence on antenna gain.

To begin, I shall introduce the formal definitions of beam solid angle, the antenna temperature and the system temperature as defined in §2-5.5b, §12-1.8, and §12-2.1 of Kraus and Marhefka (2002). The first of these — the beam solid angle — is defined as

$$\begin{aligned}\Omega_A &= \int_0^{2\pi} \int_0^\pi R_n(\theta, \phi) \sin\theta d\theta d\phi \\ &= \iint_{4\pi} P_n(\theta, \phi) d\Omega.\end{aligned}$$

where $R_n(\theta, \phi)$ is the normalized antenna power pattern; and $d\Omega = \sin\theta d\theta d\phi$ is an infinitesimal element of solid angle. This is approximately given by the half power beamwidths in the two principle planes, $\Omega_A \approx \theta_{HP}\phi_{HP}$. The antenna temperature depends on the sky:

$$T_A = \frac{1}{\Omega_A} \iint_{4\pi} T_{sky}(\theta, \phi, \nu) R_n(\theta, \phi, \nu) d\Omega, \quad (6.5)$$

where $T_{sky}(\theta, \phi)$ is the sky brightness temperature distribution. For a homogeneous sky with temperature T_{BG} , the integral simplifies to

$$T_A = T_{BG}. \quad (6.6)$$

On the other hand, if our sky consists only of a point source (a source of angular extent Ω_S much smaller than the beam angle), then

$$T_A \approx \frac{\Omega_S}{\Omega_A} T_S. \quad (6.7)$$

The overall system temperature may then be found using

$$T_{sys} = T_A + T_{AP}\left(\frac{1}{\varepsilon_1} - 1\right) + T_{LP}\left(\frac{1}{\varepsilon_2} - 1\right) + \frac{1}{\varepsilon_2} T_{RX}, \quad (6.8)$$

as was discussed in §3.7.3. Absorbing the T_{AP} and T_{LP} factors into T_{rec} (or assuming perfect efficiency), we can rewrite this as

$$T_{sys} = T_A + T_{RX}. \quad (6.9)$$

If we point the telescope at a point source with temperature T_S and assume a static background $T_{BG} \sim 5K$ elsewhere, the system temperature on source is given by

$$T_{A-ON} = \frac{1}{\Omega_A} \iint_{4\pi} T_{sky}(T_S + T_{BG})P_n(\theta, \phi, \nu)d\Omega, \quad (6.10)$$

$$= T_{BG} + \frac{\Omega_S}{\Omega_A} T_S \quad (6.11)$$

Similarly, when off source we find $T_{A-OFF} = T_{BG}$.

Position switching works by taking an on source and off source power measurement. The system temperature for on and off source are given by

$$T_{ON} = T_{A-ON} + T_{RX} \quad (6.12)$$

$$T_{OFF} = T_{A-OFF} + T_{RX} \quad (6.13)$$

and correspondingly the power measured in a single channel of bandwidth $\Delta\nu$ is given by

$$P_{ON} = G_{REC}k\Delta\nu T_{ON} \quad (6.14)$$

$$P_{OFF} = G_{REC}k\Delta\nu T_{OFF} \quad (6.15)$$

where G_{REC} is the total gain of the receiver system. The difference between these two measurements is

$$P_{ON} - P_{OFF} = G_{REC}k\Delta\nu(T_{ON} - T_{OFF}) \quad (6.16)$$

$$= G_{REC}k\Delta\nu \frac{\Omega_S}{\Omega_A} T_S. \quad (6.17)$$

Dividing through by P_{OFF} yields

$$\frac{P_{ON} - P_{OFF}}{P_{OFF}} = \frac{T_{A-ON} - T_{A-OFF}}{T_{A-OFF} + T_{RX}} \quad (6.18)$$

$$= \frac{\Omega_S T_S}{\Omega_A T_{BG} + T_{RX}} \quad (6.19)$$

$$= \frac{\Omega_S T_S}{\Omega_A T_{sys}}. \quad (6.20)$$

To relate this to antenna gain, we note that directivity is given by

$$D = \frac{4\pi}{\Omega_A} \quad (6.21)$$

and antenna gain G_A is given by $G_A = \epsilon_A D$, where $0 \leq \epsilon_A \leq 1$ is the antenna efficiency factor. Noting that these factors change as a function of frequency, we may rewrite Equation 6.20 above as

$$\frac{P_{ON} - P_{OFF}}{P_{OFF}} = \frac{\Omega_S G_A(\nu)}{4\pi \epsilon_A(\nu)} \frac{T_S(\nu)}{T_{sys}(\nu)} \equiv \beta(\nu) \frac{T_S(\nu)}{T_{sys}(\nu)}. \quad (6.22)$$

The result usually encountered (e.g. O'Neil (2002) and Winkel et al. (2012)) neglects the $\beta(\nu)$ term of 6.22, although the authors note that standing wave ripples are not removed when this technique is applied to continuum sources. For narrowband receivers, the gain and efficiency are likely to be reasonably constant, so this additional factor may not be significant. However, this is not the case with wide bandwidth backends, such as HIPSR. This is seen in the next section, which introduces the flux scale calibration.

6.7.2 FLUX SCALE CALIBRATION

For flux scale calibration, I observed Hydra A (3C218). Hydra A is a continuum source located at right ascension and declination ($09^h 18^m 06^s$, $-12^\circ 05' 45''$). The flux from continuum sources such as Hydra A, can be modelled by the power law relationship $S \propto \nu^\alpha$ between the flux density S and a quantity known as the spectral index, α . For Hydra A the flux density at 1415 MHz is 40.6 Jy, and the spectral index $\alpha = -0.83$ (Lane et al., 2004; Reynolds, 2012).

For each of the 13 beams, on-source and an off-source data were taken, using 10 s of on- and off-source data. The position switching technique detailed above was then used to compute the quantity $(P_{ON} - P_{OFF})/P_{OFF}$ for each frequency channel. An overall spectrum




Figure 6.10: Hydra A measured power (black, left axis); compared to Hydra A model (red, right axis).


was then formed from the average of all 13 position switched spectra; this technique is known as beam switching and is commonly employed at Parkes. The shown in Figure 6.10, along with a theoretical model of Hydra A computed using the spectral index and flux above. The second y-axis (Hydra A flux) has been positioned such that the 1415 MHz flux value of 40.6 Jy is coincident with the HIPSr data. The required flux calibration coefficient for each channel are then found by dividing the model by the measured Hydra A power ratio. That is,

$$\text{cal} = \frac{\text{Hydra A model flux}}{\text{HIPSr measured (ON - OFF)/OFF}} \quad (6.23)$$


These calibration coefficients are applied to an observation of a HI massive galaxy, HIZOA J0836-43, in the next section.

6.7.3 HIZOA J0836-43

The first extragalactic HI source to be observed with HIPSr was HIZOA J0836-43. HIZOA J0836-43 is a HI-massive disc galaxy lying in the Zone of Avoidance, at right ascension and declination ($8^h36^m51^s$, $-43^\circ37'41''$). It has a redshift $z = 0.036$ and correspondingly lies at




(a) Noise as a function of time for the J0836-43 observation. The dotted lines delineate the contribution from each sequential beam. A small increase in noise is seen when beam 11 is added.




(b) Plot of noise for all 13 beams (pol A). Beam 11A exhibits worse noise performance than other beams; such an outlier is not seen in polarization B.

