

# Digital Signal Processing Methods for Large- $N$ , Low-Frequency Radio Telescopes



Jack Hickish  
Balliol College  
University of Oxford

A thesis submitted for the degree of  
*Doctor of Philosophy*

Hilary Term 2014



For my grandfather, Cyril.

(With my apologies for the lack of chess-related content.)



## Acknowledgements

I am indebted to many people for their companionship, advice, and support over the last four years. Firstly, I express my gratitude to my supervisors Mike Jones and Kris Zarb Adami. Mike, for sharing your knowledge and experience, and helping to get this thesis completed. Kris, for your endless enthusiasm and desire to build new instruments for every telescope on the planet. I thank you both for your encouragement and assistance during my time working in Oxford.

I have been lucky enough to work with many excellent colleagues and friends around the world. I thank the Maltesers for bringing sophistication and culture to my time as a graduate student, and in particular Alessio for being both a fantastic character and an exceptional recipient of packets. To the staff of the Medicina Observatory I offer my warmest gratitude for your unending hospitality, assistance, and provision of fine food during my stays in Italy. To Max and his merry band of Omniscopeers, thank you for being so welcoming when I visited MIT. I am also hugely grateful to the members of the CASPER collaboration, especially those who have offered advice via the mailing list, and the developers in UC Berkeley and South Africa, whose work has benefited my research enormously.

Griffin, Charles, and Danny – I am certain I would not have survived without you. Thank you Griffin, for introducing me to the joys of NPR, candying the yams, and helping correlate stuff; Charles, for catalysing my motorcycling, your endless drilling, and general-RoachCommand NDIODEXXXX,1,0; Danny, for your appreciation of fine typesetting, your intimate knowledge of the Oxford branding guidelines, and for not complaining when I came to stay with you for a few weeks and left 7 months later. For sharing in the joy, hilarity, triumph, and despair of the last four years, you all have my endless appreciation.

To Richard, I offer my gratitude for being so welcoming and helpful when I first encountered the Oxford digital radio astronomy group as an undergraduate, and for many years of fun since. I thank Ian for sharing with me the delights of Cumnor; Aris, for finding me a beer in Belgium that tasted like an apple martini; and Sascha, for bringing a touch of Australia to the lab. I shall always have fond memories of Steve, who deserves no small amount of credit for gathering such an excellent bunch of astro-nerds in Oxford.

To Ben, Fred, and Stef, for sneaking me into the OeRC rounders match and tolerating my unannounced and disruptive coffee visits, I can only offer my thanks and apologies.

Lastly, I thank those who have enriched my life beyond the brutalist walls of the Denys Wilkinson Building and made the last four years livable. Lindall, the first out of the fire; Zoe, who almost drowned trying to save my sail; Nicole, who tamed the Griffin; and Eli, who might be a latecomer to the party but should be celebrated for her hilarious taste in men (I can't believe you dealt with those gammy feet). Laura, for your unwavering support and tolerance, I will always be grateful. To my parents, who have been so supportive of all my endeavours, thank you for the last 26 years kindness and love. To the Stevensons, my surrogate family of the last decade and more, I'm sorry I climbed on the roof (but it was Paul's idea).



# Digital Signal Processing Methods for Large- $N$ , Low-Frequency Radio Telescopes

Jack Hickish

Balliol College, University of Oxford

*A thesis submitted for the degree of Doctor of Philosophy in Hilary Term 2014*

## Abstract

Current attempts to make precision measurements of the HI power spectrum at high redshifts have led to the construction of several low-frequency, large- $N$ , interferometric arrays. The computational demands of digital correlators required by these arrays present a significant challenge. These demands stem from the treatment of radio telescopes as collections of two-element interferometers, which results in the need to multiply  $O(N^2)$  pairs of antenna signals in an  $N$ -element array. Given the unparalleled flexibility offered by modern digital processing systems, it is apt to consider whether a different way of treating the signals from antennas in an array might be fruitful in current and future radio telescopes. Such methods potentially avoid the unfavourable  $N^2$  scaling of computation rate with array size.

In this thesis I examine the prospect of using direct-imaging methods to map the sky without first generating correlation matrices. These methods potentially provide great computational savings by creating images using efficient, FFT-based algorithms. This thesis details the design and deployment of such a system for the Basic Element of SKA Training II (BEST-2) array in Medicina, Italy. Here the 32-antenna BEST-2 array is used as a test bed for comparison of FX correlation and direct-imaging systems, and to provide a frontend for a real-time transient event detection pipeline.

Even in the case of traditional  $O(N^2)$  correlation methods, signal processing algorithms can be significantly optimized to deliver large performance gains. In this thesis I present a new mechanism for optimizing the cross-correlation operation on Field Programmable Gate Array (FPGA) hardware. This implementation is shown to achieve a 75% reduction in multiplier usage, and has a variety of benefits over existing optimization strategies.

Finally, this thesis turns its focus towards The Square Kilometre Array (SKA). When constructed, the SKA will be the world's largest radio telescope and will comprise a variety of arrays targeting different observing frequencies and science goals. The low-frequency component of the SKA (SKA-low) will feature  $\sim 250\,000$  individual antennas, sub-divided into a number of stations. This thesis explores the impact of the station size on the computational requirements of SKA-low, investigating the optimal array configuration and signal processing realizations.



# Nomenclature

ADC Analog to Digital Converter

FIR Finite Impulse Response (Filter)

ASICs Application Specific Integrated Circuit

BAOs Baryonic Acoustic Oscillations

CASPER Collaboration for Astronomy Signal Processing and Electronics Research

CMAC Complex Multiply-Accumulate

CMB Cosmic Microwave Background

CPU Central Processing Unit

DM Dispersion Measure

FFT Fast Fourier Transform

FoV Field of View

FPGA Field Programmable Gate Array

FrFFT Fractional Fast Fourier Transform

FrFT Fractional Fourier Transform

FT Fourier Transform

GBE Gigabit Ethernet

GPU Graphical Processing Unit

IC Integrated Circuit

IF Intermediate Frequency

LO Local Oscillator

LSB Least Significant Bit

LUT Look Up Table

LVDS Low Voltage Differential Signaling

MSB Most Significant Bit

MSSGE Matlab Simulink System-Generator Environment

NRE Non-recurring engineering

PDF Probability density function

PFB Polyphase Filter Bank

PSF Point Spread Function

RFI Radio-frequency Interference

RF Radio Frequency

SIMD Single Instruction Multiple Data

SKA Square Kilometre Array

SNR Signal-to-noise ratio

SPEAD Streaming Protocol for Exchanging Astronomical Data

XAUI Ten Gigabit Attachment Unit Interface

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Information in the radio sky . . . . .	3
1.1.1	The Evolving Universe and Radio Astronomy . . . . .	3
1.2	Probing Global structure with Hydrogen Observatories . . . . .	8
1.2.1	Observational Challenges . . . . .	12
1.3	Coherent Manipulation of Radio Signals . . . . .	16
1.3.1	Beamforming . . . . .	16
1.3.2	Interferometry . . . . .	18
1.4	Interferometric Radio Telescopes . . . . .	20
1.4.1	The Age of Digital Radio Astronomy . . . . .	20
1.5	Current 21-cm Observatories . . . . .	21
1.5.1	EoR Observatories . . . . .	22
1.5.2	BAO Observatories . . . . .	23
1.6	Thesis Layout . . . . .	24
<b>2</b>	<b>Spatial Processing with Large-N Arrays</b>	<b>26</b>
2.1	Properties of astronomical radio signals . . . . .	26
2.1.1	Statistical properties . . . . .	26
2.1.2	Quantization . . . . .	27
2.1.3	Spatial coherence . . . . .	30
2.2	Interferometry and synthesis imaging . . . . .	30
2.2.1	Jones matrices and signal propagation . . . . .	31
2.2.2	Map making and the van Cittert–Zernike Theorem . . . . .	32
2.3	Calibration . . . . .	34
2.3.1	Sky Model Fitting . . . . .	34
2.3.2	Self-Calibration . . . . .	35
2.3.3	Baseline Redundancy . . . . .	35
2.4	Correlators . . . . .	35

2.4.1	The FX Correlator . . . . .	36
2.4.2	The XF Correlator . . . . .	37
2.4.3	Gridding . . . . .	37
2.4.4	The Direct-Imaging Correlator . . . . .	40
2.5	Choosing a correlation method . . . . .	43
2.5.1	Array configuration . . . . .	44
2.5.2	Calibratability . . . . .	48
2.5.3	Autocorrelations . . . . .	53
2.5.4	Phase Centre Tracking with Interpolation Gridding . . . . .	54
2.5.5	Computational Demands . . . . .	55
2.6	Relative Merits of Direct-Imaging Systems for Existing Telescopes . . . . .	58
2.7	Conclusions . . . . .	59
<b>3</b>	<b>An Optimized, Low-Bitwidth Correlator for FPGAs</b>	<b>62</b>
3.1	FPGAs in Radio Astronomy DSP Applications . . . . .	63
3.2	The CASPER X-engine . . . . .	63
3.2.1	Module Description . . . . .	64
3.2.2	Resource Utilization . . . . .	68
3.2.3	Successes of the CASPER X-engine . . . . .	73
3.2.4	Limitations of the CASPER X-engine . . . . .	74
3.3	Optimized CASPER X-engine . . . . .	76
3.3.1	An existing implementation of optimized complex multiplication . . . . .	76
3.3.2	Re-factoring the Correlation Algorithm . . . . .	77
3.3.3	Implementation . . . . .	79
3.3.4	Proposed X-engine Design . . . . .	81
3.3.5	Design Verification . . . . .	81
3.3.6	Quantitative Comparison of Resource Usage . . . . .	83
3.4	Conclusions . . . . .	84
<b>4</b>	<b>Implementation of a Direct-Imaging Telescope with BEST-2</b>	<b>87</b>
4.1	Motivation . . . . .	87
4.2	Introducing BEST-2 . . . . .	88
4.3	System Requirements . . . . .	90
4.4	System Architecture . . . . .	90
4.4.1	Analogue Frontend . . . . .	90
4.4.2	Digital Back-End . . . . .	90
4.4.3	“F” Engine Firmware . . . . .	93

4.4.4	“S” Engine Firmware . . . . .	101
4.4.5	“X” Engine Firmware . . . . .	105
4.5	Software (Control & Calibration) Framework . . . . .	106
4.5.1	Equalization Settings . . . . .	107
4.5.2	Data Collection . . . . .	107
4.5.3	Object tracking . . . . .	108
4.6	Conclusions . . . . .	109
<b>5</b>	<b>Results from the BEST-2 Direct Imaging Telescope</b>	<b>110</b>
5.1	System Commissioning . . . . .	110
5.2	Equivalence of Spatial FFT and Correlator Data . . . . .	111
5.3	Calibration Strategies . . . . .	118
5.3.1	Sky-Model fitting . . . . .	119
5.3.2	Calibration Stability . . . . .	120
5.3.3	Calibration Artifacts . . . . .	120
5.4	Direct Imaging Results . . . . .	122
5.4.1	Post-calibrating Direct Images . . . . .	124
5.4.2	Real-time Calibrated Images . . . . .	125
5.5	Beamforming . . . . .	127
5.5.1	Pulsar Detection . . . . .	128
5.6	Conclusions . . . . .	130
<b>6</b>	<b>Digital Signal Processing in the SKA-low</b>	<b>133</b>
6.1	Beamforming . . . . .	133
6.1.1	Beamforming Architectures . . . . .	134
6.2	SKA-low-1 Baseline Specifications . . . . .	139
6.2.1	Specification Motivations . . . . .	139
6.3	A Computational Argument for Station Configuration . . . . .	143
6.3.1	Channelization . . . . .	147
6.3.2	Summary . . . . .	151
6.4	An FFT-based correlator . . . . .	153
6.5	An FFT-based central beamformer . . . . .	154
6.5.1	The Fractional Fourier Transform . . . . .	156
6.6	Data Transport . . . . .	158
6.6.1	Internal Station Data Rates . . . . .	159
6.7	Implementation Platforms and Upgrade cycles . . . . .	160
6.7.1	Upgrading Hardware . . . . .	165

6.8	Conclusions . . . . .	167
<b>7</b>	<b>Conclusions</b>	<b>169</b>
7.1	Direct-Imaging as an FX Correlator Alternative . . . . .	169
7.2	Future Developments . . . . .	171
7.3	FX Correlator Developments . . . . .	172
7.4	The Square Kilometre Array . . . . .	172
7.5	Concluding Remarks . . . . .	173
	<b>References</b>	<b>174</b>
<b>A</b>	<b>Convolution Theorem</b>	<b>185</b>
<b>B</b>	<b>Two's Complement Number Representation</b>	<b>186</b>
<b>C</b>	<b>The CASPER X-Engine Architecture</b>	<b>188</b>

# Chapter 1

## Introduction

I look upon Karl Jansky's paper, "Electrical Disturbances Apparently of Extraterrestrial Origin", as a classic. It is in effect a wedding ceremony. It weds the science of astronomy and the science of radio and electrical engineering. It ties these sciences together by inseparable bonds.

— *Cyril Methodius Jansky, Jr.*, 1956

When Karl Jansky, working at Bell Telephone Laboratories, observed the first extra-terrestrial radio emissions originating from the centre of the Milky Way (Jansky, 1933) his results sparked a field of research which has grown into modern radio astronomy. The tools of radio astronomers have progressed significantly since this period, owing in part to a significant injection of expertise from the radar engineers of World War II. Today, radio telescopes represent some of the largest scientific instruments ever constructed. This development is exemplified by Figure 1.1, which shows a replica of Jansky's 20.5 MHz antenna as well as the central core of the Low Frequency Array (LOFAR), which was constructed some eight decades later. LOFAR is a telescope comprising from many thousands of individual antennas and represents the state-of-the-art in modern low-frequency astronomical observatories.

Radio astronomy shares its goals with its microwave, infra-red, optical, x-ray and gamma-ray counterparts. These are to observe the radiation incident from distant sources and so infer what these sources are, the processes which give rise to their emission, and the properties of the Universe which lead to such events. Progress towards these goals is made by continual improvements in the sensitivity, field-of-view, and resolution of telescopes. In the case of radio astronomy, the demand for sensitivity has driven the construction of telescopes with ever larger collecting areas, whilst high resolutions have been achieved by coherently combining signals received by multiple separated detectors.

The striving of radio astronomers to build ever-superior instruments is embodied by the Square Kilometre Array (SKA) project. This is an international, multi-billion Euro



(a) A full-size replica of “Jansky’s merry-go-round”, a rotating antenna designed to measure 20.5MHz interference sources which could degrade transatlantic voice communications. It was with an array like this that Jansky first detected extra-terrestrial radio emissions.



(b) The “Superterp” of the Low Frequency Array, which operates between 10MHz and 240MHz. This photo shows only a fraction of the total telescope, with the circular Superterp island itself covering a diameter of  $\sim 250$  m. LOFAR currently comprises an international network of 54 sub-arrays, with signals combined from more than 5,000 individual antennas. Photo reproduced from van Haarlem et al. (2013).

Figure 1.1: Eighty years of radio astronomy: (a) Jansky’s antenna, which was not originally designed as an astronomical instrument. (b) The core of the Low Frequency Array (LOFAR), constructed some eight decades later.

effort to build a telescope with over one hundred times the collecting area of any existing instrument. The SKA will utilize a variety of novel technologies to maximize its ability to enable science using observations between 50 MHz and 10 GHz. Specifically concerning the work in this thesis, at mid–low ( $\lesssim 1400$  MHz) frequencies, the SKA will utilize aperture arrays, comprising many thousands of small  $O(\text{m}^2)$  receiving elements. These can be used to synthesize large, steerable effective apertures, with no mechanical moving parts.

This thesis is concerned with the way in which signals from such distributions of receiving elements are combined, both in terms of the algorithms employed to reach final output data products, and the practical implementations of such systems.

In this introductory chapter of this thesis, the science opportunities which motivate the construction of the low-frequency SKA—and large, multi-element radio telescopes in general—will be outlined. The technical requirements of these science cases, which drive the technological innovation present in the SKA’s design, will be elaborated upon. The question addressed in this thesis is posed: what is the most effective way to process signals in large, low-frequency, interferometric arrays? Finally, an outline of the progression of work presented in this thesis is given.

## 1.1 Information in the radio sky

The brightness distribution of radio-frequency radiation on the celestial sphere encodes information that enables us to better understand the Universe in which we live. When radio astronomers attempt to measure this distribution, experiments fall into two distinct categories. In the first case, measurements of the properties of the extended features apparent in the radio sky allow inferences to be made about the cosmos at large. In the second case, targeted observations of spatially confined structures (many so small as to effectively appear as point sources) provide insight into specific, isolated phenomena.

It is the first of these scenarios which most strongly motivates the work presented in this thesis. A brief summary of some of these cosmological radio features, and their physical origins, is given in the remainder of this section.

### 1.1.1 The Evolving Universe and Radio Astronomy

Current cosmological predictions age the Universe at approximately 13.8 billion years (Bennett et al., 2012), and predict a number of distinct phases in its evolution, which radio astronomy is uniquely placed to observe.

The early Universe existed in a very hot, dense state, far too energetic for atomic matter to form. With a vast population of free electrons, photons at this time were frequently scattered, and the Universe was opaque. It was not until the Universe expanded and cooled




Figure 1.2: The all-sky map of temperature fluctuations created from nine years of data from the Wilkinson Microwave Anisotropy Probe. Reproduced from Bennett et al. (2012).

to a temperature of  $\sim 3000$  K that protons and electrons combined to form atoms. At this time, the photons and baryons became decoupled and the Universe transitioned to an effectively transparent state. The relic radiation from the epoch of decoupling has been propagating freely through space ever since. Today, having been redshifted by a factor of  $\sim 1100$  by the expanding Universe, it is observable as the Cosmic Microwave Background (CMB), a near-perfect black-body spectrum at a temperature of 2.7 K.

### The Cosmic Microwave Background

The CMB signature was discovered serendipitously in 1964 (Penzias & Wilson, 1965; Dicke et al., 1965), and has since been measured with exquisite detail (Bennett et al., 1996; Netterfield et al., 2002; Bennett et al., 2012; Ade et al., 2013). Anisotropies in the CMB signal at a level of 1 part in  $10^5$  have been extensively mapped (Figure 1.2). Observations show characteristic structure resulting from acoustic waves propagating in the photon-baryon content of the Universe before decoupling. These features have characteristic scale sizes, and are shown in Figure 1.3.




Figure 1.3: Angular power spectrum of seven-year WMAP data. The solid curve represents a best-fit  $\Lambda$ CDM cosmological model. The grey band represents uncertainty due to cosmic variance. For a detailed discussion, see Larson et al. (2011), from which this figure has been reproduced.

### The Epoch of Reionization

Much later in the history of the Universe, as the first stars and other radiating objects formed, high-energy UV photons from these sources caused the reionization of the hydrogen atoms in the universal medium. As the number of radiating sources increased, the Universe underwent a transition into an almost completely ionized state. This transition marks a major transition in the evolution of the Universe, and is referred to as the Epoch of Reionization (EoR).

Periodic reviews of EoR science exist in the literature (Furlanetto et al., 2006; Morales & Wyithe, 2010). It is generally accepted that reionization occurred in bubbles of gas, which grew in number and size as the number of radiating objects increased. Numerous simulations of this era of cosmological history exist. An example by Trac & Cen (2007), which shows the increase in the fraction of hydrogen which is ionized as the Universe evolves through redshifts 9–6, is shown in Figure 1.4.

Few strong constraints on the details of the EoR can be made with current observational data. Measurements of the free-electron scattering of CMB photons gives information on the ionization state of the medium between the earth and the surface of last scattering, with the three-year WMAP results (Spergel et al., 2007) giving a maximum likelihood for the age of the reionized Universe of 365 million years, i.e.  $z = 11.3$ . However, such measurements only give information on the integrated line of sight to the surface of last scattering, and leave the details about the sources of reionization and the duration of the reionization transition




Figure 1.4: Evolution of HII (i.e. ionized hydrogen) density through  $z = 9, 8, 7, 6$  (top left, top right, bottom left, bottom right, respectively) in a  $50 \times 50 \times 1 h^{-1}\text{Mpc}$  simulation by Trac & Cen (2007). Brighter regions indicate higher HII density. Volume-weighted HI (neutral hydrogen) fractions are 0.67, 0.46, 0.12 and  $1.5 \times 10^{-4}$  respectively, though these are not shown. Image and accompanying video animation available at <http://www.cfa.harvard.edu/~htrac/Reionization.html>.

largely unspecified.

A further constraint on the redshift at which reionization occurred comes from the study of absorption features in the spectra of distant galaxies. The Lyman-alpha forest, which is a series of absorption lines caused by neutral hydrogen, can be used to infer properties of the ionization state of the Universe. Such observations suggests that the Universe is already mostly devoid of neutral hydrogen by  $z = 6.5$  (Fan et al., 2006). However, without galactic spectra available at higher redshift, little can be learned about the earlier evolution of the reionization process.

The atomic hydrogen emission line at 21 cm (1.4 GHz) can provide a direct trace of the neutral hydrogen content of the Universe, and is thus an ideal feature with which to characterize the transition of the Universe from a neutral to ionized state. The 21-cm line has long been used to study the interstellar medium of our own and other galaxies, but only recently have attempts been made to measure the small variations in 21-cm brightness temperature from the diffuse, high-redshift sources which characterize the reionization transition.

The spatial power-spectrum of the HI signal can constrain not only the period over which reionization took place, but also the sources responsible for it. Furlanetto et al. (2004) show that the power spectrum of the HI signal evolves throughout the reionization process, and in particular has a feature corresponding to growing bubbles of ionized gas, the characteristics of which are dependent on the ionization history. A range of power spectra corresponding to different reionization histories, from Furlanetto et al. (2006), are shown in Figure 1.5. The potential to directly measure the features of this spectra is one of the great promises of 21 cm observatories.

### **Baryon Acoustic Oscillations**

The acoustic waves which give rise to the power spectrum of the CMB can also be observed at later cosmological epochs, manifesting themselves as preferential large-scale clustering of galaxies at specific length scales. These oscillations in the matter power spectrum are known as baryonic acoustic oscillations (BAOs). BAOs provide a standard ruler, which grows with the expanding Universe and allows astronomers to measure distances over cosmic time. Critically, since the BAO ruler expands with the Universe, tracking it through time gives an insight into the dark energy which drives expansion. BAO structure has already been observed in optical surveys of luminous galaxies at redshift  $z \approx 0.35$ , but the opportunity remains to observe BAOs at earlier times using the observed power spectrum of 21 cm hydrogen emission. Baryon acoustic oscillations, expressed as the relative matter power spectrum compared with a spectrum in the absence of baryons, are shown in Figure 1.6. Also shown are the error bars with which such oscillations could be measured by a proposed




Figure 1.5: RMS variation in the HI brightness temperature on different length scales. The five curves shown represent different possible ionizing efficiencies giving rise to different ionized hydrogen fractions at  $z = 10$ . Reproduced from Furlanetto et al. (Figure 20, 2006).

observing campaign with a  $200 \times 200$  m telescope targeting the  $1 \lesssim z \lesssim 2$  redshift range (Chang et al., 2008). Wyithe et al. (2007) find that observations of redshifted 21 cm radiation with a large-scale facility such as the SKA could effectively identify the BAO peak at redshifts  $3.5 \lesssim z \lesssim 12$ , with a precision in some ranges significantly exceeding current measurements.

## 1.2 Probing Global structure with Hydrogen Observatories

So redshifted 21-cm radiation, and its distribution across the sky, can give insight into the evolution of the Universe – but what sort of telescopes do we need to measure it?

The cases of both EoR and BAO studies demand observations of the 21-cm hydrogen emission line out to high redshifts. Designing a telescope to detect this radiation means determining what combination of parameters—frequency range, resolution, sensitivity, field of view—give the greatest opportunity to observe the signal we seek. These are examined in turn, below.

### Frequency Range

Deciding on a frequency range to observe is perhaps the easiest telescope parameter to fix. Depending on the redshift of the epoch we wish to investigate, the HI line, with a rest frequency of 1.4 GHz, is observed at lower frequencies. Figure 1.7 shows the observed frequencies and wavelengths of the HI line at different redshifts. Thus, an experiment designed to detect the signature of reionization prior to  $z = 6$  should operate at frequencies




Figure 1.6: Predicted baryon acoustic oscillation amplitudes, relative to a power spectrum with no baryons. A smooth curve has been removed to emphasize the oscillations. Error bars represent sensitivity predictions of a plausible  $200 \times 200$  m telescope (specified by Chang et al., 2008) designed to detect BAOs. Reproduced from Chang et al. (Figure 2, 2008).




Figure 1.7: Observed frequency and wavelength of 21 cm (1.4GHz) hydrogen emission at redshift ranges  $z \leq 12$ .

$< 200$  MHz. Similarly, BAO observations in the region  $3.5 \lesssim z \lesssim 12$  require observation bands covering  $110 \lesssim f \lesssim 315$  MHz. BAO observations in the more recent  $1 \lesssim z \lesssim 2$  range must operate at  $450 \lesssim f \lesssim 700$  MHz.

## Resolution

The resolution needed by a telescope wishing to observe features in the HI power spectrum due to reionization and BAOs is simply set by the minimum scale of these features. If a feature is predicted at a specific angular scale,  $l$ , a telescope should have a resolution of  $\sim \frac{l}{2}$  in order to Nyquist sample this variation.

In the case of BAOs, oscillations are damped at small scales by non-linear evolution (see Bassett & Hlozek (2010) for an overview of BAO physics). In essence, galactic structure which forms in the acoustic peaks of BAOs are subject to a net gravitational force from their local matter distribution. Averaging over many galaxies, these forces cause a reduction in the size of the observed power spectrum peaks, attenuating the BAO signal with increasing impact at small scales. Thorough treatment of non-linear effects (Crocce & Scoccimarro, 2008) suggests that the smallest scale of interest is  $35 \text{ h}^{-1} \text{ Mpc}$  (where  $\text{h}^{-1}$  is the dimensionless Hubble constant, approximately equal to 0.7). This implies a telescope with resolution  $17.5 \text{ h}^{-1} \text{ Mpc}$  and corresponds to an angular resolution on the sky of approximately 20 arcminutes. Since the angular resolution of a telescope (in radians) is approximately equal to the ratio of its observing wavelength and diameter, this requires a telescope approximately 200 wavelengths in diameter. At redshift  $z = 2$  (where the HI line is observed at  $\sim 470$  MHz) this telescope would need to have a size  $\sim 130$  m.

For the EoR, since ionized pockets of hydrogen evolve throughout the reionization process, a wide range of scales at various epochs are of interest. However, most of these can be captured by telescopes of kilometre-scale, operating in the 50–200 MHz range.

## Sensitivity

The sensitivity of a telescope quantifies its ability to reach a certain signal-to-noise level given an observation over a finite bandwidth and time. The signal-to-noise of a measurement is determined by the observation bandwidth,  $\Delta\nu$ , integration time,  $\Delta\tau$ , and the system temperature,  $T_{\text{sys}}$ , of the telescope.  $T_{\text{sys}}$  is a measure of the instantaneous noise contribution to a measurement from analogue electronics, the sky (which is at a finite temperature itself) and potentially any other sources of noise in the telescope receiver system. An ideal resistor at temperature  $T$  will generate a bandwidth limited noise power of  $k_B\Delta\nu T$ , where  $k_B$  is the Boltzmann constant, equal to  $1.38 \times 10^{-23} \text{ J/K}$ . (Johnson, 1928). This relationship can be

used to define  $T_{\text{sys}}$ , by equating it to the physical temperature of a resistive noise source which would generate the same noise power,  $P$ . Explicitly:

$$T_{\text{sys}} = \frac{P}{k_B \Delta\nu}. \quad (1.1)$$

Using this definition, the noise on a measurement from a telescope is then determined by the ideal radiometer equation (see, for example, Chapter 3, Burke & Graham-Smith, 2010):

$$\Delta T = \frac{T_{\text{sys}}}{\sqrt{\Delta\nu \Delta\tau}}. \quad (1.2)$$

The sensitivity to a point source can be expressed by considering the effective flux collecting area of a telescope,  $A_{\text{eff}}$ , which yields an expression for the noise on a measurement of flux,  $\Delta S$ , from a point source:

$$\Delta S = \frac{2k_B T_{\text{sys}}}{A_{\text{eff}} \sqrt{\Delta\nu \Delta\tau}}. \quad (1.3)$$

The factor of two emerges because only half the flux from an unpolarized source is measured by a single-polarization radiometer. Equation 1.3 shows the critical dependence of point source sensitivity to  $\frac{A_{\text{eff}}}{T_{\text{sys}}}$ , and thus this ratio is often quoted when specifying the top-level requirements of a telescope. It is important to recognize that  $A_{\text{eff}}$  is not the physical collecting area of a telescope, but rather is related to the gain (or directivity<sup>1</sup>),  $G$ , of an antenna (see Chapter 2, Johnson & Jasik, 1984), with

$$A_{\text{eff}} = \frac{\lambda^2 G}{4\pi}. \quad (1.4)$$

For a hypothetical isotropic radiator,  $G = 1$ , whilst for a half-wave dipole  $G = 1.64$ . The sensitivity of a telescope to a point source can always be improved by building a more directive antenna (or antenna array) with a larger effective area.

When the goal is to detect large-scale structure in the sky this is not the case. Since large-scale structure is, by definition, spread over large areas of the sky, sensitivity to its presence cannot be improved by increasing antenna directivity. For large-scale brightness distribution experiments the only reason to increase the effective area of an array is to resolve a particular spatial component of interest of this distribution. Rather, for achieving high surface-brightness sensitivity, the important quantity to maximize is the filling factor,  $f$ , of a telescope (see for example, Chapter 1, Price, 2012). This is defined by

$$f = \frac{A_{\text{eff}}}{A_{\text{phys}}}, \quad (1.5)$$

---

<sup>1</sup>These quantities differ only by a factor relating to antenna efficiency, with Antenna Gain = Antenna Efficiency  $\times$  Directivity. Antenna gain can be measured, while directivity is a theoretical quantity.

where  $A_{\text{phys}}$  is the telescope’s total physical extent. This quantity is maximized for filled aperture telescopes, such as single parabolic dishes, for which  $f \approx 1$ . The resulting sensitivity of a telescope to surface brightness temperature is (Section 9, Furlanetto et al., 2006)

$$\Delta T = \frac{T_{\text{sys}}}{f \sqrt{\Delta\nu \Delta\tau}}. \quad (1.6)$$

### Field of view

The field of view required for HI intensity mapping experiments needs to be large enough to obtain a representative sample of the Universe. Mellema et al. (2012) find that a field of view of 25 square degrees is sufficient to limit sample variance errors to approximately 3% on the most interesting scales for EoR science. To further reduce this error to the 1% level would require observing an order of magnitude more sky.

### The Ideal Telescope

The ideal telescope for HI intensity mapping at low frequencies (in particular, for EoR and high-redshift BAO science goals) has:

- A large field of view (at least  $\sim 10$  sq. deg).
- A resolution on  $\sim 0.1^\circ$  scales.
- A highly filled aperture.

For radio telescopes operating at  $\sim 100$  MHz these requirements cannot be simultaneously fulfilled by single aperture telescopes. To achieve the necessary resolution at low frequencies requires extremely large dishes, many hundreds of metres across. These are both prohibitively expensive, and have relatively small fields of view. Instead, a potentially attractive solution is a highly-filled array of antennas, which has an effective area and resolution similar to that of a single large aperture, but achieves a far wider field of view. The combination of demanding a large collecting area, high filling factor and high resolution simultaneously necessarily results in the requirement of telescopes with very large numbers of individual receiving elements.

#### 1.2.1 Observational Challenges

Whilst the signal of interest drives the requirements of Hydrogen observatories as described above, there are further observational challenges which must be overcome in order to measure the global 21cm Hydrogen signal to the precision required by BAO and EoR science goals. These are briefly outlined in the remainder of this section.

## Foregrounds

The small signatures which Hydrogen observatories aim to detect are masked by various larger “foreground” signals. These foregrounds can be divided into emission originating in our own galaxy, and emission from more distant sources.

When considering EoR (and other low frequency) observations, galactic foregrounds represent the largest contribution to detected signals. This galactic emission can further be categorized into three main components. In order of decreasing magnitude, these are Galactic diffuse synchrotron emission (GDSE), synchrotron emission from discrete sources, and free-free (bremsstrahlung) emission in diffuse ionized gas.

Galactic diffuse synchrotron emission (GDSE) originates from free electrons moving in the Galactic magnetic field, with a brightness temperature which has a power law spectrum  $T_b(\nu) \approx \nu^\beta$ . At high galactic latitudes, at frequencies of 100 MHz,  $\beta \approx -2.55$ , with steepening at higher frequencies (Banday & Wolfendale, 1991; Platania et al., 1998). The absolute temperatures amount to  $O(100\text{ K})$  at high galactic latitudes and EoR frequencies (Landecker & Wielebinski (1970), Figure 1.8). At higher frequencies, relevant for BAO studies, the GDSE temperature is somewhat lower;  $\sim 6\text{ K}$  at 408 MHz and  $4\text{ K}$  at 820 MHz (Banday & Wolfendale, 1991). The spectral indices themselves are uncertain and vary over position and frequency, with standard deviation  $\Delta\beta \approx 0.15$  (Tegmark et al., 2000).

Synchrotron emission is also generated by discreet supernova remnant sources in the galaxy. These are mostly located in the galactic plane, with only a handful present in high galactic latitude EoR fields (Green, 2009). As with the GDSE, emission from SNRs is spectrally smooth, with variation amongst sources. Diffuse free-free emission further contributes a small (about one part in a hundred) contribution to foreground temperatures. Like the other components, it has a smooth spectral dependence at low frequencies, with a spectral index of  $-2.1$ .

Further emission is due to extragalactic sources. The integrated effect of these sources is to add a thermal background of  $\sim 50\text{ K}$  at 150 MHz, with a spectral index of  $-2.65$  (Shaver et al., 1999) and variation  $\approx 0.2$  (Cohen et al., 2004).

## Foreground Removal and Detection Feasibility

With the EoR global signal appearing with a temperature of only a few mK (Furlanetto et al., 2006) at  $\sim 150\text{ MHz}$ , and the BAO signal  $\sim 100\ \mu\text{K}$  at  $\sim 500\text{ MHz}$ , it is reasonable to ask how feasible it is to detect these signatures in the presence of the astronomical foregrounds described above. There are two separate considerations to make in answering this question.




Figure 1.8: The brightness temperature of the sky at 150 MHz, with contours shown at 180 (dashed), 270, 360, 540, 1100, 2200, 3300, 4400, and 5500 K. Reproduced from Furlanetto et al. (2006), with original data from Landecker & Wielebinski (1970). The observing fields of three EoR telescopes — LOFAR, PAST and MWA, are shown, with bold, solid lines indicating locations of constant declination.

The first is to ask whether, given the temperatures of the foregrounds, is it possible to build a telescope with sufficient sensitivity to detect the signals of interest, which are many orders of magnitude smaller. This question can be easily answered using Equation 1.6.

For a BAO telescope operating at 500 MHz, the system temperature is likely to be dominated by the temperature of the receiver electronics, rather than the sky signal. A reasonable estimate of  $T_{\text{sys}}$  in this regime is 50 K, in which case an observing bandwidth-time product of  $\Delta\nu\Delta\tau = 2.5 \times 10^{11}$ . For  $18h^{-1}$  Mpc scales of interest, the frequency resolution at  $z = 1.5$  need be about 3 MHz, implying an integration time of  $\sim 10^5$  seconds ( $\sim 1$  day) per pixel. It is thus reasonable to survey small patches of sky with a large filled aperture telescope, or larger patches using a multi-receiver design to increase field of view. In the latter case, assuming a  $200 \text{ m} \times 200 \text{ m}$  telescope comprising 256  $12.5 \text{ m} \times 12.5 \text{ m}$  apertures, we arrive at the sensitivity limits, after 100 days observing, shown in Figure 1.6 (Chang et al., 2008).

A similar calculation can be made for EoR observations. At these low frequencies, the sky temperature (rather than the temperature of the telescope receiver itself) dominates  $T_{\text{sys}}$ . With  $T_{\text{sys}} = 150 \text{ K}$ , and a target sensitivity of 5 mK, the required observing bandwidth-time product is  $\Delta\nu\Delta\tau = 10^9$ . With plausible bandwidths of a few MHz, the sensitivity requirement only demands an integration time of several minutes.

The reality, of course, is not so simple. Whilst the above sensitivity calculations treat foregrounds solely as a contributor to system noise, which can be unendingly integrated

down, the reality is that power-spectrum experiments seek to identify a particular spectral signature in the presence of complex background noise, which has its own direction and frequency dependencies.

The problem of foreground contamination is one frequently addressed by EoR science groups (see, for example, Shaver et al., 1999; Morales et al., 2006; Liu et al., 2009; Parsons et al., 2012). While the exact techniques used to isolate the cosmological signal of interest from the galactic and extragalactic foregrounds varies between observatories, the basic methods employed use exploit the differences in symmetry between the cosmological signal and the foreground. Foregrounds are spectrally smooth, with power-law frequency dependence, whilst the EoR signature is expected to be a power spectrum with features that are present across both the angular and frequency dimensions. Similarly, the BAO signal can be distinguished from foregrounds by considering its frequency dependence. Unlike the smooth varying contaminant sources, the BAO signature is largely uncorrelated between redshift slices wider than 3 MHz (Chang et al., 2008).

If the modeling of foreground spectra with power laws were perfect, these unwanted signals could be easily subtracted. However, the spectral index representation is dependent on both frequency and direction. Galactic foregrounds have spectral indices which are measured to steepen with increasing frequency (Banday & Wolfendale, 1991), whilst extra-galactic foregrounds are the sums of multiple sources with different power-law characteristics (Green, 2009). Furthermore, instrumentation effects can conspire to mix spatial and frequency modes, so that variation of foregrounds as a function of direction may appear as spectral structure.

In the absence of instrument-induced corruptions, it is possible to calculate how the uncertainty in spectral index of foregrounds impacts the ability to reach necessary sensitivity levels. Chang et al. (2008) calculate that with a conservative estimate on spectral index uncertainty,  $\Delta\alpha = 1$ , the first three peaks of the BAO spectrum remain comfortably observable out to at least  $z = 2.5$ . Furlanetto et al. (2006) make an estimation of the extent to which foregrounds can be suppressed and find that plausible uncertainties in spectral index still allow foregrounds to be suppressed by a factor of more than  $10^6$ . This is sufficient for EoR detection in the high galactic latitudes, where the foregrounds dominate the EoR signature by “only” five orders of magnitude.

A more challenging limitation of foregrounds occurs because of instrumental response, which can damage the fidelity of the obtained sky maps if not known precisely. The frequency dependence of a telescopes response pattern — i.e., the frequency dependence of side-lobe positions — will tend to couple variation in the angular spectrum of a foreground signal into frequency space, undermining the assumption of smooth spectral shape. The

effects of galactic magnetic fields can also couple spatial features in the polarized galactic signal into spectral variation in the unpolarized total-power signal via Faraday rotation. For these reasons, as well as sensitivity and resolution, 21-cm observatories need place a large emphasis on calibration and stability. It is extremely important that the response of instruments is well-characterized and predictable. Indeed, it is the case that foreground removal requirements require EoR instruments to have much more stringent specifications than their science goals alone. Before any of the foreground removal steps above are performed, first the brightest point sources detected must be subtracted such that their residuals are below mK levels. In order to achieve this subtraction with sufficient accuracy — and to resolve and detect all sources of sufficient brightness — the resolution of an EoR array must be significantly greater than that required by power spectrum measurements. This requirement gives rise to tapered arrays, where a central, highly-filled core is required for EoR measurements, whilst smaller numbers of long baselines are required to aid point-source removal. Processing data from such multi-purpose arrays is discussed more in the context of the SKA in Chapter 6.

## 1.3 Coherent Manipulation of Radio Signals

Coherent manipulation of signals from multiple, independent apertures effectively decouples the relationship between field of view, resolution, and effective area, which are fixed in simple single-dish, single-pixel telescope systems. This allows greater flexibility to tailor a specific telescope to a given science case. In this section, we briefly explore the mechanisms employed in aperture-synthesis telescopes and describe how these set the overall characteristics of a telescope, and impact its suitability for different observational tasks.

### 1.3.1 Beamforming

In an array of  $N$  receiving elements, a radiation wavefront from a particular direction will reach element  $n$  at time  $t_n$ . The beamforming process relies on the fact that for a given array configuration the relative delays between arrival times  $t_{0\dots N-1}$  are a function of the direction of propagation of the incident wave and can be calculated. Artificial delays (in the time domain) or phase shifts (in the frequency domain) can then be applied to the received signals from each antenna to cause signals from a desired direction to add constructively when element signals are summed. Figure 1.9 illustrates this process for the simple case of a one dimensional array of three randomly spaced antennas. The beamforming process is recognizably equivalent to that of lens in an optical telescope, or a parabolic reflector in a radio dish, with only a subset of the collecting area being illuminated. Beamforming




Figure 1.9: An example of a simple 3-antenna beamformer synthesizing a beam in direction  $\theta$ . Signals from spatially separated antennas are delayed and summed, such that the response is greatest in the desired direction; in this case, towards radiation with unit wavevector  $\hat{\mathbf{k}}$ .

multiple antennas increases the collecting area and resolution of a telescope, whilst reducing the information output by the array to a single signal stream.

The response,  $B(\mathbf{k}, \mathbf{k}_o)$ , of a beam pointing in direction  $\mathbf{k}_o$  to radiation with other wavevectors  $\mathbf{k}(\theta, \phi)$  is given by summing signals contributions from individual antennas. If each antenna in an array has an individual response to radiation given by  $A_n(\mathbf{k})$ , this is given by

$$B(\mathbf{k}, \mathbf{k}_o) = \sum_{n=0}^{N-1} A_n(\mathbf{k}) e^{i(\mathbf{k}-\mathbf{k}_o) \cdot \mathbf{r}_n}, \quad (1.7)$$

where  $\mathbf{r}_n$  is the position vector of antenna  $n$ . In other words, the beam response is given by the Fourier transform of the distribution of receiving elements, weighted by their individual antenna responses. In the case that these responses do not vary with antenna, the field of view of the synthesized beam is approximately given by  $\frac{\lambda}{r'_{max}}$ , where  $r'_{max}$  is the maximum separation of antennas projected towards the beam direction.

Prior to beamforming, signals from individual antennas in an array see a portion of sky determined by their response  $A_n(\mathbf{k})$ . It is therefore possible to effectively form beams in any direction where  $A_n(\mathbf{k})$  is significantly non-zero. By making copies of antenna signals and adding them with different time shifts or phase weightings it is also possible to simultaneously form multiple beams. In the frequency domain, this can be expressed as a matrix-vector multiplication of a matrix,  $\mathbf{M}$ , of phase coefficients and a vector,  $\mathbf{x}$ , of narrowband signals,  $x_n$ , from each of the array elements. Each row of  $\mathbf{M}$  represents a set of coefficients for a specific beam. After multiplication, the vector of beam signals is given by

$$\mathbf{b} = \mathbf{M}\mathbf{x}. \quad (1.8)$$

Choosing the number of beams to form amounts to balancing the field of view observed by an array (which increases approximately linearly with each beam added), with the processing capacity required to evaluate Equation 1.8.

### 1.3.2 Interferometry

Interferometry has a long history in astronomy. It was first used in optical observations by Michelson & Pease (1921) to obtain sufficiently fine resolution to measure the diameters of nearby stars. Similarly to beamforming (which is in itself one type of interferometry), interferometry allows signals from separated apertures to be combined with a spatial resolution that depends on their maximum separation distance. The field of view of an interferometer is determined by the individual element sizes. The cornerstone of interferometry is

the van Cittert-Zernike theorem, which states that, in the flat sky approximation, the visibility measured by a two-element interferometer,  $V(u, v)$ , is given by a two-dimensional Fourier transform of the apparent sky brightness distribution seen by the pair of antennas. Explicitly:

$$V(u, v) = \int A(l, m)I(l, m)e^{2\pi i(ul+vm)}dldm. \quad (1.9)$$

In this equation,  $I(l, m)$ ,  $A(l, m)$  are the sky intensity distribution and the response of the individual antennas, and  $u$  and  $v$  are the spacings of the two antennas projected towards the “phase center”,  $l, m = 0$ , of the interferometer and measured in units of wavelength.