Figure 6.11: Noise levels for the J0836-43 observation (standard deviation from average for over an interference free 512 channels).



(a) Calibrated HIPSR data with HIZOA archival data overlaid.



(b) HI Profile of HIZOA J0836-43

Figure 6.12: HIPSR first light observations of HIZOA J0836-43.

a distance of 148 Mpc (Donley et al., 2006). It was first detected (unresolved) in the HI Zone of Avoidance Survey (Kraan-Korteweg et al., 2005), and remains one of the most HI-massive galaxies detected to date. The HIZOA survey was conducted using the Parkes multibeam receiver and MBCORR, and the original archival data is still readily available¹⁰. As such, it seemed a fitting source to test the HIPSR system.

The source was observed for 100 s in all 13 beams. The bandpass curve was removed by using position switching between on and off source pointings of each beam. The Hydra A based calibration detailed above was used to further calibrate the data and to set the flux scale. The data from all 13 beams was then combined to form an overall spectrum.

Higher signal to noise ratios are achieved by longer integration times. Delhaize and Staveley-Smith (2012) confirmed that the signal-to-noise ratio for decreased proportional to \sqrt{t} for up to 60 days using the multibeam receiver in conjunction with MBCORR. To confirm that the signal-to-noise ratio decreases as \sqrt{t} for HIPSR, an interference free 512 channels was selected from the data and the average and standard deviation were computed as each beam was added sequentially. The resulting noise levels are plotted in Figure 6.11a. A small increase in noise can be seen at $t = 1100$ s, when beam 11 data are added to the integrated spectrum. This was traced to unusual behaviour of polarization A of beam 11, as shown in Figure 6.11b. This is likely due to an issue with the analogue chain or the ADC response and will be investigated further. For the data reduction presented here, beam 11A was not included.

In a final calibration step, the average power level of the featureless (flat) areas of the spectrum was computed and subtracted to bring the average to 0 Jy. The final calibrated HI profile of HIZOA J0836-43 is shown in Figure 6.12a and 6.12b, along with archival data from the HIZOA catalogue. The peak flux of the HIPSR data is 55 mJy, while the HIZOA archival data has a peak flux of 47 mJy. The noise level in HIPSR spectrum is 7.9 mJy, and the HIZOA spectrum exhibits a 7.85 mJy noise level. The coefficient of determination between the two lines is $R^2 = 0.74$, which indicates a high degree of correlation between the observations.

One would not expect the flux for the two observations to match exactly, due to different systematics and flux scales used in the calibration method. The HIZOA data reduction removes baseline distortion caused by the Sun and other continuum source by fitting and subtracting a

¹⁰<http://www.atnf.csiro.au/research/multibeam/release/>

‘template distortion spectrum’ (Kraan-Korteweg et al., 2005; Barnes et al., 2001). In contrast, the HIPSR data is calibrated only off a single Hydra A observation coupled with a secondary power calibration based on the average power of featureless areas of the spectrum. Without this secondary calibration, the HIPSR data would exhibit an offset of 30 mJy; a similar offset would likely be seen in the HIZOA data before template fitting. This notwithstanding, it is clear that the HIPSR system produces high quality spectra for faint HI sources, boding well for redshifted hydrogen surveys.

6.8 FUTURE WORK

The ATNF have identified HIPSR as a possible future national facility, and there is much work which must be done to attain this status. This section identifies areas of future work that would extend the HIPSR system to its full potential.

IDENTIFYING PERFORMANCE ISSUES As mentioned above, there are a few minor issues with the HIPSR system that must be addressed. Firstly, the input power levels to the digitizer cards needs to be optimized. Secondly, the source of beam 11A’s performance issue must be identified and fixed.

SPECTROMETER ENHANCEMENTS The multibeam receiver has a calibrated noise source which can be injected after the orthomode transducer (part of the antenna feed). This calibrated source can be switched on and off at ~ 100 Hz to inject a known signal on top of the sky signal. The calibration signal is then used as a noise adding radiometer (NAR), to accurately measure the system temperature by monitoring the power difference between on and off states (Brunzic, 1988). NAR techniques allow for greater calibration accuracy to be achieved, at the expense of the system temperature (increasing T_{sys} by 5–10%).

To support NAR modes, the calibration signal must be synchronized with the ROACH boards. This could be achieved by using the GPIO pins of a master ROACH to control the cal signal, and using the 1PPS signal to synchronize all other ROACHes to within one clock cycle (5ns). In a similar fashion a frequency switching scheme could also be employed using the GPIO voltages to select between mixer frequencies.

REPLACING MBCORR FUNCTIONALITY MBCORR has several modes of operation with different numbers of channels and different bandwidths. To reproduce similar modes with HIPSR would require new firmware to be written, or for a set of analogue filters to be installed. HIPSR currently digitizes the 0–400 MHz, this can be lowered by simply decreasing the clock frequency, but anti-aliasing filters would be required to match the decrease in bandwidth: a set of 26 filters for each band.

A more practical option is to filter the signal digitally. With a few modifications to the firmware, one could apply a decimating FIR filter before the filterbank to decrease bandwidth and consequently increase spectral resolution. An alternative implementation would be to apply a coarse FFT to the incoming data, then select a sub-band to further channelize; this is known as a ‘zooming’, or ‘zoom’ spectrometer.

HETEROGENEOUS SIGNAL PROCESSING As FPGAs excel at high data bandwidth signal processing, the HIPSR wideband spectrometer does not require any of the >16 Tflops of processing power possessed by the CPU/GPU server. For lower bandwidth applications, the flexibility afforded by HIPSR’s heterogenous architecture may be leveraged. An attractive implementation of the zooming spectrometer would see coarse channelization conducted on the ROACH boards, before the data is sent to the CPU/GPU servers via 10GbE for further channelization. It is likely that sub-Hz resolution can be achieved over >10 MHz bandwidths using such an architecture. RFI excision algorithms — for example, Nita and Gary (2010), Kocz et al. (2012) — may be run on the data in real-time, decreasing the amount of data lost to interference.

6.9 CONCLUSIONS

This chapter presents preliminary results from the HIPSR wide band spectrometer. HIPSR has been tested on the astrophysical source HIZOA J0836–43, and the result is in agreement with archival data from the HIZOA survey. HIPSR has been selected to become a national facility instrument, meaning that it will be supported and maintained by the ATNF. The flexible, heterogeneous design of HIPSR, coupled with its increased processing power over existing systems, will allow it to replace these aging systems and for new modes of operation to be realized.

CHAPTER 7

CONCLUSIONS

The HI line is a wonderful tracer of matter in the Universe. While we have detected thousands of galaxies with the HI line, we have not yet sampled enough to detect the baryon acoustic oscillation signature in the spatial power spectrum of the Universe’s matter distribution. With the Square Kilometre Array scheduled for full operation in 2024, we are only a decade away from having an instrument capable of directly detecting billions of galaxies through HI emission. Such a large survey would be able to detect the BAO signature over a large redshift range, and from this data we will be able to constrain the cosmology of the Universe.

In scientific terms, a decade is a long time. In the interim, 21-cm intensity mapping instruments will likely make the first BAO detection. Highly redshifted HI arrays such as MWA, PAPER and LOFAR are poised to detect the EoR within a few years, but no dedicated HI BAO arrays at lower redshift have been built. This thesis presents instrumentation towards this effort.

In addition to this primary motivation, the instrumentation developed in this thesis has more general applications for radio astronomy. Research into wide-bandwidth signal processing systems is a vital step toward next-generation telescopes, and many existing instruments are hampered by their existing signal processing systems. For example, this is true for all of the current SKA pathfinder instruments — LOFAR, MWA, ASKAP, and MeerKAT — none of these instruments are capable of processing the full bandwidth of their receiver elements. The work presented here is a step forward in the quest for ever-increasing bandwidth DSP systems, with which scientific discovery may be advanced.

7.1 DISCUSSION

The instrumentation effort in this thesis is motivated by HI line science. The sparse aperture array research within the D-PAD project dovetails with the wide-bandwidth spectrometer work motivated by the HIPSR project. As discussed in Chapter 1, a 21-cm intensity mapping instrument can either sample the entire (u, v) plane to maximize surface brightness sensitivity, or targeted (u, v) modes can be observed with an interferometric array.

LOW REDSHIFT REGIME At lower redshift ($z < 0.5$), single dish instruments such as the Effelsberg 100 m and Parkes 64 m telescopes have enough resolution to detect the BAO using a filled aperture intensity mapping technique. The 21-cm multibeam receiver of the Parkes 64 m telescope has 13 fields of view, so is afforded a fast survey speed. The HISPEC spectrometer (Chapter 4) running on the HIPSR backend (Chapter 6), will more than double this survey speed by increasing the instantaneous bandwidth. This upgrade makes the Parkes 64 m telescope a very attractive 21-cm intensity mapping instrument, although its receiver limits observations to $z < 0.2$. I discuss this further in the ‘Future Work’ section below.

HIGH REDSHIFT REGIME At high redshift, it is prohibitively expensive to build a filled aperture instrument like a single-dish telescope, so our only option is interferometric measurement. We are fortunate to be able to build upon the legacy of the EoR arrays: using similar calibration techniques, antenna configurations, and data reduction pipelines. So far, BAOBAB is the only proposal which seeks to use an EoR array inspired design. BAOBAB plans to leverage the design and calibration techniques of the PAPER array. However, a very large N - correlator would be required. The D-PAD project (Chapter 3 and Chapter 5) represents a BX-type approach, similar to the MWA.

RFI AND GPS CONSIDERATIONS In the D-PAD project, I worked on instrumentation for sparse aperture arrays at 1.0–1.5 GHz, which corresponds to redshifts of $0 < z < 0.42$. In the UK, the radio spectrum below 1.0 GHz has many particularly strong sources of RFI, such as mobile phone communication and television broadcasting. The tests of the D-PAD station on Oxford physics building’s rooftop garden reveal a highly occupied spectrum over the 1.0–1.5 GHz band. Fortu-

nately, the PFB structure of HISPEC spectrometer has significantly better interchannel isolation than FTF filterbanks, so portions of the band are still usable. The rooftop testing campaign also showed that the L5 signal from GPS satellites can be used to map out the beam pattern of an array, which may be useful for beam pattern measurements of future BAO arrays. While the L5 signal is indeed a source of interference, having a constellation of calibration sources may prove to be useful, so it should not be rejected hastily.

7.2 FUTURE WORK

The HIPSR and D-PAD projects are both steps toward HI BAO detection. In this section, I discuss future work which could build upon the work presented in this thesis.

7.2.1 IMPROVEMENTS TO THE HISPEC SPECTROMETER

As discussed in Chapter 6, the main improvement for the HISPEC spectrometer is to add support for noise adding radiometer modes of operation. This is the only additional functionality required for the 21-cm intensity mapping proposal by Li et al. (2012). The HIPSR system has a significant amount of GPU based processing power available that HISPEC does not leverage. An exciting area of further research is exploring the possible ways in which the GPU processing can improve data quality (e.g. RFI rejection algorithms), or produce new data products (e.g. zooming spectrometer modes).

Another interesting possibility is modifying the HISPEC firmware so that a portion of the spectrum is passed to the GPU cluster and run through a transient search algorithm. This could allow transient searches in a commensal manner, which would be advantageous for detecting low probability transient phenomena, where pointing direction is unknown until the event.

7.2.2 BAO DETECTION WITH HIPSR

Now that the entire bandwidth of the Parkes 21-cm multibeam receiver can be processed, a 21-cm intensity mapping experiment with the Parkes 64 m telescope is now viable. In October 2012, Li et al. (2012) performed a large scale structure survey of the 2 Degree Field (2dF) using HIPSR. These data will be cross-correlated with optical data from the 2dF galaxy redshift survey

(Colless et al., 2001), in order to make a statistical detection of HI large scale structure. If this cross-correlation approach proves successful, the next step will be to perform a longer survey (to decrease noise), and make the detection using only radio data. Li et al. (2012) suggest that with 1500 hours of observations, Parkes — with HIPSR — could perform a volume limited survey over 15,000 square degrees of the southern sky, giving it a similar sensitivity to large scale structure as the Sloan Digital Sky Survey. In principle, such a survey could make a 3.5% measurement of the BAO distance scale.

7.2.3 A D-PAD BASED POWER SPECTRUM INTERFEROMETER

Given RFI considerations, the D-PAD components were designed for the 1.0–1.5 GHz band. At a remote site — such as the two SKA sites in Western Australia and South Africa — RFI is less of a concern, so a larger portion of frequency space may be explored. If the D-PAD antenna was frequency scaled from a central frequency of 1.25 GHz to a central frequency of 1.0 GHz, it would then span 0.8–1.2 GHz (maintaining the same fractional bandwidth). This band is of more interest, as it probes a higher redshift range ($0.8 < z < 0.2$) than that which is possible with the Parkes 21-cm multibeam system ($z < 0.2$). The sky temperature remains low, and the HI brightness temperature is comparable over this modified frequency range. The D-PAD receiver board already works well over this range (although could be further optimized), as do the second stage amplifiers. The bandpass filter would of course have to be changed, but the same design procedure could be used. A multi-beam Wilkinson power divider beamformer could be constructed at reasonable cost by extending the work of Hickling (2010).

A D-PAD based power spectrum interferometer (henceforth BAO-PAD) would be similar in scope and implementation to the BAOBAB proposal. If the aperture array station size was kept small, the array would be reconfigurable and could employ the same interchangeable minimum- and maximum-redundancy array configuration approach. BAO-PAD would have a larger collecting area per station than a BAOBAB antenna, while maintaining the same field of view, which would increase survey speed. The correlator's computational requirements scale with the number of antennas, so a BX-type array with an analogue beamformer is potentially much cheaper than the equivalent correlator-only array with the same overall number of antennas. The CASPER hardware detailed in Chapter 4 would be an ideal platform upon which to deploy the BAO-PAD

correlator, as the CASPER collaboration maintains a scalable packetized correlator architecture (Parsons et al., 2008).

7.2.4 BAO-PAD

The Oxford astrophysics radio astronomy DSP group has enough CASPER hardware with which to build a 16-input, 400 MHz bandwidth FX correlator. There are eight iADC boards left over from the 2-PAD project (Armstrong, 2012), which is enough to digitize the signals from eight dual polarization stations. There are also eight ROACH boards available, and one ROACH-2. Each ROACH board can channelize and packetize the signal of one iADC, before outputting them on their 10GbE port. The data output rate of each ROACH is then

$$800 \text{ Msamples/input/s} \times 2 \text{ inputs} \times 8 \text{ bits/sample} = 6.4 \text{ Gb/s},$$

assuming 4-bit complex data output. The ROACH-2 has eight 10GbE ports, so a switch is not required; the total input data rate is 51.2Gb/s. Data output would be over the 1GbE port of the ROACH-2.