The visibility itself can be measured directly. In the case that each antenna signal is a single, narrowband channel, it is the time-averaged product of signals from the two antennas comprising the interferometer. Using the van Cittert-Zernike relationship, one is able to build up an image of the apparent sky brightness distribution by measuring various Fourier modes of this distribution with different two element interferometers in an array.

Another, perhaps more intuitive way to view an interferometer, is within the context of beamforming. Taking the power of a single beam,  $b$ , formed by Equation 1.8 amounts to measuring

$$\begin{aligned} b^*b &= (M_0x_0 + M_1x_1 + \cdots + M_{N-1}x_{N-1})^* (M_0x_0 + M_1x_1 + \cdots + M_{N-1}x_{N-1}) \\ &= M_0^2x_0^2 + M_1^2x_1^2 + \cdots + M_{N-1}^2x_{N-1}^2 \\ &\quad + 2\Re (M_0M_1x_0x_1 + M_0M_2x_0x_2 + \cdots + M_{N-2}M_{N-1}x_{N-2}x_{N-1}). \end{aligned} \quad (1.10)$$

In other words, the power of a beam is simply a linear combination of the cross products of antenna signals,  $x_ix_j$ , with different weightings corresponding to different sky directions. As such, an interferometer can be considered to be the most general case of a beamformer. By explicitly measuring antenna cross products, an interferometer allows beams to be synthesized *post-hoc* in as many arbitrary directions as are desired. When presented in this form, an important observation is that a beamformer necessarily includes the autocorrelations,  $x_n^2$ , when calculating the power in a beam. In contrast, beams synthesized using measured visibilities need not include these. As a result, synthesized beams constructed piecemeal using a correlator can be far less noisy and more stable than the equivalent beamformed data product. We will discuss this further in Chapter 2.

## 1.4 Interferometric Radio Telescopes

Early interferometric radio telescopes were completely analogue in nature. The first was a sea-cliff interferometer which used a single receiver placed on a cliff top to perform the radio equivalent of the Lloyd’s mirror experiment (Bolton & Slee, 1953). Radiation from extra-terrestrial sources was interfered with its reflection from the sea to obtain an interference pattern. Later, Ryle constructed a two element interferometer at Cambridge University (Ryle, 1952). Ryle’s interferometer used a  $\frac{\lambda}{2}$  length of cable inserted into the signal path of one receiver to periodically switch the phase of one of the antenna signals. A power detector was then used to obtain the squares of the antenna signal sums and differences. These two measurements obtained by “phase switching” could then be differenced to obtain the multiplication of the two signals, avoiding the autocorrelation terms in Equation 1.10. Later still, Ryle developed an interferometer whose elements could be moved to sample multiple spatial modes (Ryle & Hewish, 1960). It wasn’t long before interferometers were able to capitalize on advances in wideband, analogue multiplying circuitry (Allen & Frater, 1970), which enabled direct computation of antenna pair correlation products.

Today the equipment developed by the early pioneers of interferometry has largely been superseded by digital computers, though the techniques developed over previous decades remain very similar. Namely, this involves deconstructing a potentially large array of  $N$  antennas into a collection of  $O(N^2)$  two-element interferometers. Multiplying their signals yields a set of visibilities which can be used to obtain measurements of different Fourier modes of the sky brightness distribution.

### 1.4.1 The Age of Digital Radio Astronomy

The number of transistors on integrated circuits (ICs) is observed to have approximately doubled every two years for the last three decades. This exponential progression, and the closely tied improvements in IC performance, is known as Moore’s law, after first being noted in 1965 by Intel co-founder Gordon Moore (Moore, 1965). With the Moore’s law progression of digital processing power, or the equivalent exponential reduction in processing cost, the requirements of modern telescopes are usually far more economically fulfilled by digital, rather than analogue, instruments. In fact, the analogue processing systems used by early radio astronomers have largely been made obsolete by their digital counterparts, except in the case of very large bandwidth (see, for example, Holler et al., 2011).

Since the first digital multi-channel spectrometer, which used 1-bit signal digitization (Weinreb, 1963), the allure of the accuracy and flexibility of digital processing techniques has

been clear. Today, modern telescopes frequently employ digital converters offering  $\gtrsim 8$  bits of precision and sampling rates of many gigahertz to perform large bandwidth experiments.

In the case of synthetic aperture arrays, digital radio telescopes offer an escape from the signal routing complexities of analogue correlators, which quickly become insurmountable for arrays with large numbers of elements or where large numbers of frequency channels are required. In contrast, resolution bandwidth in a digital telescope is solely limited by processing capabilities, and instruments with millions of frequency channels, such as Serendip IV, have been operating for over a decade (Werthimer et al., 2000). For arrays with large numbers of elements, digitization offers the capability to flexibly copy and route signals to different processing nodes using relatively cheap data transport technologies with no danger of data degradation. In particular, high-bandwidth international fibre-optic links have allowed telescopes to be distributed over very large geographical areas whilst still allowing real-time processing of data.

The flexibility afforded by digital signal processing techniques leads naturally to a question: what is the best way to process signals from an array of antennas? Despite the progress digital radio astronomy has made, in many cases signal processing methodology has changed very little. In the case of digital aperture synthesis arrays, whilst the number of elements in arrays has grown with the increase in computational capacity, the paradigm of treating large aperture synthesis arrays as collections of two-element interferometers has not changed. As a result, the unfavourable  $O(N^2)$  scaling of the number of interferometers in an array remains. This limits the feasibility of extremely large- $N$  telescopes. This is especially the case for low-frequency telescopes, where analogue front-end receivers may be constructed relatively cheaply. A possible alternative to this processing paradigm is to form images of the sky by directly synthesizing large numbers of synthesized beams via Equation 1.8. In such an instrument, these beams can be efficiently formed by Fourier transform algorithms that scale in computational cost with  $N \log N$ . These algorithms, and their theoretical and practical use, form the bulk of the work presented in this thesis.

## 1.5 Current 21-cm Observatories

We conclude this introduction with a brief summary of some of the currently ongoing or planned low frequency HI observatories. In this brief overview, we aim to establish the scales (both in terms of physical size, and number of receiving elements) associated with modern state-of-the-art low frequency astronomy, in order to provide some context for the work of this thesis.

### 1.5.1 EoR Observatories

Whilst efforts have been made to detect the global HI reionization signal (see, for example, Bowman & Rogers, 2011), the majority of current EoR experiments utilize multi-aperture arrays to try to make statistical detections of the EoR signal at multiple different modes. Some of the most promising of these experiments are The Precision Array to Probe the Epoch of Reionization (PAPER, Backer & PAPER Collaboration, 2009, The Murchison Widefield Array (MWA, Tingay et al., 2012) and The Low Frequency Array (LOFAR, van Haarlem et al., 2013). The Square Kilometre Array, which is still in design stage, is not included here, but is discussed extensively in Chapter 6.

**PAPER** is an interferometric array comprising 64 dipole antennas observing between 100 and 200 MHz with near full-hemisphere field-of-view. PAPER emphasizes calibration and stability, opting for single dipoles with smooth predictable response patterns, rather than beamforming multiple elements. Deployed in the Karoo desert in South Africa (with a 32-element array, primarily used for testing, also present in the USA), the PAPER team have observed in multiple array configurations, using both an array formed with random spacings up to  $\sim 250$  m, and a regular array of antennas spanning an area approximately  $12 \text{ m} \times 100 \text{ m}$ . The former has been used for characterization of sources for flux scale calibration (Jacobs et al., 2013). The latter allows maximum sensitivity to particular target scales in the EoR spectrum. Latest PAPER results from observations in the high-redundancy mode using 32 antennas have so far yielded the most stringent limits on the EoR power spectrum (Parsons et al., 2013), with  $2\sigma$  upper limits within an order of magnitude of predicted 21 cm EoR signals.

**The MWA** (Tingay et al., 2013) is an array comprising 128 tiles, each formed from 16 dipole antennas. The array operates in the frequency range 80–300 MHz, with 30 MHz of instantaneous processed bandwidth. In order to maximize sensitivity at EoR scales whilst still providing high resolution capabilities, the MWA features 50 tiles uniformly occupying a central core of 100 m diameter. A further 62 tiles are spread over an area with 1.5 km diameter, with an additional 16 tiles with up to 3 km separation. Current measurements from a prototype 32-tile MWA deployment have not yet provided EoR constraints to rival PAPER, though the recent 128-tile deployment is predicted to obtain a  $14\sigma$  detection of the EoR signal after 1600 hours observing (Hewitt & Murchison Widefield Array Collaboration, 2013; Beardsley et al., 2013).

**LOFAR** is an international, multi-purpose low-frequency telescope covering the 10–250 MHz band with separate low- and high-band antennas (van Haarlem et al., 2013).

LOFAR has stations throughout Europe, though the EoR campaign uses mainly the central 2 km of the array, and uses observations from the high-band antennas which operate from 110–250 MHz. Observations use beams formed from stations of several hundred dipole elements (see Section 1.3), and correspondingly the field of view of LOFAR is significantly smaller than that of either the MWA or PAPER. LOFAR observations are particularly complicated, owing to the complex response of the station beams. Deep images of the North Celestial Pole with a noise level of  $100 \mu\text{Jy}/\text{beam}$  have been demonstrated, which are within a factor of 1.4 of the theoretical minimum (Yatawatta et al., 2013).

### 1.5.2 BAO Observatories

A similarly large number of BAO intensity mapping experiments are also currently planned or under way. Here we very briefly summarize the Green Bank Telescope (GBT), Canadian Hydrogen Intensity Mapping Experiment (CHIME) and BINGO arrays.

**The Robert C. Byrd GBT** is the world’s largest fully steerable radio telescope, located at NRAO Green Bank, West Virginia. It is a 100 m offset paraboloid reflector, effective up to frequencies in excess of 100 GHz. The GBT has been used to conduct a wide, shallow survey of the HI line at  $z \sim 0.8$ , and currently provides the most precise constraints on the Hydrogen density fluctuations available (Switzer et al., 2013). Though undeniably powerful, the limited field of view of the GBT makes surveys up to large scales relatively slow.

**CHIME**<sup>2</sup> is a planned fully-filled 100 m  $\times$  100 m array of cylindrical receivers, with an instantaneous sky coverage of  $\sim 100$  square degrees, achieved using interferometric combination of 1280 feeds. Currently still in development, this array aims to observe from  $z \sim 0.8$  to  $z \sim 2.5$  (i.e. 400–800 MHz) using a drift scan of  $\sim 20\,000$  square degrees each day, and expects to be able to detect the first 3 BAO peaks.

**BINGO** is a single dish BAO experiment utilizing a  $\sim 40$  m reflector fitted with an array of  $\sim 50$  feed horns (Battye et al., 2013). The numerous feeds allow an instantaneous field of view of  $\sim 100$  square degrees, designed to overcome some of the limitations of single-dish experiments. Still in the design phase, BINGO is planned to observe between 960–1260 MHz and make significant detections of acoustic fluctuations up to  $z \sim 0.48$ .

## 1.6 Thesis Layout

This thesis works towards informing an answer to the question: what is the best way to process signals from antennas in low-frequency, interferometric arrays? In it, I present both an in-depth comparison of the potential opportunities and challenges presented by newly-proposed imaging systems, and results from the implementation of such a telescope. I will also explore some of the ways existing processing methodologies can be optimized, and in particular discuss processing requirements and technologies within the context of the SKA.

In Chapter 2, I will give a detailed description of how information from spatially separated antennas may be processed, elaborating on the mathematical underpinnings of beamforming and interferometry already presented. Multiple algorithms for obtaining images of the radio sky are detailed, and the equivalence of different algorithms is shown. I also discuss the computational and practical implications of the choosing a particular imaging technique. The calibration issues associated with each processing mechanism are discussed, with particular emphasis on the need for—and the effects of errors in—real-time calibration in direct imaging.

In Chapter 3, I present the Field Programmable Gate Array (FPGA) as a platform for signal processing in radio astronomy. I describe the cross-multiplication “X-engine”, designed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER), used to build FPGA-based correlators for multi-element radio telescopes. I demonstrate how dramatic efficiency improvements in FPGA-based correlators may be obtained by optimizing the multiplication algorithm for low-bitwidth inputs. I present a new X-engine design featuring drag-and-drop compatibility with the original CASPER module, but offering a 75% reduction in FPGA multiplier resource usage, up to 50% reduction in RAM usage and significant other processing savings for large arrays.

In Chapter 4, I describe the design and implementation of a spatial processing system at the Basic Element for SKA Training II (BEST-2) array. This system utilizes both traditional cross-correlation processing, as well as imaging using the computationally preferable spatial FFT algorithm presented in Chapter 2. The BEST-2 system aims to give a real-world comparison of the two image forming mechanisms. It is also capable of outputting unaccumulated beamformed signals, for use in time-domain science applications. Results from this deployment are presented in Chapter 5.

Chapter 6 explores the digital signal processing requirements of the Square Kilometre Array. These requirements directly inform important design decisions such as the number of antennas in each SKA station, and the number of stations in the array. I present a computational argument for the optimal design parameters, and argue that the current

SKA specifications could be significantly improved upon. The direct-imaging technique as a mechanism for generating large numbers of beams is also presented within the context of the SKA. Finally, I briefly consider the implications of carrying out SKA processing on different platforms, in particular emphasizing the importance of considering the potential for cost-saving via rapid upgrade cycles.

Chapter 7 draws some conclusions on the work presented in this thesis.

## Chapter 2

# Spatial Processing with Large-N Arrays

When radiation from a source is incident on an array of spatially-separated receivers, the signals detected may be processed in a variety of ways. In this chapter the mathematics and general principles of interferometry are presented. Interferometry presents a powerful method of making measurements of the radio-sky using large arrays of receivers. However, there are several potential paths from detected voltages to sky images, each with its own set of costs and benefits. These strategies are discussed in Section 2.4, with their relative merits outlined in Section 2.5. We highlight some of the subtleties of the different processing methods, which are often not discussed in the literature, and attempt to provide an unbiased comparison relevant to modern radio astronomy. In particular, we aim to find the scenarios where direct-imaging techniques may be best exploited.

### 2.1 Properties of astronomical radio signals

Before introducing the methods by which signals received at different antennas in an array are processed, it is useful first to consider some properties of such signals.

#### 2.1.1 Statistical properties

The vast majority of radio emitters in the sky generate radiation which takes the form of Gaussian random noise<sup>1</sup>. Following Zernike (1938), and considering at first only quasi-monochromatic radiation, we might consider radiation from some distant source to take the form

---

<sup>1</sup>A notable exception are coherent maser sources (Schawlow & Townes, 1958) which can give rise to—and were first discovered by their—transition lines with extraordinarily large intensity (Weaver et al., 1965; Cheung et al., 1969; Davies et al., 1967).

$$E = ae^{i\phi}e^{i\omega t}, \quad (2.1)$$

where  $e^{i\omega t}$  represents a wave with frequency  $f = \frac{\omega}{2\pi}$  at time  $t$ , and  $a$  and  $\phi$  represent, respectively, the amplitude of the emission, and its phase offset at time  $t = 0$ .

The implication of  $E$  having the form of random noise is that the instantaneous values of  $\phi$  and  $a$  are only well-defined over a period  $\tau \sim \frac{1}{\Delta f}$ , where  $\Delta f$  is the signal bandwidth. In such a case, the emission  $E$  can only be described over longer periods of time by its statistical properties, namely the mean intensity of the signal, given by

$$\langle E^* E \rangle = \langle a^2 \rangle, \quad (2.2)$$

where an asterisk denotes complex conjugation, and  $\langle \dots \rangle$  denotes a time average.

Over many such independent emitters, we expect the received signal to take the form of zero-mean Gaussian noise. The probability density function (PDF),  $P(v)$ , of such noise quantifies the probability that the signal  $v(t)$  will have a value between  $v$  and  $v + dv$  and is given by

$$P(v) = \frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{1}{2}\left(\frac{v}{\sigma}\right)^2}, \quad (2.3)$$

where  $\sigma^2 = \langle a^2 \rangle$  is the signal variance.

### 2.1.2 Quantization

In this age of digital astronomy, signals from antennas must be digitized with finite precision before any digital computation can be performed. It is therefore fortuitous that the random nature of cosmological sky signals makes them susceptible to relatively coarse quantization with limited impact on their statistical properties.

In multi-element radio telescopes, as we shall see shortly, the quantity of interest is the correlation between the signal received at two separate antennas. For reasonably large ( $\gtrsim 8$ ) quantization levels, the signal-to-noise ratio (SNR) of the correlation of antenna signals is inversely proportional to the increase in variance of each of the antenna signals due to quantization (Thompson, 1998). Calculating this increase in variance due to quantization errors can thus be used to infer the loss of correlation SNR.

Given a Gaussian signal with variance  $\sigma^2$ , a regular quantization scheme can be used to represent this signal by  $L$  levels spaced regularly in increments of  $\epsilon\sigma$  (Figure 2.1). The fractional decrease in SNR (the quantization efficiency,  $\eta$ ) resulting from such a scheme is given by (Eq. 8.67, 8.68, Thompson, 1998):




Figure 2.1: Two quantization strategies of a Gaussian distribution. Using nine bins regularly spaced in intervals of width  $\epsilon = 0.55\sigma$  (left) and eight bins of width  $\epsilon = \sigma$  (right). In either case, data values at the edges of the distribution are saturated into the nearest bin.

$$\eta(L, \epsilon) = \begin{cases} \left\{ 1 + \frac{1}{3} \left(\frac{\epsilon}{2}\right)^2 \operatorname{erf}\left(\frac{L\epsilon}{2\sqrt{2}}\right) + \frac{\epsilon^2}{2\sqrt{2\pi}} \sum_{m=L}^{L+20} (m-L+2)^2 \left[ e^{-(m-1)^2 \frac{\epsilon^2}{8}} + e^{-(m+1)^2 \frac{\epsilon^2}{8}} \right] \right\}^{-1} & \text{Even } L \\ \left\{ 1 + \frac{1}{3} \left(\frac{\epsilon}{2}\right)^2 \operatorname{erf}\left(\frac{L\epsilon}{2\sqrt{2}}\right) + \frac{\epsilon^2}{2\sqrt{2\pi}} \sum_{m=L+1}^{L+20} (m-L+1)^2 \left[ e^{-(m-1)^2 \frac{\epsilon^2}{8}} + e^{-(m+1)^2 \frac{\epsilon^2}{8}} \right] \right\}^{-1} & \text{Odd } L, \end{cases} \quad (2.4)$$

where the summation limit  $(L + 20)$  is chosen to be high enough that the contribution of further terms is arbitrarily small.

These equations are valid under the assumption that the correlation of two signals is small (i.e. the signals are dominated by uncorrelated noise) and ignores the effects of quantization errors which are correlated with input signal level. The first assumption is automatically satisfied in most astronomical observations. The second is taken to be valid for multi-level quantization in which error levels vary cyclically between maxima and minima as quantization thresholds are crossed. This behaviour does not occur in the saturation regions of the quantizer, where errors grow linearly with  $\sigma$ , and thus Equation 2.4 should be considered an approximation which only holds when saturation of the quantized signal is rare. In practice, this means that  $\epsilon$  should be large enough so that the saturation point of the quantizer occurs at  $\gtrsim 2\sigma$  and  $\lesssim -2\sigma$ , or, equivalently,  $\epsilon \gtrsim \frac{4}{L}$ .

Figure 2.2 shows Equation 2.4 evaluated for a number of possible values of  $L$ , corresponding to the even- and odd-leveled representations using 3, 4 and 5 bit quantization.




Figure 2.2: The quantization efficiency,  $\eta$ , of quantizers with  $L$  levels, spaced symmetrically around zero with spacing  $\epsilon\sigma$ . Dotted lines show simulated efficiencies based on Gaussian noise signals with a correlation coefficient of  $\rho = 0.03$ . Solid lines show efficiencies as calculated with Equation 2.4, which are only shown for  $\epsilon \gtrsim \frac{4}{L}$ .

$L = 1024$  is also plotted, which approximately represents the ideal case of a quantizer with an infinite number of levels. Simulated efficiencies for  $L = \{7, 15, 31, 1023\}$  are also shown, based on the correlation of two independent Gaussian noise sources with correlation coefficient  $\rho = 0.03$ . A number of features are apparent. At increasingly large values of  $\epsilon$  all quantization strategies become equivalent, essentially due to the extra levels of high- $L$  systems lying in low-probability regions of the input signal PDF. However, too small a value of  $\epsilon$  and much of the input signal distribution lies in the saturation zone of the quantizer, with relatively rapid impact on quantization efficiency. When  $L\epsilon \ll 1$ , the quantizer is almost always in a positive or negative saturated state, in which case the efficiency tends to 0.64, which is that of a simple two-level system (Weinreb, 1963).

The optimal values of  $\epsilon$  for each quantization scheme are shown in Table 2.1, along with the necessary increase in observing time,  $\Delta t$ , required to compensate for the noise introduced by quantization in the optimal case. In practice, most modern radio telescopes use 4-bit ( $L = 15$  or  $L = 16$ ) quantization schemes, which offer around 99% efficiency, and are reasonably tolerant to non-optimal configuration of  $\epsilon$ . More precise quantization schemes are possible, but efficiency gains are small and come at the cost of significantly increasing data rates.

$L$	$\epsilon_{opt}$	$\eta(L, \epsilon_{opt})$	$\Delta t$
2	-	0.637	2.46
3	0.612	0.810	1.52
7	0.674	0.950	1.11
8	0.575	0.966	1.07
15	0.358	0.987	1.03
16	0.334	0.989	1.02
31	0.194	0.996	1.01
32	0.188	0.997	1.01
1024	0.007	1.000	1.00

Table 2.1: Quantization efficiencies for  $L$ -leveled systems at optimal values of  $\epsilon = \epsilon_{opt}$ , as calculated using Equation 2.4. Efficiency of  $L = 2$  and  $L = 3$  systems are quoted from the derivations by Thompson (Chapter 8, 1998).  $\Delta t = \eta(L, \epsilon_{opt})^{-2}$  is the increase in observing time required to compensate for the increase in noise attributed to quantization.

### 2.1.3 Spatial coherence

In addition to short temporal coherence times, in the case of almost all cosmic emission, there is no physical reason to expect that emissions from two different parts of the sky should be statistically dependent. This allows a great simplification of the process of mapping the radio-sources in the sky, whereby individual sources of radio emission may be considered independently. In such a case, a sky image may be constructed by adding the emission intensities associated with multiple sources. This is a key principal which we shall use in the mathematical discussion of interferometry to follow.

It is worth noting that there are cases where spatial coherence can arise, resulting in emission distributions which are not amenable to mapping via the relationship to be presented in Section 2.2.2. One such example is the situation of a scattering surface between source and observer, which may cause emissions from a single point in space to appear as coherent structure on the celestial sphere. The interested reader is directed to Cornwell et al. (1989); Wolszczan & Cordes (1987) and Chapter 14 of Thompson et al. (2004) for a discussion of the mitigation of—and potential opportunities presented by—such effects.

## 2.2 Interferometry and synthesis imaging

We now move to discuss the principles of interferometry presented in Chapter 1 in more depth. The fundamental principle of interferometry is that signals from two separated apertures can be combined to give information about the spatial distribution of emitters with a resolution that depends on their separation distance. In radio astronomy, this separation is known as a baseline. This spatial resolution can be much higher than that achieved by either single aperture alone, whilst the field of view covered is larger than that covered by a single

receiver with comparable resolution. Many comprehensive treatments of interferometry can be found in the literature, for example, Thompson et al. (2004) and Taylor et al. (1999). Here we highlight the key concepts necessary for comparing imaging techniques.

In this section I will introduce a mathematical framework with which to consider interferometric techniques, following the radio interferometer measurement equation formulation, first proposed by Hamaker et al. (1996). In particular, I shall follow the formalism presented in the thorough and illuminating series by Smirnov (2011a,b,c,d).

### 2.2.1 Jones matrices and signal propagation

Following Smirnov (2011a), we consider a single, monochromatic source in an otherwise empty sky. The signal from such a source can be described at any space and time by a single complex vector,  $\mathbf{e}$ , which embodies the three polarization components of the signal. Choosing a Cartesian co-ordinate system, oriented such that the  $z$ -axis lies in the direction of propagation of radiation,  $\mathbf{e}$  takes the simple form

$$\mathbf{e} = \begin{pmatrix} e_x \\ e_y \end{pmatrix}. \quad (2.5)$$

In order to express the signal detected far from the source, we require a mechanism to map the vector  $\mathbf{e}$ , the signal at the source, to the signal at the detector,  $\mathbf{e}'$ . Under the assumption that all effects on the signal as it propagates from source to detector are linear with respect to  $\mathbf{e}$ , we may choose to use matrix multiplication to relate  $\mathbf{e}'$  to  $\mathbf{e}$ :

$$\mathbf{e}' = \mathbf{J}_{\mathbf{e} \rightarrow \mathbf{e}'} \mathbf{e}. \quad (2.6)$$

In such a formulation, the  $2 \times 2$  complex matrix  $\mathbf{J}$  is known as a Jones matrix, after Robert Clark Jones, who first used such a calculus for describing the state of light propagating through linear optical systems (Jones, 1941). We may consider  $\mathbf{J}_{\mathbf{e} \rightarrow \mathbf{e}'}$  to be the product of a chain of effects experienced by a signal on its path from source to detector, which may be explicitly deconstructed as

$$\mathbf{J}_{\mathbf{e} \rightarrow \mathbf{e}'} = \mathbf{J}_n \mathbf{J}_{n-1} \cdots \mathbf{J}_0, \quad (2.7)$$

where the individual terms are interpreted as separate effects on the signal, with  $\mathbf{J}_0$  describing a transformation early in the signal path, and  $\mathbf{J}_n$  representing an effect immediately before detection. Since what is measured at the detector is ultimately a pair of voltages which are related in some linear way to the received signal, we can add this term as an effect

in the Jones chain as well. This leads to a compact description of the relationship between detector voltages and source signal, namely<sup>2</sup>:

$$\mathbf{v} = \begin{pmatrix} v_a \\ v_b \end{pmatrix} = \mathbf{J}_{\mathbf{e} \rightarrow \mathbf{v}} \mathbf{e} = \mathbf{J}_{\mathbf{e}' \rightarrow \mathbf{v}} \mathbf{J}_{\mathbf{e} \rightarrow \mathbf{e}'} \mathbf{e}. \quad (2.8)$$

Some, none, or all of the transformations present in such a Jones chain may commute, depending on their exact form, though not on the co-ordinate system in which they are expressed.

### 2.2.2 Map making and the van Cittert–Zernike Theorem

Having laid the basic foundations of Jones matrix formalism, we are now ready to introduce the visibility matrix, that is, the quantity explicitly measured by an interferometer. For two antennas,  $p$  and  $q$ , separated by a vector  $\mathbf{b} = \mathbf{p} - \mathbf{q}$  (Figure 2.3), and respectively measuring voltages  $v_p$  and  $v_q$ , the visibility matrix is, explicitly

$$V_{pq} = \begin{pmatrix} \langle v_{pa} v_{qa}^* \rangle & \langle v_{pa} v_{qb}^* \rangle \\ \langle v_{pb} v_{qa}^* \rangle & \langle v_{pb} v_{qb}^* \rangle \end{pmatrix} = \langle \mathbf{v}_p \mathbf{v}_q^H \rangle = \mathbf{J}_p \langle \mathbf{e} \mathbf{e}^H \rangle \mathbf{J}_q^H, \quad (2.9)$$

where the final equality of Equation 2.9 follows only when  $\mathbf{J}_p$  and  $\mathbf{J}_q$  are constant over the period of averaging implied by the  $\langle \dots \rangle$  operator. When the  $xy$  polarization basis is used, Equation 2.9 can be expressed directly in terms of the Stokes parameters,  $I$ ,  $Q$ ,  $U$ ,  $V$ , which fully describe partially polarized radiation (see, for example, Hamaker & Bregman, 1996)

$$V_{pq} = \frac{1}{2} \mathbf{J}_p \begin{pmatrix} I + Q & U + iV \\ U - iV & I - Q \end{pmatrix} \mathbf{J}_q^H.$$

This is the radio interferometer measurement equation for a single point source, more commonly expressed by defining the matrix of Stokes parameters as  $\mathbf{B}$ , the sky brightness matrix:

$$V_{pq} = \frac{1}{2} \mathbf{J}_p \mathbf{B} \mathbf{J}_q^H. \quad (2.10)$$

When considering a continuous distribution of spatially incoherent sources over the celestial sphere, the visibility matrix  $V_{pq}$  is determined by integrating Equation 2.10 over the sky observed by an interferometer. Expressing the surface integral over the  $(l, m)$  plane tangential to the celestial sphere at the centre of the interferometry field yields

---

<sup>2</sup>We explicitly follow Smirnov (2011a) in the definition of the components of  $\mathbf{v}$ , and use  $v_a$  and  $v_b$  (rather than  $v_x$  and  $v_y$ ) to emphasize the fact that the detection polarization states need not be  $x$  and  $y$ . Frequently, for example, left and right circular polarization states are used instead.




Figure 2.3: An elementary interferometer, consisting of two receivers, separated by a baseline  $\mathbf{b}$ , observing a distant source with signal  $\mathbf{e}_0$ . As the signal propagates from source to detector, it is modified by the media through which it propagates. At the point of digitization, the signal from each of the two receivers,  $\mathbf{v}_p$  and  $\mathbf{v}_q$ , will likely have experienced a series of different corruptions.

$$V_{pq} = \frac{1}{2} \iint \mathbf{J}_p(l, m) \mathbf{B}(l, m) \mathbf{J}_q^H(l, m) \frac{dl dm}{n}, \quad (2.11)$$

where  $n = \sqrt{1 - l^2 - m^2}$ . Smirnov (2011a) shows that by separating the Jones matrices in this equation into a direction independent part,  $\mathbf{G}$ , a direction dependent part,  $\mathbf{E}$ , and scalar phase terms (which are associated with free-space propagation), Equation 2.11 can be recast in the form traditionally derived in radio-interferometry literature:

$$V_{pq} = \frac{1}{2} \mathbf{G}_p \left( \iint \mathbf{E}_p \mathbf{B} \mathbf{E}_q^H e^{-2\pi i(u_{pq}l + v_{pq}m)} dl dm \right) \mathbf{G}_q^H. \quad (2.12)$$

In this equation,  $u$  and  $v$  are the components of the baseline between antenna pair  $pq$  projected onto the  $lm$  coordinate system and measured in units of wavelength. In this form, the all-sky measurement equation expresses the van Cittert-Zernike theorem (VCZ). This states that the correlation matrix associated with an interferometer measures the two-dimensional Fourier transform of the sky brightness. Each unique baseline, that is, every unique  $uv$  pair, measures a different Fourier component of the brightness. Equation 2.12 generalizes VCZ by modifying the sky brightness,  $\mathbf{B}$ , with direction and antenna-dependent effects,  $\mathbf{E}_i$ . The job of the ambitious radio astronomer is then to “simply” invert Equation 2.12 to recover the sky brightness distribution, and all the physics it encodes.

## 2.3 Calibration

An inversion of Equation 2.12 allows recovery of the sky brightness distribution, up to the limits determined by the number of  $uv$  co-ordinates sampled by any given interferometer. Direction-independent corrections,  $\mathbf{G}$ , which take the form of Jones chains (Equation 2.7), may be accounted for by matrix multiplication with their inverses,  $\mathbf{G}^{-1}$ . Similarly, the direction-dependent effects may either be inverted in the image plane by multiplication, or equivalently removed in the  $uv$  plane by convolution.

Given this recipe for obtaining the true sky brightness, image-making relies upon being able to obtain, and invert, the Jones matrices which corrupt signals between emission and detection. There are a number of ways by which this can be achieved, three of the most important of which—simple sky model fitting, self-calibration and baseline redundancy calibration—are briefly summarized here. The interested reader is directed to Cornwell & Fomalont (1999) for a more comprehensive overview.

### 2.3.1 Sky Model Fitting

The method of sky model fitting relies on *a priori* information about the area of sky being observed with a telescope. The general principle of fitting a sky model begins with a model

sky,  $\hat{I}$ , and a corresponding set of model visibilities,  $\hat{V}_{pq}$ . The set of corrupted<sup>3</sup> visibilities,  $\mathbf{G}_p \hat{V}_{pq} \mathbf{G}_q^H$ , should reproduce the observed visibilities, within the noise of the observed data. Given this equivalence relationship, the Jones terms,  $\mathbf{G}$ , can be computed by a least squares minimization of the residual difference between model and observed visibilities. In particular, such a calibration is frequently performed by observing a single, bright, isolated point-source. In this case the model visibilities take the particularly simple form of identical, constant, real terms.

### 2.3.2 Self-Calibration

The method of self-calibration is similar to that of sky model fitting, yet requires far less prior knowledge of the area of sky being observed with a telescope. Starting only with the assumption that the sky being observed is positive-valued, and made up of discrete point sources, an initial calibration can be obtained by minimizing the difference between the observed sky and a model containing only the brightest points. Further calibration improvements can then be obtained by iterating the self-calibration procedure, allowing the model sky to be improved at each cycle using information from the observed, and partially calibrated, data.

### 2.3.3 Baseline Redundancy

Redundant calibration requires no *a priori* information, and instead relies on the fact that different two-element interferometers in an array with the same baseline vector measure the same component of the sky brightness Fourier decomposition. Provided there are at least  $N$  such redundancies in the array, the system of simultaneous equations formed by equating all redundant visibilities can be used to solve for the  $N$  direction-independent  $\mathbf{G}$  terms. A significant limitation of the redundancy calibration method is its inability, without a prior sky model, to compute direction-dependent effects. Furthermore, where direction-dependent effects vary across antennas the assumption that pairs of redundant interferometers measure the same quantity breaks down. In this case, the quality of baseline redundancy as a calibration estimator may become limited (Noorishad et al., 2012).

## 2.4 Correlators

Regardless of how an interferometer is calibrated, the visibilities themselves ( $V_{pq}$  in Equation 2.12) must be measured. This can be accomplished in a variety of ways. A device used to compute visibilities, or correlations, is known as a correlator. In this section we summarize

---

<sup>3</sup>for simplicity (and to represent the state of most current calibration software packages) we consider only direction-independent effects.

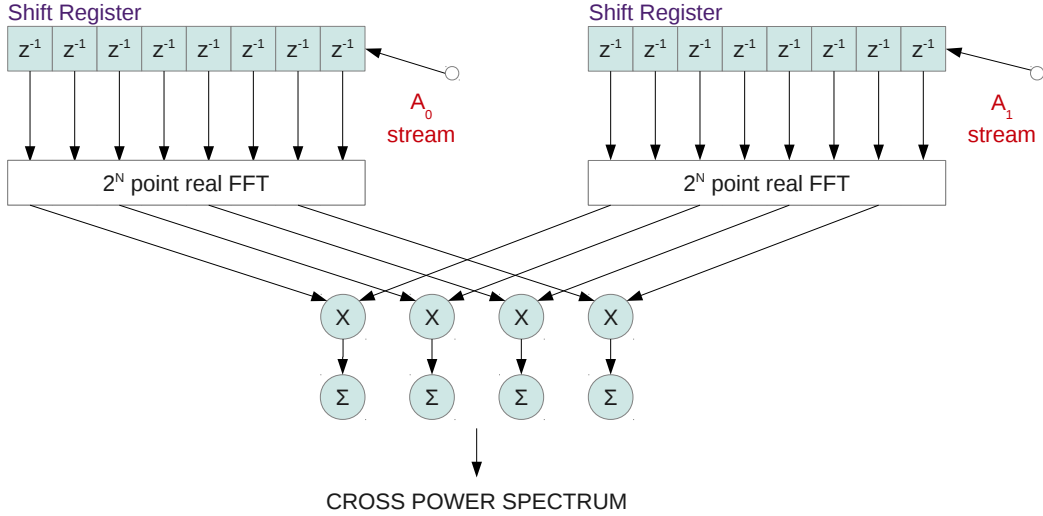


Figure 2.4: Schematic depicting FX correlation of two antennas. Digitized signals from two antennas are first channelized into  $2^{N-1}$  channels by Fourier transform, before multiplying each matching channel pair together to form the cross power spectrum. This spectrum is then accumulated to reduce output data rate and increase SNR. In real-world implementations the initial Fourier transform is usually pipelined so that number of operations per clock cycle is uniform.

the main approaches to building correlators proposed in radio astronomy literature and used in practice. We conclude by introducing the concept of direct imaging, which breaks the paradigm of treating an antenna array as a set of two element interferometers and instead generates sky images directly using beamforming techniques.

### 2.4.1 The FX Correlator

An FX correlator is a correlator which performs a frequency domain channelization of the time domain voltage stream from each station in an interferometer (the “F” operation), before carrying out the complex multiplications of pairs of streams (the “X” operation) and accumulating the resultant visibilities. The FX correlation of a pair of antenna signals is depicted schematically in Figure 2.4.

For an interferometer with  $N$  elements there are  $N(N - 1)/2$  different visibilities to be calculated. If spectral decomposition of the real, time-domain signals into  $M$  channels is required, the total number of complex multiply-accumulate operations (CMACs) performed by the correlator is of order  $NM \log_2(M) + MN^2$ . The first contribution is from the fast Fourier transform (FFT, Cooley & Tukey, 1965) of  $N$  data streams. The second contribution is from the  $O(N^2)$  visibilities which need to be formed for each of  $M$  frequency channels. Since these operations need be carried out only once per spectra, that is, once per  $M$  samples, normalizing to unit bandwidth yields a computational cost in CMACs per Hz of

order  $N \log_2(M) + N^2$ .

### 2.4.2 The XF Correlator

The convolution theorem states

$$\mathcal{F}\{f * g\} = \mathcal{F}\{f\} \mathcal{F}\{g\}, \quad (2.13)$$

where  $\mathcal{F}\{f\}$  is the Fourier transform of a function  $f$ , and  $f * g$  is the convolution of two functions  $f$  and  $g$ . Since much use is made of the convolution theorem in radio astronomy, and in this thesis, the standard proof is replicated in Appendix A.

FX correlation—that is, the Fourier transform of two data vectors, followed by element-wise multiplication—is equivalent to the right hand side of Equation 2.13. The left hand side suggests that there exists another route to the desired end product of cross power spectra, namely by convolution of the two antenna signal streams, and Fourier transform of the result. Such an implementation of a correlator is called an “XF” or lag correlator, as the delay and multiply stage precedes the following Fourier transform from time to frequency space. The XF correlation of a pair of antenna signals is depicted schematically in Figure 2.5.

For an interferometer with  $N(N - 1)/2$  antenna pairs and  $M$  channels,  $MN(N - 1)$  real multiplications are required to perform the convolution. The Fourier transform requires  $O(M \log(M))$  operations, but need be performed only once per integration.

### 2.4.3 Gridding

Assuming that visibilities can be measured and calibrated, Equation 2.12 allows recovery of the sky brightness (or rather the apparent sky brightness,  $\mathbf{E}_p \mathbf{B} \mathbf{E}_q^H$ , seen by an antenna pair) by inverse Fourier transform. This transform can be computed in its general discrete form, using the visibilities as measured (and calibrated). However, where images with large numbers of pixels are required—that is images formed from interferometers whose resolution is far greater than their field of view—it is usually worth interpolating the measured visibilities onto a regular grid of locations in the  $uv$  plane, allowing the discrete Fourier transform to be efficiently implemented with an FFT. This interpolation process is usually referred to as *gridding*. Gridding is accomplished by means of a convolution of measured visibilities with an appropriate kernel, producing a continuous function which can be regularly resampled as required. The convolutional gridding process is covered in depth in Thompson et al. (Chapter 10, 2004) and Taylor et al. (Chapter 7, 1999), here we simply lay out the fundamental choices which must be made when gridding.

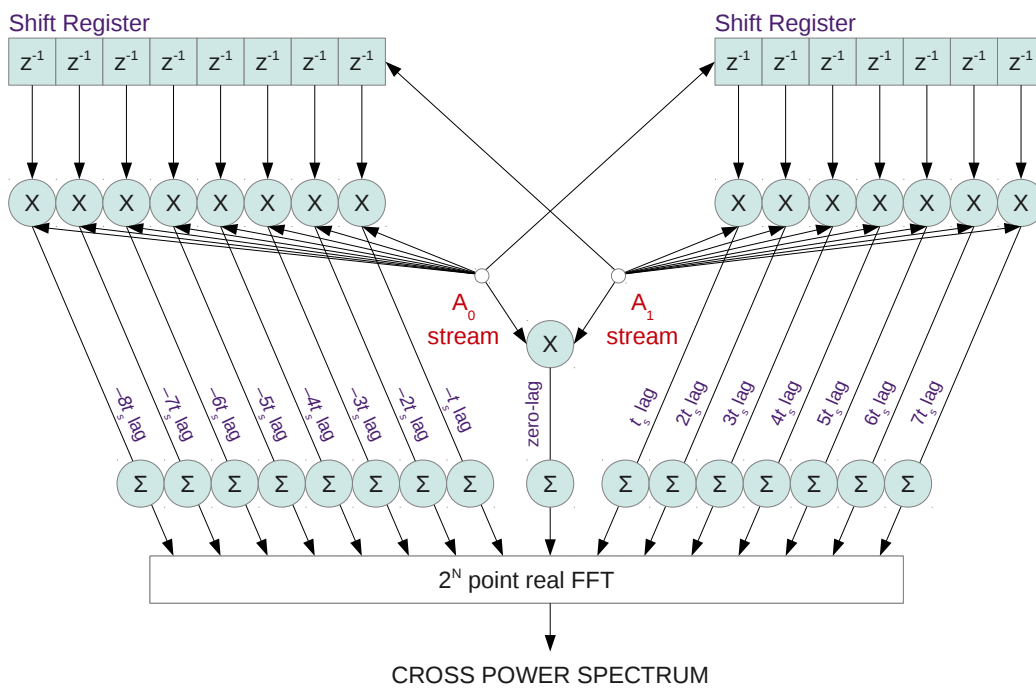


Figure 2.5: Schematic depicting XF, or “lag” correlation of two antennas. Digitized signals from two antennas are convolved with  $2^M$  lags, before Fourier transforming the resulting convolution to form a cross power spectrum with  $2^{M-1}$  channels. In an XF correlator, accumulation takes place before the Fourier transform stage, with the transform required only once per spectrum integration.

## Grid Size

The grid configuration used in the gridding process defines two main properties in the image produced. Firstly, the spacing between grid points defines the size of the patch of sky imaged. Secondly, the extent of  $uv$  grid points generated determines the area of each pixel. Typically, an area of sky is imaged that is approximately the size of a telescopes field of view, whilst the pixel size is chosen so that each pixel corresponds to an area a few times smaller than the resolution of the telescope being used. Explicitly, the area of sky to be imaged is defined by the image edges at  $l = \pm l_{\max}$ ,  $m = \pm m_{\max}$ , given by

$$\begin{aligned} l_{\max} &= \pm \frac{1}{2\Delta u} \\ m_{\max} &= \pm \frac{1}{2\Delta v}. \end{aligned} \tag{2.14}$$

A similar relationship determines the pixel areas,  $\Delta l \Delta m$ , which are set by the extent of the  $uv$  grid. With  $u$  and  $v$  extending between  $\pm u_{\max}$  and  $\pm v_{\max}$ , the pixel sizes are given by

$$\begin{aligned} \Delta l &= \pm \frac{1}{2u_{\max}} \\ \Delta m &= \pm \frac{1}{2v_{\max}}. \end{aligned} \tag{2.15}$$

## Gridding Kernel

The operation of resampling the measured visibilities,  $V(u, v)$  on a regular grid with spacing  $\Delta u$  and  $\Delta v$  can be implemented by computing, for each grid point  $(u_c, v_c)$

$$\begin{aligned} V^R(u_c, v_c) &= \sum_{i=1}^n C(u_c - u_i, v_c - v_i) V(u_i, v_i) \\ \text{where } u_c &= 0, \pm\Delta u, \pm 2\Delta u, \dots \\ v_c &= 0, \pm\Delta v, \pm 2\Delta v, \dots, \end{aligned} \tag{2.16}$$

and the sum is over all the visibilities measured by the interferometer. This includes conjugates of these measured values, calculated as  $V(-u, -v) = V^*(u, v)$ .  $C(u, v)$  is a convolution function which can be arbitrarily specified.