The ROACH-2's FPGA has 2016 DSP48E multipliers, and can comfortably run at FPGA clocks speeds of up to 250 MHz. Recent work by Hickish et al. (2012) allows a single DSP48E block to compute a 4-bit complex multiply; this has been incorporated into the CASPER 'X-engine' (see Parsons et al., 2008). The total number of DSP48Es required can be estimated by

$$N_{ops} = N_{XE} N_{taps} N_{stokes}, \quad (7.1)$$

where $N_{XE} = 10$ is the number of X-engine kernels running in parallel (based on input data rate); $N_{taps} = N_{ant}/2 = 16$ is the number of taps in each X-engine, N_{ant} is the number of single polarization inputs, and $N_{stokes} = 4$ is the number of Stokes parameters computed (Foster, 2012). Using these values, a total of 640 DSP48Es are required: under 33% of the available resources.

If another ROACH-2 was procured, one could build a 32-input, 400 MHz bandwidth FX system in a similar fashion. A very capable correlator could therefore be designed with minimal outlay (\sim £2000). Based on the experience of D-PAD, the analogue components for a B=16, X=16 sta-

tion array could be built for under £50,000: a very modest price for a pathfinder 21-cm intensity mapping instrument.

7.2.5 BAO-PAD SENSITIVITY

The BFoM of BAO-PAD would depend upon the station configuration, however a compact configuration would be required to maintain brightness sensitivity. From Equation 1.35, the BFoM of an instrument is

$$BFoM = 2n_p \Delta\nu \Omega_{FoV} \left(\frac{f}{T_{sys}} \right)^2.$$

As was shown in Sections 1.6.4 and 1.6.3, the minimum filling factor for an intensity mapping instrument is $f \approx 0.036$. To reach this filling factor requires ~ 1800 antennas of 9 dBi directivity within a 200λ diameter.

The BAO-PAD pathfinder would have only a fraction of this number of antennas. If the initial BAO-PAD array was kept within a 50λ diameter, then ~ 113 antennas are required to maintain this filling factor. With sixteen stations, each with sixteen antennas, this filling factor requirement would be fulfilled (256 antennas over a $50\lambda \times 50\lambda$ area).

A 50λ diameter array would have resolution of about 1 degree, which is insufficient to Nyquist sample the first peak of the power spectrum. As such, the BAO signature could not be measured. Nevertheless, BAO-PAD would be an ideal pathfinder instrument with which to test the viability of a sparse aperture approach to 21-cm intensity mapping.

7.3 CONCLUDING REMARKS

It is an exciting time for radio astronomy: we are on the edge of detecting the EoR through redshifted hydrogen, and construction on the SKA will begin within 5 years. The two SKA precursors, ASKAP and MeerKAT, are well into the construction phase, and at higher frequencies, construction on the Atacama Large Millimetre Array (ALMA) is scheduled to be completed in 2013.

Nevertheless, the performance of a radio telescope is almost always limited by its signal processing systems. The work presented here provides a specific example of how increasing the processed bandwidth of a receiver facilitates science. More generally, this work will inform the design of the next-generation of digital signal processors for radio astronomy.

We are in a strong technological position to detect the baryonic acoustic oscillations embued in the matter distribution of the Universe, through 21-cm intensity mapping. This could be done at a fraction of the cost of the SKA, ALMA, and even the SKA precursors. The instrumentation research presented in this thesis brings us a step closer to this BAO detection, which will be a giant leap forward in our understanding of this enigmatic Universe.

REFERENCES

A

- Abdalla, F. B., Blake, C., and Rawlings, S. (2010). Forecasts for dark energy measurements with future HI surveys. *Monthly Notices of the Royal Astronomical Society*, 401:743.
- Abdalla, F. B. and Rawlings, S. (2005). Probing dark energy with baryonic oscillations and future radio surveys of neutral hydrogen. *Monthly Notices of the Royal Astronomical Society*, 360:27.
- Anderson, K., Alexov, A., Baehren, L., et al. (2010). LOFAR and HDF5: Toward a New Radio Data Standard. *arXiv, astro-ph.IM*.
- Anderson, L., Aubourg, E., Bailey, S., et al. (2012). The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in the Data Release 9 Spectroscopic Galaxy Sample. *eprint arXiv, 1203:6594*.
- Armstrong, R. P. (2012). High-Performance Signal Processing Architectures for Digital Aperture Array Telescopes. PhD thesis, St. Peters College, University of Oxford.
- Atmel (2009). AT84AD001C Dual 8-bit 1 Gsps ADC. Technical report, Atmel.
- Avago (2010a). MGA-30889: Flat Gain High Linearity Gain Block. Technical report, Avago.
- Avago (2010b). MGA-633P8: Ultra Low Noise, High Linearity Active Bias Low Noise Amplifier. Technical report, Avago.

B

- Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., and Witzel, A. (1977). The absolute spectrum of CAS A - an accurate flux density scale and a set of secondary calibrators. *Astronomy and Astrophysics*, 61:99.
- Bailes, M., Bates, S. D., Bhalerao, V., et al. (2011). Transformation of a Star into a Planet in a Millisecond Pulsar Binary. *Science*, 333:1717.
- Barkana, R. and Loeb, A. (2007). The physics and early history of the intergalactic medium. *Reports on Progress in Physics*, 70:627.

- Barnes, D. G., Staveley-Smith, L., de Blok, W. J. G., et al. (2001). The Hi Parkes All Sky Survey: southern observations, calibration and robust imaging. *Monthly Notices of the Royal Astronomical Society*, 322:486.
- Bates, S. D., Bailes, M., Bhat, N. D. R., et al. (2011). The High Time Resolution Universe Pulsar Survey - II. Discovery of five millisecond pulsars. *Monthly Notices of the Royal Astronomical Society*, 416:2455.
- Battye, R. A., Brown, M. L., Browne, I. W. A., et al. (2012). BINGO: a single dish approach to 21cm intensity mapping. *arXiv, astro-ph.CO*.
- Beardsley, A. P., Hazelton, B. J., Morales, M. F., et al. (2012). The EoR Sensitivity of the 128 Antenna Murchison Widefield Array. *arXiv, astro-ph.IM*.
- Bellanger, M., Bonnerot, G., and Coudreuse, M. (1976). Digital filtering by polyphase network: Application to sample-rate alteration and filter banks. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 24(2):109 – 114.
- Bennett, C. L. (2006). Cosmology from start to finish. *Nature*, 440:1126.
- Bergman, J. E. S., Åhlén, L., Stål, O., et al. (2005). ELVIS - ELectromagnetic Vector Information Sensor. *arXiv, astro-ph*.
- Bergman, J. E. S. and Carozzi, T. D. (2008). Canonical Electromagnetic Observables for Systematic Characterization of Electric and Magnetic Wave Field Data on board Spacecraft. *eprint arXiv, 0804:2092*.
- Blake, C. and Glazebrook, K. (2003). Probing Dark Energy Using Baryonic Oscillations in the Galaxy Power Spectrum as a Cosmological Ruler. *The Astrophysical Journal*, 594:665.
- Braun, R. and van Cappellen, W. (2006). Aperture Arrays for the SKA: Dense or Sparse? *SKA Memorandum series*, (87).
- Brigham, E. O. (1988). *The FFT and its applications*. Prentice-Hall Inc.
- Brunzie, T. J. (1988). A noise-adding radiometer for the Parkes antenna. In *its The Telecommunications and Data Acquisition Report p 117-122 (SEE N88-18774 11-32)*, page 117.
- Bunton, J. (2000). An Improved FX correlator. *Alma Memorandum Series*, (342).
- Burke-Spolaor, S., Bailes, M., Johnston, S., et al. (2011). The High Time Resolution Universe Pulsar Survey - III. Single-pulse searches and preliminary analysis. *Monthly Notices of the Royal Astronomical Society*, 416:2465.