Since the  $uv$  and image planes are related via a Fourier transform, a convolution in  $uv$  space amounts to a multiplication of the resultant image. An obvious choice for the gridding kernel,  $C(u, v)$ , is a sinc function. Such a function, appropriately scaled, has a Fourier transform consisting of a two-dimensional top hat function of the same dimensions as the desired image, and unit magnitude. In this case the convolution amounts to a multiplication

of the image by unity. For computational reasons, it is desirable to be able to truncate the sum in Equation 2.16 so that only  $uv$  samples in close proximity to a grid point contribute to its value. This is at odds with the choice of a sinc function for the convolution kernel, which has significant non-zero components even at large arguments and cannot be truncated without significant artifacts in the resultant image. Rather, a convolution kernel which falls naturally to zero is usually chosen, often by simply multiplying the ideal sinc function by a window—often Gaussian—which falls rapidly and smoothly to zero beyond a few grid spacings. The effects of such a window can then be compensated for in the resultant image via multiplication.

#### 2.4.4 The Direct-Imaging Correlator

An FX correlator computes the Fourier transform of antenna signals and then forms  $uv$  samples by cross-multiplication of all antenna pairs. This is equivalent to an autocorrelation of channelized antenna signals in antenna-space. This correlation can be implemented by performing a spatial Fourier transform on the antennas in an array, taking the element-wise power of the result, and inverse Fourier transforming to recover the cross-power spectrum. The direct-imaging correlator is a correlator that performs precisely these steps, and is so named because, prior to the final inverse Fourier transform, the intermediate data product is an image of the sky. In fact, telescopes that function in this way are frequently not thought of as interferometers at all, but rather as beamformers that generate enough simultaneous beams to obtain a complete image of the sky.

A number of variants of the direct imaging correlator have appeared relatively recently in literature (Tegmark & Zaldarriaga, 2009, 2010; Morales, 2011), all of which aim to exploit the ability, in certain circumstances, to implement Fourier transforms using logarithmically scaling “fast” algorithms. Direct imaging telescopes naturally fall into two categories, depending on whether the configuration of antennas in an array places them on regular grid points.

#### Regularly Gridded Arrays

Where antennas naturally lie on a regular grid, a fast Fourier transform can be used to calculate their cross power spectrum. This is entirely equivalent to the observation that when antennas lie on a regular grid, their beamforming coefficients become equal to FFT coefficients (Williams, 1968). A regular array has many repeated baselines, and duplicated measurements made by these are indistinguishable in a sky image. It follows that the cross-correlations of different interferometers separated by the same baseline are inseparably averaged in the image-making process. In this way, the output data product of such a

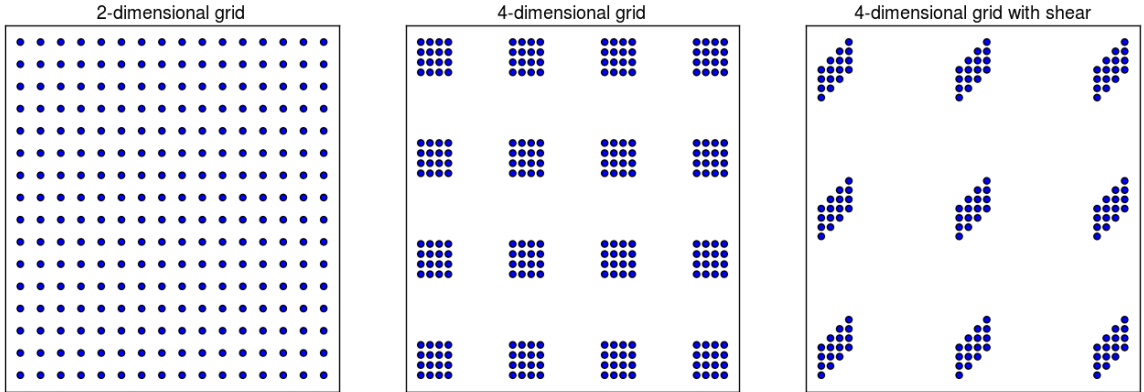


Figure 2.6: Examples of antenna layouts susceptible to FFT-based interferometers. Grids may contain arbitrary levels of hierarchy, and may feature rotation or shear.

telescope is, in the visibility domain, a single measurement of each unique Fourier mode measured by the array.

The largest telescope featuring in radio astronomy literature employing direct imaging is located at Waseda University and comprises an  $8 \times 8$  grid of dishes (Otoabe et al., 1994). More recently, Tegmark & Zaldarriaga (2009), proposed a similar design but with zero-padding of FFTs used to allow complete recovery of visibilities after averaging beam powers. Such telescopes, referred to by Tegmark & Zaldarriaga as “Fast Fourier Transform Telescopes (FFTTs)”, all feature regular two-dimensional grids of antennas. It is further noted by Tegmark & Zaldarriaga (2010) that arrays with hierarchical grid structure—dubbed “Omniscopes”—are also amenable to FFT-based interferometry. Examples of possible grid-based telescope layouts are shown in Figure 2.6.

The zero-padding required by FFTTs and Omniscopes is worthy of some explanation, particularly since this zero-padding is essentially the only way in which these types of telescopes differ from previous direct-imaging implementations. We explicitly illustrate the FFTT process using a toy example, similar to that concocted by Tegmark & Zaldarriaga (2010).

First we consider a toy array, consisting of a one-dimensional line of 4 antennas located at  $x = 0, 1, 2, 3$ . If these antennas measure, respectively, complex voltages  $v(x) = v_0, v_1, v_2, v_3$ , then in the visibility plane, after summing together redundant baselines,  $V(u)$  is given by

$$V(u) = \int v(x)v^*(x+u) dx \equiv (v \star v)(u), \quad (2.17)$$

where  $v \star v$  defines the correlation operator. The relationship between  $v(x)$  and  $V(u) \equiv (v \star v)(u)$  is shown pictorially in Figure 2.7, which illustrates the 7 samples of  $V(u)$  measured

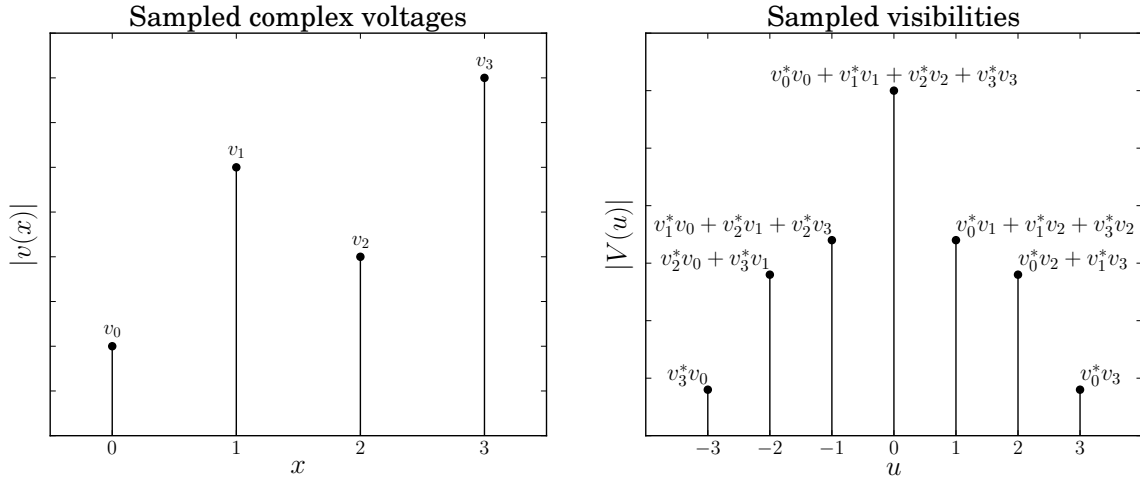


Figure 2.7: The complex voltages measured by 4 antennas (left), and the corresponding  $u$  samples (right).

by our toy array. By a proof almost identical to that for the convolution theorem (Appendix A) the correlation operation satisfies

$$\mathcal{F}\{a \star b\} = \mathcal{F}\{a\}^* \mathcal{F}\{b\}, \quad (2.18)$$

or equivalently, for our autocorrelation

$$\mathcal{F}\{v \star v\} = \mathcal{F}\{v\}^* \mathcal{F}\{v\} = |\mathcal{F}\{v\}|^2. \quad (2.19)$$

$\mathcal{F}\{v \star v\} = \mathcal{F}\{V(u)\}$  defines the one dimensional sky brightness distribution via VCZ, but Equation 2.19 shows that we can equally well reach this distribution directly via a Fourier transform of complex voltage distribution measured by the array. We again emphasize that this is *always* true for any co-planar array, no matter what the layout of antennas. However, in the case of regularly sampled  $v(x)$ , this transform can be efficiently computed with an FFT.

For our array of 4 antennas  $V(u)$  has 7 unique real components (Fig 2.7). Correspondingly, to make a sky image directly from  $v(x)$  which encodes all the information in the sampled  $V(u)$  requires the construction of at least 7 regularly spaced beams on the sky. In practice, this can be accomplished with a power-of-two FFT by zero-padding the measured voltage vector to a length of 8, and then performing a standard FFT.

This example simply serves as an illustration of the well known fact that a discrete correlation (or convolution) can be implemented using the discrete form of Equation 2.18, provided adequate zero-padding is included (see, for example, Chapter 13, Press et al., 1992). In general, a one-dimensional array comprising  $N$  antennas requires a minimum of

$N - 1$  padded zeros for correlation to be calculable with an FFT. However, as the array reaches higher dimensionality, zero-padding must be applied in each dimension. An  $m$ -dimensional array of  $N$  antennas in any arrangement requires  $\sim 2^{m-1}N$  padded zeros, from which an image can be constructed via a  $\sim 2^m N$  point FFT. In practice, since many of the inputs to this FFT are zero, the cost is reduced by a factor  $\frac{1}{m} \sum_{i=1}^m 2^{i-m}$ . In the case of a two-dimensional  $\sqrt{N} \times \sqrt{N}$  array, this factor is  $\frac{3}{4}$ . This results in a two-dimensional,  $N$ -point, zero-padded FFT cost of  $3 \log_2(4N)$ . In the limit of large numbers of dimensions this reduction factor is approximately  $\frac{2}{m}$ , in which case the multi-dimensional FFT of a large- $N$  array requires approximately  $\frac{2^{m+1}}{m}$  times more computation than a single dimensional array with the same number of antennas.

### Irregular Arrays

Where antennas are irregularly positioned in an array, Fourier transform imaging can still be implemented using FFTs provided the antenna signals are first interpolated onto a regular virtual grid using an appropriate convolution kernel. This paradigm, suggested by Morales (2011) under the name Modular Optimal Frequency Fourier (MOFF) correlator<sup>4</sup>, allows the efficiency of the FFT algorithm to be exploited, whilst still allowing complete flexibility in antenna placement. After initial generation of the virtual grid of antennas the processing performed is essentially identical to that of an FFTT, though the number and spacing of points in the virtual grid may be chosen independently of the number of antennas.

Choosing grid positions for virtual antenna locations is equivalent to choosing grid positions in the  $uv$  plane. This is because the  $uv$  grid points correspond to the baseline vectors between antenna positions, projected towards a chosen sky-direction. The parameters defining the antenna grid can thus be set using the same arguments as for  $uv$  grids, as laid out in Section 2.4.3. To date, no telescope featuring a MOFF correlator has appeared in radio astronomy literature.

## 2.5 Choosing a correlation method

A pictorial representation of the possible routes from sampling sky signals to forming a sky image is shown in Figure 2.8. In the remainder of this chapter we consider the domains in which the various available architectures are most beneficial. We also consider the fidelity of the images produced by each of these architectures and assess their unique limitations and benefits. Before evaluating implementation efficiencies, that is, required computing power

---

<sup>4</sup>“Optimal” in this context means that the information encoded within the generated image is equivalent to that which would be present in an FX correlation matrix. This is proven by Morales (2011) only when an appropriate convolution function is used, though the image-making mechanism itself is independent of this choice.

for each processing method, we start by considering the mathematical differences between the approaches, and the resultant differences in their capabilities.

### 2.5.1 Array configuration

The first issue which must be considered is the impact each processing mechanism has on the allowable array configurations. MOFF and FX/XF-based imagers effectively place no limitations on allowable array configurations, whilst FFTT or Omniscope-based arrays require that elements be on a regular grid. This forces Omniscopes to suffer a number of drawbacks.

1. Regular arrays suffer from high sidelobes in their point-spread function (PSF).
2. If the grid spacing does not achieve Nyquist sampling of incoming wavefronts, grating lobes will appear in the array’s PSF.

Both of these effects are liable to reduce the ability of a telescope to produce high dynamic range images. However, they can be somewhat mitigated by appropriate array design. Side-lobe response levels can be improved by array apodization or “tapering”, that is, weighting antenna contributions such that the array beam quickly falls to zero away from the pointing centre. The effects of a selection of weighting schemes are shown in Figure 2.9, for an example array consisting of 32  $\frac{\lambda}{2}$ -spaced antennas. These weights, which can be applied directly to antenna signals (as in Figure 2.9), or applied to visibility matrices after correlation or direct imaging, invariably reduce the signal to noise of observations, by downweighting certain measurements.

Another way to address both the PSF and aliasing issues experienced by regularly gridded direct imagers is to emulate a random array by placing antennas on a fine-pitched, but sparsely-populated, grid. This scenario is shown in Figure 2.10 for two example grid sizes. Here we take the positions of a LOFAR LBA station as an example array and consider the PSF that would result if imaging were conducted in one of three ways:

1. By FX/XF correlator, using the true antenna positions.
2. By MOFF correlator, using the true positions interpolated onto a regular grid with a sinc-function convolution kernel, truncated at 3 m.
3. By FFTT, where we take the PSF of a hypothetical LOFAR LBA which had been constructed on a grid of regularly spaced, but sparsely populated points.

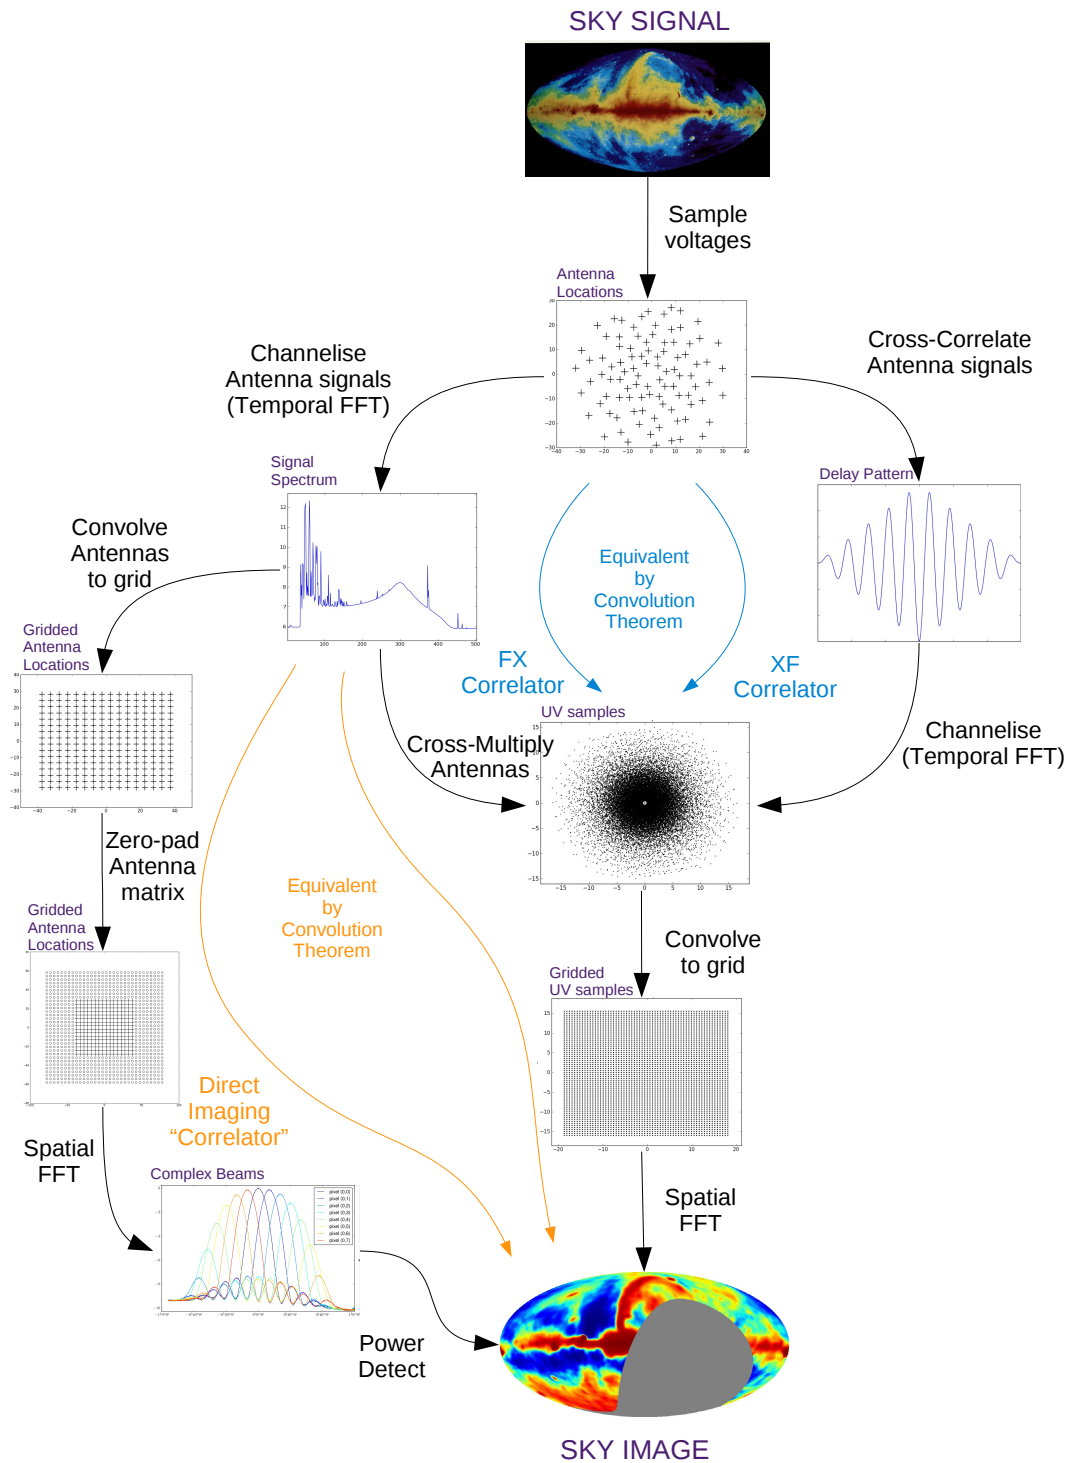


Figure 2.8: Various routes from voltages sampled by a distributed set of antennas to a synthesized sky image. (Sky signal graphic from Haslam et al. (1982). Sky image graphic from Foster (2013))

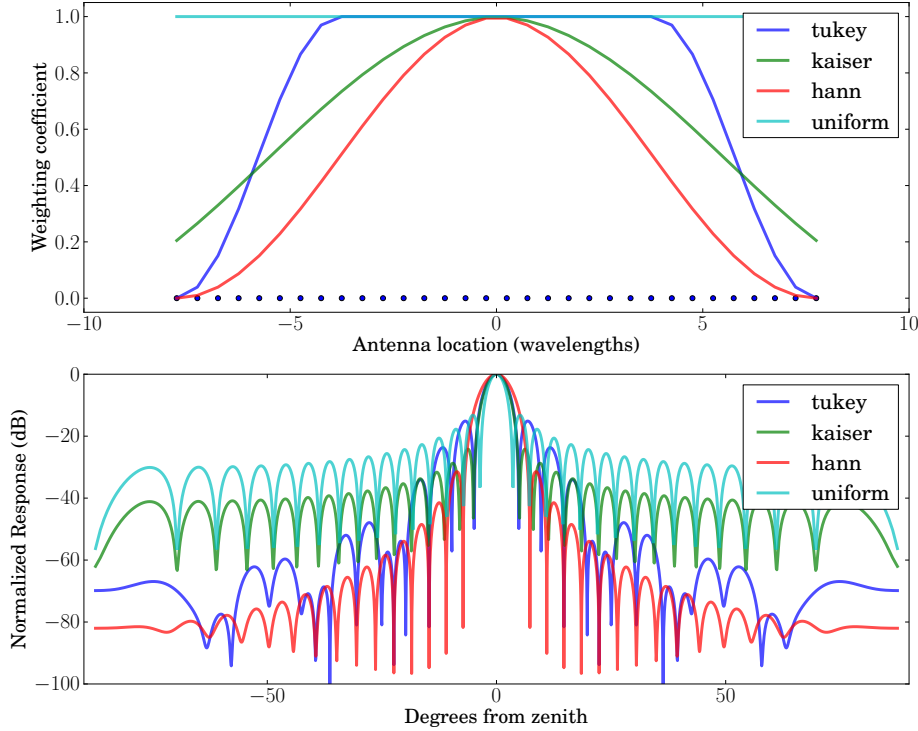


Figure 2.9: Responses of a 32 element,  $\frac{\lambda}{2}$ -spaced array, with Tukey, Kaiser, Hann, and uniform apodization schemes.

In the last case, we model the hypothetical array by moving each antenna in the LOFAR station to the nearest grid point. Also computed, for reference, is the PSF corresponding to an untapered array with antennas lying on each of the grid points.

Examining Figure 2.10, we see that PSFs based on both true and interpolated antenna positions are almost identical. In this case, the number of grid points need not necessarily reflect the wavelength of the sampled radiation in order to produce a high quality PSF (though the spacing and quantity of grid points needs to be appropriately chosen for the end-result image to have desired dimensions and resolution). In the third case, the gridded nature of the array mandates that at least some antennas have spacings of  $\frac{\lambda}{2}$  to avoid aliasing. However, providing this criteria is met, an array based on a sparse grid can be created which has a PSF of comparable quality to a random array.

The number of grid points required by a sparse-regular or interpolated-random array dictates the amount of processing required in a digital system. This quantity is worthy of brief investigation.

Direct-imaging systems require that the field of view of an array be completely sampled by synthesized beams. The field of view of an antenna is approximately given by  $\frac{\lambda}{d}$ , for antenna elements of size  $d$ , whilst the synthesized beam resolution is approximately  $\frac{\lambda}{D}$  for

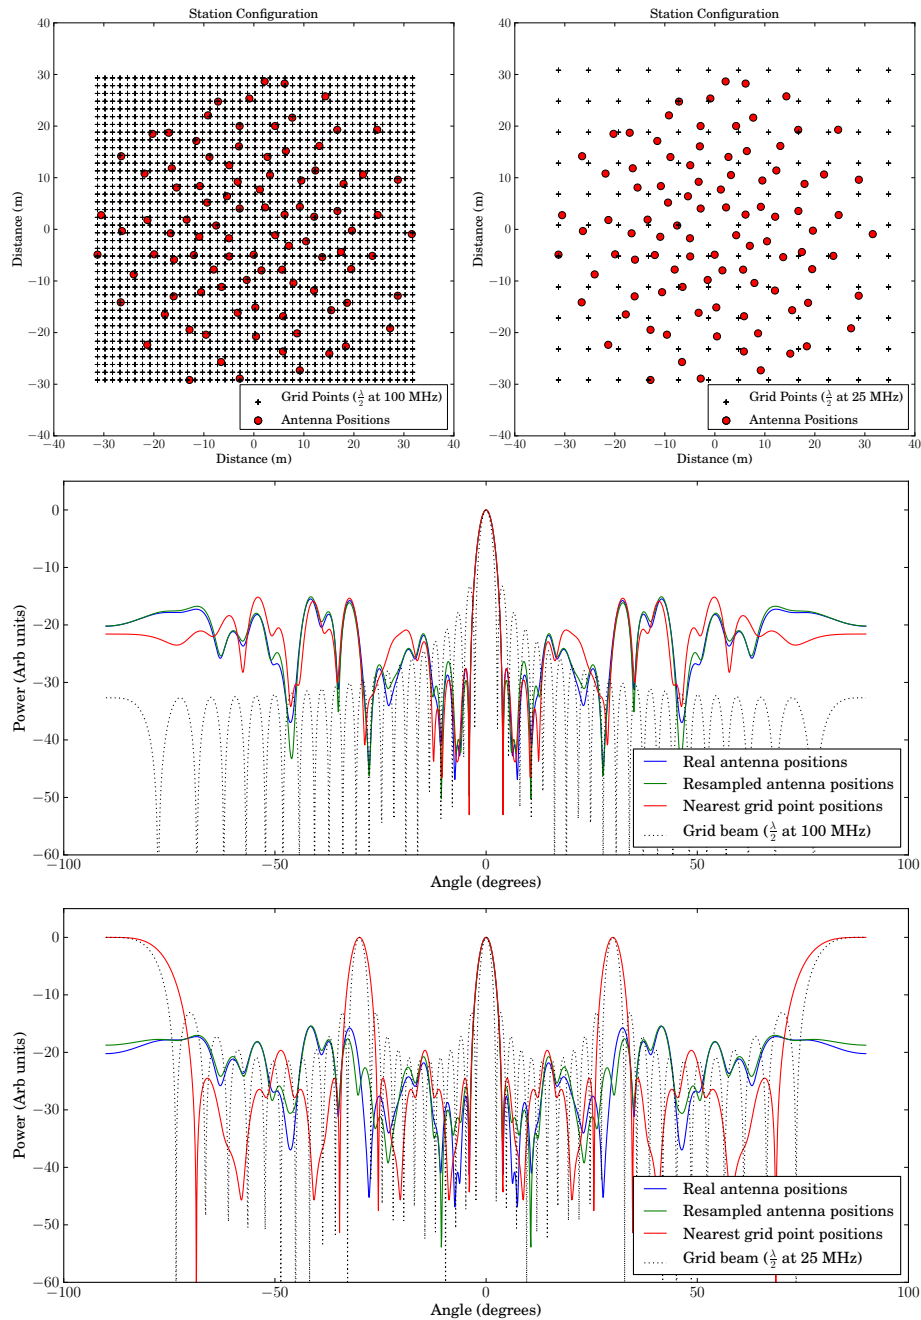


Figure 2.10: Antenna distributions of an international LOFAR station with potential gridding locations for grid points  $\frac{\lambda}{2}$ -spaced at 100 MHz (top-left), and 25 MHz (top-right). Array-beam slices from a such a station are shown at 100 MHz. Shown in blue is the true pattern obtained by DFT. The green line represents the pattern obtained by re-sampling antennas to a regular grid. The red line shows the hypothetical response of a LOFAR station with antennas moved such that they all sparsely populate a regular grid. The grid spacing used is  $\frac{\lambda}{2}$  at 100 MHz (centre), and 25 MHz (bottom).

an array of size  $D$ . By Equations 2.14 and 2.15, this implies a number of grid points,  $G$ , given by

$$G \sim \left(\frac{D}{d}\right)^2. \quad (2.20)$$

This requirement can only be relaxed if we are willing to tolerate grating lobes (as in 2.10), which may be acceptable for power-spectrum science goals.

More strictly, since it is the array field of view that is of interest rather than the explicit size of the array elements, a more accurate definition of  $G$  can be made by considering the effective area of each antenna,  $A_{\text{eff}}^{\text{el}}$ . In this case  $G$  is given by

$$G = \frac{A_{\text{physical}}}{A_{\text{eff}}^{\text{el}}} \sim \frac{N}{f}.$$

It is therefore the filling factor,  $f$ , of an array which plays a large part in determining the efficiency of a direct-imaging system. For many types of antenna, this is liable to change significantly with frequency (see Equation 1.4). In this case, if the antennas in an array are physically located on a grid,  $G$  is set by smallest effective area per element in the observation band. If antennas are interpolated onto a grid, then a different grid spacing may be chosen for every frequency range, potentially resulting in a more efficient (if more complex) digital system.

### 2.5.2 Calibratability

The ease with which the the antenna based corruptions  $\mathbf{G}_i$ ,  $\mathbf{E}_j$  (Equation 2.12) can be compensated for in the procedure of estimating the sky brightness distribution depends on the imaging mechanism used. In this sense, FX and XF correlators represent something of a gold standard for calibratability. Since such instruments measure and record every individual visibility formed by an array, providing the calibration coefficients  $\mathbf{G}^{-1}$  and  $\mathbf{E}^{-1}$  can be computed, all antenna-based effects can be accounted for. Of course, the problem of obtaining such calibration coefficients, which may be time, direction, and frequency dependent is non-trivial. Nonetheless, assuming adequate calibration algorithms exist, FX and XF correlators should produce images unaffected by antenna gain variation.

With FFTT and MOFF imaging systems, this is not the case. Such systems compute and record images of the sky. These images can be used to obtain the familiar spatial Fourier modes measured by a set of two-element interferometers, but only provide one value for each unique baseline in an array (or, in the case of a MOFF correlator, each unique virtual baseline). These “visibilities” are equal to the sum of identical baselines as measured in an FX or XF correlator, as illustrated in Figure 2.7. This simple fact cannot be stressed

enough, as it underpins all arguments made about calibration of direct-imaging correlators. *Direct imaging correlators measure, once, the sum of redundant baselines in an array.*

### Real-time Calibration

The summing effect of a direct-imaging correlator means that, for optimal use, pairs of antennas (or grid points) contributing to such a sum should yield the same correlation value. An explicit example can be made by considering a single redundant baseline of the toy interferometer in Section 2.4.4. Considering the baseline  $u = 1$ , which has 3 contributing antenna pairs  $\{p, q\} = \{0, 1\}, \{1, 2\}, \{2, 3\}, \{3, 4\}$ . If the calibrated visibility from any of these two-element interferometers is  $V_{p,q} = \tilde{V}_{u=1}$ , then the mean visibility as measured by an FX correlator after calibration is

$$\begin{aligned} V_{u=1}^{FX} &= \frac{1}{3} \sum_{\{p,q\}} \left\{ \mathbf{E}_p^{-1} \mathbf{G}_p^{-1} \left[ \mathbf{G}_p \mathbf{E}_p \tilde{V}_{u=1} \mathbf{E}_q^H \mathbf{G}_q^H + n_{pq} \right] (\mathbf{G}_q^H)^{-1} (\mathbf{E}_q^H)^{-1} \right\} \\ &= \tilde{V}_{u=1} + n'_{0,1} + n'_{1,2} + n'_{2,3}, \end{aligned} \quad (2.21)$$

where  $n_{pq}$  represents Gaussian noise in the  $p, q$  baseline measurement, and  $n'_{pq} = \mathbf{E}_p^{-1} \mathbf{G}_p^{-1} n_{pq} (\mathbf{G}_q^H)^{-1} (\mathbf{E}_q^H)^{-1}$ . An FX correlator permits calibration of each redundant baseline individually. This allows recovery of  $\tilde{V}_{u=1}$ , up to Gaussian noise terms, independent of the gains  $\mathbf{G}$ ,  $\mathbf{E}$  (providing they are invertible). Summing redundant visibilities obtained with an FX correlator also permits inverse noise weighting of measurements, to ensure that the overall measurement of  $V_u^{FX}$  has optimal SNR.

The corresponding measurement made with a direct imager is

$$V_{u=1}^{IM} = \frac{1}{3} \mathbf{E}_{u=1}^{-1} \mathbf{G}_{u=1}^{-1} \left( \sum_{\{p,q\}} \left\{ \mathbf{G}_p \mathbf{E}_p \tilde{V}_{u=1} \mathbf{E}_q^H \mathbf{G}_q^H + n_{pq} \right\} \right) (\mathbf{G}_{u=1}^H)^{-1} (\mathbf{E}_{u=1}^H)^{-1} \quad (2.22)$$

Since a direct imager measures only one value for each unique visibility, any calibration of this value cannot act separately on its contributing parts. We express this by permitting a single calibration of  $\mathbf{E}$  and  $\mathbf{G}$ , with coefficients denoted by  $\mathbf{E}_{u=1}^{-1}$  and  $\mathbf{G}_{u=1}^{-1}$ . In general, only if  $\mathbf{E}_p = \mathbf{E}_{u=1}$  and  $\mathbf{G}_p = \mathbf{G}_{u=1}$  for all antenna pairs will the FX and direct imaging implementations be equivalent with  $V_{u=1}^{IM} = V_{u=1}^{FX}$ . This condition can be satisfied by calibrating antenna signals before correlation takes place. In literature, such calibration is likened to that required by beamformers. In beamforming, it is assumed that the beams form the data products of interest, in which case, any gain artifacts—even those common to redundant baselines—may intolerably alter the formed beam's shape or pointing direction.

However, if it is assumed that the visibilities are the data product of interest computed by direct imagers, then post-imaging calibration can remove *common* gain calibrations from redundant baselines with no impact on the final measured visibilities. Gain errors which are not common to redundant baselines can also be tolerated, though these impact the SNR of final visibility measurements.


The impact of phase and amplitude gain errors for the case of 1000 redundant baseline measurements, with each individual measurement having equal SNR, is shown in Figure 2.11. Here SNR is obtained by simple Monte Carlo simulation, assuming Gaussian phase and amplitude error distributions. SNR is expressed as a fraction relative to the measurement which could be made with an optimally calibrated array. As expected, where gains differ dramatically, measuring a single visibility with 1000 redundant baselines effectively fails to increase SNR above that achieved with a single measurement. Conversely, moderate calibration errors (phase  $\lesssim 5^\circ$ , amplitude  $\lesssim \frac{1}{2}$  dB) can be tolerated with minimal SNR loss. In this example case, the SNR loss with such measurements could be compensated for with an increase in observing time of approximately 3%. Similar calculations for SNR loss are shown for 2, 10, 100 and 1000 redundant visibility measurements for phase-only errors in Figure 2.12.

The case of amplitude-only errors is shown in Figure 2.13. Here, the SNR in the summed visibility measurement depends on the ability to accurately inverse-noise weight each contributing visibility.

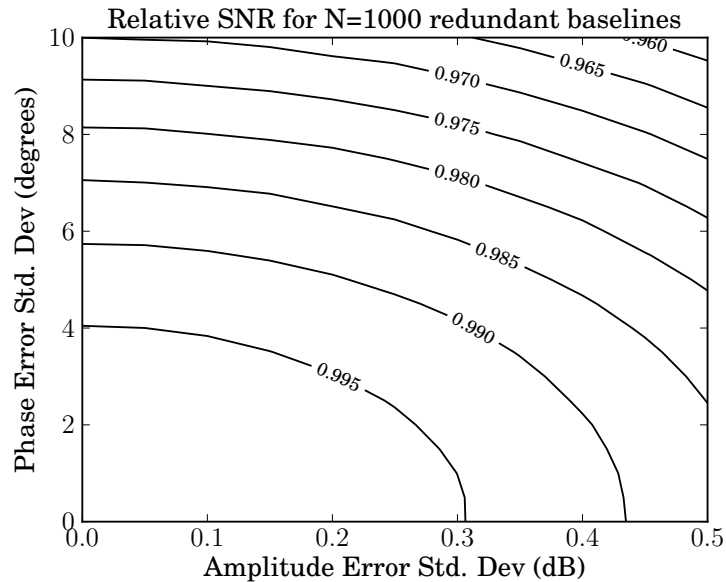
The simple results in this section give an idea of the level of accuracy at which real-time calibration must be achieved for direct-imagers to be feasible. Namely, phase and amplitude (or SNR deviation) residuals must be no greater than  $\sim 5^\circ$  and 0.5 dB. However, again we emphasize that this is the tolerable level of gain variation amongst redundant baseline measurements. It is not the variation of measurements from the true value. To emphasize this, two cases which may give rise to larger errors which are common to redundant baselines are explicitly stated.

**Pointing errors:** Errors in the direction of the interferometer phase centre give rise to identical errors amongst redundant baselines.

**Ionospheric corrections:** The ionosphere, which acts as to distort the phase of radiation from different directions, may give rise to direction-dependent gain errors. However, in the case of array sizes which are smaller than the characteristic scale sizes of ionospheric variations ( $O(\text{km})$ ), these errors should largely be the same among redundant baselines.



(a) The effect on SNR ratio per-baseline, when 100 redundant baselines are averaged with residual phase and amplitude errors. At large phase errors, where measurements sum approximately incoherently SNR falls to approximately 0.03 of the calibrated value. This relative SNR of  $\left(\frac{1}{\sqrt{1000}}\right)$  illustrates the case where adding 1000 incoherent measurements does not improve on the SNR of a single measurement.



(b) A zoomed plot showing SNR reduction with relatively small residual antenna complex gain errors. A relative SNR of 0.995 represents a required increase in observing time of  $\sim 1\%$  to reach the SNR per-baseline of a perfectly calibrated array. A relative SNR of 0.975 represents a required increase in observing time of  $\sim 5\%$ .

Figure 2.11: The reduction in SNR of 1000 redundant visibility measurements due to errors in phase and amplitude calibration.

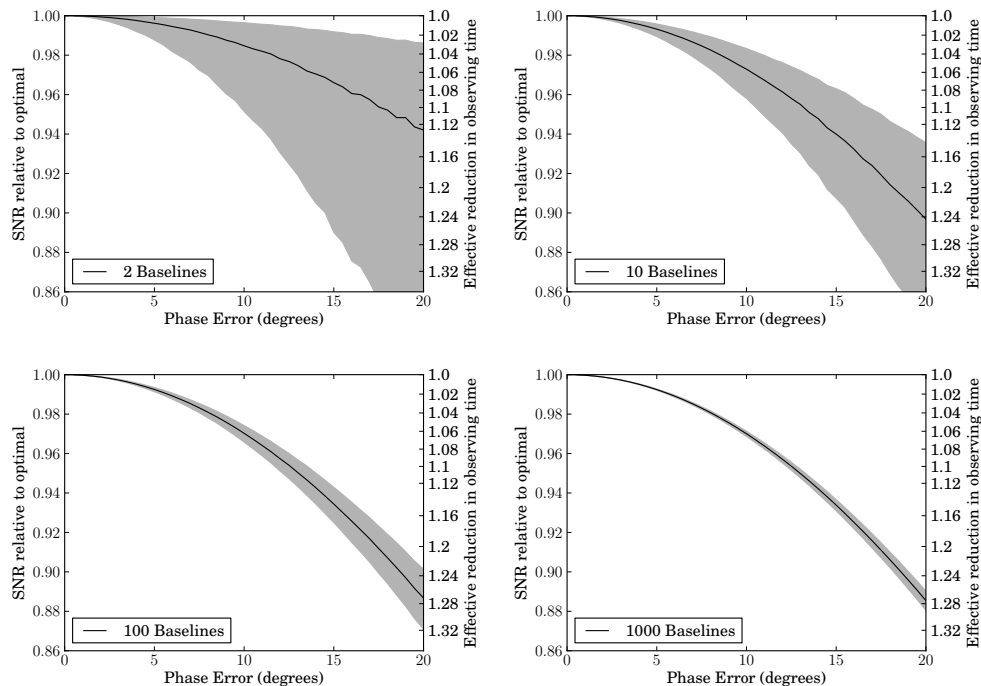


Figure 2.12: Reduction in SNR due to adding redundant baseline visibilities formed from antennas with phase errors. Shaded regions represent 1 standard deviation uncertainty.

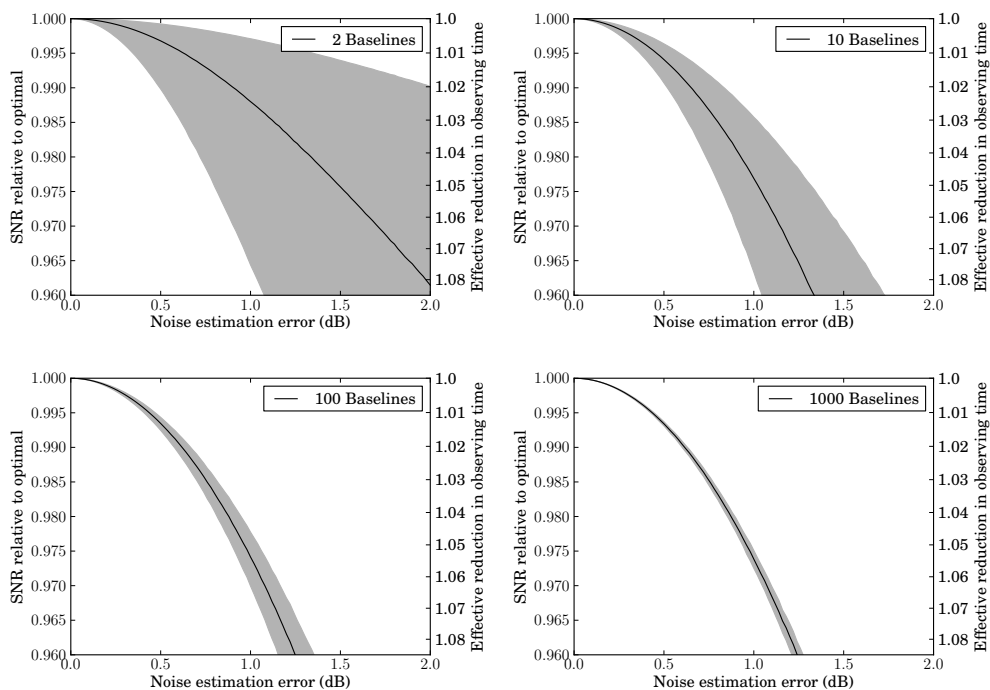


Figure 2.13: Reduction in SNR due to adding redundant baseline visibilities formed from antennas with incorrect noise weighting. Shaded regions represent 1 standard deviation uncertainty.

## Obtaining Calibration Coefficients

The results above suggest that gain errors should not pose a great obstacle to the use of direct-imaging systems. However, real-time calibration, as well as a mechanism for calculating the per-baseline corrections present in Equation 2.22, is required. The real-time calibration element can simply be implemented by means of any of the calibration methods described in Section 2.3, provided a subset of the full correlation matrix is recorded in parallel to direct imaging. Since the full correlation matrix renders direct images redundant, such a system is assumed to instantaneously compute only a subset of the total baselines of the array, possibly over a limited bandwidth. Provided calibration solutions can be computed quickly, relative to the rate at which the antenna gains fluctuate, this information can be used to compute the gain corrections which need be applied in real-time. In fact, even if calibration cannot take place at a rate adequate to match all gain variations, provided a coarse calibration can be computed in good time, correlation matrix data can later be used for per-baseline calibration. In Chapter 5 we discuss the deployment of a 32 antenna direct-imaging system, which was found to perform almost optimally even for calibrations which were only updated daily.

### 2.5.3 Autocorrelations

So far, our discussion has been based on the premise that images formed by correlators and direct imagers are mathematically equivalent. However, an important difference between images formed directly and those formed by synthesis of two-element interferometers is the presence of autocorrelation data in the final image product. Autocorrelation measurements are invariably noisier than non-zero baseline correlation measurements, because noise between different antennas is uncorrelated. Furthermore, since the autocorrelations are frequently many orders of magnitude larger than the cross-correlations, small gain fluctuations in antenna receiver chains can induce variations in the autocorrelations which are substantially larger than the cross-correlation signals. Because of these properties, autocorrelation data is usually not included when synthesizing images from interferometers. This entails a crucial difference between indirect and direct image-forming systems. In the former, individual baselines are available before image creation, and any autocorrelations can simply be discarded before any further processing is performed. In the latter, the image is immediately formed, and encodes measurements from *all* baselines in an array, including autocorrelations. This simple fact breaks the true equivalence of FFT-based and correlation-based image forming. Only if the unwanted autocorrelations can be removed from the image in both cases are the techniques equivalent.

In the case of a regular gridded array, autocorrelations can be isolated and removed by transforming the directly-obtained image back into the visibility plane, and then reforming it with the autocorrelations removed. This process is identical to simply subtracting the mean value from the directly obtained images. For an irregular array with a MOFF correlator, the situation is not so simple. Here, the gridding process of the MOFF correlator is equivalent to the  $uv$  gridding performed in traditional synthesis imaging, but without first excluding autocorrelations. The convolutional gridding process has the effect of spreading the contributions of autocorrelations into all  $uv$  grid points that lie within a convolution kernel width of the origin. Thus, not just the mean offset in the synthesized image is made more noisy, but also all other small  $k$  modes.