C

-
- Carozzi, T. D. and Woan, G. (2009). A generalized measurement equation and van Cittert–Zernike theorem for wide-field radio astronomical interferometry. *Monthly Notices of the Royal Astronomical Society*, 395(3):1558–1568.
- Chang, T.-C., Pen, U.-L., Bandura, K., and Peterson, J. B. (2010). An intensity map of hydrogen 21-cm emission at redshift $z = 0.8$. *Nature*, 466:463.
- Chang, T.-C., Pen, U.-L., Peterson, J. B., and McDonald, P. (2008). Baryon Acoustic Oscillation Intensity Mapping of Dark Energy. *Physical Review Letters*, 100:91303.
- Chikada, Y., Ishiguro, M., Hirabayashi, H., Morimoto, M., and Morita, K.-I. (1987). A 6 x 320-MHz 1024-channel FFT cross-spectrum analyzer for radio astronomy. *IEEE*, 75:1203.
- Colless, M., Dalton, G., Maddox, S., et al. (2001). The 2dF Galaxy Redshift Survey: spectra and redshifts. *Monthly Notices of the Royal Astronomical Society*, 328:1039.
- Condon, J. and Ransom, S. (2010). *Essential Radio Astronomy*. Technical report, National Radio Astronomy Observatory.
- Cordes, J. (2009). *Survey Metrics*. SKA Memos, (109).
- Cross, S., Crida, R., Bennett, T., and Welz, M. (2009). *Guidelines for Communication with Devices*. National Research Foundation (South Africa).
- CSIRO (2011). <http://www.csiro.au/science/Parkes-five-decades-of-discovery>.

D

-
- Davis, M., Gerke, B. F., Newman, J. A., and Team, D. (2005). Constraining Dark Energy with the DEEP2 Redshift Survey. *Observing Dark Energy*, 339:128.
- de Lera Acedo, E. (2011). SKALA: SKALog-periodic AntennaA candidate for the SKA AA-low. In *The Path to SKA-low Workshop*.
- de Vos, M., Gunst, A., and Nijboer, R. (2009). The LOFAR Telescope: System Architecture and Signal Processing. *Proceedings of the IEEE*, 97(8):1431 – 1437.
- DeBoer, D., Gough, R., Bunton, J., et al. (2009). Australian SKA Pathfinder: A High-Dynamic Range Wide-Field of View Survey Telescope. *Proceedings of the IEEE*, 97(8):1507 – 1521.
- Delhaize, J. and Staveley-Smith, L. (2012). Studying neutral hydrogen in galaxies to $z = 0.14$ with the Parkes Telescope and the 2dF Galaxy Redshift Survey. in prep.
- Dewdney, P., Hall, P., Schilizzi, R., and Lazio, T. (2009). The Square Kilometre Array. *Proceedings of the IEEE*, 97(8):1482 – 1496.

Dicke, R. H., Peebles, P. J. E., Roll, P. G., and Wilkinson, D. T. (1965). Cosmic Black-Body Radiation. *Astrophysical Journal*, 142:414.

Dobbs, M. (2009). The Canadian Hydrogen Intensity Mapping Experiment (CHIME). In Canadian Institute for Advanced Research (CfAR) Cosmology and Gravity Group Annual General Meeting.

Donley, J. L., Koribalski, B. S., Staveley-Smith, L., et al. (2006). A massive spiral galaxy in the Zone of Avoidance. *Monthly Notices of the Royal Astronomical Society*, 369:1741.

Dunn, M. J. (2011). Global Positioning System Directorate Systems Engineering & Interface Specification IS-GPS-705. Technical report, Navstar GPS Space Segment.

E

Eisenstein, D. J., Zehavi, I., Hogg, D. W., et al. (2005). Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *The Astrophysical Journal*, 633:560.

Ellingson, S. W., Taylor, G. B., Craig, J., et al. (2012). The LWA1 Radio Telescope. arXiv, astro-ph.IM.

Engargiola, G. (2002). Non-planar log-periodic antenna feed for integration with a cryogenic microwave amplifier. *ATA Memos*, (45):1–4.

Erker, S., Tholert, S., Furthner, J., and Meurer, M. (2009). L5–The New GPS Signal. Proceedings of the 13th IAIN World Congress and Exhibition, Stockholm, October 27th - 30th 2009.

F

Faria, J. A. B. (2002). On the transmission matrix of $2n$ -port reciprocal networks. *Microwave and Optical Technology Letters*, 33(3):151–154.

Faulkner, A. (2011). Dense aperture arrays for the Square Kilometre Array. In General Assembly and Scientific Symposium, 2011 XXXth URSI, pages 1 –4.

Foster, G. (2012). personal communication.

Foster, G., Hickish, J., Price, D., and Adami, K. Z. (2012). A Spatial FFT and FX Correlator for the BEST-2 Array. in prep.

Furlanetto, S. R., Oh, S. P., and Briggs, F. H. (2006). Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe. *Physics Reports*, 433:181.

G

Gade, S. and Herlufsen, H. (1987). Use of weighting functions in DFT/FFT analysis (Part I). Brüel & Kjær, Windows to FFT Analysis (Part I) Technical Review, (3).

GBT (2011). Observing with the Green Bank Telescope. observer guide, NRAO.

Giovanelli, R. (2005). ALFALFA: a Legacy Extragalactic HI Survey. American Astronomical Society Meeting 207, 207:1488.

gps.gov (2012). <http://www.gps.gov>.

Griffiths, D. J. (1999). Introduction to Electrodynamics. Prentice-Hall Inc.

H

Hamaker, J. P. (2000). Understanding radio polarimetry. IV. The full-coherency analogue of scalar self-calibration: Self-alignment, dynamic range and polarimetric fidelity. *Astronomy and Astrophysics Supplement*, 143:515.

Hamaker, J. P. (2006). Understanding radio polarimetry. V. Making matrix self-calibration work: processing of a simulated observation. *Astronomy and Astrophysics*, 456:395.

Hamaker, J. P. and Bregman, J. D. (1996). Understanding radio polarimetry. III. Interpreting the IAU/IEEE definitions of the Stokes parameters. *Astronomy and Astrophysics Supplement*, 117:161.

Hamaker, J. P., Bregman, J. D., and Sault, R. J. (1996). Understanding radio polarimetry. I. Mathematical foundations. *Astronomy and Astrophysics Supplement*, 117:137.

Harris, F. (2004). *Multirate Signal Processing for Communication System*. Prentice-Hall Inc.

Hickish, J., Foster, G., Adami, K. Z., and Jones, M. (2012). Multipliers, multipliers, multipliers. in prep.