This issue is not addressed by Morales (2011), though it can be combated in two ways:

1. Ignore all contaminated  $uv$  measurements. In this case, any correlations corresponding to baselines within a convolution-kernel distance of one another are simply rejected. If these short baseline measurements are required, they can be computed independently using standard correlation methods.
2. Record directly each antenna's autocorrelation, in parallel with MOFF processing. Knowledge of the gridding kernel used to resample the array allows the contribution of autocorrelations to each  $uv$  grid point to be calculated. If autocorrelations are known, these contributions can be removed in post processing.

The former offers the simplest method of correction, though with a potentially significant increase in real-time infrastructure. The latter requires relatively little dedicated real-time processing, but requires significant post-processing, especially if the resampled grid locations are varying functions of time. In either case, the increase in complexity and material cost of a MOFF system to avoid autocorrelation contamination needs to be considered in any potential deployment. More generally, it should always be remembered that direct and indirect imaging processes are not entirely equivalent because of their different treatment of autocorrelations.

#### 2.5.4 Phase Centre Tracking with Interpolation Gridding

As the sky rotates over an array, the projected baselines of pairs of antennas towards a given sky-coordinate changes. The resultant cyclic variation of the phase of a given visibility is known as a fringe. In order that time-averaging of data does not wash out these fringes, it is necessary to compensate for the rotation of the earth prior to the averaging process. This compensation is usually termed “fringe-stopping” or “delay tracking”.

In the case of any array where processing is performed on signals from fixed antennas, some variety of fringe-stopping or phase rotation is always necessary. Worse still, Equation 2.12 shows that the fringe rate of a visibility is a function of sky co-ordinates. Hence, when compensation is applied, it is only effective for a relatively small patch of sky. For all-sky imaging with a large array, it is therefore necessary to process multiple fields of view, each with a different phase centre. Since the MOFF correlator correlates virtual grid points, which can be constructed in arbitrary locations, it provides a mechanism for a novel implementation of phase-centre tracking. Such an array can, in principal, use a time varying virtual grid with positions fixed relative to the sky, rather than the ground. Such a system is the digital equivalent of an Equatorial mount. An initial phase delay would be required to phase the array to the local Celestial Pole, whilst a rotating virtual grid would be constructed in this plane to track the sky.

The advantage of such a system over traditional fringe stopping is that the rotation of the entire sky is compensated for in a single operation. Since the MOFF correlator requires gridding regardless of any tracking requirements, a time-variant grid need not necessarily be an overwhelming computational overhead, particularly since the relevant gridding coefficients can be pre-computed prior to observation.

### 2.5.5 Computational Demands

For correlators with  $N$  antenna inputs generating  $M$  spectral channels, the computational demands of FX, XF and direct imaging with and without gridding (to  $G$  grid points) are summarized in Table 2.2. In each case processing requirements have been normalized to unit bandwidth, and it is assumed that accumulation of visibilities occurs over long enough timescales that post-accumulation computations are negligible. To aid comparison, we assume that in both the FFTT and MOFF cases, images are generated which have twice the number of pixels as grid points in each of two dimensions. That is, we assume that zero-padding is used in both cases.

The total costs for the processing paradigms in Table 2.2 are shown for different array sizes in Figure 2.14, for  $M = 1024$ , and the two cases  $G = N$  and  $G = 4N$ . In the case of the MOFF correlator, we assume that  $K = 4$ , i.e. that each antenna signal contributes only to its 4 nearest virtual grid points. Figure 2.14 exemplifies the driving force for interest in FFT-based direct-imaging systems, which is the reduction in computation costs for large- $N$  arrays. Since XF correlators are clearly undesirable for large- $N$  arrays, the baseline against which direct imaging systems should be judged is the FX correlator. In this case, assuming that any required gridding operations make up a small fraction of the overall computational demands of an imager, direct imaging is more efficient than FX correlation when

Processing Paradigm	Channelization	Multiplication	Gridding
FX Correlator	$CN \log 2M$	$C \times \frac{1}{2}N(N+1)$	1
XF Correlator	1	$2M \times \frac{1}{2} \frac{C}{4}N(N+1)$	1
Direct Imaging (via FFT)	$CN \log 2M$	$C(3G \log 4G) + 2CG$	1
Direct Imaging (with gridding)	$CN \log 2M$	$C(3G \log 4G) + 2CG$	$CNK$

Table 2.2: Computation cycles for correlators with  $N$  antenna inputs generating  $M$  spectral channels. The factor  $C$  relates to the number of real operations required for a complex multiply-accumulate, and is either 8 if real multiplications and additions are both counted, or 4 if only real multiplications are counted. Processes given a computational cost order 1 are those which may be applied after data rates have been lowered by accumulation.




Figure 2.14: Computational comparison of FX, FFTT, MOFF and XF architectures, for  $G = N$ ,  $M = 1024$ ,  $K = 4$  (left) and  $G = 4N$ ,  $K = 4$  (right).

$$\begin{aligned}
C(3G \log 4G) + 2CG &\leq \frac{1}{2}CN(N+1) \\
G \log(3G) &\lesssim \frac{N^2}{6}.
\end{aligned}$$

## Operation Bitwidth

An important consideration when comparing the relative efficiency of interferometer architectures is to appropriately weight the variety of different operations required. In particular, the bit-widths required by different operations can vastly change the efficiency with which an architecture can be implemented on any given platform. In Section 2.1.2 it was shown that FX correlators can be very effectively implemented with 4-bit input signals, and 4-bit complex multiplies. In contrast, to maintain SNR in a large FFT requires significantly larger computational word sizes, even if the input signal is of only 4-bits. A detailed analysis of noise resulting from FFT word sizes can be found in Qadeer et al. (2011), here it suffices to

Processing Paradigm	Output Data Quantity (words per integration)
FX Correlator	$\frac{1}{2}MN(N + 1)$
XF Correlator	$\frac{1}{2}MN(N + 1)$
Direct Imaging (via FFT)	$4GM$
Direct Imaging (with gridding)	$4GM$

Table 2.3: Output data rate per integration for different processing mechanisms.

note that for large FFTs to accurately produce expected array-beams, phase factors should be accurate to greater than  $1^\circ$ . This implies significantly more precision than 4-bits in the FFT twiddle factors, resulting in multiplies far wider than the  $4 \times 4$  bits required by FX correlators. In Chapter 3 it is explicitly shown that Field Programmable Gate Array (FPGA) hardware, used extensively in modern radio astronomy, can perform approximately 4 times as efficiently if multiplies are restricted to fewer than  $5 \times 5$  bits. These savings would not be possible in FFT-based interferometers. With custom computing hardware, based on Application Specific Integrated Circuits (ASICs) the chip area required for a multiplier scales approximately with the product of the multiplicand bit-widths (Chapter 15, Smith, 1997). Again, this means that FFT-based processors are likely to be more costly per-operation to implement than simpler FX correlators. It is likely to be only on floating-point hardware, such as commodity CPU and GPU systems that the increased bit-width of FFT-based operations will not impact processing costs. However, in these systems, it is frequently the input data rate, rather than the computation requirements, that limit processing capacity. In this case a more efficient signal-processing method may not necessarily improve the overall system performance.

### Data Output Rate

Of arguably greater importance than the potential computational savings of direct-imaging systems is the substantial reduction in output data rates, summarized in Table 2.3. In the case of direct-imaging systems, the output data rate is directly proportional to the number of pixels (or beams) forming the image, rather than the  $O(N^2)$  visibilities. As discussed in Chapter 6, for the case of the Square Kilometre Array, output data rates for large arrays can total many Terabytes/second. In this case, regardless of whether individual visibilities, rather than sky images, are available as data products, it quickly becomes unfeasible to store them in their native form.

Array	$N$ (dual-polarization)	Frequency Range (MHz)	$A_{\text{phys}}$ ( $\text{m}^2$ )	$A_{\text{eff}}$ ( $\text{m}^2$ )
LOFAR LBA SEPCAM	96	15–75	3 310	3 310–512
LOFAR LBA AARTFAAC	288	15–75	96 200	9 930–1 536
LEDA	256	30–88	7 850	7 850–992
MWA	512	80–300	1 770 000	10 700–6 700

Table 2.4: Specifications of a selection of current large- $N$ , low-frequency interferometers. In the case of the LOFAR and LEDA arrays, which are based on dipole antennas, per-element effective area is taken to be  $\min\left(\frac{\lambda^2}{3}, \frac{\pi D^2}{4N}\right)$ . In the case of the MWA, per antenna effective area is based on information from the MWA Collaboration (Figure 10, 2009).  $A_{\text{eff}}$  is computed as  $N \times$  Effective area per element.

## 2.6 Relative Merits of Direct-Imaging Systems for Existing Telescopes

To put the analysis presented above into perspective, it is useful to consider the potential merits of direct imaging telescope systems based on the specifications of real arrays used in modern radio astronomy. Given that the main driver for direct imaging systems is their computational efficiency when applied to large, relatively dense arrays, we consider here some of the largest- $N$  interferometric systems currently under development. These are the LOFAR single station correlator (SEPCAM, Foster, 2013), the LOFAR “superterp” correlator (AARTFAAC, Prasad & Wijnholds, 2012), the Murchison Widefield Array (MWA, Lonsdale et al., 2009) and the Large Aperture Experiment to Detect the Dark Ages (LEDA, Greenhill et al., 2012). The relevant specifications of these arrays are given in Table 2.4. In the case of LOFAR and LEDA, which utilize dipole antennas, effective areas per element have been calculated as  $\min\left(\frac{\lambda^2}{3}, \frac{\pi D^2}{4N}\right)$ . The effective area of the MWA is estimated from the MWA Collaboration (Figure 10, 2009) from which a two-part linear spline has been approximated. Figure 2.15 shows the filling factors of these arrays, which are all less than 1 over the majority of their observing ranges.

Given these array specifications, Tables 2.2 and 2.3 can be used to estimate the relative efficiency of a direct-imaging, versus traditional FX correlator. Using  $G = \frac{N}{f}$ , and ignoring potential gridding costs, this relative efficiency is plotted in Figure 2.16a. The low filling factors of these arrays lead to computationally inefficient application of direct-imaging. LEDA and LOFAR both have bands of operation where direct imaging is more efficient, but these are outweighed by significant region where direct imaging is several times more costly than traditional FX-based interferometry.

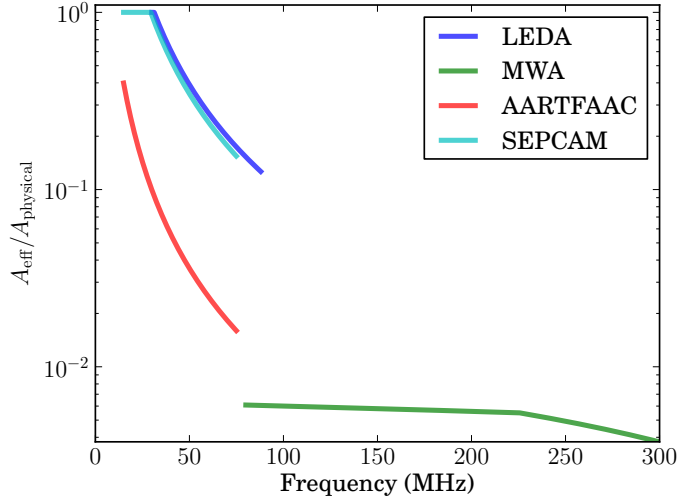


Figure 2.15: Filling factors, that is, the ratio of effective to geometric area, of the four example arrays. In the case of LEDA, AARTFAAC and SEPCAM

In terms of data output rate, the case for direct imaging systems is much more compelling. Output data rates from all arrays but the MWA are significantly reduced by FFT imaging at most frequencies. Though most literature focuses on the potential computational improvements offered by FFT imagers, in this analysis it is clear that it may well be output data rate reductions which drive adoption of such instruments.

We finally note that in the case of FFT-based imagers, filling factor issues could be mitigated by simply demanding that the antennas in these arrays be placed on a densely populated, but widely spaced grid. In this case, computational savings can be significant, and depend only on the total number of antennas in the arrays. However, the ease of imaging in this case comes at the expense of synthesized beam quality, and significant grating lobe structure will exist if this strategy is adopted.

## 2.7 Conclusions

This Chapter has explored the theoretical underpinnings of interferometry and synthesis imaging, including the relatively recent appearance of direct-imaging techniques as replacements for traditional correlators. Much of that discussed is extensively covered in substantial texts on interferometry, but a number of subtleties with regard to direct-imaging which are frequently not recognized in recent literature are worthy of emphasis.

Firstly, the difference of treatment in autocorrelations in direct-imaging and traditional interferometry systems requires careful consideration when proposing new synthesis-imaging systems. The removal of autocorrelation data in the traditional image-making process breaks the equivalence with direct FT-based imaging. The requirement that autocorrelations must

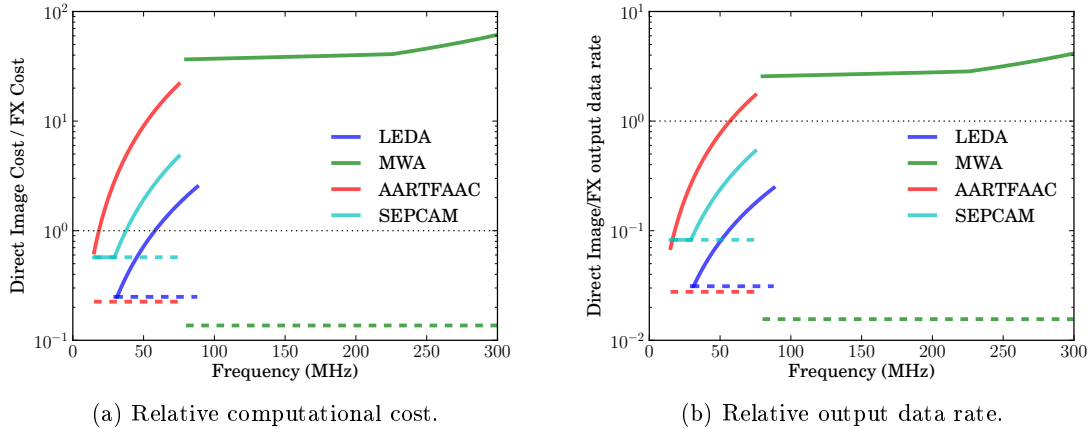


Figure 2.16: Relative computational cost and output data rate of direct-imaging and FX correlation with the LEDA, MWA, AARTFAAC, and SEPCAM arrays. Solid lines denote cost for  $G = \frac{N}{f}$  grid points. Dashed lines show relative cost for  $G = N$  grid points.

be similarly removable in directly formed images places extra requirements on direct-imaging systems. Whilst this ability is “built in” to the FFTT process, the more advanced MOFF correlator requires additional data to be gathered from an array to facilitate autocorrelation removal.

Secondly, the emphasis on the post-calibration opportunities presented by direct-imaging systems has been discussed. Beamformers, which place emphasis on the beam shape as the product of interest, suffer from any failure in real-time calibration. Direct-imagers, however, can tolerate reasonable ( $\sim 5^\circ$ , 1 dB) phase and gain errors with no systematic corruption of the final image product. This is the case because the summed, redundant baselines present in the directly-obtained image can be calibrated like traditional visibility measurements, albeit with a small reduction in SNR. The idea that direct-imaging observations do not require perfect real-time calibration is potentially of great practical importance, but also may help to lessen the concerns of the some opponents to direct-imaging systems. Further investigation into such calibration, including what information in addition to a raw image (if any) best aids computation of gain corrections will no doubt prove an important part of the future of direct-imaging interferometers.

It is probably reasonable to say that the immediate usefulness of direct-imaging systems as replacements for traditional FX correlators has been overstated in recent literature. Whilst the potential computational cost savings of such telescopes are undoubtedly huge, the specifications of the largest arrays currently under development suggest that they will not appear in use unless a significant change in design philosophy emerges. Currently, the strong desire of astronomers to avoid grating lobe structure, whilst demanding arrays

with high resolution and fractional bandwidth but limited cost, mandates the use of arrays with low filling factors. These factors can quickly overwhelm the efficiency of direct imaging, and undermine the main appeal of these systems over their traditional interferometry counterparts. In addition, and unconsidered in current arguments proposing such systems, bit-width-dependent penalties potentially increase the computational cost of direct-imaging operations relative to low-bitwidth cross-multiplication.

With all these potential limitations in mind, it may seem that direct-imagers will remain absent from modern radio telescopes. However, whilst current arrays may not favour their deployment, the constant demand for ever more capable telescopes will doubtless push the scale of “large- $N$ ” arrays to the point that more efficient image-forming algorithms are required. With arrays currently deployed, it is already the case that data output rate reductions offered by direct imaging systems may drive their adoption. It is also the case that direct-imaging systems can already be made to be far more efficient than FX correlators, providing that low filling factor arrays are arranged on densely populated but widely spaced grid points. Whilst in the imaging sense the grating lobes resulting from this configuration may be intolerable, for many scientific uses the Fourier modes measured by such an array are satisfactory, and spatial degeneracies can be mitigated by other means.

Finally we note that merely having access to the all-sky electric field distribution, as produced by direct imagers, is potentially very useful. All the direct-imaging systems outlined in this chapter can equally be used as efficient FFT-based beamformers. This use is further discussed, in the context of the SKA, in Chapter 6.

## Chapter 3

# An Optimized, Low-Bitwidth Correlator for FPGAs

FX correlators, which are the predominant type seen in newly-constructed interferometric arrays, have computational complexity which scales as  $O(N^2)$  with the number,  $N$ , of antennas in an array. This scaling is driven by the need to multiply all possible pairs of antennas in the array, with one complex multiply required to calculate each entry in the array's  $N \times N$  correlation matrix. The interconnection of data streams required to feed the processors performing the correlation calculations also has an unfavourable  $O(N^2)$  scaling, and leads to data routing infrastructure being a significant contributor to the overall cost of an FX correlation system. For this reason, correlators invariably take advantage of the noise-like nature of the received sky signal. This allows quantization to a low (usually  $\lesssim 4$ ) number of bits, reducing required data routing bandwidth without significantly impacting the fidelity of data.

This reduction of bits carried out to aid signal transport has further implications for downstream processing. Since, other than completely custom circuitry, all processors calculate products of numbers using multiplier cores with fixed input bitwidths, multiple operations performed on small bitwidth inputs can often be combined into single large-bitwidth calculations<sup>1</sup>. In correlators built for platforms that allow low-level control of signal processing algorithms, significant performance increases can be gained by such optimizations, with few computational resources required to arrange data in optimal form. In particular, we single out devices that are capable of performing bit shifting, concatenation, and masking operations with no processing overhead as especially susceptible to such optimizations.

In this Chapter, we examine the potential optimizations which allow all four real multiplications, and one of the two additions, required for a complete 4-bit complex multiply

---

<sup>1</sup>This should not be confused with the concept of “Single Instruction Multiple Data” (SIMD) processes, which efficiently implement multiple identical operations on elements of an input data vector using single instructions.

to be combined into a single  $18 \times 18$  bit multiplication operation. We present an existing implementation of such a solution, and demonstrate a novel alternative which offers further advantages when used on current Field Programmable Gate Array (FPGA) based platforms. We then proceed to use this multiplier implementation to optimize an existing correlator design, and examine the resulting performance improvements.

### 3.1 FPGAs in Radio Astronomy DSP Applications

Field Programmable Gate Arrays (FPGAs) are reconfigurable devices comprising a repeated array of logical, memory and arithmetic elements, with programmable interconnect. These can be used to perform logical functions on digital signals. Large computational and interconnect requirements have become commonplace in the field of radio astronomy. FPGAs, which are capable of higher data transfer bandwidths and compute rates than general purpose Central Processing Units (CPUs) at a fraction of their power consumption, are an attractive processing solution.

The ease with which FPGAs can be interfaced with analogue to digital converters (ADCs), using readily-customized, high-bandwidth interfaces also makes them ideally suited for use in radio astronomy. Even where more flexible CPU platforms are chosen to carry out data processing in a radio astronomy instrument, some variety of FPGA is usually still used to facilitate ADC to CPU communication.

The low-level control of signals in an FPGA afforded to a firmware designer allows some algorithms to be implemented very efficiently. However, this low-level control can also limit the portability of firmware between different FPGA models, which may have different types and quantities of underlying circuit elements. This contrasts with the task of creating software for CPUs, where software is relatively portable between CPU models, even through multiple generations of technology.

### 3.2 The CASPER X-engine

The Collaboration for Astronomy Signal Processing and Electronics Research group (CASPER) is an international collaboration whose primary goal is to simplify and accelerate the design of astronomy instrumentation through the reuse of open source software and hardware resources (Parsons et al., 2006). Such accelerated design aims to minimize the so-called *time-to-science* of an instrument, the time between conception of an experiment and deployment of a science-capable instrument. Minimizing this delay allows the Moore's-law progression in technology to be leveraged as effectively as possible, and avoids the common

scenario where an instrument, taking several years to design, is based on technology which is obsolete long before it is deployed.

Whilst CASPER developments encompass a broad range of radio (and more general) astronomy systems, the flagship design of CASPER is the Packetized Correlator (Parsons et al., 2008). This is an implementation of a highly-flexible, highly-scalable, FX correlator architecture, based on FPGAs for data processing and commercial Ethernet switches for signal transport and routing. The subject of this section is the X-engine<sup>2</sup>, the sub-module of the Packetized Correlator responsible for calculating cross products between pairs of antenna signals. In this chapter, I will use the CASPER X-engine as a platform on which an optimized cross-multiplying core will be based.

I will begin with a brief functional overview of the CASPER X-engine core in its original form and describe both its successes and limitations. I will then discuss possible optimizations to the X-engine design, and describe the implementation of such a module.

### 3.2.1 Module Description

The CASPER X-engine is a parameterized module which processes data in fixed-length windows. Each window comprises blocks of  $M$  time samples of a single frequency channel, serially from  $N$  dual-polarization antennas,  $A_{0,\dots,N-1}$ . A graphical representation of a window of data is shown in Figure 3.1. Consecutive windows of data usually (but not necessarily) cycle through a collection of different frequency channels. Windows of  $MN$  samples must be fed into the X-engine in an uninterrupted block, whilst gaps between blocks are allowed, as long as they are an integer multiple of  $MN$  clock cycles long. This requirement allows the X-engine to operate with no internal clock-enable logic.

The X-engine, depicted in schematic form in Figure 3.2 and further described in Appendix C, comprises a series of “taps”. Each tap calculates and accumulates the correlation between a particular pair of dual-polarization, complex-valued antenna streams by taking as inputs signals propagated through a series of  $M$ -sample delays.

The first tap of the delay chain calculates the autocorrelations of each antenna, whilst tap  $p$  calculates the correlation between antennas  $A_i$  and  $A_{i+p}$  for  $0 \leq i < N - p$ . Tap  $p$  is also responsible for calculating the correlation of  $A_j$  and  $A_{j+N-p}$  for  $0 \leq j < p$ , which occur as tap inputs precisely at times when  $A_i, A_{i+p}$  pairs are not valid because  $i > N - p$ . These times represent the occasions where the multiplier inputs, were they to be separated by a delay chain  $M \times p$  samples long, would be associated with different windows of data that don’t require correlation.

---

<sup>2</sup>designed by Lynn Urry and implemented and maintained by Jason Manley, Aaron Parsons and Terry Filiba. See [https://casper.berkeley.edu/wiki/Win\\_x\\_engine](https://casper.berkeley.edu/wiki/Win_x_engine)

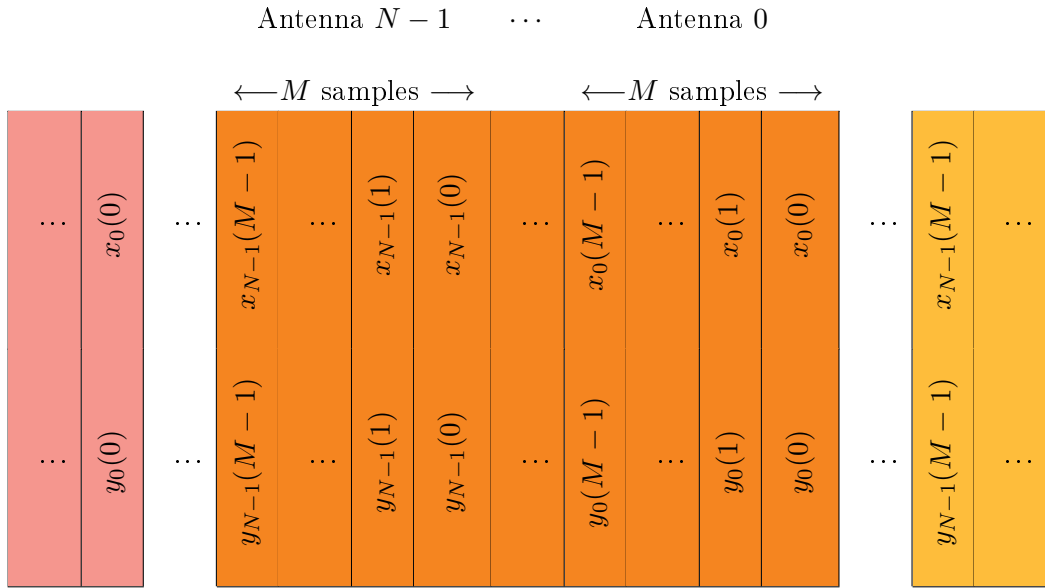


Figure 3.1: Data is input into the CASPER X-engine in a series of windows. An example of the organization of data in such a window is shown above, where  $x_m(s)$ ,  $y_m(s)$  represent the  $s^{\text{th}}$  time sample from the  $x$  and  $y$  polarizations of the  $m^{\text{th}}$  antenna, respectively. Each window comprises  $M$  time samples from a single frequency channel for a series of  $N$  dual-polarization antennas.

The allocation of different antenna pair calculations between different taps is shown in Figure 3.3, for the example of an 8-antenna system. Using the method described, all independent correlations between  $N$  antennas can be calculated with an X-engine comprising  $\lfloor N/2 + 1 \rfloor$  taps. Such a system utilizes all taps with 100% efficiency for odd values of  $N$ , and all but the last tap with 100% efficiency in the case of even  $N$ . In this latter scenario, the last tap operates with 50% efficiency, since the correlations  $A_i^* A_{i+p}$ ,  $0 \leq i < N - p$  and  $A_j^* A_{j+N-p}$ ,  $0 \leq j < p$  are identical copies.

The natural readout order of the X-engine tap chain is in the form depicted in Figure 3.4. Readout occurs in a series of continuous blocks, with each block containing a result from each tap in the X-engine tap chain. Even though windows of data are input sequentially into the X-engine, the architecture employed results in overlapping blocks of output results. These overlapping blocks must be separated and reordered prior to further downstream accumulation of the correlation matrix. In the CASPER X-engine design, this reordering is performed in on-chip FPGA memory, which is used to buffer an entire correlation matrix containing  $N(N + 1)/2$  values.

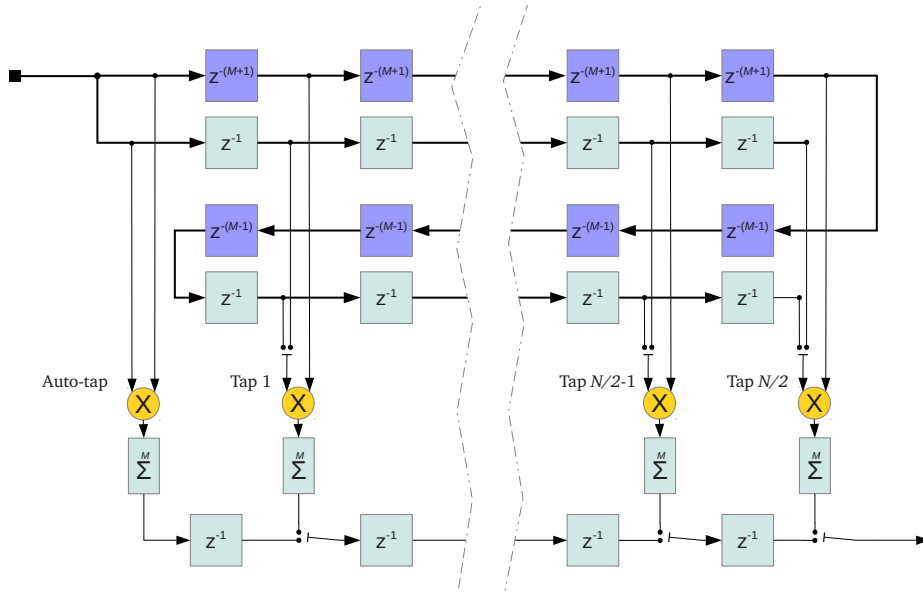


Figure 3.2: Schematic depicting the flow of data through a windowed X-engine module, processing  $N$  antennas with  $M$  accumulated samples per correlation. Based on a similar block diagram from Parsons et al. (2008).




Figure 3.3: The distribution of correlation calculations among different taps for an 8 antenna system. We represent the correlation of a pair of antennas  $A_i$  and  $A_j$  as  $V_{i,j}$ . All taps are 100% utilized for odd numbers of antennas, whilst the last tap is 50% utilized when processing even numbers of antennas.

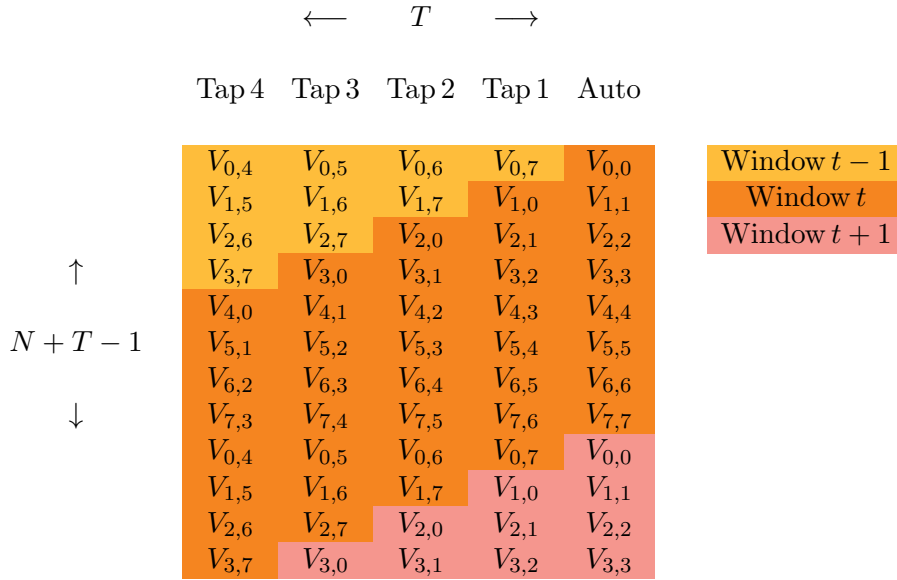


Figure 3.4: The natural readout order of data from an X-engine tap chain configured for  $N = 8$  antenna correlation (with  $T = 5$  taps). We represent the correlation of a pair of antennas  $A_i$  and  $A_j$  as  $V_{i,j}$ . Readout is from top-left to bottom right. For an X-engine configured to integrate over  $M$  samples, a row of  $T$  results (one calculated by each tap) is output in a continuous block every  $M$  clock cycles, resulting in the condition that  $M \geq T$  for a system to be viable. The natural output order of the tap chain gives rise to windows of data which overlap, with the results from three windows,  $t - 1$ ,  $t$  and  $t + 1$  mixed in the same block of  $T \times (N + T - 1)$  outputs. This non-trivial ordering must be accounted for in downstream processing before further accumulation of the correlation matrix can take place. Note also the duplication of visibilities in the column calculated by tap 4. The presence of pairs of visibilities  $V_{j,i} = (V_{i,j})^*$  is a symptom of the non-optimal nature of X-engines processing even numbers of antennas. Such duplicated results are masked and ignored in downstream processing.

### 3.2.2 Resource Utilization

Simple analysis of the CASPER X-engine allows quantification of the amount of FPGA resources which are required for the CASPER X-engine. This offers some insight into where optimizations may be implemented most effectively. In examining the X-engine, we parameterize the number of antennas and integration length as  $N$ , and  $M$ , respectively. We also introduce a new parameter,  $B$ , for the number of bits of input signal, and  $T$ , for the number of taps in the engine. As noted in the previous section,  $T$  is not an independent X-engine parameter, and is fixed by the number of antennas needing to be correlated, with  $T = \lfloor N/2 + 1 \rfloor$ . In the following discussion we assume that the dual-polarization, complex inputs to the X-engine are quantized to real/imaginary 4-bit precision pairs (i.e.  $B = 16$ ). Such a quantization scheme is common in modern radio astronomy, with 4-bits representing a compromise between the reduction in signal-to-noise caused by coarse quantization, and the processing and infrastructure costs incurred by fine quantization.

#### DSP Slice Usage

The X-engine can be configured so that the underlying complex multiplies are implemented in a variety of ways. “Soft-logic” may be used to implement multipliers using the most abundant and flexible logic elements of the FPGA. RAM-based look up tables (LUTs) can also be used to perform multiplication. However, the preferred configuration uses FPGA DSP slices. These are dedicated areas of the FPGA designed for efficiently implementing arithmetic, and in particular, multiplication, functions.

The current generation of FPGAs manufactured by Xilinx feature “DSP48E1” cells, each of which comprises a  $25 \times 18$ -bit multiplier and a 48-bit adder/subtractor, capable of adding or subtracting results output from neighbouring DSP slices (Xilinx Inc., 2012a). Recent generations of Xilinx FPGAs, as well as those produced by other manufacturers, feature DSP cells with similar functionality, although multiplier input width may be limited to  $18 \times 18$  bits. More advanced features, including pre-adder and pattern detect capabilities, which are also present in a DSP48E1 slice (depicted in full in Figure 3.5) are not used in the correlator development presented here (and vary amongst different manufacturers and generations of FPGAs) and will not be discussed further.

Using the simple implementation shown in Figure 3.6, four DSP48 slices are capable of computing a complex multiply, with a total of four multipliers, one adder, and one subtractor utilized. The current CASPER X-engine is configured for dual-polarization antennas and requires four complex multiplies (for each of the polarization combinations,  $A_x B_x$ ,  $A_x B_y$ ,  $A_y B_x$ ,  $A_y B_y$ , of two antennas  $A$  and  $B$ ) per tap. The number of DSP slices,  $n_{DSP}$ , used in

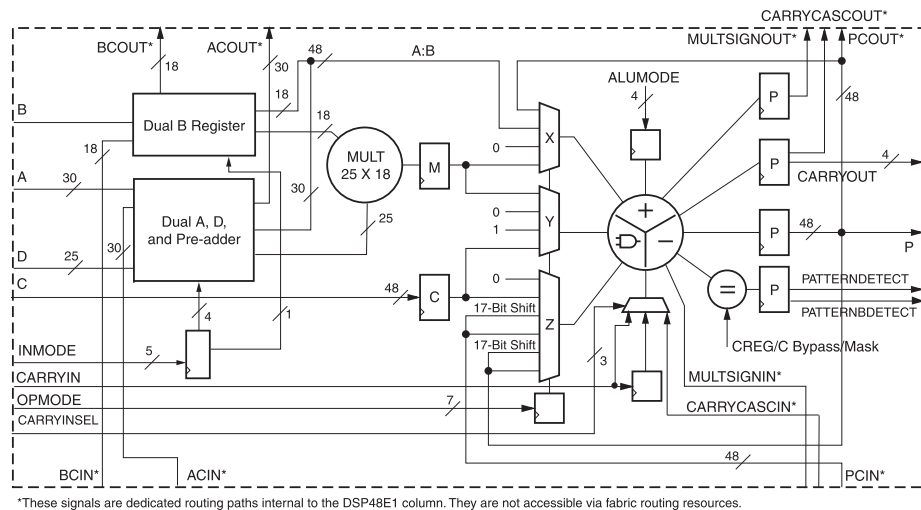


Figure 3.5: Xilinx Virtex-7 DSP48E1 Slice, reproduced from Xilinx documentation (Xilinx Inc. (2012a)). Core features are a  $25 \times 18$  bit multiplier core, and 48 bit post-multiplication arithmetic unit, capable of addition and subtraction, as well as more advanced features. Of particular importance, and present on all modern high-performance FPGAs, are the “PCIN” and “PCOUT” signal paths, which allow computation results from one DSP slice to be passed as inputs to neighbouring slices.

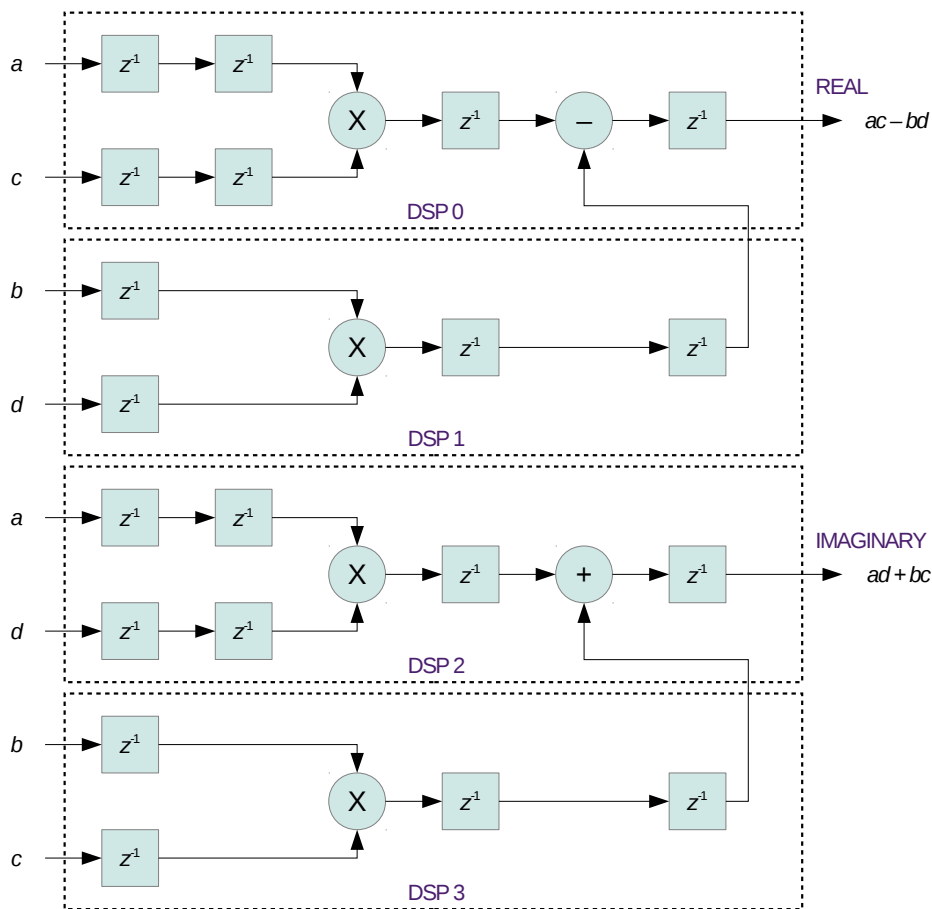


Figure 3.6: An implementation of a complex multiply, implemented using four DSP slices. The arrangement shown calculates the product of two complex numbers  $a + ib$  and  $c + id$ . There are various possible pipelining configurations, with that shown giving rise to a multiply with four clock cycles latency.

an X-engine is, therefore:

$$n_{DSP} = 4 \times 4 \times T = 16 \times \lfloor N/2 + 1 \rfloor . \quad (3.1)$$

The number of multiplications in the X-engine cannot be significantly reduced<sup>3</sup>. However, since each  $4 \times 4$ -bit multiply uses only a fraction of the computational power of a DSP slice, it is the case that the total number of DSP slices required by the X-engine design can be decreased. This will be shown shortly.

### Block RAM Usage

Inter-tap delays in the X-engine are usually implemented in block RAMs (or BRAMs) on the FPGA by writing an incoming data stream to memory continuously, and reading it out some number of clock cycles later. In the X-engine, there are two categories of delays.  $(M + 1)$ -cycle delays fall between each of the  $T$  multiplier stages, whilst  $T$  repeated  $(M - 1)$ -cycle delays form a single  $T \times (M - 1)$  delay unit between the last and first taps. These two types of delays are implemented differently on the FPGA, owing to the physical constraints on input/output bandwidth and depth to each block RAM element. In current generations of Xilinx FPGAs, block RAM elements are fundamentally 18 kb in size<sup>4</sup>. In most circumstances they are configured in the X-engine as 18-bit wide RAMs, with  $2^{10}$  addressable memory locations, though a BRAM can also function as a 36-bit wide, 512-word deep unit, by utilizing dual 18-bit read and write ports.

Provided that  $M \lesssim 2^{10}$  and  $B < 18$  (or  $M \lesssim 512$ ,  $B < 36$ , in the dual-ported case) the  $(M + 1)$ -cycle delays are each implemented as one block RAM and a small number of single-cycle register delays. The  $T \times (M - 1)$ -cycle delay requires multiple BRAMs, owing to its required depth of over  $2^{10} = 1024$  words. Here, the number of BRAMs required is  $\lceil T(M - 1)/1024 \rceil$ . Thus, in total, the number of 18 kb BRAMs used in an X-engine when  $B < 18$ ,  $n_{BRAM-18}$ , is:

$$n_{BRAM-18} = T \times \lceil M/1024 \rceil + \lceil T(M - 1)/1024 \rceil . \quad (3.2)$$

In practice, usually  $M \propto N$  so that the increase in data rate associated with cross multiplication is nullified by increased accumulation length, and at the very least  $M \geq T$  to ensure sufficient readout time of X-engine correlation results.

<sup>3</sup>Except by representation of input data in some multiplier friendly fashion using either logarithms (which allow multiplications to be replaced by additions), or magnitude-argument specification of complex numbers (which can be multiplied with only a single real multiplication). In either case, the processing required to encode and decode the representation of the numbers usually outweighs any potential performance benefits. For this reason, such endeavours are not considered in this thesis.

<sup>4</sup>In actual fact, each Block RAM is fundamentally a 36 kb region of storage, though 18 kb is the smallest usable unit of memory supported

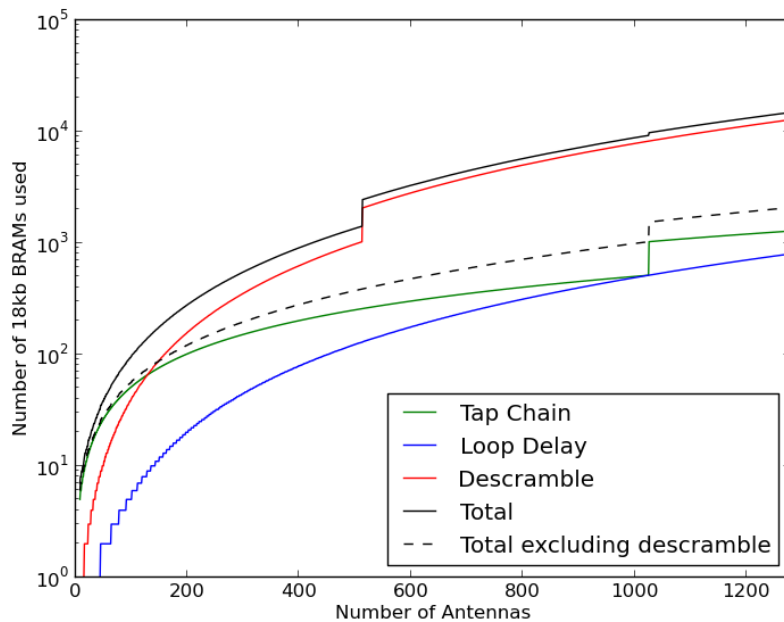


Figure 3.7: Number of 18 kb block RAMs required by an X-engine accumulating  $M = N$  samples per visibility, as a function of number of antennas. Contribution to the total number of BRAMs from each of the tap, loop and descramble subsystems is shown. For this system, descrambling is the dominant BRAM source for large ( $N \gtrsim 128$ ) numbers of antennas.

A complete X-engine also uses BRAM for reordering (“descrambling”, in CASPER parlance) windows of  $O(N^2)$  visibilities before outputting. However, since output data are usually immediately accumulated in off-chip memory, the functionality of this reorder can be bundled into a downstream vector accumulator. This functionality does not currently exist, but is currently under development as part of this work.