Hickling, A. M. (2010). Assessment and Improvement of the Beam Quality of the D-PAD Stations and Tooling to Aid Investigations Involving Multiple Stations. *Astrophysics Summer Project*.

Hong, J.-S. and Lancaster, M. J. (2001). *Microstrip filters for RF/microwave applications*. John Wiley and Sons Inc.

J

Jarosik, N., Bennett, C. L., Dunkley, J., et al. (2011). Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results. *The Astrophysical Journal Supplement*, 192:14.

Jones, D. H., Saunders, W., Colless, M., et al. (2004). The 6dF Galaxy Survey: samples, observational techniques and the first data release. *Monthly Notices of the Royal Astronomical Society*, 355:747.

Jones, R. C. (1941). A New Calculus for the Treatment of Optical Systems. I. Description and Discussion of the Calculus. *J. Opt. Soc. Am.*, 31:1–6.

K

Keith, M. J., Jameson, A., van Straten, W., et al. (2010). The High Time Resolution Universe Pulsar Survey – I. System configuration and initial discoveries. *Monthly Notices of the Royal Astronomical Society*, 409:619.

Keith, M. J., Johnston, S., Bailes, M., et al. (2012). The High Time Resolution Universe Pulsar Survey – IV. Discovery and polarimetry of millisecond pulsars. *Monthly Notices of the Royal Astronomical Society*, 419:1752.

Kestur, S., Davis, J. D., and Williams, O. (2010). BLAS Comparison on FPGA, CPU and GPU. In *Proceedings of the 2010 IEEE Annual Symposium on VLSI, ISVLSI '10*, pages 288–293, Washington, DC, USA. IEEE Computer Society.

King, O. (2009). C-BASS: The C-Band All Sky Survey. PhD thesis, The University of Oxford.

Klein, B., Hochgürtel, S., Krämer, I., et al. (2012). High-resolution wide-band fast Fourier transform spectrometers. *A & A*, 542.

Klein, B., Philipp, S. D., Gusten, R., Kramer, I., and Samtleben, D. (2006). A new generation of spectrometers for radio astronomy: Fast Fourier Transform Spectrometer. In Mather, J. C., MacEwen, H. A., and de Graauw, M. W. M., editors, *Space Telescopes and Instrumentation I: Optical, Infrared, and Millimeter*, volume 6275, page 627511. SPIE.

Kocz, J., Bailes, M., Barnes, D., Burke-Spolaor, S., and Levin, L. (2012). Enhanced pulsar and single pulse detection via automated radio frequency interference detection in multipixel feeds. *Monthly Notices of the Royal Astronomical Society*, 420:271.

Kraan-Korteweg, R. C., Staveley-Smith, L., Donley, J., Koribalski, B., and Henning, P. A. (2005). The Universe Behind the Southern Milky Way. *Maps of the Cosmos*, 216:203.

Kraus, J. D. and Marhefka, R. J. (2002). *Antennas for All Applications*, pages 439–441. McGraw-Hill.

Kuon, I. and Rose, J. (2006). Measuring the gap between FPGAs and ASICs. In *Proceedings of the 2006 ACM/SIGDA 14th international symposium on Field programmable gate arrays, FPGA '06*, pages 21–30, New York, NY, USA. ACM.

L

-
- Lane, W. M., Clarke, T. E., Taylor, G. B., Perley, R. A., and Kassim, N. E. (2004). Hydra A at Low Radio Frequencies. *The Astronomical Journal*, 127:48.
- Larson, D., Dunkley, J., Hinshaw, G., et al. (2011). Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Power Spectra and WMAP-derived Parameters. *The Astrophysical Journal Supplement*, 192:16.
- Li, Y.-C., Pen, U.-L., Staveley-Smith, L., et al. (2012). HI Intensity Mapping: Parkes-2dFGRS and BAO science. ATNF proposal P641.
- Loeb, A. (2006). The Dark Ages of the Universe. *Scientific American*.
- Loeb, A. and Wyithe, J. S. B. (2008). Possibility of Precise Measurement of the Cosmological Power Spectrum with a Dedicated Survey of 21 cm Emission after Reionization. *Physical Review Letters*, 100:161301.
- Lonsdale, C., Cappallo, R., Morales, M., et al. (2009). The Murchison Widefield Array: Design Overview. *Proceedings of the IEEE*, 97(8):1497 – 1506.
- Ludwig, A. (1973). The definition of cross polarization. *Antennas and Propagation, IEEE Transactions on*, 21(1):116–119.
- Lyons, R. G. (2004). *Understanding digital signal processing*. Prentice Hall, 2nd edition.

M

-
- Mailloux, R. J. (2005). *Phased Array Antenna Handbook*. Artech House Inc., 685 Canton Street, Norwood, MA 02062, 2nd edition.
- Mandel, L. and Wolf, E. (1995). *Optical Coherence and Quantum Optics*. Cambridge University Press, 1 edition.
- Matthaei, G. L., Young, L., and Jones, E. M. T. (1980). *Microwave Filters, Impedance-Matching Networks, and Coupling Structures*, pages 421–519. Artech House.
- McMahon, P. L. (2008). *Adventures in Radio Astronomy Instrumentation and Signal Processing*. CASPER Theses, pages 1–135.
- Meyer, M. J., Zwaan, M. A., Webster, R. L., et al. (2004). The HIPASS catalogue – I. Data presentation. *Monthly Notices of the Royal Astronomical Society*, 350:1195.
- Morales, M. F. (2011). Enabling Next-Generation Dark Energy and Epoch of Reionization Radio Observatories with the MOFF Correlator. *Publications of the Astronomical Society of the Pacific*, 123:1265.

- Morales, M. F. and Wyithe, J. S. B. (2010). Reionization and Cosmology with 21-cm Fluctuations. *Annual Review of Astronomy and Astrophysics*, 48:127.
- Mort (2010). OSKAR: Simulating digital beamforming for the SKA aperture array. *Phased Array Systems and Technology (ARRAY)*, 2010 IEEE International Symposium on, pages 690 – 694.
- Mort, B., Dulwich, F., and Salvini, S. (2012). OSKAR2 [computer software].
- Mucci, R. (1984). A comparison of efficient beamforming algorithms. *Acoustics, Speech and Signal Processing*, IEEE Transactions on, 32(3):548 – 558.
- Mueller, H. (1948). The foundation of optics. *J. Opt. Soc. Am.*

N

- Navtech (2012). <http://www.navtechgps.com/extra/GNSSfacts.asp>. online.
- Nita, G. M. and Gary, D. E. (2010). The generalized spectral kurtosis estimator. *Monthly Notices of the Royal Astronomical Society: Letters*, 406:L60. (c) Journal compilation © 2010 RAS.
- Noordam, J. E. and Smirnov, O. M. (2010). The MeqTrees software system and its use for third-generation calibration of radio interferometers. *A&A*, 524:A61.
- NRAO (2011). CASA 3.3.0. [computer software].

O

- OFCOM (2010). United Kingdom Frequency Allocation Table. Technical report, National Frequency Planning Group.
- Olofsson, H. A. O., Torchinsky, S. A., Chemin, L., et al. (2009). Profiling the EMBRACE tile beam using GPS satellite carriers. *Proceedings of Wide Field Astronomy & Technology for the Square Kilometre Array (SKADS 2009)*. 4-6 November 2009. Chateau de Limelette, page 42.
- O’Neil, K. (2002). Single-Dish Calibration Techniques at Radio Wavelengths. *Single-Dish Radio Astronomy: Techniques and Applications*, 278:293. ISBN: 1-58381-120-6.

P

- Pantoja (1986). Analysis of a wideband microwave log-periodic antenna feeding technique. *Antennas and Propagation Society International Symposium*, 1986, 24:757 – 760.
- Pantoja (1987). A microwave printed planar log-periodic dipole array antenna. *Antennas and Propagation*, IEEE Transactions on, 35(10):1176 – 1178.
- Parsons, A. (2012). The BAO Broadband and Broad-beam (BAOBAB) Array. Unpublished.

- Parsons, A., Backer, D., Bradley, R., and ..., J. A. (2009). The Precision Array for Probing the Epoch of Reionization: 8 Station Results. arXiv, 0904.2334.
- Parsons, A., Backer, D., Chang, C., and Chapman, D. (2006). Petaflop/second FPGA signal processing for SETI and radio astronomy. *Signals*.
- Parsons, A., Backer, D., Chen, H., et al. (2008). A Scalable Correlator Architecture Based on Modular FPGA Hardware, Reuseable Gateway, and Data Packetization. arXiv.
- Parsons, A., Pober, J., McQuinn, M., Jacobs, D., and Aguirre, J. (2011). A Sensitivity and Array-Configuration Study for Measuring the Power Spectrum of 21 cm Emission from Reionization. eprint arXiv, 1103:2135.
- Parsons, A. R. (2009). Low-Frequency Interferometry: Design, Calibration, and Analysis Towards Detecting the Epoch of Reionization. *CASPER Theses*, pages 1–127.
- Peixeiro (1988). Design of log-periodic dipole antennas. *Microwaves, Antennas and Propagation, IEE Proceedings H*, 135(2):98 – 102.
- Pen, U.-L., Staveley-Smith, L., Peterson, J. B., and Chang, T.-C. (2009). First detection of cosmic structure in the 21-cm intensity field. *Monthly Notices of the Royal Astronomical Society: Letters*, 394:L6.
- Penzias, A. A. and Wilson, R. W. (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophysical Journal*, 142:419.
- Percival, W. J., Reid, B. A., Eisenstein, D. J., et al. (2010). Baryon acoustic oscillations in the Sloan Digital Sky Survey Data Release 7 galaxy sample. *Monthly Notices of the Royal Astronomical Society*, 401:2148.
- Peterson, J. B., Aleksan, R., Ansari, R., et al. (2009). 21-cm Intensity Mapping. *Astro2010: The Astronomy and Astrophysics Decadal Survey*, 2010:234.
- Pozar, D. M. (2005). *Microwave Engineering*. John Wiley and Sons Inc., 3rd edition.

R

- Rawlings, S., Abdalla, F. B., Bridle, S. L., et al. (2004). Galaxy evolution, cosmology and dark energy with the Square Kilometer Array. *New Astronomy Reviews*, 48:1013.
- Reynolds, J. (2012). Personal communication.
- Ries, P. (2007). Beam Patterns of Low Frequency Radio Antennas using LEO Satellites. *PAPER memo 15*.

S

-
- Sault, R. J., Hamaker, J. P., and Bregman, J. D. (1996). Understanding radio polarimetry. II. Instrumental calibration of an interferometer array. *Astronomy and Astrophysics Supplement*, 117:149.
- Schediwy, S. (2010). SANDGROPER ARC Proposal. ARC Research Proposal.
- Smirnov, O. M. (2011a). Revisiting the radio interferometer measurement equation. I. A full-sky Jones formalism. *A&A*, 527:A106.
- Smirnov, O. M. (2011b). Revisiting the radio interferometer measurement equation. IV. A generalized tensor formalism. *A&A*, 531:A159.
- Smith, S. W. (2002). *Digital Signal Processing: A Practical Guide for Engineers and Scientists*. Newnes.
- Smoot, G. F., Bennett, C. L., Kogut, A., et al. (1992). Structure in the COBE differential microwave radiometer first-year maps. *Astrophysical Journal*, 396:L1.
- Staveley-Smith, L., Wilson, W. E., Bird, T. S., et al. (1996). The Parkes 21 CM multibeam receiver. *Publications Astronomical Society of Australia*, 13:243.
- Straw, R. D., editor (2007). *The ARRL Antenna Book*, chapter 10. The ARRL Inc., 21st edition.
- Szomoru, A. (2010). The UniBoard. In 10th European VLBI Network Symposium and EVN Users Meeting: VLBI and the New Generation of Radio Arrays.

T

-
- Taylor, G. B., Carilli, C. L., and Perley, R. A., editors (1999). *Synthesis Imaging in Radio Astronomy II*, volume 180. Astronomical Society of the Pacific.
- Tegmark, M. and Zaldarriaga, M. (2009a). Fast Fourier transform telescope. *Physical Review D*, 79:83530.
- Tegmark, M. and Zaldarriaga, M. (2009b). Omniscope: Large Area Telescope Arrays with only $N \log N$ Computational Cost. arXiv, 0909.0001.
- Thompson, A. R., Emerson, D. T., and Schwab, F. R. (2007). Convenient formulas for quantization efficiency. *Radio Science*, 42:3022.
- Thompson, A. R., Moran, J. M., and Jr., G. W. S. (2004). *Interferometry and Synthesis in Radio Astronomy*. WILEY-VCH Verlag, second edition.
- Tinbergen, J. (2005). *Astronomical Polarimetry*. Cambridge University Press.
- TTi (2006). TGR2050 Instruction Manual. Technical report, Thurlby Thandar Instruments.

V

Vaidyanathan (1990). Multirate digital filters, filter banks, polyphase networks, and applications: a tutorial. *Proceedings of the IEEE*, 78(1):56 – 93.

W

Weinberg, S. (2008). *Cosmology*. Cosmology. OUP Oxford.

Weinreb, S. (1963). A digital spectral analysis technique and its application to radio astronomy. MIT Research Laboratory of Electronics Technical report.

Weinreb, S., Barrett, A. H., Meeks, M. L., and Henry, J. C. (1963). Radio Observations of OH in the Interstellar Medium. *Nature*, 200:829.

Werthimer, D. (2012). personal communication.

Wikimedia (2010). Circular Polarization and Linear Polarized light comparison [image]. <http://commons.wikimedia.org>.

Winkel, B., Kraus, A., and Bach, U. (2012). Unbiased flux calibration methods for spectral-line radio observations. eprint arXiv, 1203:741.

Wolf, E. (1954). *Nuovo Cimento*, 12(884).

Wyithe, J. S. B., Loeb, A., and Geil, P. M. (2008). Baryonic acoustic oscillations in 21-cm emission: a probe of dark energy out to high redshifts. *Monthly Notices of the Royal Astronomical Society*, 383:1195.

X

Xilinx (2012). Virtex-5 FPGA XtremeDSP Design Considerations. Technical report, Xilinx Incorporated.

Z

Zarb-Adami, K. (2011). personal communication.

Zaroubi, S., de Bruyn, A. G., Harker, G., et al. (2012). Imaging neutral hydrogen on large-scales during the Epoch of Reionization with LOFAR. arXiv, astro-ph.CO.

Zernicke, F. (1938). *Physica*, 5(758).