Figure 3.7 shows a breakdown of the total BRAM use of X-engines of various sizes. In this figure, we assume that the integration length is equal to the number of antennas in the system, i.e.  $M = N$ . Several features of Figure 3.7 are striking. Firstly, it is clear that for large X-engines (with  $N \gtrsim 128$ ), the dominant contributor to RAM usage is the post-correlation reordering of the correlation matrix. The dominance is exacerbated at  $M > 512$ , at which point the values in the accumulated correlation matrix become larger than 18 bits, and require twice as much memory to store them. Secondly, far more BRAMs are required in the X-engine tap chain than the loop delay. Although the amount of memory required by the chain is approximately equal to that needed by the loop delay, the division of the tap chain (which has to be broken into  $T$  discrete blocks, each using at least one BRAM) leads to a much larger requirement of discrete BRAM blocks. Note that at  $M \approx 1024$ , the BRAMs in the tap chain are optimally filled, and the contribution to BRAM use from the tap chain and loop delay is equivalent. However, at  $M > 1024$ , a step function occurs in tap delay BRAM usage, as individual delays become larger than 1024 words, and require implementation in 2 BRAMS.

Figure 3.7 leads us to conclude that, providing data can be descrambled without requiring on-chip RAM, the BRAM use of the X-engine is best optimized by considering the tap chain used in the architecture. To this end we recall that for  $M < 512$ , it is possible to implement taps delays up to 36 bits wide in a single block RAM, suggesting that increasing the input data rate to the tap chain will result in more efficient RAM use. This represents a possible route to significant resource savings, which we shall take advantage of in Section 3.3.4.

### 3.2.3 Successes of the CASPER X-engine

The main achievement of the CASPER X-engine, and the packetized correlator in general, is the flexibility of operation. The parameterized nature of the X-engine, combined with the use of industry-standard Ethernet connections in the packetized correlator makes the architecture abundantly scalable. The system is currently used by numerous radio telescopes in a variety of science applications, a subset of which are shown in Table 3.1. The CASPER correlator can be adapted to any new telescope system extremely rapidly, with deployment of a digital backend possible on a timescale between several days and several weeks, depending on the exact architecture of the new system. In addition to the libraries of firmware available

	Polarizations	Antennas	Frequency Coverage
PAPER	2 (1)	32 (64)	100 – 200 MHz
GMRT	2	30	50 – 1450 MHz (various bands)
BEST-2, Medicina	1	32	400 – 416 MHz
KAT-7	2	7	1200 – 1950 MHz

Table 3.1: Radio telescopes currently using the CASPER packetized correlator.

from the CASPER group, this flexibility is heavily reliant on a substantial set of open source software utilities, provided by the South African KAT team<sup>5</sup>.

### 3.2.4 Limitations of the CASPER X-engine

Inefficient use of both BRAM resources, which are not optimally used in the tap chain for most values of  $M$ , and DSP-slice-based multiplier resources, which use large (at least  $18 \times 18$ -bit) multipliers to compute the products of 4 bit inputs, represent areas of the current design suitable for optimization.

A more general area for improvement of the current X-engine is the language used for its development. The CASPER X-engine is currently implemented in Xilinx System Generator, a graphical, block-diagram-based FPGA programming environment. This operates on top of MATLAB’s Simulink design and model-simulation platform. Such an environment lends itself to designing streaming-based processing architectures, such as the CASPER X-engine, which are often most readily described in schematic form. The intuitive nature of Simulink as a design tool also makes the CASPER X-engine, and the packetized X-engine system in its general form, relatively easy to use for designers with little or no experience with more traditional FPGA programming languages, such as Verilog or VHDL. A screenshot showing the hierarchical, schematic description of the X-engine as used in Simulink is shown in Figure 3.8.

Though intuitive to use, System Generator is not without its drawbacks. Parameterization of large modules, even with only moderately complicated substructure, can prove cumbersome. Furthermore, these modules—or “blocks”—are inconvenient to use in systems designed using other tools, and cannot be simply ported to FPGAs designed by manufacturers other than Xilinx. In addition, MATLAB, and the Xilinx tools necessary to use the CASPER X-engine, are expensive. This represents a significant hurdle to using the design for those who do not have access to the relevant software via large institutions, which either have the resources to purchase such tools, or are eligible for academic donations.

<sup>5</sup>See the correlator control python package “corr” at <https://github.com/ska-sa/corr>

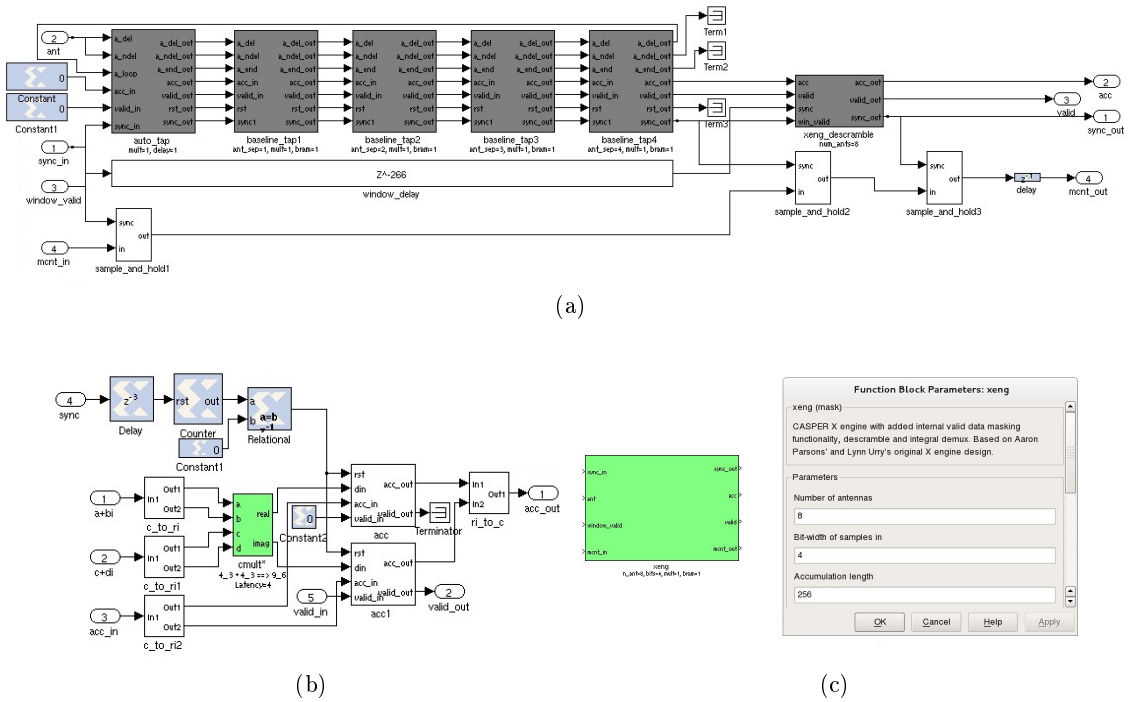


Figure 3.8: (a) The structure of the CASPER X-engine as viewed in the graphical MATLAB Simulink environment. The five, similar grey blocks represent the five taps in an X-engine configured for processing eight antennas. The grey block on the right is the module responsible for reordering overlapping windows of output data. (b) The sub-module of the X-engine forming a complex multiplication and accumulation circuit. Blue blocks in this diagram represent Xilinx FPGA primitives, which perform basic functionality such as integer clock cycle delays, counters, and simple logical functions. White and green blocks represent user defined blocks, constructed using a hierarchy of modules, ultimately described by primitives. (c) The X-engine module as exposed to the end user. The top-level block masks any underlying structure from the user, and allows configuration of X-engine parameters using a simple graphical user interface.

To combat these issues, the CASPER community has been working to consider other FPGA programming platforms, which seek to maintain the ease-of-use of the Simulink-System Generator combination but alleviate the aforementioned limitations. As yet, all such efforts are still in early stages, however, it is my belief that any future CASPER toolflow should be based on modules described with traditional Hardware Description Languages. Whilst it is possible to build a user-friendly layer on top of an HDL platform which allows users to manipulate and build their own modules<sup>6</sup>, the lesson learned from System Generator is that it is sensible not to solely commit all development efforts into a high level language.

### 3.3 Optimized CASPER X-engine

In the remainder of this chapter, I will demonstrate a number of optimizations and improvements which have been applied to the CASPER X-engine. In Section 3.3.1, a method for optimizing low-bitwidth multiplies, demonstrated by de Souza et al. (2007), is described. Central to this optimization is the principle that forcing one component of a complex number to be positive allows two 4-bit components to be concatenated into a single 18-bit multiplier input, with only one of the components requiring sign extension. In Section 3.3.2 I present another possible, and extremely simple, way of alleviating the need to sign extend *any* of the components in a  $4 \times 4$  bit complex multiply. I will show how this gives rise to a multiplier implementation which can be even more efficient, when used as a component of a correlator, than that of de Souza et al. (2007). Readers unfamiliar with two's complement signed binary number representation may wish to refer to Appendix C for a brief overview.

#### 3.3.1 An existing implementation of optimized complex multiplication

Parsons et al. (2008) note that an  $18 \times 18$  bit multiplier is capable of calculating all component cross multiplication products of 3-bit complex numbers, but conclude that 4-bit complex multiplies are not possible. de Souza et al. (2007) take the optimizations further, and manage to implement a 4-bit complex multiply in a  $18 \times 18$  bit multiplier by first generating a new complex number with an extracted sign bit. Given two complex numbers,  $a + ib$  and  $c + id$  to multiply, they first generate new complex numbers:

$$\begin{aligned} S_a(A + iB) &= a + ib & A > 0 \\ S_c(C + iD) &= c + id & C > 0, \end{aligned}$$

where  $S_a$  and  $S_b$  are sign bits equal to  $\pm 1$ . With the sign bits extracted, and working under the assumption that the initial complex number components  $a$ ,  $b$ ,  $c$ , and  $d$  only take values in

---

<sup>6</sup>The free, open-source, python-based myHDL project (<http://www.myhdl.org>) is one such possibility, currently under active investigation by CASPER developers

the range  $[-7,+7]$  (i.e. the value -8 is not allowed),  $A$  and  $C$  need only be represented by 3 bit values. The full multiplication  $(a + ib)(c + id) = S_a S_c ((AC - BD) + i(AD + BC))$  can then be implemented as shown in Figure 3.9. Minimal spillage of logic outside of the multiplier is required, with sign choice of the output being absorbed into a downstream accumulator, capable of either adding or subtracting the complex number  $((AC - BD) + i(AD + BC))$  depending on the sign  $S_a S_c$ .

This implementation of 4-bit complex multiplication is clearly significantly more efficient than a solution requiring four separate multipliers. However, before considering other possible implementations, we note that the given algorithm does not utilize the adder/subtractor module available in the FPGA's DSP slice, nor does it allow parallel outputs to be easily summed, without either explicitly converting the signs of each output or implementing more complex logic to detect the cases where the parallel samples may have different combinations of signs. As will be shown in the following section, in my own optimizations of the multiply operations, these defects are avoided, by adopting a different implementation strategy.

### 3.3.2 Re-factoring the Correlation Algorithm

The correlation algorithm involves calculating the time-average of the quantity  $X^*Y$ , where an asterisk indicates complex conjugation, and  $X$  and  $Y$  are complex values with  $X = a + ib$  and  $Y = c + id$ . Quantities  $a', b', c', d'$  can be defined, such that  $a' = a + 1$ ,  $b' = b + 1$ ,  $c' = c + 1$  and  $d' = d + 1$ . Similarly, we define  $X'$  and  $Y'$  such that  $X' = a' + ib'$  and  $Y' = c' + id'$ .

We now consider the time average of the quantity  $X'^*Y'$  which can be expressed as follows:

$$\begin{aligned}
\langle X'^*Y' \rangle &= \langle (a' - ib')(c' + id') \rangle \\
&= \langle [(a + 1) - i(b + 1)][(c + 1) + i(d + 1)] \rangle \\
&= \langle X^*Y \rangle + \langle (a + b + c + d) - i(a - b - c + d) + 1 \rangle \\
&= \langle X^*Y \rangle + \langle (a' + b + c' + d) - i(a' - b - c' + d) \rangle. \tag{3.3}
\end{aligned}$$

We see that the average visibility  $\langle X'^*Y' \rangle$  is equal to  $\langle X^*Y \rangle$  up to a correction which takes the form of a linear combination of the real and imaginary parts of  $X$ ,  $Y$ ,  $X'$ , and  $Y'$ . Where  $a, b, c, d$  are fixed-point, signed,  $m$ -bit binary numbers lying in the range  $[-1, 1)$ , the transformation  $a, b, c, d \mapsto a', b', c', d'$  is merely the transformation between signed-binary and offset-binary representations of  $2^m$  values. This can be trivially accomplished by inverting the MSB of each of  $a, b, c$ , and  $d$ . After such a transformation  $X'$  and  $Y'$  are fixed-point *unsigned*  $m$ -bit numbers whose real and imaginary components lie in the range




Figure 3.9: Schematic depicting 4 bit signed complex multiply implementation in an  $18 \times 18$  bit multiplier core. In this implementation, sign extraction of the inputs allows the real parts of the operands to be treated as unsigned values. The imaginary parts, which remain signed, are sign-extended to 7 bits before being passed to the multiplier. Compensation for sign extraction, as well as calculation of the MSB of the imaginary part of the multiplier output is performed in non-DSP logic. In this method, operands are required to fall in the symmetric range  $[-7,7]$ , such that the real parts of the two input values can both be represented as 3-bit, positive values with a corresponding sign bit. Reproduced (with minor corrections) from de Souza et al. (2007).

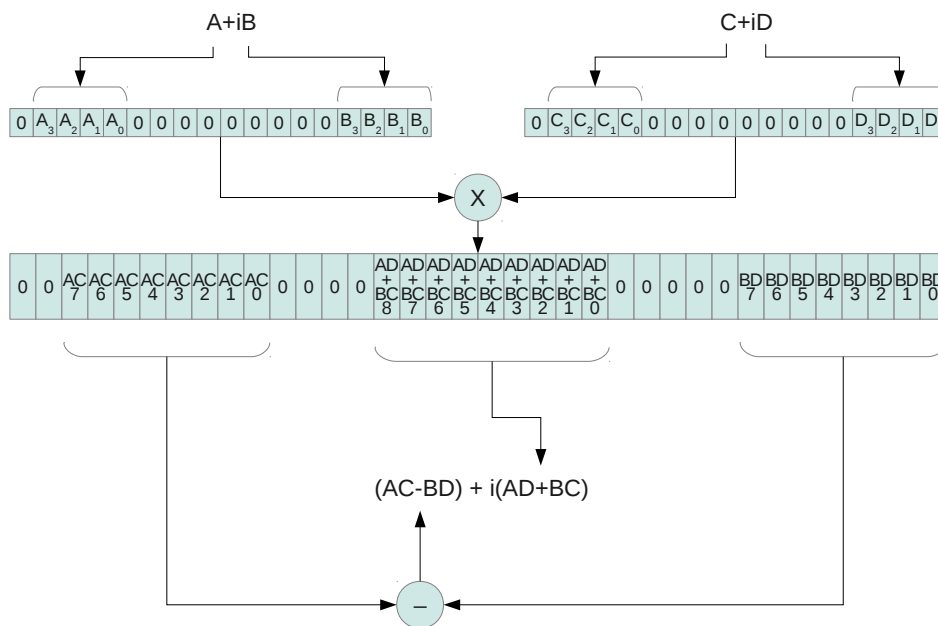


Figure 3.10: Schematic depicting a 4 bit unsigned complex multiply implementation in an  $18 \times 18$  bit multiplier core. This implementation requires that both real and imaginary components of each operand be unsigned, but directly computes the result of multiplication with no intervention of non-DSP logic (save the subtraction necessary to obtain the real part of the output).

$[0, 2)$ . Being able to re-factor the correlation so that the multiplication of signed numbers can be replaced by the multiplication of unsigned numbers offers significant implementation advantages, which we will now explore.

### 3.3.3 Implementation

Figure 3.10 shows how, using a single  $18 \times 18$ -bit multiplication, a 4-bit complex multiply of two unsigned complex numbers can be performed, with only a single additional subtraction operation required after the multiplication. Unlike the implementation shown in Figure 3.9, the output bit vector contains a minimum of 4 bits zero padding between each component of the multiplier output. This allows up to  $2^4$  multiplication results to be accumulated in the same single 18-bit word container without the risk of cross-contamination of components. The same method can also be used with  $5 \times 5$  bit multiplications, in which case  $2^2$  multiplication results can be accumulated without overflow.

The ability to sum multiple complex products without needing to first explicitly com-

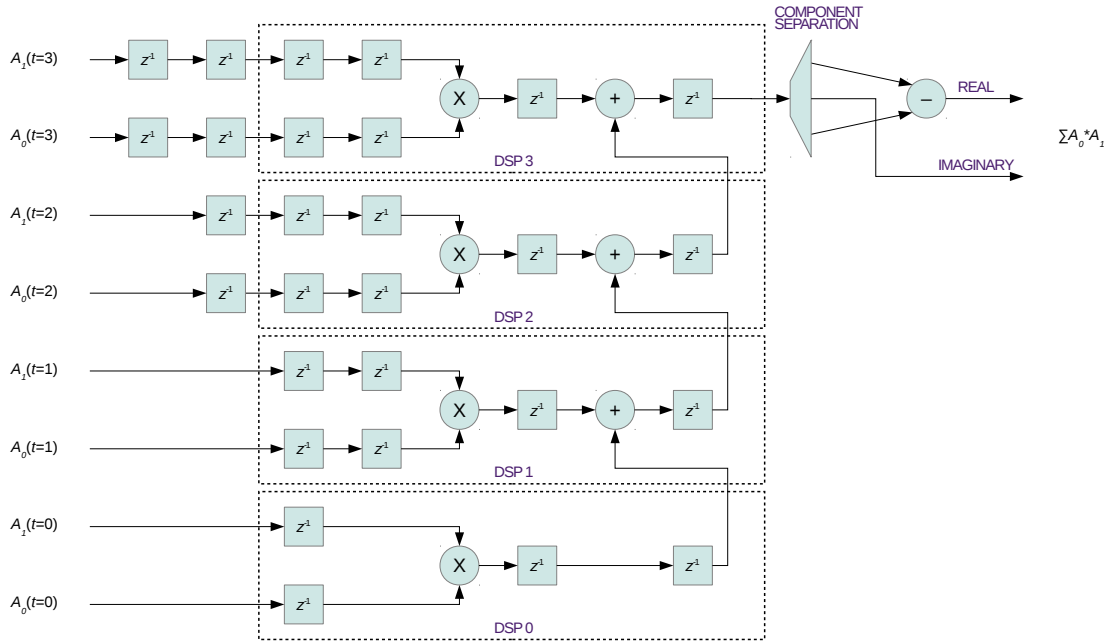


Figure 3.11: Schematic depicting the use of four DSP slices to calculate the sum of the correlation between two antennas,  $A_0$  and  $A_1$ , at four time samples. This implementation requires only a single subtractor to be implemented in non-DSP resources. Note, however, that pipelining of each different time stream is necessary to maintain correct alignment of different time samples. For a chain of fewer than three DSP slices, all pipelining requirements can be fulfilled by DSP resources. For a chain of three or more DSP slices, as depicted here, it is necessary to implement extra pipeline registering externally. Where inputs are 4-bit real, 4-bit imaginary unsigned numbers, and the DSP multipliers accept  $18 \times 18$  bit inputs, it is possible to chain together 16 DSP slices in a configuration similar to that shown without risk of overflow corrupting results.

pute and separate real and imaginary parts opens up new possibilities when we consider implementation of the multiplication algorithm in FPGA DSP slices such as that shown in Figure 3.5. Unsigned multiplication allows multiple DSP slices to be chained together, with the results of the parallel multiplications added in DSP resources. Such a configuration is shown schematically in Figure 3.11, for the example case of four grouped multiplications, in which each multiplication calculates the correlation between different time samples of a pair of antennas.

The advantages of using chained DSP slices in this way are multifold. First, each addition inside a DSP saves resources in the main FPGA fabric. Second, chaining DSP slices means that only a single FPGA-fabric subtractor is required to generate the imaginary (or real, if one complex input is conjugated during multiplication) component of the complex multiply.

As well as these advantages of DSP chaining, Figure 3.11 also demonstrates some undesirable side-effects of such a data configuration. Pipeline registers, which are required to maintain correct alignment of data streams are liable to spill outside DSP slices when chains

greater than two long are created. It should also be remembered that the architecture is only possible at all because of the unsigned nature of the inputs. The correlation of these inputs may only represent the correlation of corresponding signed inputs if the relationship in Equation 3.3 is properly taken into account.

Both of these potential issues can be mitigated when chains of unsigned multipliers are implemented in a windowed X-engine architecture such as that used by CASPER. Pipelining the parallel data streams, though necessary, need only be carried out once, before data flows through the series of baseline taps. Correction of the unsigned results can be accomplished relatively cheaply by only correcting the  $O(N^2)$  visibilities after accumulation. To achieve this, the  $O(N)$  real and imaginary components— $a'$ ,  $b$ ,  $c'$  and  $d$  in Equation 3.3—must be summed alongside the visibilities. Relevant correction factors can then be calculated and subtracted at the end of the averaging process. The order of this added computational burden,  $\Delta$ , of this correction process is

$$O(\Delta) = O(N) + O(N^2/M), \quad (3.4)$$

where the first term comes from summing the input component parts, and the second term comes from correcting  $O(N^2)$  visibilities after integrating  $O(M)$  samples. In practice, to keep data rates at a manageable level, the input and output data rates of a correlator are often balanced, so that the minimum accumulation length  $M \approx N$ , in which case  $\Delta$  has  $O(N)$  complexity. In this case, the overhead may be neglected when compared to the  $O(N^2)$  complexity of visibility calculation.

### 3.3.4 Proposed X-engine Design

Figure 3.12 gives a schematic for an X-engine architecture based on the CASPER tap-chain architecture, but enabling correlations to be calculated by considering only unsigned inputs. The design proposed is capable of accepting multiple parallel signal samples simultaneously, allowing the correlator to be constructed using optimized DSP chains as shown in Figure 3.11. The inequality between the visibilities calculated with unsigned integers and the true correlations is compensated for by accumulating individual real and imaginary components of input signals, and applying the relevant correction to the output correlation matrices.

### 3.3.5 Design Verification

To avoid the platform lock-in associated with creating FPGA code using Xilinx System Generator, the optimized X-engine design presented in this chapter was implemented in Verilog. Wherever possible, manufacturer-specific cores were avoided in favour of vendor-independent behavioural code. In practice, this has meant that Xilinx-specific code is only

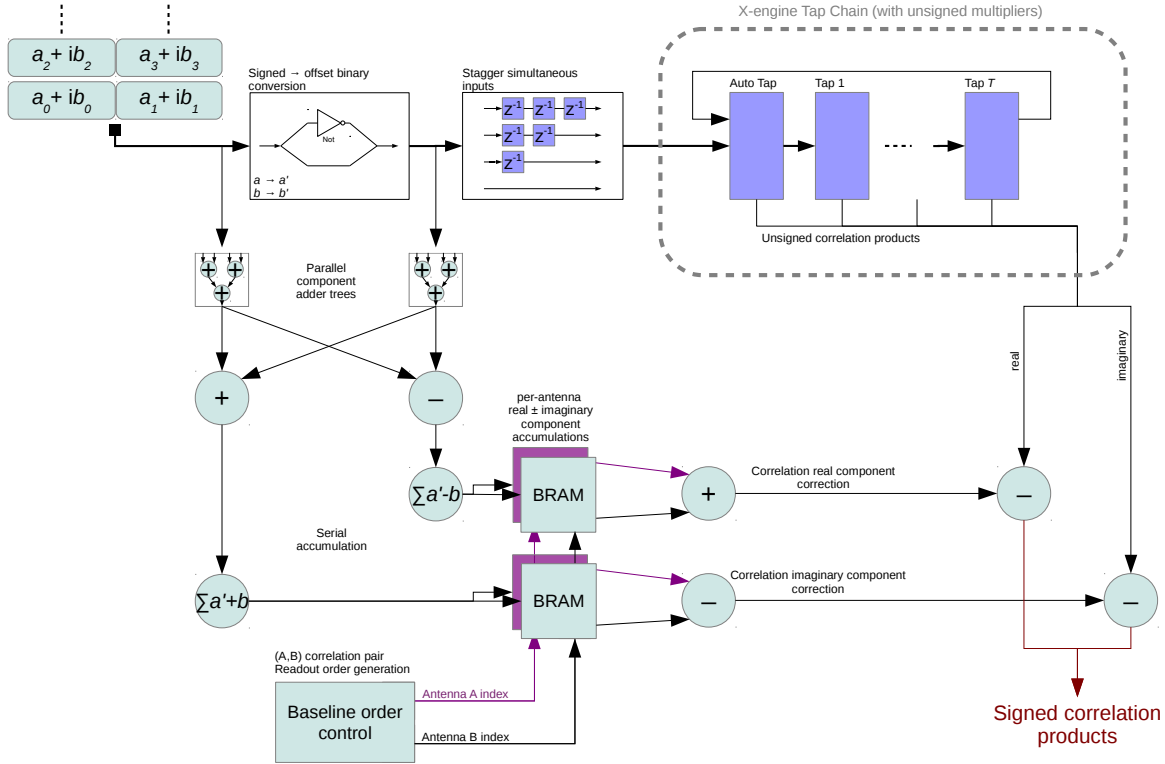


Figure 3.12: A top level block diagram of an X-engine built around unsigned multiplication. Input data order may be identical to the CASPER X-engine, though multiple samples to be accumulated may be input in parallel, as well as serially. Inputs to the X-engine are first converted to unsigned integer form, by inversion of the MSB of all parallel components, and then staggered in time, to compensate for delays imposed by downstream adding chains. Alongside the main correlation substructure, sums and differences of components are accumulated and stored for each antenna in the correlation system. This is marginally more efficient than merely summing components, as fewer large bitwidth additions are required to generate the correlation corrections post-accumulation. When accumulation of correlations is complete, the necessary corrections to the real and imaginary parts of this visibility are applied, making the output of this module identical to that of the CASPER X-engine. In this diagram only one polarization is shown, though the tap chain and surrounding infrastructure are fundamentally dual-polarized. This means input real/imaginary component summing is necessary for two polarizations, and correction of unsigned correlation products must occur for all four polarization cross products.

user for low-level implementation of the DSP slice chaining, which can be difficult to infer if not coded explicitly.

Various testing processes were used to verify the code that was produced. Initial “golden vector” simulation test benches were created and run using the free, open-source, Icarus Verilog<sup>7</sup>. Xilinx ISE Simulator (ISim<sup>8</sup>) was also used for further verification. Final hardware tests were carried out on a Xilinx Virtex-5 based “ROACH” board, by comparison of the optimized X-engine output with the original CASPER design, with both subject to identical pseudo-random noise inputs. In all cases, the outputs of the optimized X-engine have proved bitwise-identical to the original CASPER module, as expected.

Verilog code for the X-engine described here, as well as Simulink wrappers, are available via a public online repository<sup>9</sup>. More detailed implementation details than those presented in this chapter are given in Appendix C.

### 3.3.6 Quantitative Comparison of Resource Usage

Figure 3.13 shows the FPGA resource usage of an  $N = 64$ ,  $M = 256$  X-engine based on standard complex multiplication, multiplication as prescribed by de Souza et al. (2007), and unsigned multiplication as proposed in Section 3.3.2. The most obvious feature of such a comparison is the 75% reduction of DSP slices required to implement the cross-multiplication in the design, when using either of the multiplier optimizations presented in this chapter. Either implementation comes with a small increase in logic use, with either extra logic required to implement multiplier DSP spillover and downstream management of sign bits in the case of signed multiplies, or logic required to manage, and subtract offsets created by unsigned multiplies. In the latter case, a small amount of RAM is also required to store accumulations of individual complex components before subtracting them from correlator results.

More significantly, Figure 3.14 demonstrates the resources required when the bandwidth required to be processed is increased by factors of 2 and 4. In the case of the CASPER X-engine, such an increase in processing can currently only be achieved by creating 2 and 4 instances, respectively, of the module. In the case of the X-engine implementation presented as part of this thesis, increased processing bandwidth can be achieved by increasing the number of samples simultaneously input into the module, and using the efficient multiply-sum architecture depicted in Figure 3.11. Having a single module process increased bandwidth, rather than instantiating multiple modules, has huge advantages in logic use as infrastructure is not required to be duplicated. Furthermore, RAM use is much improved as (for

---

<sup>7</sup>See <http://iverilog.icarus.com>

<sup>8</sup>See <http://www.xilinx.com/tools/isim.htm>

<sup>9</sup>[https://github.com/oxfork/mlib\\_devel/tree/master/ox\\_library/hdl\\_lib](https://github.com/oxfork/mlib_devel/tree/master/ox_library/hdl_lib)




Figure 3.13: The resources required to implement an  $N = 64$ ,  $M = 256$ , windowed X-engine using the original CASPER design, original design with modified complex multiplication as prescribed in Section 3.3.1, and implementation using unsigned multiplication as prescribed in Section 3.3.2. Resources are given as a percentage of those available on the Xilinx SX95T FPGA, used in CASPER’s popular ROACH platform.

integration lengths less than 512 samples) tap delays can be implemented using 32-bit wide  $\times$  512 deep BRAM blocks, using dual-port functionality. RAM required by post-correlation descrambling is also reduced, as the buffers required to reorder the correlation matrix are not required to be duplicated.

### 3.4 Conclusions

A new X-engine implementation has been described here which obtains a significant reduction in FPGA resources when compared to the original CASPER design. In particular a 75% saving in required DSP slices. The design compares favourably with the alternative approach to minimizing DSP use advocated by de Souza et al. (2007). Furthermore, the added functionality of the new X-engine, which allows increased processing bandwidth by inputting multiple samples for parallel accumulation, can give rise to compelling improvements in non-DSP resource use.

In addition to the resource use improvements seen, I believe the implementation platform of the optimized X-engine should also be seen as a significant improvement over the




Figure 3.14: The resources required to implement a correlator processing twice (left) and four times (right) the bandwidth of a single X-engine. In the case of the original CASPER design, processing more bandwidth simply means creating multiple instances of identical X-engines, with a linear increase in resource use. The X-engine presented in this thesis not only immediately quarters DSP requirements, but also exhibits much better than linear scaling in logic and BRAMs use, owing to more efficient use of BRAMs where input data bus widths exceed 18 bits, and far less duplication of peripheral infrastructure required by the X-engine. Resources are given as a percentage of those available on the Xilinx SX95T FPGA, used in CASPER’s popular ROACH platform.

CASPER Simulink block. Coded almost entirely in behavioural Verilog, the module is much more readily portable between FPGA platforms and development environments than the original implementation, whilst a Simulink wrapper for the module has been created to allow seamless integration into the current CASPER toolflow. The new module, when used with the Simulink wrapper, represents a direct drop-in replacement for the original X-engine, potentially allowing existing CASPER correlator backends to be upgraded with no hardware costs and minimal firmware development time.

Lastly, I hope that the general-purpose Verilog submodules which have been created in the course of this work can be used to speed porting of other CASPER library blocks to some more flexible HDL. An open source library of astronomical processing modules that can be easily used by astronomers regardless of their chosen software and hardware development platforms will be a great benefit to the those working to construct new instruments.

## Chapter 4

# Implementation of a Direct-Imaging Telescope with BEST-2

We have seen in the previous chapter how low-level optimizations can significantly affect the efficiency with which the correlation stage of an FX correlator can be implemented. In this chapter, however, rather than considering optimizations to the classical FX architecture, we instead move to consider another type of processing mechanism altogether; the *Direct-Imaging Correlator*. Described in depth in Section 2.4.4, in this chapter we present an implementation of one of the simplest varieties of direct-imaging correlators, applicable to arrays of receivers positioned regularly on a two-dimensional grid. The requirement for such an array configuration is met by the Basic Element for SKA Training II (BEST-2), in Medicina, Italy. We begin by laying out the motivations for deploying such an instrument, before detailing the design and operation of the new system and in particular the digital backend, which was newly constructed as part of the work presented in this thesis. This backend facilitates both FX correlation and direct-imaging functionality, and also serves as a beamforming front-end for time-domain science applications.

### 4.1 Motivation

The equivalence of correlation matrices obtained by either cross-multiplication of antenna signals, or Fourier transform-based direct-imaging, was discussed at length in Chapter 2 (see, in particular, Figure 2.8). However, the quality of data obtained with a direct-imaging system need to be calibrated in real time. In Chapter 2 the effects of deficiencies in calibration of direct-imaging systems were assessed using simulation and theoretical calculation. In the construction and deployment of the instrument described in this chapter, we aim to directly address such issues, and compare the performance of a spatial-FFT-based imager to a traditional FX correlator, whilst experimenting with a variety of real-time and non-real-time calibration techniques. We also wish to demonstrate the use of the intermediate data




Figure 4.1: (a) A schematic of the Northern Cross Telescope, indicating the subset of eight reflectors of the N-S arm which form the BEST-2 array. (b) A photo of the BEST-2 array. The focal line of each cylindrical concentrator is critically sampled by 64 dipoles at 408MHz. Signals from these 64 dipoles are summed into 4 groups, which form the analogue feeds available for downstream processing.

products of the spatial FFT—beams on the sky—for time-domain science applications, and in particular, pulsar monitoring and detection.

## 4.2 Introducing BEST-2

The Basic Element for SKA Training II (BEST-2) is a subset of the Northern Cross cylindrical array (Figure 4.1), at the Medicina observatory in Italy. The array comprises eight East-West oriented cylindrical concentrators, each with 64 dipole receivers spaced such that the cylinder focal line is critically sampled at 408 MHz. Signals from the 64 dipoles are combined in groups of 16 using analogue circuitry, resulting in four analogue channels per cylinder. This results in a total of 32 effective receiving elements positioned regularly on a  $4 \times 8$  grid, shown in Figure 4.2a. The top-level specifications of the BEST-2 array, which provides  $\sim 1^\circ$  resolution over a  $\sim 38$  square degree field of view, are summarized in Table 4.1.

Such a receiver configuration leads to a set of highly redundant baselines, which allows efficient spatial processing with an FFT. In the case of BEST-2, the array possesses 496 possible antenna pairings (discounting pairings of an antenna with itself), forming the 52 unique baselines shown in Figure 4.2b.

### BEST-2 Specifications

Effective Area	1002m <sup>2</sup>
System Temperature	86K
Longest baseline (N-S)	70m
Longest baseline (E-W)	17m
RF band	400 – 416MHz
Total analogue channels	32
Primary Field of View (Dec, 408MHz)	5.7°
Primary Field of View (RA, 408MHz)	6.6°
Synthesized beam size (Dec, 408MHz, at zenith)	0.52°
Synthesized beam size (RA, 408MHz)	1.73°

Table 4.1: The top-level specifications of the BEST-2 array (Bartolini, 2010).




Figure 4.2: The relative locations of the 32 effective receiving elements of BEST-2, indicated by crosses, lying on a  $4 \times 8$  rectangular grid (a). This configuration of receivers leads to a highly redundant array with 496 non-zero baselines, only 52 of which are unique (b). Figure credit: G. Foster.

### 4.3 System Requirements

The desire to design a BEST-2 backend suitable for comparing direct and interferometric image forming results in a number of top-level requirements. First, the backend should be capable of performing simultaneous and effective imaging by both methods. Such an ability allows most effective comparison of data sets, and also allows potentially different experiments with both systems to be conducted independently, should this ability be required.

Second, the system should be capable of outputting one or more sky-beams at full data rate—i.e. with no power accumulation—to downstream processors. This ability allows experimentation with processing systems suitable for examining transient cosmic events, in particular using algorithms implemented on graphical processing units (GPUs) such as those developed by Magro et al. (2011); Armour et al. (2012); Barsdell et al. (2012). Thirdly, the system developed should not require modification of the existing BEST-2 analogue frontend, and should be suitable for implementation mainly on hardware already available at the Medicina Observatory. In this way deployment difficulty and cost, as well as interference with telescope use, are minimized.

### 4.4 System Architecture

In this section we describe the system developed as part of this thesis, which meets the requirements set out above.

#### 4.4.1 Analogue Frontend

The BEST-2 analogue frontend is extensively described in Perini (2009); Perini et al. (2009); Montebugnoli et al. (2009). The key components of the system are an initial 16-dipole analogue combiner, used to reduce 64 dipole receivers on each cylinder focal line to 4 analogue channels; RF-over-optical-fibre links from the telescope front-end to a temperature-controlled receiver room; and a mixing system designed to convert the 408 MHz RF sky signal to a 16 MHz IF, centred at 30 MHz. The chain also features per-channel, digitally controlled attenuators allowing control of analogue signal equalization. A high-level schematic of the BEST-2 analogue chain is shown in Figure 4.3.

#### 4.4.2 Digital Back-End

The CASPER-designed Reconfigurable Open-Architecture Computing Hardware (ROACH) board, pictured in Figure 4.5, is a digital signal processing platform based around a Xilinx Virtex 5 SX95T FPGA, enhanced by a number of peripheral elements. The components of a ROACH board are shown in block diagram schematic form in Figure 4.4. Peripherals

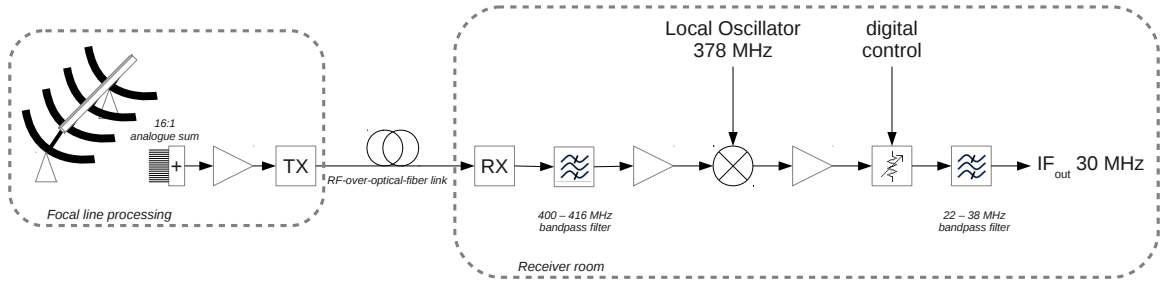


Figure 4.3: The BEST-2 analogue front end, featuring RF-over-fiber links from the focal line processing modules to a temperature-controlled receiver room, where down-mixing, digitization and signal processing takes place. Four such chains form the frontend for each of the four BEST-2 cylinders. Schematic adapted from Montebugnoli et al. (2009).

include a single DDR RAM and pair of QDR memory modules, and a PowerPC CPU. This CPU facilitates user-interaction with shared memory elements of FPGA firmware via a standard 10/100 Mb/s Ethernet connection. The ROACH platform also includes a number of general purpose interfaces, including two 40-differential-pair Z-DOK connectors, enabling connection of a variety of ADC daughterboards. Several ROACH boards, as well as other CASPER-designed hardware, are already used in the Medicina observatory for a variety of digital backends, and can be reprogrammed to perform the roles required by the system presented here.

The newly designed BEST-2 digital backend is based on three ROACH processing platforms, and a number of software-based processing nodes, connected in the topology shown in Figure 4.6. An overview of the processing performed on each of the three ROACH boards is described in detail in Sections 4.4.3, 4.4.4 and 4.4.5, respectively. These processing nodes will henceforth be referred to as the “F” (for frequency transform), “S” (for spatial transform), and “X” (for cross-multiplication) engines, indicating their roles in the complete system. The functionality of each is described in detail in the remainder of this section. F-engine firmware was designed jointly by myself and Griffin Foster during the course of this work. S-engine firmware, and associated software processing is solely my work, whilst the X-Engine used in the backend presented here is solely the work of Griffin Foster, and is discussed in Chapter 3 of his doctoral thesis (Foster, 2013).

Data is passed between ROACH boards using the Ten Gigabit Attachment Unit Interface (XAUI) over a copper cable connection. XAUI is a lightweight, bidirectional, point-to-point transmission protocol implemented over four receive and transmit lanes, each operating at 3.125 Gb/s with 8b/10b data encoding. Downstream processing of time-domain beam data is accomplished using a GNU/Linux-based server, which hosts an Nvidia Graphical Processing Unit. Recording of data to hard disks and telescope control and monitoring is

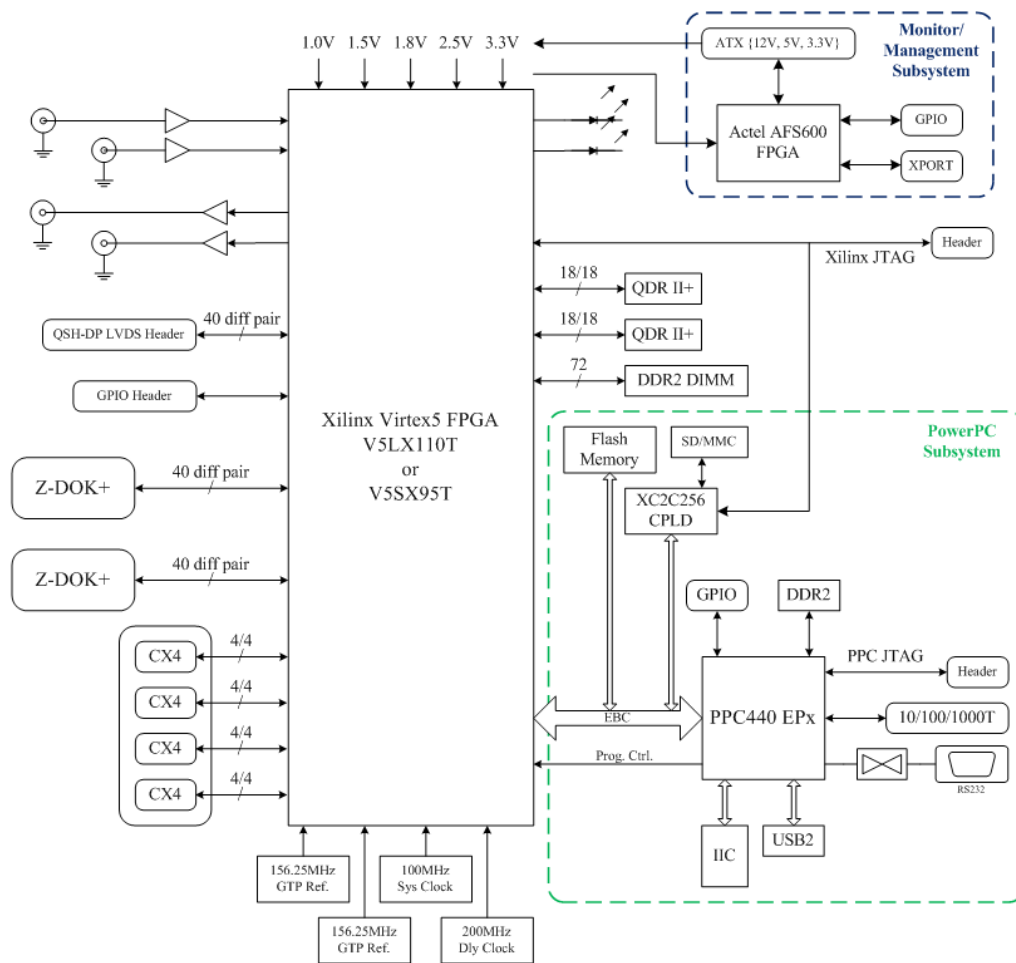


Figure 4.4: Top level diagram of the Xilinx SX95T-based ROACH board. The ROACH features dual 2Mb QDR memory chips, four CX4 10Gb/s interfaces, a single FPGA-accessible DDR2 DIMM socket and dual Z-DOK expansion ports, allowing interfacing to a variety of ADC daughter cards. The PowerPC CPU, running Linux, allows software interaction with shared memory elements on the ROACH, and facilitates remote configuration via Ethernet. Figure credit: CASPER Collaboration.




Figure 4.5: The ROACH board, housed in a 1.3U ATX case, and powered with a standard ATX power supply. On the right of the board are the twin, 40-pair Z-DOK connectors, which facilitate interfacing with a variety of ADC cards, including the 64 input ADC board used in the BEST-2 backend.

performed by standard GNU/Linux-based servers, connected using an off-the-shelf Ethernet network, with both 100 Mb and 10 Gb components.

#### 4.4.3 “F” Engine Firmware

The F-Engine FPGA-based processing node is responsible for digitization, channelization and transmission of coarsely quantized antenna signals to downstream processing boards. In this section we describe the components of the firmware designed for the F-Engine processor, forming the system depicted in Figure 4.7.

##### ADC Interface

Central to the simplicity of the BEST-2 digital backend is the ability to digitize all 32 analogue streams from BEST-2 on a single processing node. This avoids the need to synchronize multiple boards, and the requirement for any complicated cabling systems. To this end, the signals are digitized with the 64ADCx64-12 board developed by Rick Raffanti<sup>1</sup>. This board uses eight Texas Instruments ADS5272 8-channel, 12-bit ADC chips<sup>2</sup> to digitize up to 64 single-ended analogue signals at up to 65 MSa/s. It is capable of interfacing to a single

<sup>1</sup>See <https://casper.berkeley.edu/wiki/64ADCx64-12>

<sup>2</sup>See <http://www.ti.com/product/ads5272>

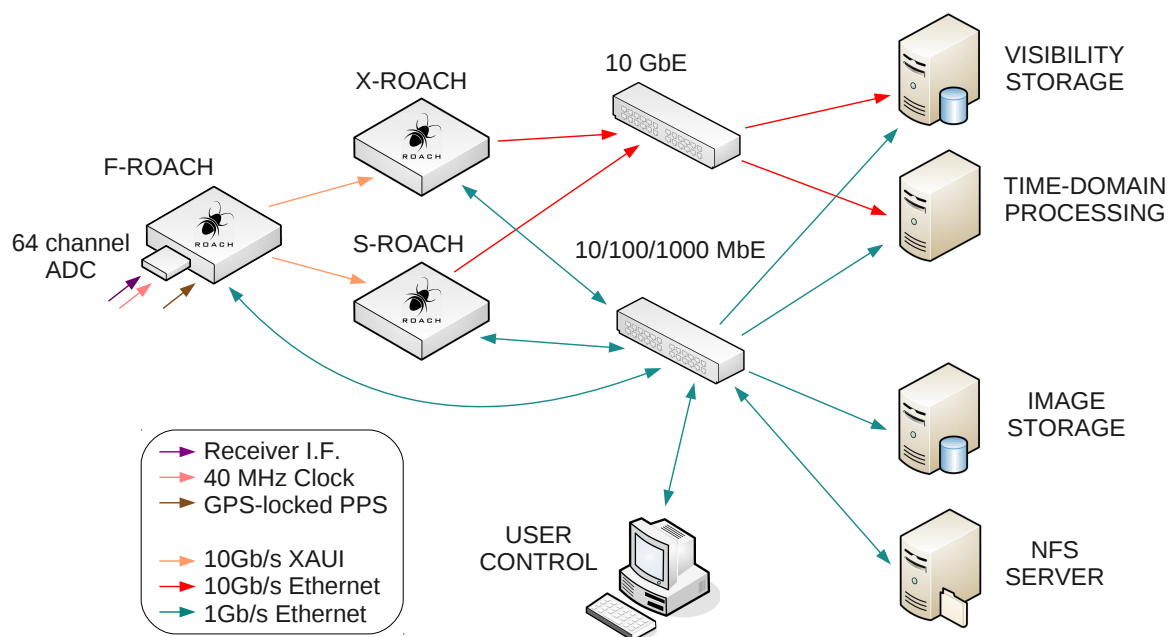


Figure 4.6: The top level topology of the newly designed BEST-2 backend, which uses both 1Gb and 10Gb Ethernet for data transport to CPU and GPU nodes. 10Gb/s XAUI is used for ROACH-ROACH connections. “F”, “X” and “S” ROACH boards are respectively responsible for channelizing, cross correlating and spatially Fourier transforming signals from the BEST-2 array. In practice, the roles performed by any of the five CPU-based nodes need not be split amongst physically distinct hosts.

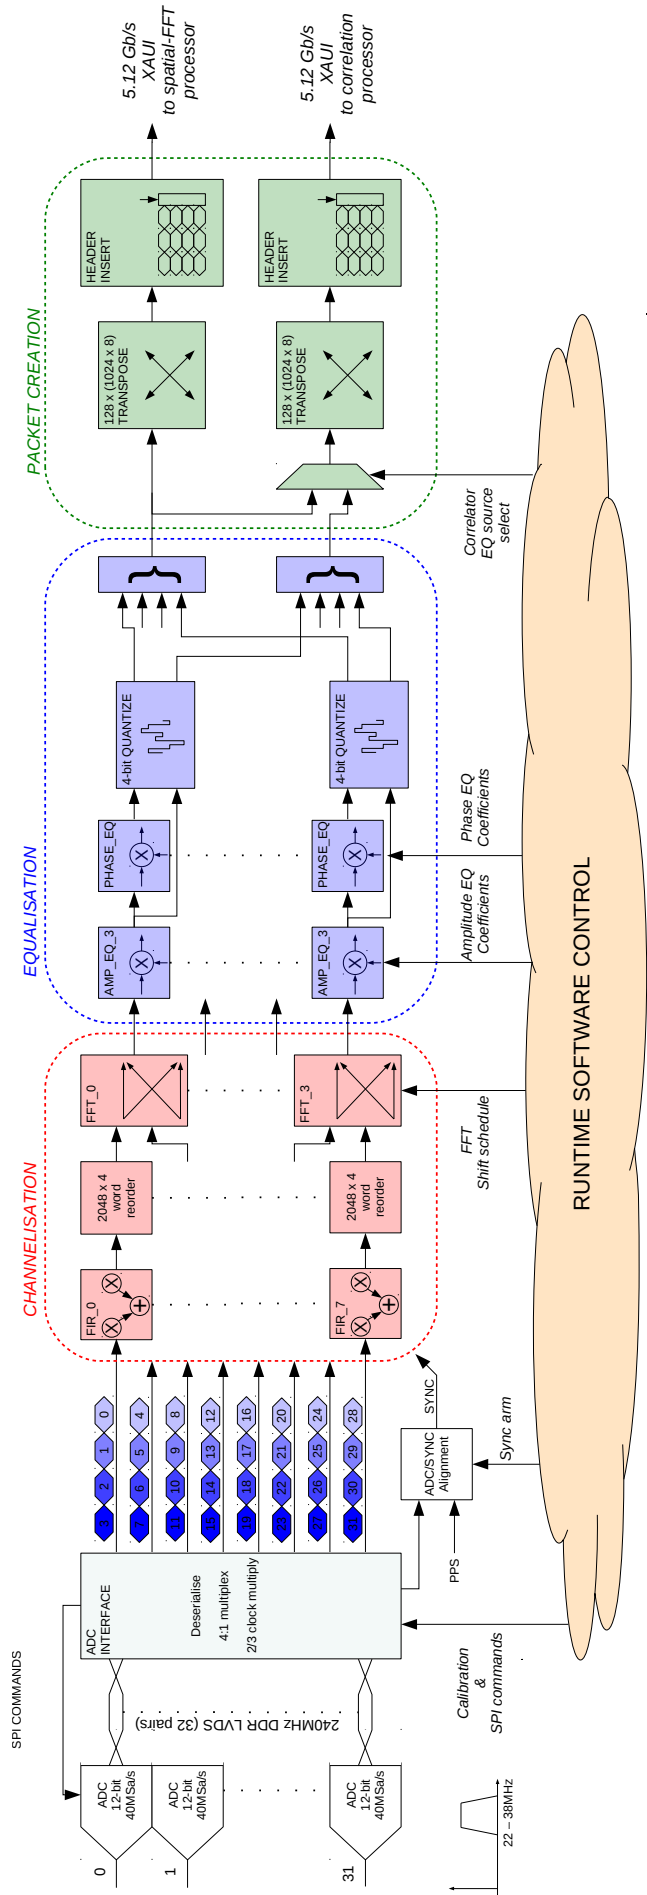


Figure 4.7: A block diagram describing the processing carried out on the "F-engine" ROACH board, which is used to channelize ADC data from an attached daughterboard. Following channelisation, software-defined amplitude and phase corrections are applied to the frequency-domain data, which is then quantized to 4-bit precision and output over XAUI for further processing.

ROACH board using a dual Z-DOK interface. The ADS5272 chip is suitable for the 20–40 MHz IF of BEST-2, when sampling in the second Nyquist zone at 40 MSa/s.

Whilst the 64ADCx64-12 is physically compatible with the ROACH board, using the ADC card in the CASPER toolflow requires that an interface be constructed, and integrated into the MATLAB Simulink System-Generator Environment (MSSGE). To this end, an interface, shown schematically in Figure 4.8, was designed in Verilog, and integrated into the CASPER toolflow as a “yellow block”. In Simulink, “yellow blocks” are blocks which can be used to interface to FPGA peripherals, such as ADCs, external memory, and Ethernet interfaces. Such blocks abstract implementation details of the underlying interfaces and controllers from the end user, details which are only integrated into a design at compile-time.

In the case of the 64ADCx64-12, the controller designed interfaces the eight separate 12-bit, 8-channel Texas Instruments ADS5272 ADC chips of the ADC card with the ROACH’s FPGA, which receives 64 digitized channels over 64 pairs of DDR LVDS signals. These are grouped in sets of four, and presented at the interface output as 16 parallel streams, each carrying four multiplexed channels at a rate of four times the ADC sampling clock. A single clock from one of the ADC chips is used to generate all required FPGA clocks. Synchronization of the eight chips is achieved using Xilinx IODELAY instances controlled via a runtime software interface. An SPI controller is also implemented, allowing control of programmable ADC functions. A Simulink yellow block abstraction has been created to allow use of the Simulink GUI to select top-level parameters such as ADC clock speed and SPI interface requirements. Timing constraints for the FPGA design, as well as external pin connections are then transparently configured by the toolflow. Having been integrated into the CASPER toolflow, a number of radio astronomy groups outside Oxford are using the 64ADCx64-12 in their own research projects.

## **Channelization**

After digitization, signals need to be channelized to enable radio-frequency interference (RFI) to be isolated and to facilitate further data processing in the frequency domain. Channels need to be well isolated, and the channelization process is not well-served by a simple Fourier transform. Instead, a polyphase filter bank (PFB) channelization scheme, originally proposed by Bellanger et al. (1976), is frequently adopted. PFBs allow creation of a set of identical, regularly-spaced, filters, by combination of a Finite-Impulse-Response (FIR) filter and a discrete Fourier Transform. PFBs enable channelizers to be constructed with far greater channel isolation than “pure” Fourier transform-based filters, whilst still allowing the efficiency afforded by the Fast Fourier Transform implementation of the discrete Fourier

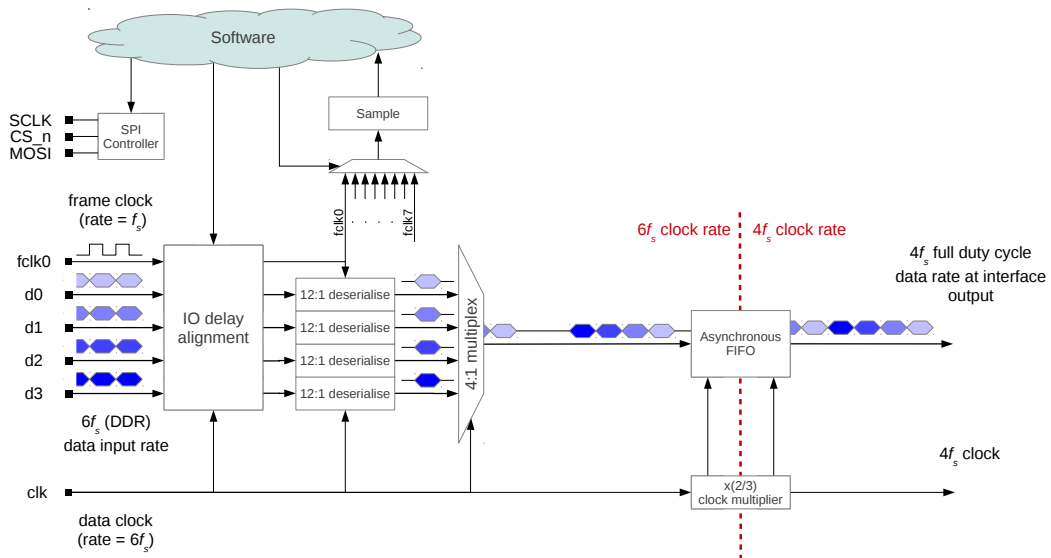


Figure 4.8: The FPGA-side interface for the 64 channel ADC card (hardware designed by Rick Raffanti). Eight ADC chips present 64 digitized channels to the FPGA in 64 DDR LVDS streams. These are grouped in sets of four, and presented at the interface output as 16 parallel streams, each carrying four multiplexed channels at a rate of four times the ADC sampling clock,  $f_s$ . The signal path making up one such stream of four channels is depicted in the schematic above. IO delays are computed by sampling and aligning the frame clocks associated with each ADC chip. An FPGA software interface controls these delays, and also allows writing to the programmable registers of the ADC chips via a lightweight, uni-directional Serial Peripheral Interface (SPI).

Transform (Cooley & Tukey, 1965) to be exploited. Harris & Haines (2011) provide a succinct and relevant overview of PFB implementations frequently encountered in the field of radio astronomy. Here we simply state that, with a suitably chosen set of coefficients,  $h_n$ , and a discretely sampled input data stream,  $x_n$ , a  $C$ -channel polyphase filter bank is constructed by taking the  $C$ -point Fourier Transform of the quantity

$$x'_i = \sum_{p=0}^{P-1} h_{Cp+i} x_{Cp+i}. \quad (4.1)$$

This is a  $P$ -tap FIR filter, with taps spaced at intervals of  $C$  time samples.

In the system presented here, channelization is achieved by means of a  $C = 2048$  point real-input FFT, with a 4-tap, Hann-windowed FIR filter front-end. This generates 1024 unique, complex-valued frequency channels in the digitized band, corresponding to an RF range of 398–418 MHz, and a channel resolution of 19.5 kHz.

The set of Hann-windowed filter coefficients are given by

$$h_n = \frac{1}{2} \left( 1 + \cos \left( \frac{2\pi n}{PC} \right) \right) \times \text{sinc} \left( \frac{\pi n}{C} \right) \quad n = \left[ -\frac{PC}{2}, \frac{PC}{2} - 1 \right], \quad (4.2)$$

and are implemented in a FIR filter on the time-multiplexed, streaming ADC data using the implementation shown in Figure 4.9. The implementation chosen is based on that provided by the CASPER collaboration libraries (Parsons et al., 2006), with added parameterization to allow time-multiplexed inputs. Since four ADC channels are multiplexed onto a single data stream, delays between taps (implemented in BRAMs) are  $4C$  samples long with coefficients changed every four clock-cycles. Eight such filters are instantiated simultaneously, to process all 32 ADC channels in eight parallel streams. In order to more efficiently utilize RAM, whilst not encountering timing issues associated with high signal fanouts, 4 sets of coefficients are shared between the 8 filters.

Following the FIR filter, a  $4 \times 2048$  word reordering of each data stream occurs, to group 2048 sequential time samples from each ADC channel. This allows the channelization process to be completed with a standard implementation of a 2048 point streaming radix-2 FFT, one of the centre-pieces of the CASPER library of DSP processing modules (Parsons et al., 2006). The response of the resulting filter, along with that obtained by a simple Fourier Transform, is shown in Figure 4.10.

The input data stream into the FFT module is real-valued, giving rise to spectra of 2048 channels, of which 1024 are independent, and 1024 are conjugated copies. Efficiency is maximized by using complex FFT cores to calculate two transforms simultaneously, with only the 1024 independent frequency channels generated.

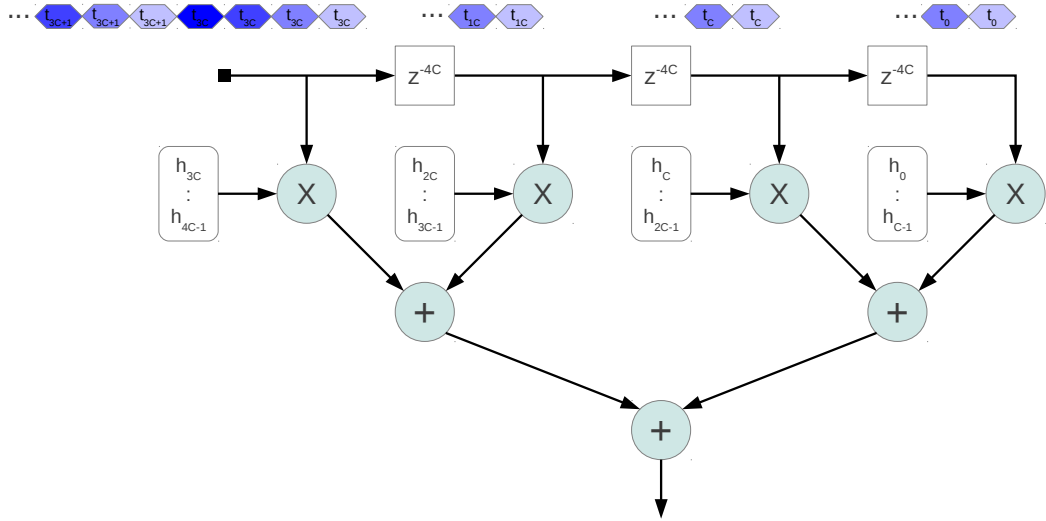


Figure 4.9: The FIR architecture used to implement 4 tap polyphase filter banks on 4 time-multiplexed data streams. Coefficients are stored in BRAM-based ROMs, and are incremented every four clock cycles.




Figure 4.10: The response of an FFT and 4-tap polyphase filter bank (based on a Hann window) as a function of input frequency,  $f$ , relative to the filter centre,  $f_0$ , for channels spaced in increments of  $\Delta_f$ . For a purely FFT-based filter the first sidelobes are at  $\sim 13$ dB, and isolation improves only slowly away from the filter centre. In contrast, the PFB-based filter has a first sidelobe at  $\sim -44$ dB with a response quickly falling to  $< -100$ dB.

A  $2^m$  channel radix-2 FFT comprises  $m$  “butterfly” stages, with two complex numbers summed at each stage. In the fixed point computation performed on an FPGA, the width of the data-bus is 18-bits, and is maintained throughout the transform. This bitwidth is chosen so as to maximize the efficiency with which data can be buffered in FPGA block RAM. Since the bus width is not allowed to grow through each summation, scaling of data by a factor of  $\frac{1}{2}$  at each butterfly stage is necessary to guarantee against potential overflow. In practice, input signals rarely occupy all 18 available bits. In any case, since the signals are largely noise, additions at each butterfly stage are mostly incoherent and scaling need not necessarily be performed at every stage. The scaling configuration, or “shift schedule” (so-called because a  $\frac{1}{2}$  scaling is equivalent to a single bit shift towards the LSB), in the Medicina channelizer can be set at runtime via software control, which also monitors the FFT for overflow events.

### **Equalization and Calibration**

Before coarse quantization of subband data streams can occur, care must first be taken to ensure that the quantization levels chosen optimally represent the signals of interest, in line with the values give in Table 2.1. In the case of the BEST-2 backend data is sampled with 12-bit resolution, channelized, and then quantized using a 15-level, 4-bit scheme. For optimal operation, this requires setting signal levels such that one LSB represents approximately  $0.36\sigma$ , for an input noise signal of variance  $\sigma^2$ . In general, the characteristics of a telescope analogue front-end mean that a compensation of the instrumental bandpass must be applied to data to ensure optimal quantization of all subbands. In the BEST-2 backend, this equalization is achieved by pre-multiplying 18-bit data streams by a set of 18-bit coefficients, with 512 independent coefficients available to represent the bandpass of each analogue channel.

Since calibration of each data stream is also required for downstream processing using a spatial FFT, gain calibration may also be incorporated into the equalization stage, by appropriate modification of coefficients. Phase calibration of signals is carried out in a dedicated post-equalization, multiplication stage using 18-bit complex coefficients. The separation of amplitude and phase calibration into real and complex multiplication stages allows precision of phase coefficients to be independent of amplitude calibration, and also provides the opportunity to feed the downstream S-Engine with phase calibrated data, without modifying the phase of X-Engine inputs. The ability to select between phase calibrated and uncalibrated X-engine inputs has been a useful diagnostic tool during development and deployment of the BEST-2 system.

## Inter-Board Communication

In order to accumulate visibilities and images in downstream processing, it is useful to order data such that multiple time samples from a single frequency channel arrive consecutively. This allows short-term integration of visibilities or beam data to occur, reducing the output data rate of the spatial processing systems. Long-term accumulation can then occur in a single, large block of memory with only limited IO requirements. This role can be filled by either the QDR or DDR peripheral memory elements provided by the ROACH platform. In the BEST-2 design, QDR is used, as it is of sufficiently large capacity, and has a significantly simpler interface.

In the BEST-2 design, the data structures processed by the S-engine and X-engine processors are different. The X-engine, which is based on the standard CASPER windowed design, requires windows of data as shown in Figure 3.1. The S-engine, described in Section 4.4.4, processes windows of the form shown in Figure 4.11. Whilst the option exists to transmit data to both processors in a common format, and re-order downstream, the choice was made in the BEST-2 design to transmit data to each processor in its native processing order. This decision allows a common packet receiver module to be used in both S and X-engines, but incurs the cost of two independent reordering systems in the F-engine before transmission may occur. Such an overhead may be prohibitive in some systems, but for the BEST-2 backend is easily accommodated, and in any case is required in order to enable S and X engine data to be sourced from different stages in the F-engine equalization pipeline.

The CASPER collaboration advocates the use of 10 GbE switched systems to enable correlators to be constructed with multiple F-engine ROACH boards. In BEST-2 we comply with the packet header formats required to decode data from these systems, but implement the F-X and F-S communication as lightweight point-to-point XAUI links. In such a configuration packet order is maintained, allowing extremely lightweight receiving modules to be used.

### 4.4.4 “S” Engine Firmware

The S-Engine FPGA-based processing node is responsible for formation of electric-field and total-power beams on the sky by spatial Fourier transform, using the algorithm prescribed in Section 2.4.4. In this section we describe the components of the firmware designed for the S-Engine processor, forming the system depicted in Figure 4.12.

#### Spatial FFT

Signals are input into the spatial FFT in the S-Engine in the format shown in Figure 4.11. The input data bus width is 32 bits, corresponding to 4 antenna samples on each FPGA

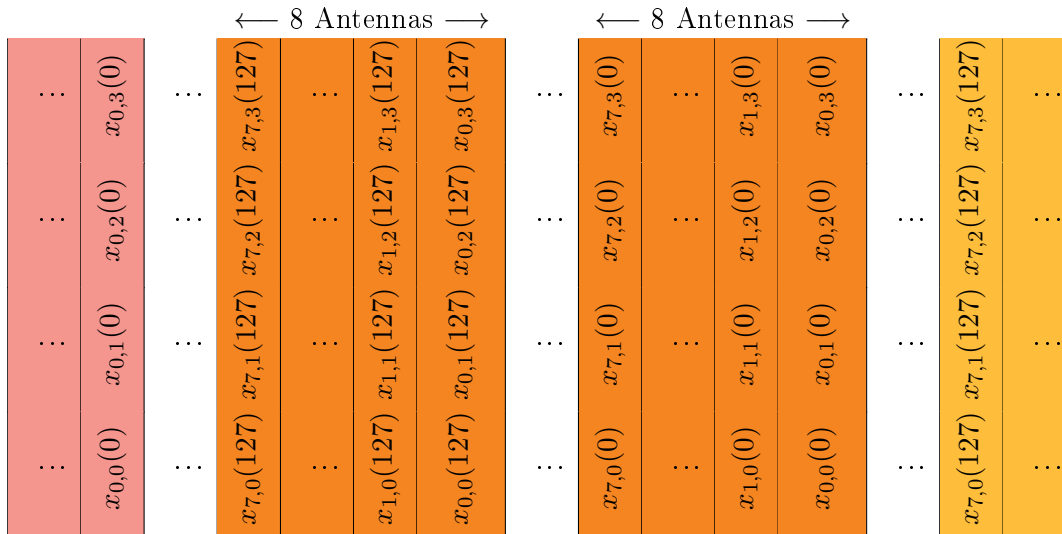


Figure 4.11: Data is sent to the S-Engine in a series of windows. An example of the organization of data in such a window is shown above, where  $x_{i,j}(s)$  represents the  $s^{\text{th}}$  time sample from the antenna located at point  $(i, j)$  on the grid of antenna locations. Each packet sent comprises 128 words, with 8 packets forming a complete window of 128 time samples from a single frequency channel of all 32 antennas. Consecutive windows of data cycle through the series of 1024 of different frequency channels.

clock. Over 8 consecutive clock cycles, a complete set of 32 antenna signals are received for a single time sample and frequency channel. This data order allows the spatial FFT in the S-Engine to be performed with two entirely parallel butterfly operations, with no serial stages, separated by a data transpose (see Figure 4.12). At each FFT, zero padding is added to the data vector as prescribed by the imaging algorithm. Consequently, data rate doubles following each FFT, and the intermediate transpose takes the form of an  $8 \times 8$  word corner turn operation. Such a corner turn can be achieved with the multiplexer-based configuration proposed by Urry (2000), and is implemented using a standard CASPER library module. Over a period of  $128 \times 8 = 1024$  clock cycles, a set of 128 time samples for a single channel of the 32-antenna array are processed. This represents the fundamental data chunk or “window” size. Consecutive windows arriving at the S-Engine contain consecutive frequency channels, with 1024 windows required to process the full digitized bandwidth of the BEST-2 array.

To avoid overflow in the FFT computation stages, data are shifted at every butterfly stage, and computations take place on an 18-bit data bus.

### Time-domain beams

After 128 beams have been created by process of spatial Fourier transform, it is possible to arbitrarily select 8 of these beams for output at 16-bit precision over a 10 Gb Ethernet

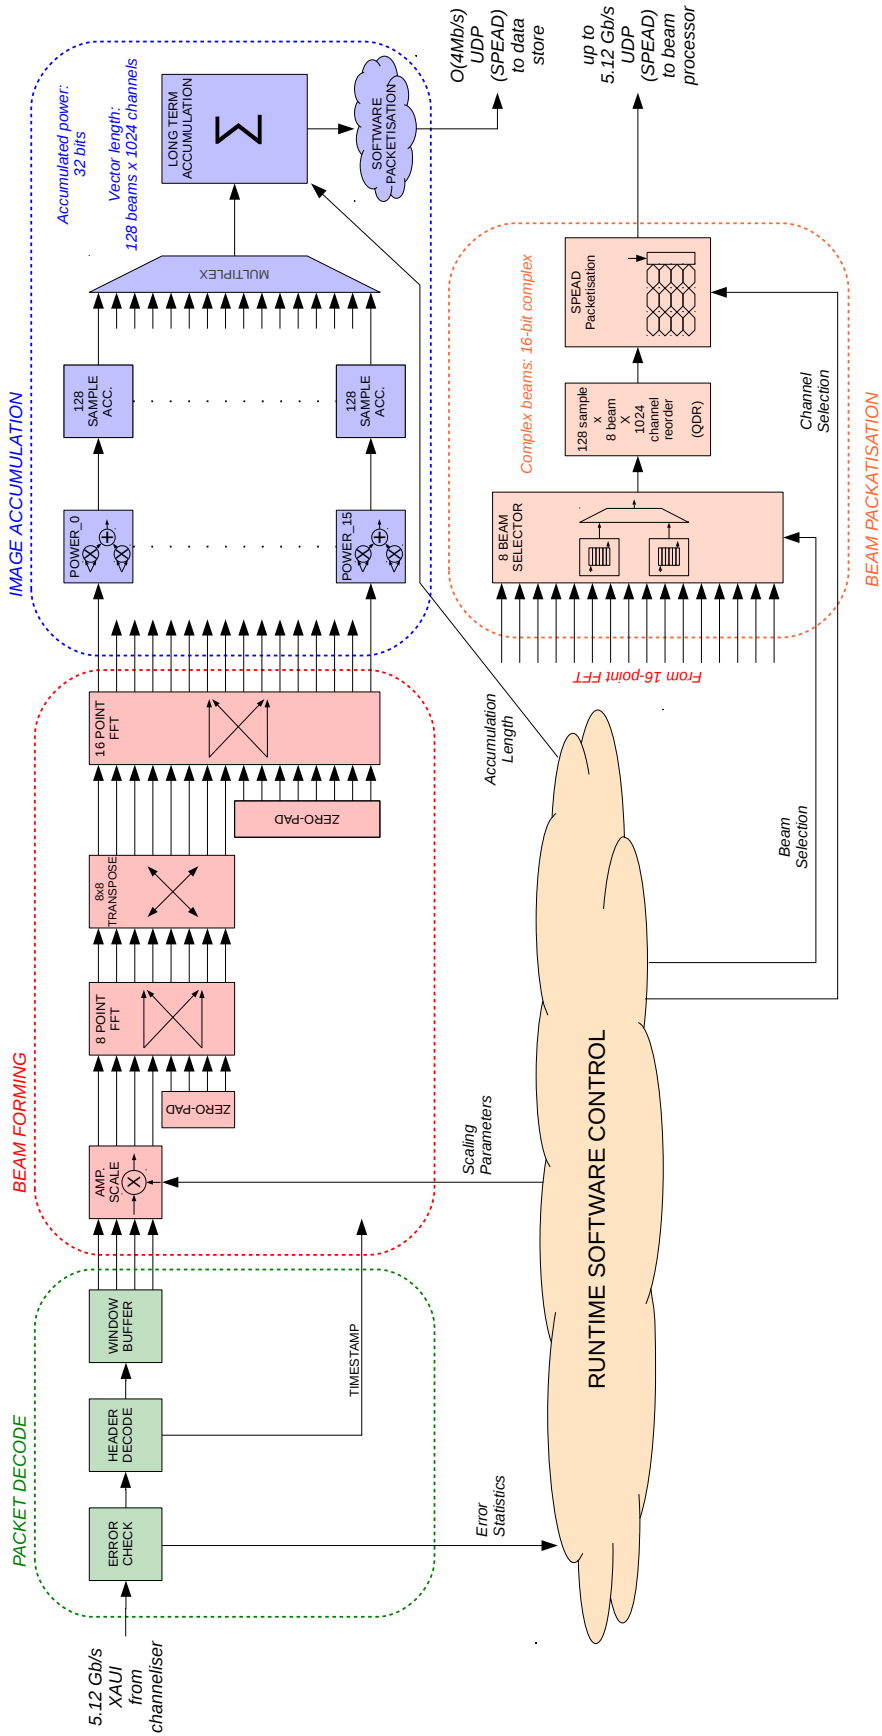


Figure 4.12: A block diagram describing the processing carried out on the “S-engine” ROACH board, which is used to perform a zero-padded spatial FFT on an incoming data stream, forming a set of beams on the sky. Following beamforming, accumulation of beam powers is performed to form low time-resolution sky images. A runtime-programmable set of eight beams, with user selectable frequency range are output via a 10GbE network, for further processing of high time-resolution data.

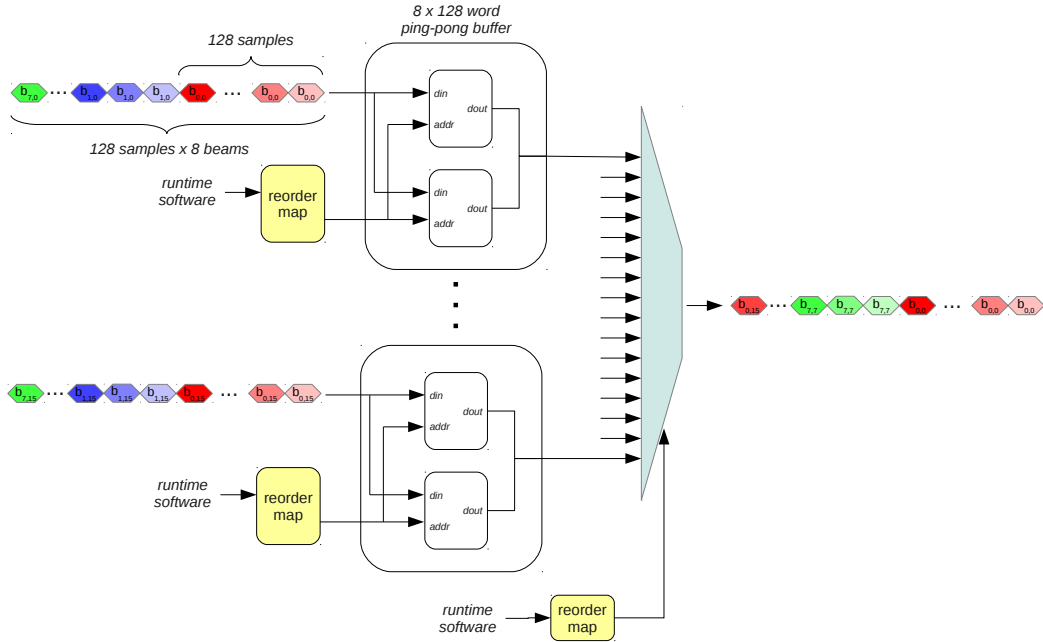


Figure 4.13: Spatial FFT output data are 16 parallel streams of 128 samples  $\times$  8 beams, cycling through 1024 frequency channels. Runtime-programmable double-buffered reordering, followed by 16 way multiplexing, is used to generate a single stream consisting of any 8 beams.

interface. Runtime-programmable reordering and multiplexing (shown in Figure 4.13) allows selection of these 8 beams from the 128 available as 16 parallel FFT output streams.

In order to maximize Ethernet data throughput, it is desirable to generate Ethernet frames with large packet sizes. In order to achieve this, whilst maintaining the possibility of parallelizing the processing of different beams amongst different host computers, packets are generated which contain data from a single beam, but multiple time samples and channels. For a data word size,  $W$ , in bytes, a packet containing  $T$  samples and  $C$  channels has payload size  $P$  bytes:

$$P = W \times T \times C \quad (4.3)$$

In the BEST-2 backend,  $W = 4$ , and  $T = 128$  is chosen to conveniently match the accumulation length of the imager.  $C = 8$  is chosen to achieve packets with payloads of 4096 kB. These are generated with a simple data transpose in QDR memory space, and headers are inserted into the data stream in compliance with the Streaming Protocol for Exchanging Astronomical Data (SPEAD<sup>3</sup>) transmission protocol in a custom-built packetizing module. SPEAD (Manley et al., 2012) is a flexible, self-describing, lightweight, application-level

<sup>3</sup>See <https://github.com/ska-sa/PySPEAD> for an open source Python-based implementation of the SPEAD protocol.

datastream format. It is designed to provide a standard output format for radio astronomy instruments outputting UDP streams. Masking of the transposed data stream also allows sub-selections of the 1024 generated frequency channels to be transmitted, reducing the transmitted data rate if desired.

The packetizing system implemented emphasizes simplicity. Whilst it exhibits flexibility when choosing subsets of generated beams and frequency channels for downstream processing, it does not provide the facility for trading number of beams against transmitted bandwidth in a completely user-defined way. In particular, a system which allows arbitrary numbers of beams and channels to be transmitted, up to the saturation point of the 10 Gb Ethernet link would represent a significant improvement in the design. Such an upgrade is under consideration, pending results from the existing time-domain backend.

### **Image accumulation**

Following generation of 128 complex-valued beams, power is calculated and 128 samples of each beam are accumulated in 16 parallel, 8-entry, distributed-RAM vector accumulators. This allows multiplexing of accumulated images into a single 32-bit wide 131072-entry vector accumulator, implemented in peripheral QDR memory. Long-term accumulation is performed in this memory for a runtime-configurable number of samples, after which data are output over 100 Mb Ethernet. This output is controlled by a routine running on the ROACH PowerPC software processor. As in the case of the unaccumulated beam outputs, these data are output following SPEAD protocol specifications.

#### **4.4.5 “X” Engine Firmware**

The X-Engine FPGA-based processing node is responsible for generating accumulated visibility matrices for the antennas in the BEST-2 array. An overview of the firmware for the X-engine, designed by Griffin Foster, is shown in Figure 4.14. Interested readers are directed to his doctoral thesis (Foster, 2013) for a full description of the design, which is based on the CASPER windowed X-engine described in Chapter 3. Since this cross-correlation module expects dual-polarization antennas, signals from the BEST-2 array are processed as 16 dual-polarization, rather than 32 single-polarization, receivers. This leads to a visibility matrix with 544 entries. The standard X-engine module is only able to process a single dual-polarization signal on each clock cycle, and two such X-engine instances are required to process the full bandwidth of the BEST-2 array.

Unlike the images formed by the S-engine, visibility matrices may be dumped to disk on a timescale of  $O(10)$  ms over a 10 Gb Ethernet connection, which has enabled the FX

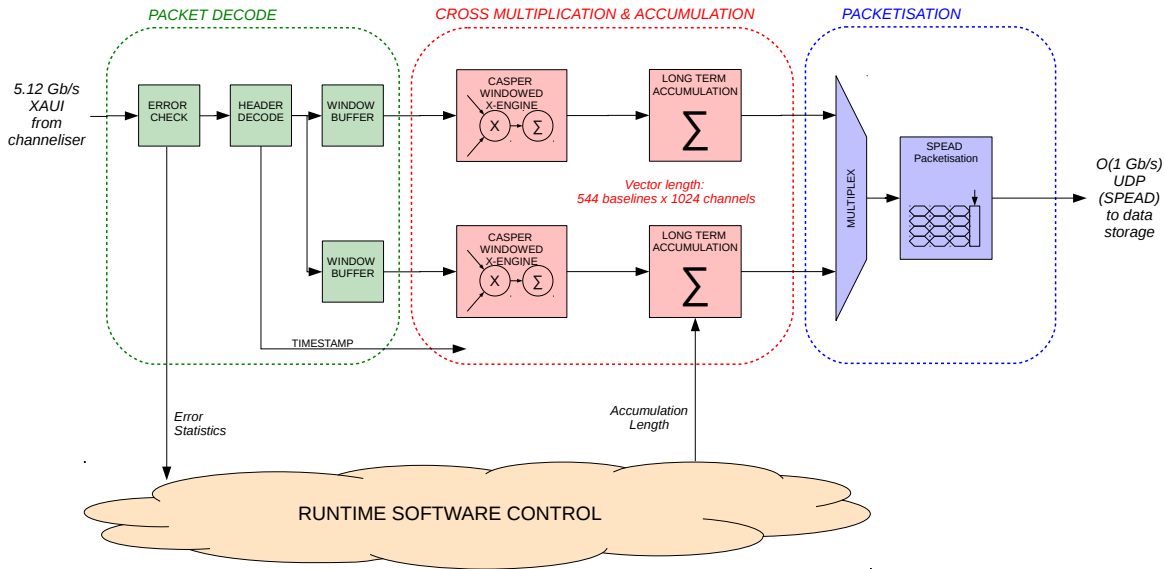


Figure 4.14: A block diagram describing the processing carried out on the “X-engine” ROACH board, which is used to perform cross multiplication and accumulation of antenna signals.

component of the BEST-2 backend to be used as a testbench for rapid imaging software systems.

## 4.5 Software (Control & Calibration) Framework

A suite of software designed for control and monitoring of the BEST-2 system has been developed by Griffin Foster and myself as part of the Medicina deployment. All code is available in the public repository at <https://github.com/jack-h/medicina.git>.

All software development was carried out in Python, due to the ability to rapidly prototype and test functionality, and the vast array of python packages available. Packages extensively used include the numerical NumPy package<sup>4</sup>, the astronomical computation package PyEphem<sup>5</sup>, and the KATCP library<sup>6</sup>, designed for simply interfacing with CASPER ROACH boards.

A variety of scripts allow configuration of the BEST-2 instrument, according to parameters (such as image accumulation length, beam selection, FFT shift schedule, etc.) specified in .xml files, which are read into python objects with a simple parser based on work by Wai Yip Tung. These scripts allow the BEST-2 backend to be controlled and monitored, both using command line utilities and a simple graphical user interface, which allows visualization of the signal paths in the F-engine.

<sup>4</sup>See <http://www.numpy.org/>

<sup>5</sup>See <http://rhodesmill.org/pyephem/>

<sup>6</sup>See <http://pypi.python.org/pypi/katcp>

### 4.5.1 Equalization Settings

Software-side equalization settings are controlled using a simple, but versatile set of python methods. These allow for management of real, complex, and arbitrary bitwidth coefficients loaded in FPGA shared memory. Currently loaded coefficient sets are persistently stored, as well as coefficient components corresponding to calibration gains, passband shapes, and constant offsets. At each update of the coefficients loaded on the FPGA a meta-data issue is triggered, and the new coefficient sets are transmitted via Ethernet, along with an upload timestamp, to the downstream data receivers. This coefficient meta-data is recorded with telescope signal data, allowing multiplications in the FPGA data streams to be compensated for in post-processing.

As discussed in Section 2.3, the complex gains of the BEST-2 antennas must be calibrated before beams can be effectively formed via spatial FFT. In the BEST-2 backend, a full speed correlator is available in parallel to any direct-imaging observations. Data from this correlator are used to calibrate the array, using an algorithm detailed by Boonstra & Van Der Veen (2001) and Section 5.3.1.

Passband coefficients may be specified using a set of polynomial coefficients, or automatically updated using autocorrelation data. Phase coefficients can also be applied to data streams to centre the fan of synthesized beams at the pointing centre of the BEST-2 cylinders.

Since the overall per-antenna scaling on the F-engine exists to ensure signal variance is optimally quantized, a further amplitude equalization stage exists on the S-engine, allowing this scaling to be removed. This allows the overall scaling to be that required by the per-antenna amplitude calibration of the array. Modifying these coefficients also allows FFT windowing functions to be applied to data, and facilitates inverse-noise scaling of FFT inputs.

### 4.5.2 Data Collection

Whilst inter-ROACH communications use the XAUI protocol, data products of both the “S” and “X” engines utilize the UDP-based SPEAD protocol to transmit data over Ethernet. Data recorders are implemented using the associated pySPEAD implementation<sup>7</sup>, whilst raw beam reception is implemented using a custom optimized reader, designed by Alessio Magro. The SPEAD protocol facilitates the transmission of meta-data, including equalization settings and instrument configuration information. Visibility and image data are recorded to disk in HDF5 format<sup>8</sup>, which supports flexible and efficient management of

---

<sup>7</sup>See <https://github.com/ska-sa/PySPEAD>

<sup>8</sup>See <http://www.hdfgroup.org/HDF5/>




Figure 4.15: Beam responses (with aliases outside the primary beam removed) for the eight beams formed by BEST-2 pointed at different azimuths at 400, 408, and 416MHz. Beam 4 is plotted twice, as its aliases give rise to equal beam response in the positive and negative azimuth directions.

large data sets. A variety of data manipulation and visualization tools have been developed which can directly process recorded HDF5 data files. Data sets can also be converted to the FITS Interferometry Data Interchange Format (FITS-IDI) using a conversion tool originally developed by Danny Price<sup>9</sup>. FITS-IDI is supported by many standard radio astronomy packages, including the Common Astronomy Software Applications (CASA).

### 4.5.3 Object tracking

Whilst the electric-field beams generated by spatial Fourier Transform are primarily suited to “drift-scan” observations, where astronomical sources are allowed to transit the beams without being tracked, software has also been written to automatically switch beams to “pseudo-track” up to 8 simultaneous targets during transit (one per beam). In this mode of operation, software will dynamically configure the FPGA outputs to transmit the nearest beam to a target source. This tracking method has the significant drawback that sources appearing between beam centres can not be observed with the full sensitivity of the ar-

<sup>9</sup>Available in the “pyfitsidi” repository at <https://github.com/telegraphic/pyfitsidi.git>

ray. As an example, Figure 4.15 shows the response of the 7 synthesized beams covering a right-ascension range of  $4^\circ$ . Since neighbouring beams overlap at  $\sim 81\%$  of peak power, signal-to-noise is reduced when tracking objects lie between pointing directions. This tracking mechanism is not a recommended operation mode, with beam interpolation being a preferable (but more computationally expensive, and not yet implemented) tracking solution.

A further complication in processing time-domain beams is the varying of beam centers with frequency. Since the fractional bandwidth of the BEST-2 analogue frontend is relatively small, this effect is limited, especially around the pointing centre of the telescope. The beams around the pointing-centre are also preferred as they suffer least from aliasing and primary beam suppression. For this reason the effect of pointing variance is neglected in current observations. It will, however, give rise to an observed variation in frequency dependence of radiation from transiting sources. This effect can in principle be compensated for in downstream processing, but may potentially represent a significant complication to some observations.