Zimmerman, G. A. and Gulkis, S. (1991). Polyphase-Discrete Fourier Transform Spectrum Analysis for the Search for Extraterrestrial Intelligence Sky Survey. The Telecommunications and Data Acquisition Progress Report, 107:141.

APPENDIX A

SCHEMATICS AND COMPONENT LISTINGS

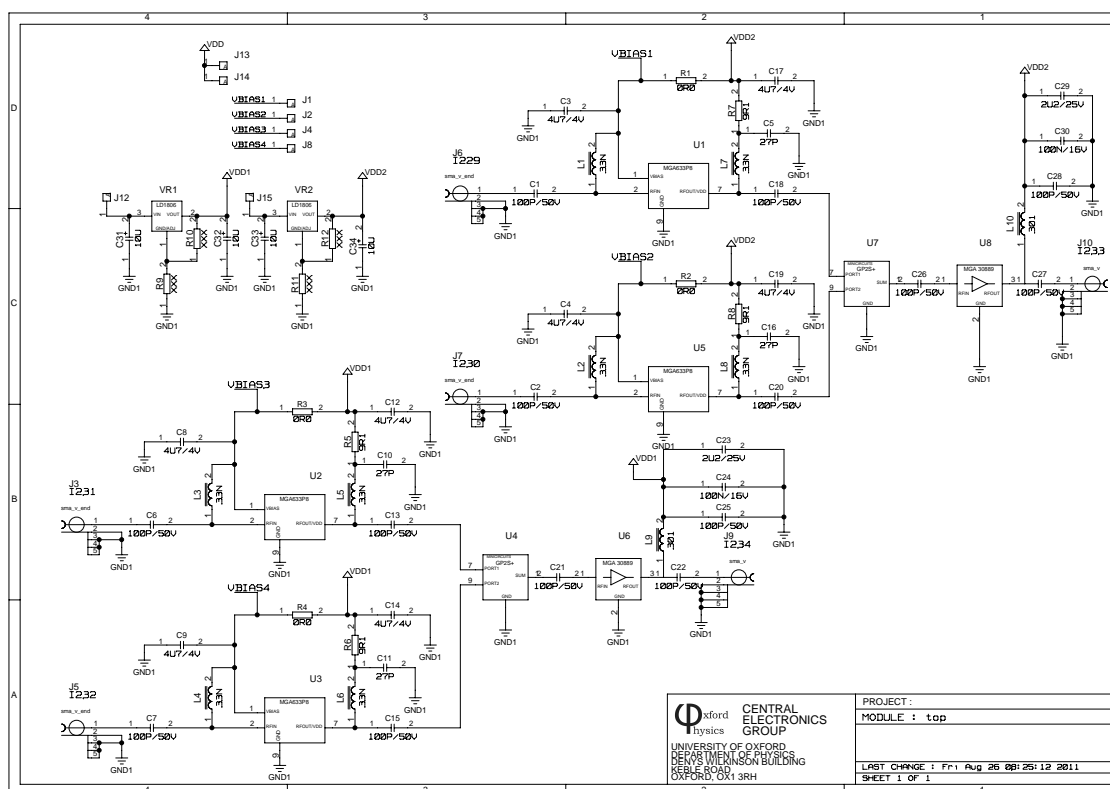

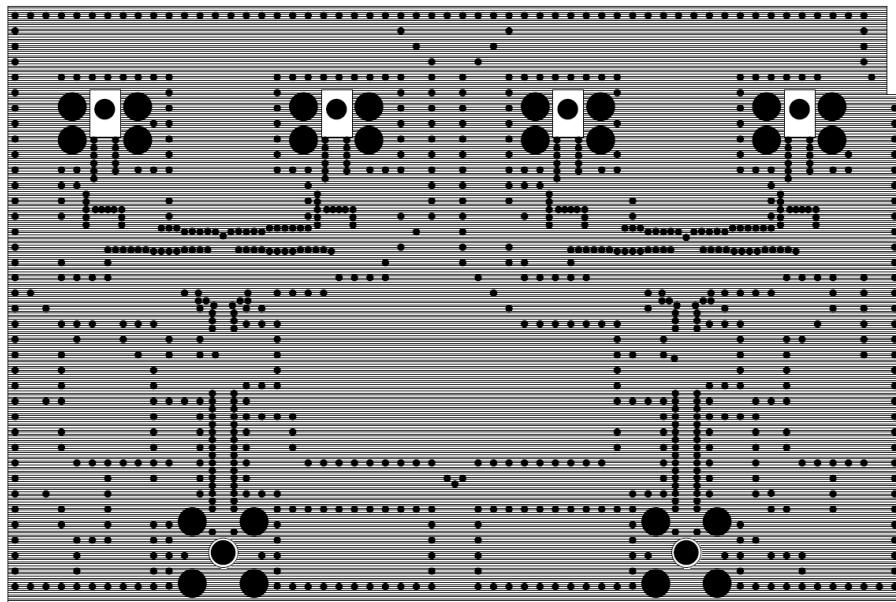


Figure A.1: Receiver board, circuit schematic.



(a) Receiver board, top



(b) Receiver board, bottom

Figure A.2: Receiver board solder resist layouts

Table A.1 : Receiver board bill of materials (BOM)

SYM_NAME	COMP_DEVICE_TYPE	COMP_CLASS	QUANTITY	REFDES
CASE_A	CAPACITOR_CASE_A-10U,+/-20%,NEC	DISCRETE	4	C31, C32, C33, C34
DPAK	LD1806	IC	2	VR1, VR2
DQ1225	MINICIRCUITS_GP2S	IC	2	U4, U7
PROBEPAD	PROBEPAD	IC	8	J1, J2, J4, J8, J12, J13, J14, J15
QFN8	MGA_633P8	IC	4	U1, U2, U3, U5
SMA_VERT	SMA_V	IC	2	J9, J10
SMA_VERT_ENDLAUNCH	SMA_V_END	IC	4	J3, J5, J6, J7
SMD0402	CAPACITOR_0402-100P/50V,+/-5%,A	DISCRETE	14	C1, C2, C6, C7, C13, C15, C18, C20, C21, C22, C25, C26, C27, C28
SMD0402	INDUCTOR_0402-33N,TOKO	DISCRETE	8	L1, L2, L3, L4, L5, L6, L7, L8
SMD0402	RESISTOR_0402-0R0,5%,NEOHM	DISCRETE	4	R1, R2, R3, R4
SMD0402	RESISTOR_0402-9R1,5%,ROHM	DISCRETE	4	R5, R6, R7, R8
SMD0603	INDUCTOR_0603-301,TDK	DISCRETE	2	L9, L10
SMD0603	RESISTOR_0603-XXX,5%,ROHM	DISCRETE	4	R9, R10, R11, R12
SMD0805RFLWCAP	CAPACITOR_0805-2U2/25V,+80/-20A	DISCRETE	2	C23, C29
SMDCAP0402	CAPACITOR_0402-4U7/4V,+/-20%,KA	DISCRETE	8	C3, C4, C8, C9, C12, C14, C17, C19
SMDCAP0402	CAPACITOR_0402-27P,+/-5%,MURATA	DISCRETE	4	C5, C10, C11, C16
SMDCAP0402	CAPACITOR_0402-100N/16V,+/-20%A	DISCRETE	2	C24, C30
SOT89	MGA_30889	IC	2	U6, U8