It should be emphasized that none of the beam pointing effects mentioned here affects the ability to make images with the BEST-2 system, as image processing can utilize the full set of 128 beams to generate pixels with arbitrary pointing centres.

## 4.6 Conclusions

The digital backend deployed on BEST-2 has been presented. All digital components to this backend were produced as part of the work presented in this thesis, using CASPER-designed FPGA hardware and a variety of custom and publicly available DSP modules. A substantial software suite has also been created, to allow flexible and user-friendly configuration of the instrument using simple .xml configuration files and a collection of python scripts. The BEST-2 backend is unique in offering simultaneous direct-imaging, correlation and beam-forming capabilities. It is also designed to be readily compatible with CASPER 10 GbE-based architectures. This should facilitate easy upgrading of the imaging system to a larger array if required in the future.

In the following Chapter we present a variety of observations made with the BEST-2 array, and analyse the performance of the direct-imaging system with respect to the expected performance outlined in Chapter 2.

## Chapter 5


# Results from the BEST-2 Direct Imaging Telescope

In Chapter 4, the digital backend comprising a direct-FFT imaging system, FX correlator and beamforming system was described. In this chapter, results from this deployment are presented. The effective operation of the direct-imaging system is verified, and both real-time and non-real-time calibration mechanisms are used to generate images of some of the brightest radio sources in the sky. The aim of these images is to assess the overall effectiveness of the imaging system as an FX-correlator alternative. Preliminary results showing pulsar detection using the BEST-2 beamforming system are also presented.

### 5.1 System Commissioning

Initially, FX correlator observations were used to verify the correct operation of the BEST-2 array. Figure 5.1 shows the autocorrelation spectra of the 32 receivers of the BEST-2 system. In order that corruptions due to 4-bit quantization are minimized, a 6<sup>th</sup> order polynomial is used to compensate for the slight slope across the band, and gain differences between receiver chains. These corrections are applied in real time before quantization using the equalization capabilities described in Section 4.5.1. Strong RFI is clear in the lower part of the band ( $\lesssim 404$  MHz), but effective analogue filtering and digital channel isolation allow the majority of BEST-2's 16 MHz of bandwidth to remain relatively clean.

Observations of the bright radio source Cassiopeia A (Cas A, 3C461), with a flux at 408 MHz of  $\sim 4000$  Jy (Baars et al., 1977), serve to test the correlator system. Given the transiting nature of the BEST-2 array and the declination of Cas A of  $+58^{\circ}53'$ , observation of Cas A is limited to approximately 50 minutes each day. Observations of Cas A as it cuts the E-plane of pairs of BEST-2 elements are shown in Figure 5.2, for a time series of 3.5 second integrations. The physical locations of the corresponding array baselines are also shown, using the antenna numbering scheme adopted by the Medicina observatory.



(a) Autocorrelation power spectrum with no band-pass correction. (b) Autocorrelation power spectrum after 6<sup>th</sup> order polynomial fit.

Figure 5.1: Autocorrelation spectra of 32 BEST-2 receivers, before and after application of bandpass corrections.



Primary beam sidelobes are clearly visible at  $\sim -15$  dB, with short baselines having greater off-axis response and more varied sidelobe patterns. Of particular note is the beam-pattern of BEST-2’s longest four redundant North-South baselines. A clear pointing offset is visible, depending on where the baseline is situated in the array. Such variance undermines the assumption that redundant baselines see an identical sky, and emphasizes the importance of electromagnetic uniformity in high-precision arrays which rely on such principles. In the case that such variance of beam shape cannot be avoided we see that direction-dependent post-calibration of observations will always be necessary if the entire array field of view is to be imaged effectively.

In addition to cross-power beam patterns, fringes recorded on six East-West oriented baselines during a transit of Cas A are shown in Figure 5.3. As expected, redundant baselines have equal fringe rates, though with different absolute offsets, caused by uncalibrated phase errors on each signal path.



## 5.2 Equivalence of Spatial FFT and Correlator Data

The first stage of verifying the BEST-2 direct FFT imaging system is to compare its outputs with data from simultaneous observations with the FX correlator. This comparison is initially performed in the visibility (rather than image) domain.



An inverse-transform of imager data can be used to obtain visibilities for each unique baseline in the BEST-2 array. Regardless of calibration, these can be directly compared with the equivalent redundant baselines, and their means, computed by the FX correlator operation. Figures 5.4 and 5.5 show the visibilities obtained in the case of a non-redundant,



(a) A selection of cross-power beam patterns for baselines with largest E-W components.



(b) Cross-power beam patterns for redundant N-S baselines located at different positions in the array.



(c) A representative selection of cross-power beam patterns for short baselines.

Figure 5.2: Cross-power beam patterns for a selection of baselines in the BEST-2 array. Data corresponds to the measured power in a single spectral channel (at 409 MHz) during a transit of Cassiopeia A.

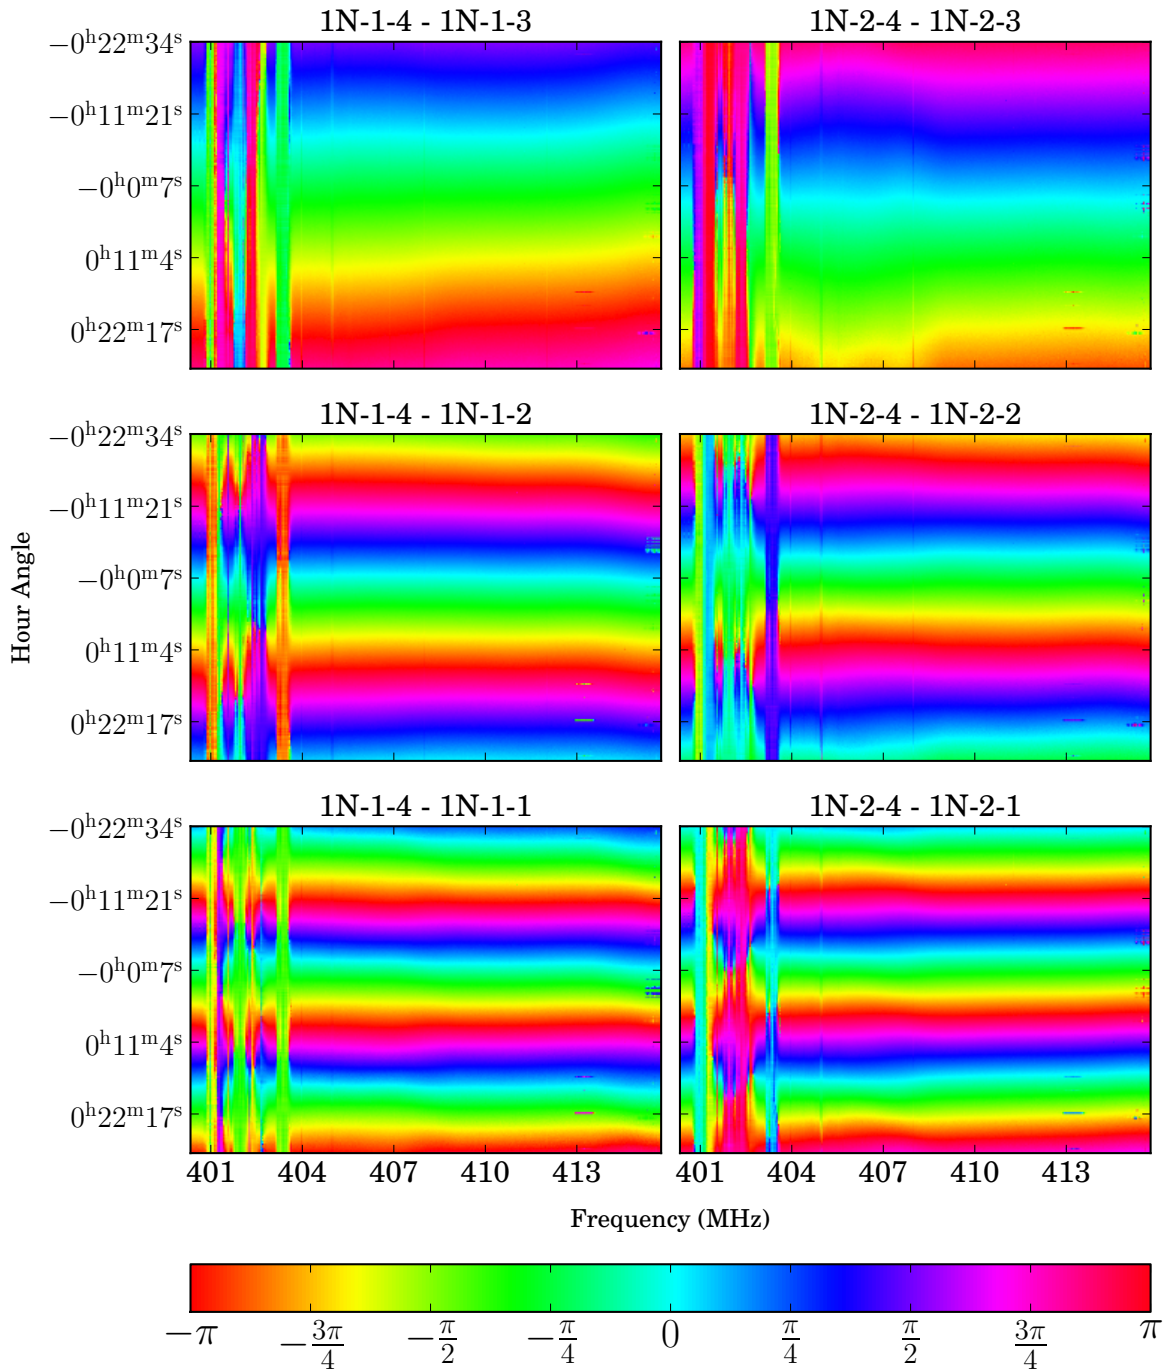


Figure 5.3: Fringes from Cassiopeia A during a transit of BEST-2. Each pair of plots represents a redundant E-W baseline, with baseline length increasing in powers of two from top to bottom. Strong, persistent RFI is visible below 404 MHz, whilst transient RFI bursts can be seen in all visibilities at 413 MHz.




Figure 5.4: Comparative measurements of the non-redundant baseline between antennas 1N-1-1 and 1N-7-4 using the FX (blue-line) and FS (red-line) systems. The polynomial fits used to calculate signal variances are also shown (black line).

and 16-times redundant baselines, respectively. Data are shown in digital backend units (DBUs), i.e., the unscaled outputs of the digital system. Noise on these data, after dividing by a 5<sup>th</sup> order polynomial fit, are shown in Figure 5.6 and 5.7 for the non-redundant and redundant cases, respectively. In the case of the non-redundant baseline 1N-1-1 – 1N-7-4 signal variance on the visibility amplitude obtained by direct imaging is approximately 1.4% higher than that measured by FX correlation. In the case of the 16 times redundant baseline between antennas 1N-1-1 and 1N-1-3, variance on the amplitude of the visibility measurement is approximately 8% lower when measured with the direct-imaging system. Given that in the absence of any calibration we expect measurements obtained by the two systems to be identical up to arithmetic noise, an explanation of these differences is required.

### Data Synchronization

Whilst data between the two observing systems are recorded simultaneously, individual integrations (of 6.7 seconds, or 131072 samples, in the case used above) are not synchronized. Since the five minutes of data used to compare noise represents relatively few integrations, this lack of synchronization is enough to introduce significant variance in the relative noise of the two systems. To illustrate this, a simple simulation was conducted. The effects of computing noise ratios based on two sets of 80 regular, 131072 sample integrations of two

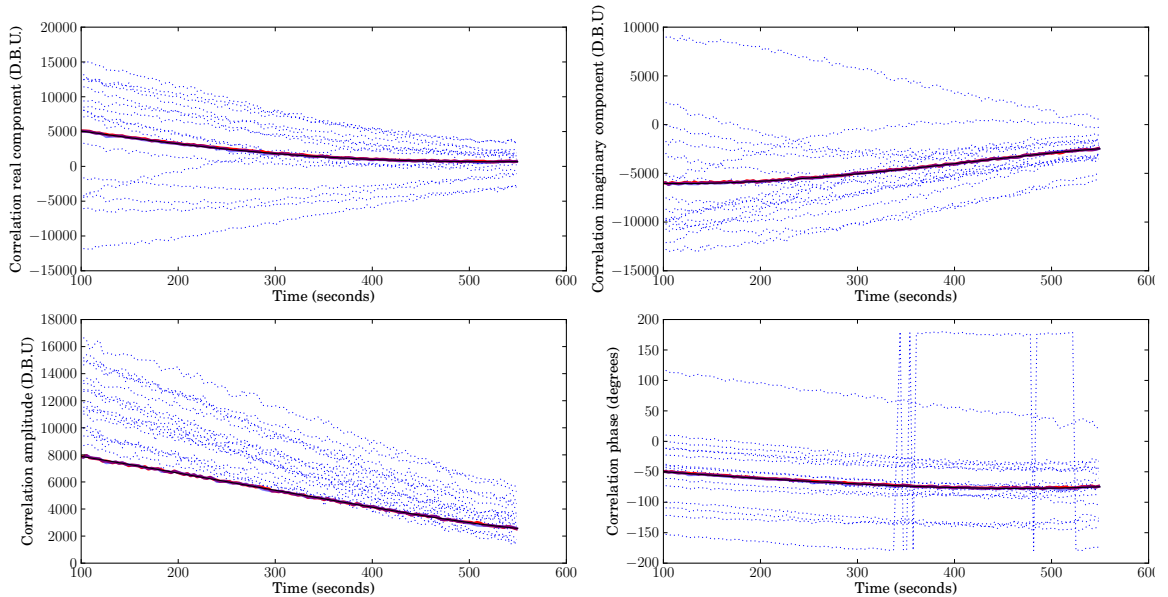


Figure 5.5: Comparative measurements of the highly-redundant baseline between antennas 1N-1-1 and 1N-1-3 (and similarly spaced pairs). Dotted blue lines represent FX measurements of each of the 16 redundant baselines. The mean FX (solid blue line) and FS (solid red line) measurements are also shown, as well as polynomial fits used to calculate signal variances (solid black line). Variances of the FX and FS system visibility amplitudes are  $9.47 \times 10^{-5}$  and  $1.03 \times 10^{-4}$  DBU, respectively.

partially-correlated Gaussian noise streams were evaluated. These simulations suggest that in a single channel, the ratio of variances in the FX and FS measurement sets has mean  $\left\langle \frac{\sigma_s^2}{\sigma_x^2} \right\rangle = 1.01 \pm 0.03$  and variance  $0.03 \pm 0.016$ . Computing the ratio of variances for all unique baselines in the data set yields  $\left\langle \frac{\sigma_s^2}{\sigma_x^2} \right\rangle = 0.99$ , with variance 0.019, in agreement with these simulations.

Using more time samples and a larger subset of frequency channels would allow the equivalent performance of the two measurement systems to be verified with higher certainty. Unfortunately, limited data and frequent RFI contamination make this difficult with the available data set (though a larger, calibrated data set is examined in Section 5.4.2). For this reason, further measurements were taken with both correlator and imaging systems with all antenna signals simulated by digitally synthesized Gaussian noise. With this arrangement in place, noise variance was analysed for approximately 12 minutes of data (in  $\sim 3$  second integrations), over 700 frequency channels. Under perfect operation, simulations of this data set result in a approximately Gaussian distributed variance ratio, with mean  $1.005 \pm 0.002$  and variance  $0.011 \pm 0.004$ . These values match with those calculated from the hardware data sets, which have variance ratios with mean 1.006 and variance 0.008. The measured values of  $\frac{\sigma_s^2}{\sigma_x^2}$ , averaged over 700 channels for each baseline, are shown in Figure 5.8.

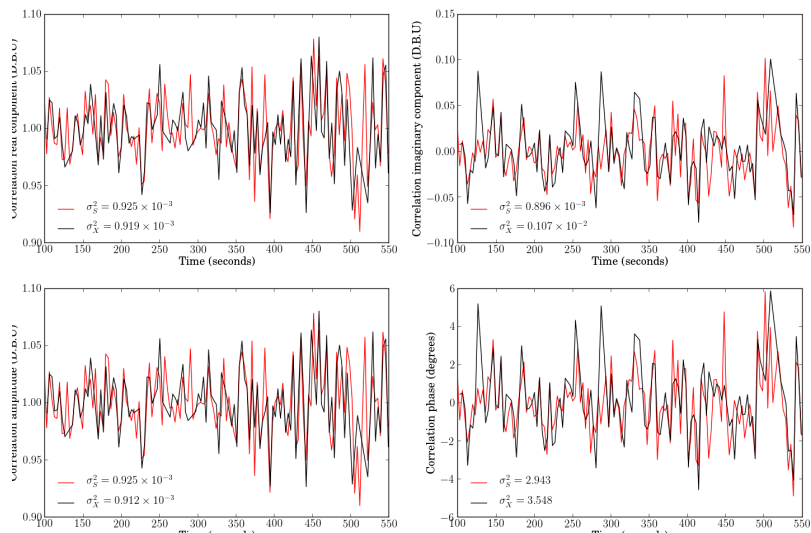


Figure 5.6: Residuals of cross correlations of the non-redundant baseline between antennas 1N-1-1 and 1N-7-4 using the FX (black-line) and FS (red-line) systems, after dividing the polynomial fits shown in Figure 5.4.




Figure 5.7: Residuals of cross correlations of the 16 times redundant baseline between antennas 1N-1-1 and 1N-1-3 using the FX (black-line) and FS (red-line) systems, after dividing the polynomial fits shown in Figure 5.5.




Figure 5.8: Ratio of variances of baseline data, as computed by the direct imaging and cross correlation systems. Data has mean 1.006 and variance 0.008. Simulations suggest under perfect operation, the variance ratios have mean  $1.005 \pm 0.002$  and variance  $0.011 \pm 0.004$ , in agreement with measured data.




Figure 5.9: A histogram of noise on real part of the three times redundant 1N-1-1 - 1N-8-2 baseline. Plotted values reflect data taken from 400 samples and 400 frequency channels. Antenna data are simulated on the FPGA by mutually uncorrelated white noise generators.

Another test of the noise levels of measurements from the two systems is to simply compute the RMS noise of all time samples and channels for a single redundant baseline in the system. This assumes that the noise in each frequency channel is equal, which is reasonable for the case of simulated white-noise inputs. In the case of the 3-times redundant baseline between antennas 1N-1-1 and 1N-8-2, using digitally generated noise sources, the distribution of noise in the real part of measured visibilities is shown in Figure 5.9. Noise levels on measurements from two systems agree to within 0.06%. Other baselines show similar results. Given these results, we conclude that both correlator and imaging systems are working as expected, and up to small arithmetic rounding effects generate equivalent data products.

### 5.3 Calibration Strategies

In order to combine baseline data in a way that preserves signal-to-noise ratio, a key requirement of the BEST-2 direct imaging system is the prior knowledge of gain corruptions in the receiving chains of each of the 32 elements in the array. Given the regular layout of antennas in this array, redundant baseline calibration strategies such as those described by Liu et al. (2010); Noorishad et al. (2012) are applicable. However, in this work I opt to use a general sky-model based algorithm, which assumes prior knowledge of the sky observed by the telescope, but does not impose the condition that redundant baselines behave identically.

This strategy is chosen for a number of reasons. The strong sidelobes of BEST-2 primary beams, which are not identical even amongst redundant baselines, make calibration challenging when there are bright sources away from the phase-center of the array. In such a situation, the assumption of redundant baselines breaks down, and redundant-baseline cali-

bration methods are likely to prove ineffective. Even without requiring baseline redundancy, varying sidelobe gains across the array makes for challenging calibration, as gain corruptions are strongly direction-dependent. To avoid this pitfall, the simplest calibration method is to wait for a transit of a bright source. In this case, the model sky is trivially a point source at the array phase center, from which gain solutions can be very easily extracted.

### 5.3.1 Sky-Model fitting

The calibration algorithm used in this chapter is the column ratio gain estimation method, described by Boonstra & Van Der Veen (2001). This is shown to be competitive with other algorithms for high signal-to-noise observations. Chosen for simplicity, and ease of implementation, the algorithm uses ratios of antenna cross-correlations to estimate the diagonal (autocorrelation) elements of a measured visibility matrix in the absence of a system noise contribution. This allows creation of a modified correlation matrix, which features these elements, and can subsequently be used, under the assumption of a model sky consisting of a point source at phase center, to obtain per-antenna gain solutions. For clarity, the algorithm is explicitly outlined (following Boonstra & Van Der Veen, 2001) below.

1. In the absence of noise, and given a model sky consisting of a point source an phase center, the observed correlation matrix,  $\mathbf{R}$ , is simply equal to  $\mathbf{g}\mathbf{g}^H$ , where  $\mathbf{g}$  is the vector of antenna gains we wish to obtain. In the presence of noise, the diagonal elements of  $\mathbf{R}$  (i.e., the autocorrelations) are subject to an additive noise term, whilst in the situation of high signal-to-noise, the off-diagonal elements are asymptotically unchanged.
2. Ratios of antenna gains,  $\alpha_{ij} = g_i/g_j$ , can be computed by considering the two rows,  $\mathbf{c}_i$  and  $\mathbf{c}_j$  of  $\mathbf{R}$ , and solving

$$\mathbf{c}_i = \alpha_{ij}\mathbf{c}_j. \quad (5.1)$$

3. Representing the  $ij$ -th element of  $\mathbf{R}$  by  $r_{ij}$ , a least squares solution of Equation 5.1, which does not include diagonal elements of  $\mathbf{R}$ , is given by

$$\alpha_{ij} = (\mathbf{c}_i^H \mathbf{c}_i)^{-1} (\mathbf{c}_i^H \mathbf{c}_j) = \frac{\sum_{k \neq i, j} r_{ki}^* r_{kj}}{\sum_{k \neq i, j} r_{ki}^* r_{ki}}. \quad (5.2)$$

This allows elements of  $\mathbf{g}$  to be estimated as  $|g_i|^2 = \mathbb{R} \left( \alpha_{ij} r_{ij}^* \right)$  for any  $j$ .

4. A modified correlation matrix,  $\mathbf{R}'$  is constructed, which is equal to  $\mathbf{R}$  but with auto correlations replaced by estimates of  $|g_i|^2$  using the calculation above. Eigenvalue

decomposition of  $\mathbf{R}'$  produces an estimate of  $\mathbf{g}$ , given by  $\mathbf{u}_1\sqrt{\lambda_1}$ , where  $\lambda_1$  is the largest eigenvalue of  $\mathbf{R}'$ , and  $\mathbf{u}_1$  its corresponding eigenvector.

The column ratio gain estimation algorithm has been implemented in a straightforward python routine, which is able to directly interface with the hdf5-format data products generated by the BEST-2 FX correlator.

### 5.3.2 Calibration Stability

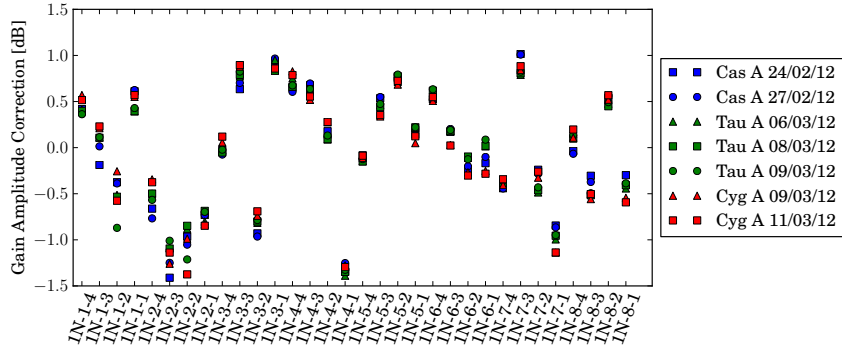
A number of calibration sets have been obtained for the BEST-2 array, using observations of the bright radio sources Cas A, Tau A and Cygnus over a number of days. The calibration phases and amplitudes computed from these observations are shown in Figure 5.10. The degree of freedom associated with defining a phase and amplitude reference has been fixed by requiring that, over each calibration observation, the mean gain correction is unity.

After removing the mean amplitude gain of each antenna over the 7 observations, the amplitude gains have a standard deviation  $\sigma = 0.092$  dB (Figure 5.10b). Phase calibrations, after similarly removing the per-antenna mean and compensating for potential, small pointing errors, have a standard deviation of  $\sigma = 1.9^\circ$  (Figure 5.10d).


Calibrations are noticeably source-dependent, which we primarily attribute to temperature fluctuations throughout the day (at the epochs observed, the transit times of Cygnus, Cas A and Tau A were approximately 0800, 1130 and 1800 UTC, respectively). Inaccuracies in the sky model, i.e. the absence of other sources, also potentially cause calibration errors, though these have not yet been investigated further. Given the calibration variances found in these preliminary observations, in the absence of systematic errors, additional noise per baseline is expected to be  $< 1\%$  (Section 2.5.2).

### 5.3.3 Calibration Artifacts


Figure 5.11a shows a  $32 \times 32$  antenna correlation matrix before and after calibration. The data in this matrix represent an observation of Cygnus lasting approximately 1 minute. After calibration, the standard deviation of the phase and power of the corrected data from the trivial sky model—that is, a  $32 \times 32$  matrix whose elements are all equal to unity—are  $1.04^\circ$  and 0.03, respectively. Clear structure in the residual errors indicate non-antenna-based gain errors, associated with coupling in the array analogue signal paths before digitization, or model-sky errors. Collapsing the two-dimensional correlation matrix into a one-dimensional vector of 52 unique, non-zero baselines shows that the shortest baselines suffer most from these effects (Figure 5.11b). This is likely to be the case if errors are due to coupling between receivers, or due to large-scale galactic structure not included in the calibration sky model. Phase errors appear largely confined to baselines formed with antennas on the




(a) Amplitude calibration coefficients.



(b) Amplitude calibration coefficients relative to each antenna's mean.  $\sigma = 0.092$  dB.



(c) Phase calibration coefficients.



(d) Phase calibration coefficients relative to each antenna's mean, after compensating for small pointing errors.  $\sigma = 1.9^\circ$ .

Figure 5.10: Amplitude and phase calibrations for the BEST-2 antennas during transits of multiple sources at multiple epochs.

same or neighbouring cylinders. Such errors, assuming they are stable and known, do not necessarily interfere with the operation of the direct imaging system. Since they affect redundant baselines largely uniformly, they can be calibrated in the end data products with no loss of sensitivity.

## 5.4 Direct Imaging Results

With the ability to calibrate the array, and data reassuringly suggesting that calibrations are relatively stable with time, BEST-2 operations with the imaging system were begun. Initial imaging with the BEST-2 array was carried out without calibration. In the simplest case, the output of the BEST-2 imaging system is a  $16 \times 8$  pixel image of the sky. In the absence of any real-time phasing (or fringe-stopping, in the interferometry sense), a typical output data product is shown in Figure 5.11a, which shows a 3.36 second snapshot during a transit of Cas A.

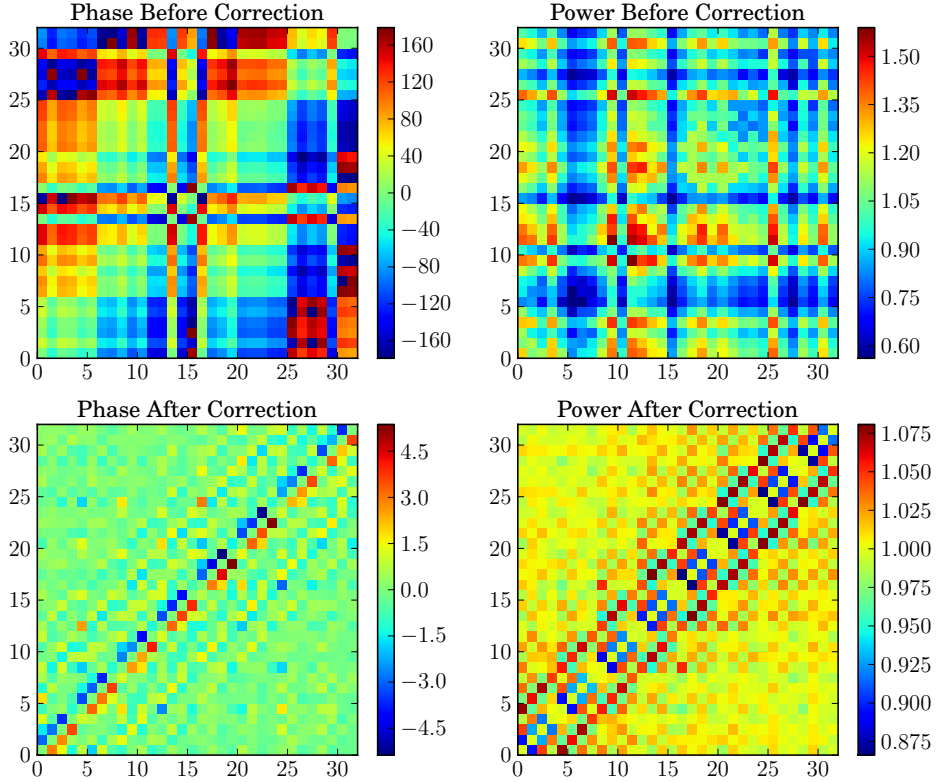
Though Figure 5.11a encodes all the information from the BEST-2 array it can be further processed by phasing the data towards Cas A, so that the position of Cas A appears in the centre of the image, and by re-imaging with more pixels. In order to leverage the considerable existing software development efforts in the area of synthesis imaging, this further processing is carried out by converting the image data sets into correlation matrices. Since only one data entry exists for each redundant baseline in the array, this correlation set contains 53 baselines. These are nominally identified as being one of the 32 correlations including antenna 0, and 21 correlations with antenna 3. Once the image data sets exist as correlations, it is simple to phase and image them using existing radio astronomy software tools. In the case of BEST-2 data, phasing is achieved using the Aipy package<sup>1</sup>, whilst imaging is performed using CASA<sup>2</sup>. In particular, the latter enables direct-imaging data to be processed using the usual tools of interferometry including rotation and frequency synthesis.

A  $5^\circ \times 5^\circ$ ,  $256 \times 256$ -pixel image generated in CASA using the same observation shown in Figure 5.11a is shown in Figure 5.11b. Here we have used a two-minute observation between 409 MHz and 410 MHz to generate the image, which is centered on the location of Cas A. Note that phasing the correlations towards Cas-A has moved the image centre, though lack of calibration means that the peak response of Cas A appears offset. This image is, up to arithmetic rounding artifacts in the spatial-FFT system, identical to an image formed using an uncalibrated FX correlation matrix.


---

<sup>1</sup>Astronomical Interferometry in Python: <https://pypi.python.org/pypi/aipy>.


<sup>2</sup>Common Astronomy Software Applications: <http://casa.nrao.edu>.



(a) Baseline phases and amplitudes before and after calibration. The data in this plot represents a  $\sim 1$  minute observation of the transit of Cygnus using a single 19.5 kHz subband.



(b) Baseline phases and amplitude residual errors after calibration. Errors are shown for each of the 52 non-zero redundant baselines in the BEST-2 array, with mean values in red.



(a) Raw data output of the BEST-2 imaging system, showing Cas A approaching transit. Colour scale is normalized linear power.

(b) An uncalibrated, multi-frequency synthesis image, formed using 2 minutes of data between 409 MHz and 410 MHz. The image is centered on the location of Cas A. Colour scale is normalized linear power.

Figure 5.11: Uncalibrated direct-images of Cas A.

### 5.4.1 Post-calibrating Direct Images

In traditional interferometric observations, the  $O(N^2)$  measurements in the correlation matrix significantly over-determine the  $O(N)$  antenna gain calibration coefficients which need to be computed in the image-making process. This allows the self-calibration procedure outlined in Section 2.3 to be carried out, in which the coefficients are obtained by iteratively minimising the difference between the measured data set and a sky model, which is determined from the telescope data itself. As discussed in Section 2.5.2, the directly obtained image can be calibrated after observation by applying one gain coefficient to each unique baseline. Since this is equivalent to applying one gain coefficient per independent pixel in the direct image, this is not overconstrained like the traditional self-calibration problem. This means that a sky-model cannot be iteratively developed from the telescope data set. Rather, if one starts with a sky model, calibration coefficients can be found which match the data to this model. Obviously, if we have a perfect sky model, then there is little point in making new images of the sky. Therefore, in the case of direct-imaging systems, post-calibration assumes that we have access to further data sources than just the direct images themselves. These need not be full time and frequency resolution correlation matrices, but must be sufficient to provide a significant level of overdetermination of the array gain calibration solutions. In the case of the BEST-2 system, we have access to the full correlation matrix from the FX correlator, which fulfils this role. In an effective large-scale deployment of a direct-imaging instrument, it is assumed that a full correlator — which

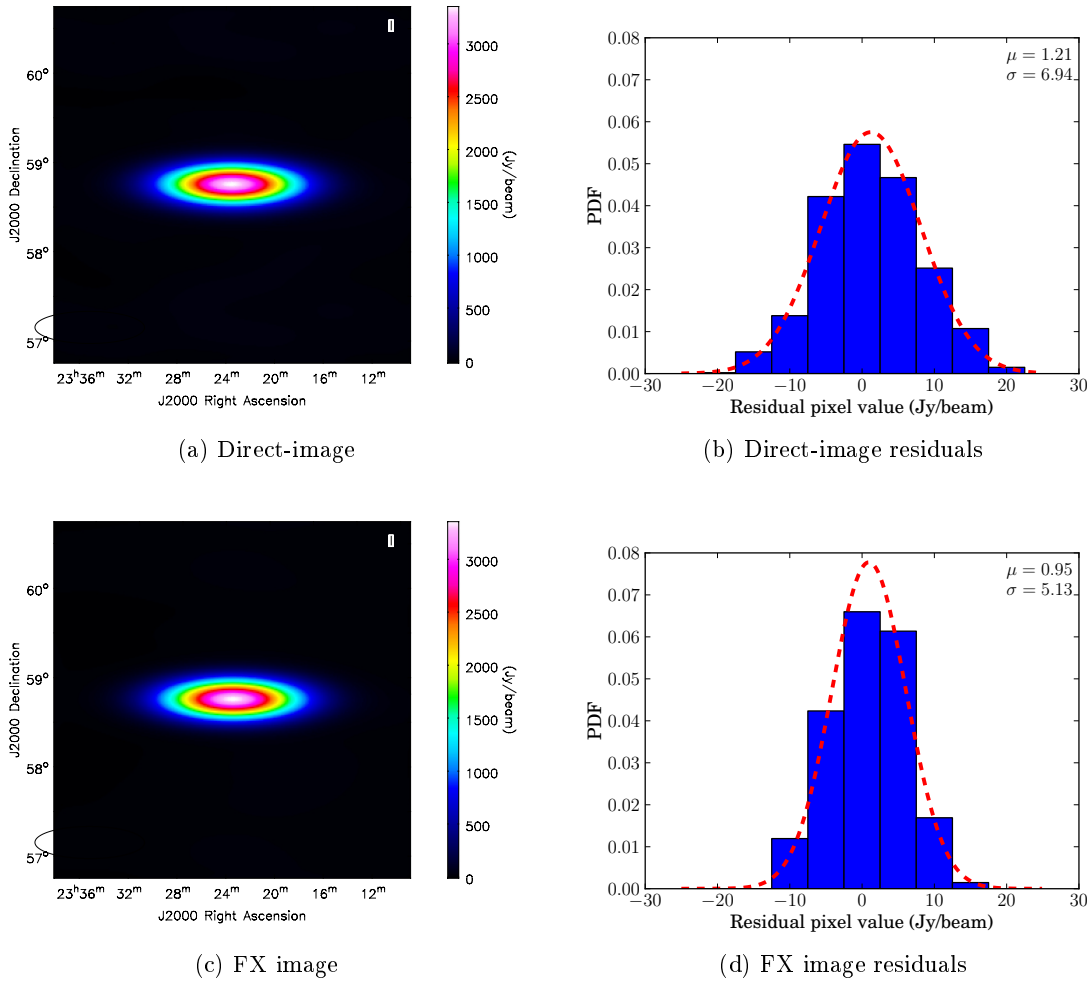


Figure 5.12: Post-calibrated, cleaned images of Cas A, with residuals after cleaning.

undermines the efficiency of the direct-imaging process — will not be present, but rather a correlator operating on a subset of baselines, frequency channels and time samples will provide calibration data.

Calculating gain corrections using the method outlined in Section 5.3.1 and applying these to the direct images yields an image that can be cleaned with standard CASA routines. Using natural weighting of visibilities, the cleaned images, and their residuals, are shown in Figure 5.12. As expected, though in both cases the cleaned images represent point sources with Gaussian noise, signal-to-noise is significantly higher (by  $\sim 35\%$ ) in the FX-imaging case. This is simply a result of the lack of real-time calibration.

#### 5.4.2 Real-time Calibrated Images

Whilst it is clearly possible to use a direct-imaging system for map making even without any real-time calibration, signal-to-noise is significantly lower than that which can be obtained





Figure 5.13: A real-time calibrated image of Tau A, generated with a single channel and integration of direct-imaging data. Colourscale is normalized power, in dB.

with a traditional FX correlator. To test the performance of the direct imager when real-time calibration is used, antenna gain solutions were obtained by observing a transit of Cygnus and applying the algorithm described in Section 5.3.1. These calibration coefficients were then used to calibrate the antenna signals in real time during a subsequent observation of Tau A, using the firmware infrastructure outlined in Section 4.5.1. The directly-formed image obtained from this observation is shown in Figure 5.13. In this case, the image is generated directly from the raw hdf5 data set, which has been phased towards Tau A and interpolated to give a higher pixel density. This dirty image exhibits the PSF structure expected from the  $4 \times 8$  antenna grid of BEST-2.

Since the real-time calibrated direct-image is approximately equivalent to an ideal FX image, it is simplest to compare signal-to-noise of the direct-imaging and FX systems in the visibility domain. The analysis of Section 5.2 has been repeated for FX and direct-imaging data sets comprising 200 seconds and 100 channels. Given the data synchronization limitations already mentioned, a simulation matching these parameters was also conducted. The simulation suggests that in the presence of a perfectly operating system, the variance ratios of the two data sets should have a mean  $1.000 \pm 0.003$  and variance  $0.0004 \pm 0.0002$ . We find that the real data has variance ratios with mean 0.9994 and variance 0.00032, in agreement with these simulations. These data are shown in Figure 5.14a.

Since we do not expect the real-time calibration applied to be perfect, the Tau A direct-imaging data set has also been compared to the FX visibility set which has been further



(a) Variance ratios when both FX and FS systems are operated with same real-time calibration. (b) Variance ratios when FX data are post-calibrated.

Figure 5.14: Ratios of variances on visibilities from the direct imaging and FX correlator systems. After further calibration, correlator noise levels improve only slightly ( $\lesssim 0.5\%$ ) relative to the real-time calibrated direct imaging data set. For the 200 second (6.7 second integrations) and 100 channels analyzed, simulations suggest that the variance ratios of the two systems with identical noise properties should have mean  $\mu_r = 1.000 \pm 0.003$ , and variance  $\sigma_r^2 = 0.0004 \pm 0.0002$ .

post-calibrated. In this case, variance ratios (shown in Figure 5.14b) have mean 1.0041 and variance 0.00036. The mean signal-to-noise of the FX data set has been improved, but only by a small amount ( $\sim 0.5\%$ ). This is consistent with the small amount of variation in gain calibrations seen over multiple observations in Section 5.3.2.

Finally, we examine the noise on visibility measurements as a function of baseline redundancy. This relationship is plotted for the Tau A data sets in Figure 5.15. In an ideal array we expect noise levels to fall with the square-root of redundancy. We observe noise  $\propto \text{redundancy}^{-0.43}$ . Though this relationship has not yet been thoroughly investigated, we attribute the scaling to variation of signal-to-noise levels on individual BEST-2 receivers.

## 5.5 Beamforming

With the ability to calibrate antenna gains in real-time in place, the unaccumulated outputs of the FFT imager are calibrated electric-field beams on the sky. These beams remain fixed in position relative to the array location. Figure 5.16 shows the power in each of eight beams in an E-W row, during a transit of Cas A. This plot emphasizes the sampling of the BEST-2 primary beam by the array's synthesized beams. Significant aliasing effects are visible for all off-centre beams, becoming more severe for beams further from the array phase centre. This is a result of the hierarchical beamformer used in BEST-2, the first layer of which is the analogue summation of 16 critically spaced dipoles. Section 6.1 shows an explicit example of




Figure 5.15: RMS noise on baselines of different redundancy, as measured by both digital backends. The dashed line is a linear fit with gradient  $-0.43$ . In the ideal case of all antenna signals having equal signal-to-noise levels, we expect the a fit with gradient  $-0.5$ . I.e. a noise level proportional to the square root of the baseline redundancy.

the difference between beams formed in a hierarchical, rather than single-stage, beamformer.

Figure 5.16 also shows the loss of signal when the source is between beams. In the imaging cases described in the previous section, manipulation of the entire set of beams (or more accurately, the correlations they represent) allows phasing to any direction. However, when the beams are processed separately, or only a subset of beams are available, in general such arbitrary pointing and tracking is not possible.

### 5.5.1 Pulsar Detection

The case of non-tracking synthesized beams has one obvious use case. This is as a front-end for a blind transient-event search. In such a search, a drift scan—where the sky drifts through a set of fixed beams—is just as effective as a targeted scan, where a single patch of sky is targeted for a long observation. The survey effectiveness is simply defined by the product of time and field of view processed, and the sensitivity of the instrument used to conduct the observations. The BEST-2 digital backend has the capability of outputting eight unaccumulated synthesized beams for further processing. Alessio Magro, working from the University of Malta, has developed a GPU pipeline for processing these data streams. This pipeline performs a search for transient events at multiple dispersion measures<sup>3</sup> (DMs), and

<sup>3</sup>A dispersion measure (DM) is the constant of proportionality which relates the arrival time of different frequency components of a pulse, which are related by  $t \propto \frac{1}{f}$ . DM is determined by the electron density integrated over the line of sight from observer to source, usually quoted in units of  $\text{pc cm}^{-3}$ . See Lorimer &




Figure 5.16: The normalized power observed in each of eight beams in an E-W row, during a transit of Cas A.

multiple periodicities (Magro et al., 2013).

Whilst a full description and results from the transient search pipeline are currently in preparation, here we show some of the early results from the pulsar system, to demonstrate the potential of the BEST-2 array for transient monitoring. Early testing was performed by pointing the array to the declination of bright pulsar B0329+54. As the pulsar transited the array, a row of eight East-West beams were processed in real-time to search for transient events. After compensating for different possible dispersion measures and summing the signal over 10 MHz of the observing band, we can measure the signal-to-noise of the pulse as a function of DM. This measurement, conducted in real-time, is shown in Figure 5.17. Peak signal-to-noise occurs at a DM of approximately  $27 \text{ pc cm}^{-3}$ , that expected given existing measurements of the pulsar.

After compensating for the DM of the pulsar, multiple pulses can be summed to obtain a profile of the average pulse shape. Pulse profiles as obtained by each of the eight beams processed are shown in Figure 5.18. These profiles were obtained by summing approximately 50 pulses, though data quality has suffered due to the need to excise significant quantities of RFI. Figure 5.19 shows a single one of these profiles archival data from a 408 MHz Lovell Kramer (2005, Chapter 1) for an introduction to pulsar astronomy.




Figure 5.17: Signal to noise ratio of B0329+54 time-series dedispersed with different DMs. The peak signal to noise is observed at the pulsars expected DM of  $27 \text{ pc cm}^3$ . Figure credit: Alessio Magro

Telescope observation overlaid (Gould & Lyne, 1998).

Work continues on the transient search pipeline to effectively mitigate RFI contamination, and to intelligently identify interesting events for follow-up observation. Here we simply present the BEST-2 backend as an example of the role FFT beamforming can play in transient survey campaigns. In particular we emphasize the ease and efficiency with which such functionality can be incorporated into a larger imaging system.

## 5.6 Conclusions

In this chapter some early results from the BEST-2 direct-imaging, correlator, and pulsar processing systems have been presented. Within expected uncertainties, we find that the direct-imaging system is capable of generating correlation datasets with equivalent signal-to-noise as those directly obtained with an FX correlator. We also demonstrated the ability to form calibrated images even in the absence of real-time calibration of the direct images. As is expected, these images have significantly lower signal-to-noise than those obtained with either a calibrated direct-imaging observation, or a full FX correlation matrix. However, we see that almost all signal-to-noise can be retained even with slow ( $\sim$ day) real-time calibration coefficient update cycles. These are likely to be far easier to implement and manage than direct-imaging systems which demand rapid and continuous real-time calibration updates.




Figure 5.18: Sampling (at zenith) of the BEST-2 primary beam by the 128 digitally generated beams. The large oval is an approximate representation of the BEST-2 primary beam. Outermost beams, which have equal response from two directions, are shown twice.




Figure 5.19: A profile of pulsar B0329+54, obtained by folding  $\sim 50$  dedispersed pulses observed a single beam from the BEST-2 imaging system. Overlaid is a profile of B0329+54 obtained using the Lovell telescope at 408 MHz (Gould & Lyne, 1998)

This suggests, as was posed in Chapter 2, that the usability of direct-imaging systems for map-making experiments is likely to be determined by the ease and efficiency with which antenna gain calibrations can be obtained. Unlike traditional FX correlation, post-calibration direct-imaging potentially decouples the observation data from the data used to compute calibrations. Further work in the field of direct-imaging should investigate the quantity of non-direct-image data required to effectively perform post-calibration of images, and the computational cost of obtaining such data.

We note that even under the assumption of perfect per-antenna real-time calibration, the assumption of redundant baseline equivalence is not quite accurate in the case of BEST-2. Primary beam shapes of each receiver are seen to vary depending on their position in the array, giving rise to direction-dependent gain calibration solutions. This implies that wide-field imaging will only ever be possible with some quantity of post-processing and post-calibration of data. This further supports the argument that the real-time component of gain calibration need only be relatively crude.

We have also shown the effectiveness of the BEST-2 direct-imaging system as a frontend for transient processing pipelines, with little implementation cost beyond that already required by the Fourier transform-based map-making process. The BEST-2 transient pipeline continues to be developed by Alessio Magro, with the BEST-2 direct-imaging and beam-forming front-end providing a useful and versatile test platform.

## Chapter 6

# Digital Signal Processing in the SKA-low

The low-frequency component of the Square Kilometre Array (SKA-low) will be rolled out in two stages, the first of which (SKA-low-1) will feature  $\sim 2.5 \times 10^5$  antennas. These will be sub-divided into stations, which will use beamformers to synthesize (likely multiple) electronically steerable fields of view on the sky. Signals from each station will then be combined by an interferometer, capable of computing the visibilities which directly measure the spatial power spectrum of the radio-sky and can be used to form images.

The main science motivation for SKA-low is the study and characterization of the Epoch of Reionization (see Section 1.1.1, Mellema et al., 2012). In 2010 a preliminary system description of SKA-low-1 was proposed by Dewdney et al. (2010), advocating the construction of an array consisting of 50 stations, each comprising 11200 small  $O(\text{m}^2)$  antennas. More recently, Dewdney (2013) have revised the specifications of SKA-low-1, producing a baseline design of the system. The evolution of specifications between those proposed by Dewdney et al. (2010) and Dewdney (2013) include some radical changes. Aside from a rescoping of the observing frequency band from 70–450 MHz to 50–350 MHz, the new specifications define an array comprising 911 stations, each with 289 antennas.

In this Chapter we explore the possible processing implications of the two suggested options for SKA-low-1, and in general the effect of station size and quantity on processing requirements. Specific processing implementations which might be employed in SKA-low are considered and discussed, including the potential for using fast Fourier transform-based beamforming and imaging systems such as that demonstrated in Chapter 5.

### 6.1 Beamforming

Initial processing in SKA-low will consist of digitizing antenna signals and forming beams at a station level. In this section we give a brief overview of relevant beamforming architectures

and antenna configurations which may be employed.

### 6.1.1 Beamforming Architectures

There are numerous ways in which beamforming systems may be implemented. First and foremost, the geometric delays (see Figure 1.9) associated with signals from different directions of arrival can be compensated for in either the time domain or the frequency domain, giving rise to two distinct categories of beamformers. In the former case, beamforming takes the form of applying programmable delays to signal streams, which may include interpolation of digitally sampled data. This is the mechanism used, for example, by the Sub Millimetre Array (SMA) beamformer (Nagpal, 2006). Alternatively, in the case of frequency domain beamforming, data streams are divided into narrow-band channels, each of which is phase-rotated by multiplication with a complex weight (see Section 1.3.1).

The requirement to channelize all data streams is a major burden of frequency domain beamforming. Nevertheless, frequency-domain methods are often chosen in radio astronomy for a number of reasons. Firstly, phase rotation by complex multiplication is trivially implemented in modern computing hardware, and can compensate for delays with precisions far greater than the system's sampling period. Where multiple beams are required, multiplying by multiple coefficient sets is straightforward and efficient. Secondly, having access to channelized data streams allows easy correction of frequency-dependent corruptions in voltage samples by appropriate modification of beamforming coefficients. Sub-sample delays and frequency-dependent effects can be compensated for in the time-domain, by use of an FIR interpolation filter with appropriately chosen kernel (namely the Fourier transform of the frequency domain corrections). However, the filtering process can be computationally demanding, with large kernels required to achieve required precision (D'Addario, 2008). In many cases, the filtering requirements of such a system may outweigh the costs of channelization. In any case, filtering techniques need to be replicated in their entirety for each required beam, rendering them inefficient when several beams are required. We assume in the remainder of this chapter that most, if not all, beamforming in SKA-low-1 will take place in the frequency domain.

In either of the time or frequency domain cases, the beamforming process for large arrays is often too large a task to be carried out by one processing node. In this case, it can be subdivided among processors in a variety of ways. Here we highlight two mathematically distinct beamforming mechanisms. These are the *hierarchical* and *partial* processes.

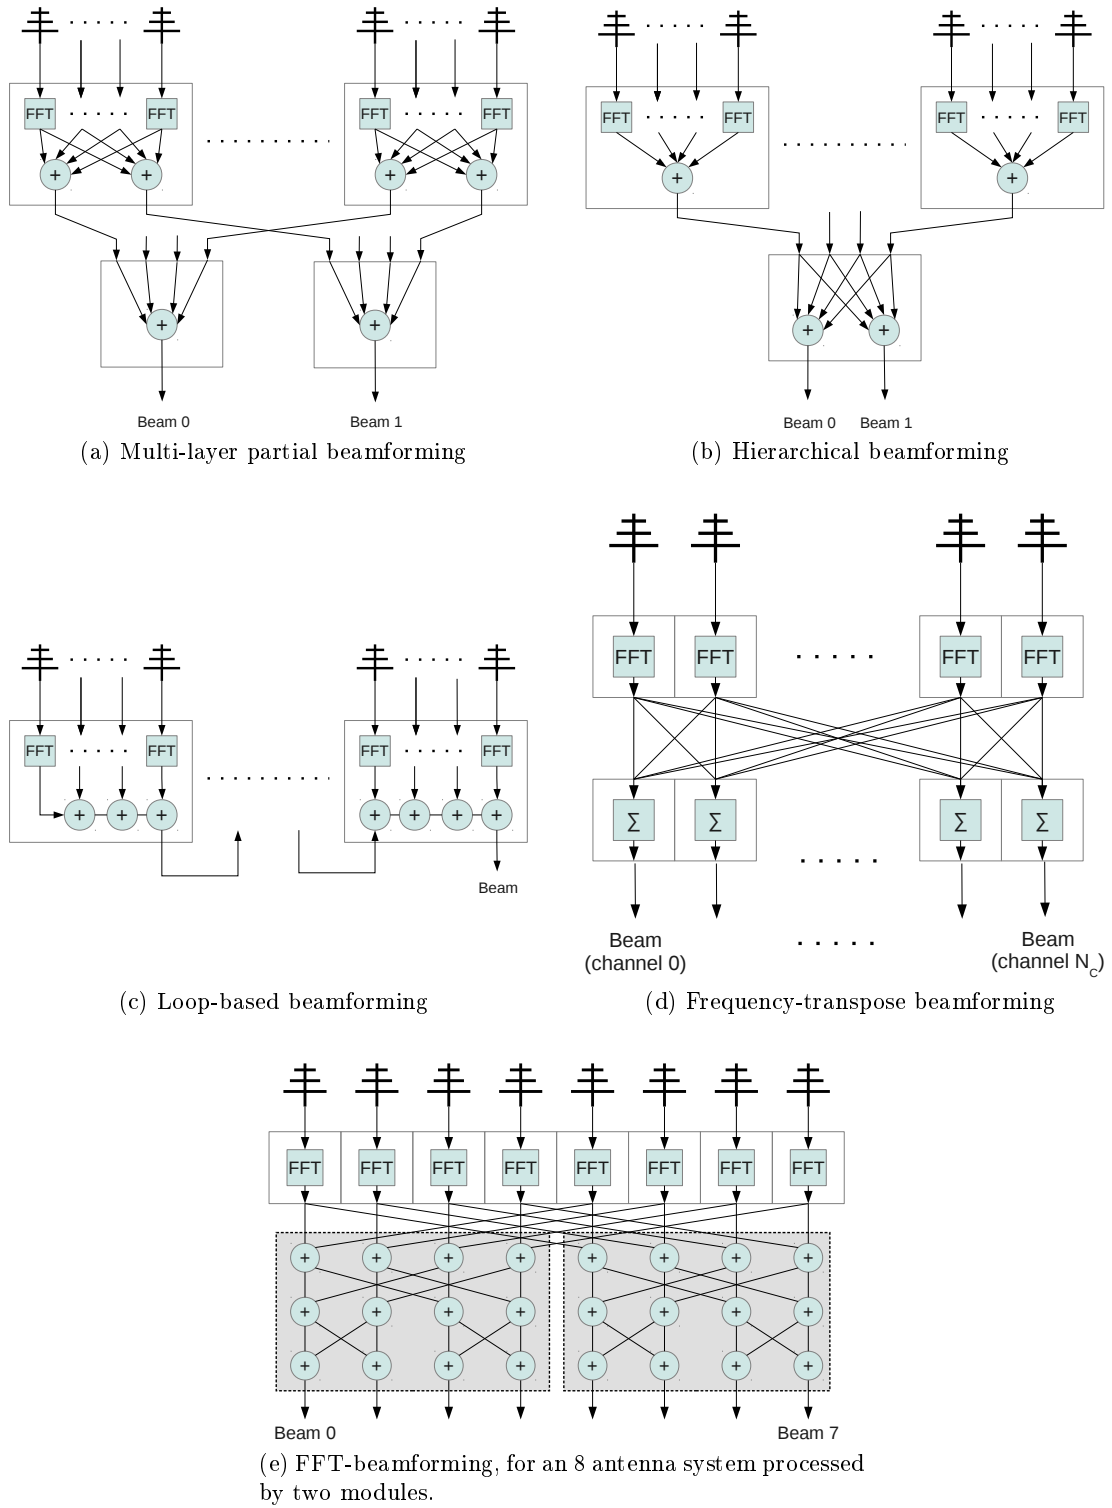


Figure 6.1: Five potential beamforming architectures.

### Multi-layer, partial beamforming

The beamforming process—that is, calculation of the matrix-vector multiplication  $\mathbf{M}\mathbf{x}$  (Equation 1.8)—involves multiplying a vector of  $N$  antennas by an  $b \times N$  matrix of coefficients to form  $b$  beams. It thus incurs a computational cost of  $N$  CMACs per beam. Partial beamforming involves splitting this process into a number of stages. At each stage, a subset of the entries of  $\mathbf{x}$  are multiplied by their respective beamforming coefficients. When the process is broken down into  $S$  stages the complete computation takes the form of  $S$  smaller multiplications of a  $b \times \frac{N}{S}$  matrix by an  $\frac{N}{S}$  length vector. By suitable choice of the coefficients in  $\mathbf{M}$ , the partial beamforming process allows  $b$  completely independent beams to be formed, and is shown for the example case of two layers of processing and two beams in Figure 6.1a.

The result of each matrix-vector multiplication in a beamforming layer is a length  $b$  vector. For this reason, partial beamforming has the significant disadvantage that if the number of beams required to be formed is greater than the number of input signals at a top-level node, the data rate output from each node at this level will be greater than the input rate. Partial beamforming in the layered form shown in Figure 6.1a is therefore unfeasible for large arrays requiring vast numbers of beams, but has been successfully used in a number of astronomical instruments, for example, the Dual-polarization All Digital Aperture Array demonstrator, 2-PAD (Armstrong et al., 2009).

### Hierarchical beamforming

Hierarchical beamforming is mathematically distinct from the partial beamforming alternative, and involves breaking the main array to be beamformed into a series of sub-arrays, usually referred to as tiles. One or more beams is formed at each sub-array, before the different sub-arrays are combined by a higher-level beamformer, as if they were single antennas. Hierarchical beamformers allow a reduction of field-of-view (and, therefore, data rate) at every level of processing by forming multiple array beams from a single tile beam. The hierarchical beamforming process, for the example case of two levels of hierarchy is shown in Figure 6.1b. In this system, two station beams are formed from a single tile beam, effectively halving the output bandwidth of the first layer of processors when compared to the partial-beamforming system. The penalty for this bandwidth reduction is the loss of field of view within which array beams may be formed. In this example the two output beams are independent, but are constrained by the field of view of the tile beams formed in the first layer of processing. An example of the effects of such an envelope function corresponding to a one-dimensional array comprising 16 regularly-spaced antennas in 4 tiles is shown in Figure 6.2. When forming a beam away from the tile pointing center, the hierarchical beamforming




Figure 6.2: Hierarchical versus partial beamforming for the case of a one dimensional array of  $16 \frac{\lambda}{2}$  spaced antennas. Two beams are formed, one at zenith, and one at an altitude of  $80^\circ$ . In the hierarchical case, the array is subdivided into 4 tiles.

process results in significant loss of forward gain, as well as degraded sidelobe performance. These issues can be somewhat alleviated by suitable distribution and gain-weighting of antennas and tiles (see Section 2.5.1) but remain a limitation of the hierarchical beamforming process.

As well as interconnection bandwidth improvements, if more than one beam is required hierarchical beamforming is also more computationally efficient than partial beamforming. The computational cost of an  $H$ -level hierarchical beamformer can be found by considering summing the costs of the beamforming processes at each level of the hierarchy. In CMACs, this cost is given by

$$b_0 N + \sum_{i=1}^{H-1} \left\{ \left( \prod_{j=0}^i b_j \right) T_{i-1} \right\}, \quad (6.1)$$

where  $b_i$  is the number of beams formed at level  $i$  of the hierarchy, and  $T_i$  is the number of tiles in level  $i$  of the hierarchy. The total number of beams produced by such a system is  $\prod_{i=0}^{H-1} b_i$ .

### Ring-based beamforming

Though not a mathematically distinct processing mechanism, a ring-based beamformer, shown in Figure 6.1c, is an implementation with no hierarchical processing structure, but rather a serial chain of addition stages whereby each processing node contributes to the

formation of a complete beam. A ring-based beamformer is used in the Low Frequency Array (LOFAR, Lubberhuizen & Kooistra, 2007; de Vos et al., 2009). The architecture offers an efficient data transport mechanism, which requires no external data switching infrastructure and minimizes required hardware. However, it relies on the fact that the chain of processing nodes have sufficient I/O bandwidth to connect to a single bus which carries the entire output bandwidth of the beamformer. For this reason, the architecture is most suited to beamformers where the output data rate is only a moderate fraction of the input bandwidth of each processing node.

### Frequency Parallelization

The beamforming process can be trivially parallelized by distributing beamforming of different frequency channels across processors. This method, depicted in Figure 6.1d, involves first channelizing input time-domain signals, and then performing a transpose, or corner-turn, operation, whereby signals from all antennas are grouped into common frequency channels. The corner-turn operation can be performed in a variety of manners. In the simplest case, a custom mesh of interconnect can be used to provide all-to-all connections between input and output nodes. Alternatively, corner-turning can be accomplished by a butterfly arrangement of multiplexers (Urry, 2000; Lutomirski et al., 2011), or by utilizing general-purpose routing hardware such as commercial off-the-shelf Ethernet switches (McMahon et al., 2007).

### FFT beamforming

FFT beamforming uses the FFT-based imaging techniques described in Chapter 2 to form beams on the sky. It incurs a computational cost of  $O(N \log N)$  operations to form  $N$  beams on the sky. The only difference between an FFT-beamformer and a FFT direct-imager is that the output of the former is complex-valued and unintegrated, whilst the output of the latter is an integrated power. FFT beamforming is mathematically equivalent to partial beamforming, under the constraint that the beamforming coefficient matrix  $\mathbf{M}$  takes the form of an  $N \times N$  matrix of Fourier coefficients. In this case, the formed beams point at regularly spaced intervals on the sky.

Large FFT beamforming systems can potentially be implemented efficiently by using the architecture shown in Figure 6.1e. Here, the butterfly architecture of the FFT is spread across multiple processing nodes, with inter-node communication limited to the early stages of the FFT. This is identical to the architecture proposed by Morales (2011) for the MOFF correlator. Potentially, subdividing processing in this way can allow large amounts of partially formed beams to be produced without the associated interconnection data-rate issues associated with partial beamforming.

The main drawback of the FFT beamforming system is the requirement for all independent beams to be formed, even if only a small number of beams are required for further processing. Furthermore, as in the direct-imaging case, FFT beamforming either requires antennas to lie, or be digitally manipulated onto, a regular grid. Where arrays are sparse, significant computational penalties result. Beam tracking is not trivially implemented in an FFT beamformer, since there is little flexibility available to modify the beamforming coefficients. In the FFT beamforming case, tracking can only be accomplished either by manipulating virtual antenna positions, interpolating between fixed beam positions or simply choosing the closest beam available to the pointing direction of interest. Choosing one of the first two strategies means incurring significant computational costs, whilst the third strategy (as adopted in Chapter 4) results in non-optimal and time-dependent sensitivity towards the direction of interest, which must be compensated for in post-processing.

### Mixed Architecture Beamforming

None of the beamforming mechanisms mentioned in this section is mutually exclusive. It is frequently the case that a combination of these methods are combined to meet the processing and interconnect limitations of an instrument. For example, at a station-level, the LOFAR beamformer utilizes (in some operation modes) two-layer hierarchical beamforming, using an initial analogue sum of antenna signals, followed by a downstream ring of processor nodes which form a serialized partial-beamforming system.

## 6.2 SKA-low-1 Baseline Specifications

We now proceed to analyse the potential computational cost of SKA-low-1, considering some of the beamforming architectures described in the previous section. The key SKA-low-1 specifications, as originally stated by Dewdney et al. (2010), and later redefined by Dewdney (2013) are shown in Table 6.1. We shall, in the remainder of this chapter, refer to these configurations as SKA-low-2010 and SKA-low-2013, respectively.

### 6.2.1 Specification Motivations

The specifications for SKA-low-2013 are motivated by the main science goal of the low frequency SKA, namely study of the highly redshifted 21cm HI line. Key design decisions have largely followed the recommendations of the SKA-EoR science working group (Mellema et al., 2012). These are (Dewdney, 2013):

**Frequency Range:** EoR study requires observations at redshifts  $25 \lesssim z \lesssim 5.5$ , setting a minimum frequency coverage of 54–215 MHz. A secondary desire for the SKA to be

		Dewdney et al. (2010) (SKA-low-2010)	Dewdney (2013) (SKA-low-2013)
Lower Frequency	$f_{low}$	70 MHz	50 MHz
Upper Frequency	$f_{high}$	450 MHz	350 MHz
Bandwidth	$B$	380 MHz	300 MHz
Antennas per Station	$N_A$	11200	289
Number of Stations	$N_S$	50	911
Total Antennas	$N_{TOT}$	560 000	263 279
Antenna Area	$A_{el}$	2.27 m <sup>2</sup>	2.25 m <sup>2</sup>
Station Diameter	$D$	180 m	35 m
Station filling factor	$f = \frac{N_A A_{el}}{\pi(\frac{D}{2})^2}$	1	0.7
Total Physical Aperture	$A_{phys}$	$1.3 \times 10^6$ m <sup>2</sup>	$0.8 \times 10^6$ m <sup>2</sup>
Channel resolution (maximum)	$RBW$	1 kHz	1 kHz
Number of channels	$N_C$	380 000	250 000
Number of Beams	$N_B$	480	1

Table 6.1: Key specification of SKA-low, as defined by Dewdney et al. (2010) and Dewdney (2013).

able to observe HI absorption from the redshifts  $0 \lesssim z \lesssim 20$  further motivates an upper frequency limit of SKA-low which meets the lower end of the SKA mid-frequency array, currently specified as 350 MHz.

**Sensitivity:** Mellema et al. (2012) recommend 108 MHz as the optimum frequency for maximum sensitivity, implying that the dense-sparse transition of the array should be at approximately this frequency. Significant detection of the EoR signature in a reasonable observing time requires a high filling factor array. At EoR resolutions, this calls for an array of  $\sim 1$  km in size, implying a total collecting area  $\sim 10^6$  m<sup>2</sup>.

**Polarization:** Dual-polarization capability is required, mainly as it aids the removal of polarized foreground sources from the unpolarized EoR signature.

**Field of View:** Spatial structure of the EoR signature should be detectable on scales up to several degrees. In order to detect structure at these scales, single beam observations with instantaneous field of view of  $\sim 5$  degrees across (at 108 MHz), have been specified. This requirement, together with the dense-sparse transition frequency, imply array stations approximately 35 m in diameter. This requirement is briefly discussed below.

**Array Configuration:** Given the array station size of 35 m, the required sensitivity determines the number of stations which are needed to reach this target. This area needs to be concentrated in a dense core to achieve sensitivity to large scale structure, whilst a significant portion of collecting area need also be spread over a core radius of  $\sim 3$  km to

facilitate studies at  $\sim 5$  arcminute scales. In addition, some collecting area is required at longer baselines, in order to aid calibration (which requires mapping of unresolved point sources) and model direction-dependent effects.

**Sky Coverage:** Beams should be steerable, in order that fields appropriate for EoR observation be accessible, and in order to facilitate secondary pulsar and HI absorption science goals.

The most radical difference between SKA-low-2013 and SKA-low-2010 is the requirement that large fields of view should be accessible with single beams. This has dramatically reduced the size of stations in the proposed design, resulting in a corresponding growth in quantity of stations to maintain the target collecting area and filling factor. The argument made for large beams is that information on scales larger than the beam size is lost, and cannot be recovered by mosaicing multiple fields of view.

VCZ (Equation 2.12) expresses the fact that an interferometer measures a Fourier mode of the apparent sky, modulated by the beam shape of the stations in the interferometer. Thus, in the  $uv$  plane, an interferometer measures a single point, with a point spread function that is equal to the Fourier transform of the beam pattern (see for example, Campbell, 2002). Since the beam shape of a station is equivalent to the power of the aperture's Fourier transform, the transform of this beam is just the autoconvolution of the aperture distribution of a station. The example case of an interferometer formed of two circular stations of diameter  $D$  and separated by a baseline  $b$  is shown in Figure 6.3. Such a configuration has maximum sensitivity to scales at  $u = b$  (which is measured by the cross correlation) and  $u = 0$  (which is measured by each station's autocorrelation). The autocorrelation of an interferometer is usually discarded, and the shortest baseline between two non-overlapping stations is  $b \leq D$ . This results in a minimum detectable spatial scale with  $u = b$ , although some sensitivity extends to scales of approximately  $u = b \pm D$ .

The argument used by Mellema et al. (2012) to justify a station size of 35 m is based on the fact that the beam size should be comfortably larger than largest scales of interest, which are around  $2^\circ$ . A station with a beam width of  $\sim 5^\circ$  at the frequency of interest was then chosen. However, we see that it is actually the shortest baseline (and thus the approximate size of a station) which needs to match this scale. With stations effectively touching, the minimum baseline and maximum station size would have a size around twice that proposed by the EoR working group. Such a configuration still has significant sensitivity to scales larger than this size.

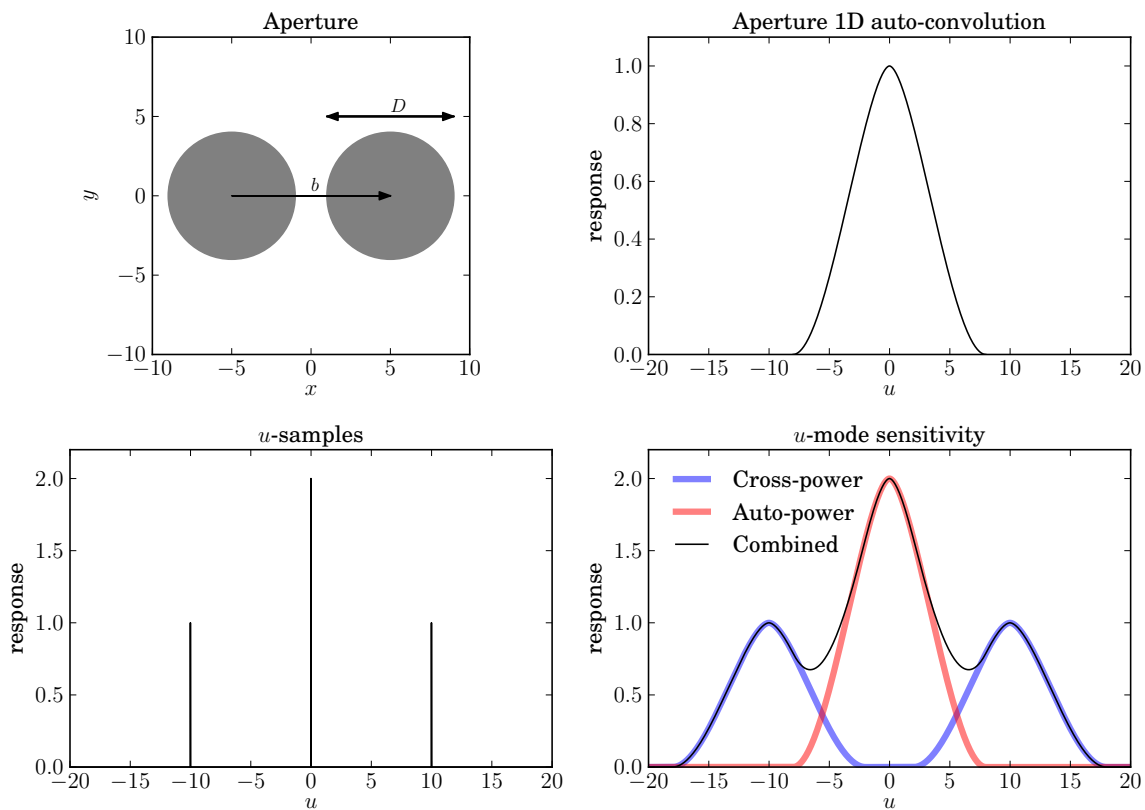


Figure 6.3: Top left: A two element interferometer, formed from two circular apertures with diameter  $D = 8$ , separated by a baseline,  $b = 10$ . Top right: The one-dimensional auto-convolution function of a circular aperture with  $D = 8$ . Bottom left: The  $u$ -coverage of a two element interferometer with baseline  $b = 10$ . Bottom right: The relative sensitivity, in  $u$ -space, of such an interferometer to different angular scales in the sky temperature distribution. Interferometers invariably discard the contribution from the auto-power spectra, causing a loss of  $u$ -sampling at  $|u| < D$ .

### 6.3 A Computational Argument for Station Configuration

A concern raised by Mellema et al. (2012) is that, whilst smaller stations may give more power spectrum sensitivity at large scales, increasing the number of stations requiring correlation may have serious downsides in terms of computational cost. In this section we examine the effect of station size on the computational demands of SKA-low-1. Given a fixed number of antennas in an array and a requirement for a certain instantaneous field of view, a straightforward argument can be made for the optimal assignment of antennas into multiple stations. We proceed to lay out this argument here, by considering the number of operations required by the station beamformers, and the central correlator in such an array. The work in this section was carried out with Danny Price, and is currently under review for publication in the SKA memo series.

We first consider an array with  $N_{TOT}$  antennas, which require division into  $N_S$  stations of  $N_A$  antennas each. We consider a general array capable of detecting  $N_P$  polarizations, where  $N_P$  may be equal to 1 or 2. We parameterize the required instantaneous field-of-view of the array as  $\Omega_A$ , the physical antenna size as  $A_{el}$ , and the station filling factor as  $f$ . For any given observing frequency, the field of view of a single beam,  $\Omega_B$ , from such a station is inversely proportional to the area of the station, and hence the number of antennas it contains. For stations of diameter  $d$ , the field of view at wavelength  $\lambda$  is (up to a numerical factor of order unity depending on the station shape, internal antenna distribution and any applied apodization)

$$\Omega_B = \frac{\pi}{4} \left( 1.3 \frac{\lambda}{d} \right)^2 \quad (6.2)$$

$$= 1.3^2 \frac{\pi^2}{16} \frac{f}{N_A A_{el}} \lambda^2, \quad (6.3)$$

where the second line follows since  $\frac{N_A A_{el}}{f}$  is the station area, and therefore  $\pi \left(\frac{d}{2}\right)^2 = \left(\frac{N_A A_{el}}{f}\right)$ . We make the definition  $\Lambda = 1.3^2 \frac{\pi^2}{16 A_{el}}$ , so that

$$\Omega_B = \Lambda \frac{\lambda}{N_A}. \quad (6.4)$$

Given the requirement that the array have an instantaneous field of view  $\Omega_A$ , each station in the array must generate  $N_B$  beams where

$$N_B = \frac{\Omega_A}{\Omega_B} = \frac{\Omega_A N_A}{\Lambda \lambda}. \quad (6.5)$$

The number of computations<sup>1</sup> required to generate these beams at  $N_S$  stations is

<sup>1</sup>Here we count in complex multiply accumulate (or CMAC) operations. A further implementation-specific factor is required to compute the real operation count.

$$C_B = N_S N_A N_P N_B = N_P N_{TOT} \frac{\Omega_A N_A}{\Lambda \lambda}, \quad (6.6)$$

where the last equality follows using Equation 6.5, and noting that  $N_A N_S = N_{TOT}$ .

Given a correlator with inputs from  $N_S$  stations, the number of computations required to calculate all visibilities (including autocorrelations) for  $N_B$  beams is, per unit bandwidth,

$$C_X = \frac{1}{2} N_B N_P^2 N_S (N_S + 1) = \frac{1}{2} \frac{\Omega_A N_A}{\Lambda \lambda} N_P^2 N_{TOT} \left( \frac{N_{TOT}}{N_A} - 1 \right). \quad (6.7)$$

For simplicity, we consider channelization operations to be independent of array configuration. In this case the total computation rate is given by summing the contributions from beamforming and correlation operations, yielding

$$C_{BX} = C_B + C_X = \frac{1}{2} \frac{\Omega_A N_A}{\Lambda \lambda} N_P^2 N_{TOT} \left( \frac{2N_A}{N_P} + \frac{N_{TOT}}{N_A} - 1 \right). \quad (6.8)$$

Differentiating Equation 6.8 reveals that  $C_{BX}$  is minimized for

$$N_A = N_A^{opt} = \sqrt{\frac{N_{TOT} N_P}{2}}. \quad (6.9)$$

For dual-polarization arrays, computational efficiency is maximized for  $N_A = N_S = \sqrt{N_{TOT}}$ , whilst for single polarization arrays minimum computational load is required when  $N_S = \sqrt{2} N_A$ .

Equation 6.8 is plotted in Figure 6.4 for SKA-low-2013-like specifications (Table 6.1), with  $\Omega_A = 20$  square degrees. The number of antennas has been fixed at  $N_T = 263\,279$ , but  $N_A$  has been allowed to vary. Results are shown for both the case where the number of beams is allowed to be fractional (or, equivalently, the top level array specifications result in a large number of beams at all frequencies), and also the case where the number of beams is required to be integral. Results have been scaled by an array bandwidth of 300 MHz to obtain an operation rate in CMACs per second, and are shown for the maximum and minimum frequencies in the SKA-low-1 band. Notably, the optimal configuration is independent of frequency, except through discontinuities in beam numbers at different frequencies.

The analysis presented in this section serves only as a starting point for considering array configurations. Non-computational considerations, for example interconnect bandwidth and memory capacity requirements, may drive final configuration specifications. It is also the case that some computations, such as low-bitwidth correlation operations, can be computed with far greater efficiency than general wide-bitwidth multiplications (see Chapter 4). Furthermore, heterogeneous computing platforms, if deployed in the SKA, might cause similar operations to have different cost weightings when performed on different hardware. All these considerations make it clear that the minimum of Equation 6.8 does not necessarily describe




Figure 6.4: The computational cost associated with an SKA consisting of  $\sim 250,000$  antennas, operating at 50 – 350 MHz and required to have a field of view of 20 sq. degrees. Left: An evaluation of Equation 6.8. Right: The same, with the number of beams forced to be an integer, yielding at least the required field of view. Dotted, dashed, and solid lines show the beamforming, correlation and total operation rates, respectively. At low beam numbers, this results in step functions in cost where more field of view is processed than is required. The number of antennas per station specified in SKA-low-2013 and SKA-low-2010 are indicated by vertical lines.

the preferred station configuration in terms of engineering constraints. However, Figure 6.4 makes it clear that to simply assume that an SKA with fewer stations and more beams is computationally easier to build is not correct. With the number of stations significantly fewer than the numbers of antennas in each, the computational cost of beamforming quickly dominates correlation. In particular, normalizing to the case of case of 263 279 antennas (as per SKA-low-2013) an SKA comprising 279 antennas per station may be almost an order of magnitude easier to build in terms of raw computation than the previous specifications of 11 200 antennas per station.

### Alternative Beamforming Methods

We can repeat the above analysis for the other beamforming architectures described above. For the hierarchical beamforming case, we consider tiles consisting of 16 antennas each. In this case, the minimum field of view of the tile beam is  $\sim 25$  square degrees even at 350 MHz. We consider the most computationally efficient but least flexible case where only a single tile beam is formed. For the FFT-based beamforming system, we use the filling factor of 0.7 quoted by Dewdney (2013)<sup>2</sup> and the arguments presented in Section 2.5.1 to determine a number of grid points required by a beamformer. We assume that no real-

<sup>2</sup>This is the geometric filling factor, based on the physical (rather than effective) antenna sizes. However, it is representative of the effective filling factor at  $\sim 111$  MHz (Dewdney, 2013), and is used here to provide a conservative costing.

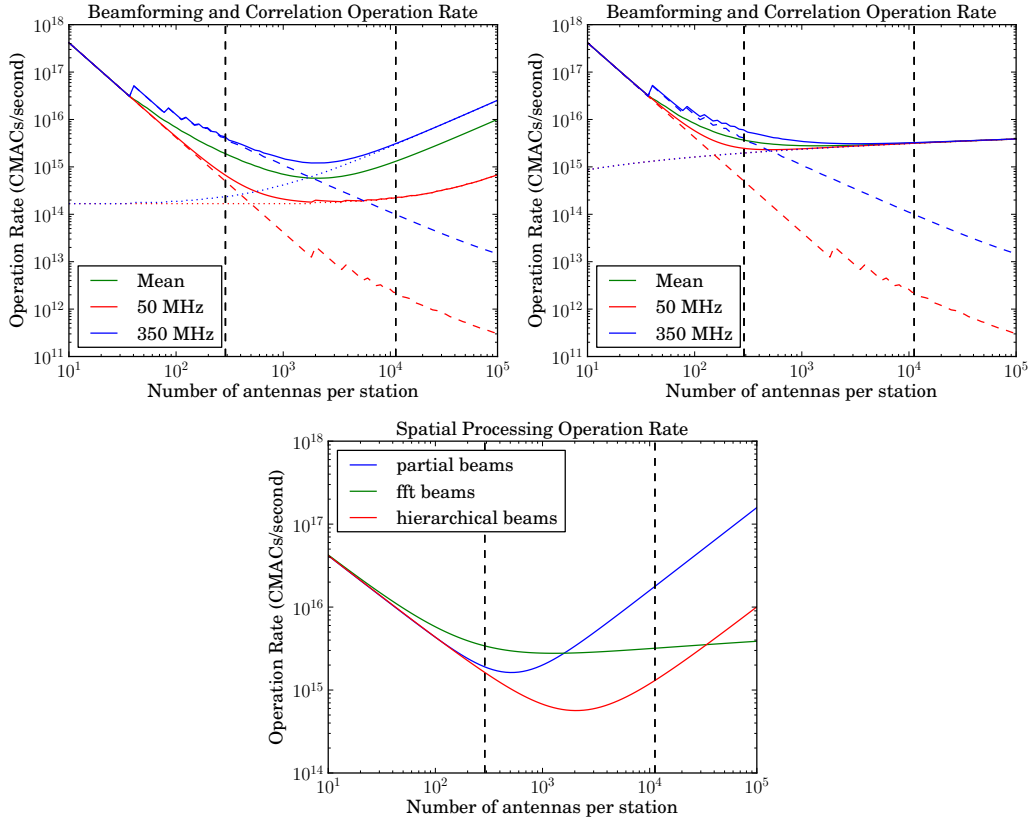


Figure 6.5: Hierarchical beamforming, with tiles of 16 antennas (top left). FFT beamforming without gridding (top right). As in Figure 6.4 dotted, dashed, and solid lines show the beamforming, correlation and total operation rates, respectively. A comparison of arrays formed with three beamforming strategies (bottom) for the mean number of beams each system generates.

time interpolation of antenna samples is required, though this will potentially increase the computed computational costs.

The computational costs of these described hierarchical and FFT-based beamforming arrays are shown in Figure 6.5, for different array configurations. Also shown is a comparison of these two systems with the originally derived partial-beamforming array, averaged over the frequencies at which the array operates.

Both hierarchical and FFT beamforming systems result in a reduced array-wide computational cost as the station size gets larger and the overall computational budget becomes beamformer dominated. Below  $\sim 1000$  antennas per station, FFT beamforming is the least efficient strategy. This corresponds to station sizes of  $\sim 70$  m which is twice that proposed by Mellema et al. and approximately the maximum appropriate for EoR science (see Section 6.2.1). It seems unlikely that stations larger than this will be accommodated in SKA-low-1 unless a strategy of variable station sizes is adopted. For this reason, and in conjunction

with the limitations of FFT-beamformers with regards to tracking and potentially gridding, FFT-beamformers are unlikely to play a useful role in SKA-low-1.

Hierarchical beamforming is always more efficient than partial beamforming whenever more than one beam is formed. For SKA-low-2013, this difference is negligible, but at stations of  $\sim 1000$  antennas, the total computational cost of a partial beamforming system is greater by a factor  $\sim 4$ . However, in this case the beams are fully independent and are not confined to a single field defined by the tile beam size.

### Single Beam Stations

We conclude this section with a comment on arrays where field of view is not required to be constant over frequency. This is, in fact, not explicitly required by SKA-low-2013 specifications, which rather specify that there will be only one beam formed per station, over the entire observing band. In this case, the computational cost is shown in Figure 6.6. Since only a single beam is processed,  $N_A$  effectively becomes a tracer of processed field of view, via Equation 6.3. At large  $N_S$  (i.e. small  $N_A$ ), computation is predictably correlator-dominated, however, a lower limit exists where processing is beamformer-dominated. This limit is given by Equation 6.6 and is simply (normalized to unit bandwidth)  $N_{TOT}N_P$ . It is not a significant function of beamforming architecture. In this sense,  $N_{TOT}N_P$  sets an absolute minimum to the amount of computation that is required in an antenna array and makes it futile to seek to reduce the computational cost of an array merely by reducing numbers of stations and processed field of view.

#### 6.3.1 Channelization

The above exploration of the relative computational cost does not take into account the requirements of channelization of input signals. This is because, in the most basic scenario (which we term “single pass channelization”), when each antenna’s data stream is split into  $N_C$  channels, the computational cost is independent of the subdivision of antenna groups into array stations. Explicitly, the number of complex operations per second required for channelization of all data streams is

$$C_C = B \times N_{TOT}N_P (\log_2(N_C) + N_{taps}) . \quad (6.10)$$

The factor  $N_{taps}$  has been introduced, representing the number of taps in a (presumed) polyphase filter-bank channelization implementation (see Section 4.4.3). In the case of SKA-low-2013, under the assumption that  $N_{taps} = 4$ , the total channelization cost<sup>3</sup> is

---

<sup>3</sup>We ignore the minor implication that the number of channels in an FFT-based channelizer cannot be arbitrarily specified, and for instance, will almost certainly be required to be a power of 2 or 4.




Figure 6.6: The computational cost associated with an SKA consisting of  $\sim 250,000$  antennas, operating at 50 – 350 MHz and processing one beam per station at all frequencies. The number of antennas per station specified in SKA-low-2013 and SKA-low-2010 are indicated by vertical lines.

$C_c = 2.9 \times 10^{15}$  CMACs/s. In this case,  $C_C$  represents a significant contribution to the overall computational requirements of the array and in fact dominates the cost below  $\sim 150$  MHz. Overall, however, the computation budget is still most impacted by the configuration-dependent processing costs at the array’s higher frequencies, even around the computationally optimal array layout.

Another channelization strategy is to coarsely channelize individual antenna streams, and then only finely channelize beam data streams, of which there are far fewer. In this case (which we term “hybrid channelization”) the computational cost of channelization can be split into a coarse, per-antenna contribution,  $C_{C-coarse}$ , and a fine, per-beam contribution,  $C_{C-fine}$ .  $C_{C-coarse}$  is given by Equation 6.10, with  $N_C$  replaced by an appropriate number of channels for coarse channelization,  $N_{C-coarse}$ :

$$C_{C-coarse} = B \times N_{TOT} N_P (\log_2(N_{C-coarse}) + N_{taps}) . \quad (6.11)$$

The beam channelization contribution is

$$\bar{C}_{C-fine} = B \times \bar{N}_B N_S N_P (\log_2(N_{C-fine}) - \log_2(N_{C-coarse}) + N_{taps}) . \quad (6.12)$$

Here, for simplicity, rather than considering  $C_{C-fine}$  as a function of frequency, we take into account the variation of number of beams at different frequencies by considering the mean number of beams across the band of interest,  $\bar{N}_B$  (shown in Figure 6.7). Throughout

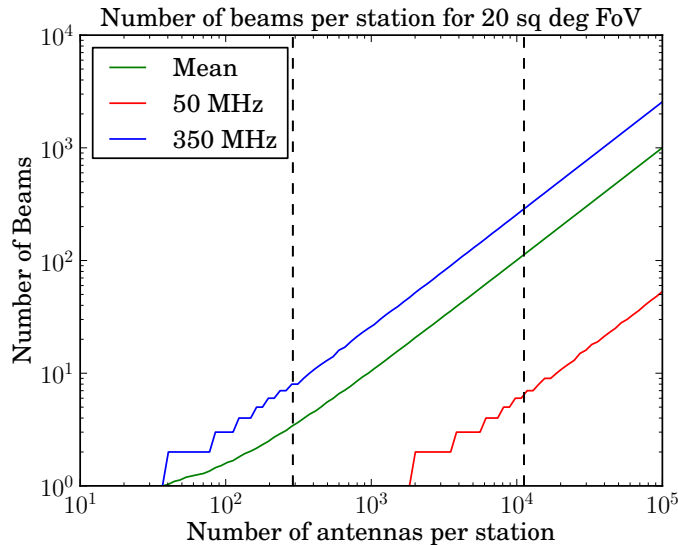


Figure 6.7: Number of beams required to cover a 20 sq. degree field of view.

this chapter we will use a bar to denote quantities which have been computed for the mean number of beams.

### Coarse channel width

A maximum bandwidth of antenna-level frequency channels,  $RBW_{coarse}$ , is set by the requirement that beamforming of individual channels is able to be achieved with a single phase multiplication. This condition can be expressed by demanding that  $\frac{1}{RBW_{coarse}}$  is shorter than the difference in arrival times of radiation at different elements of the array. Explicitly:

$$RBW_{coarse} = k \frac{c}{D}, \quad k \ll 1, \quad (6.13)$$

where  $c$  is the speed of light. For the case of SKA-low-2013,  $\frac{c}{D} \approx 8$  MHz. This, and the further requirement that frequency-dependent gain variations in antenna signal chains should be able to be calibrated accurately, leads to the requirement that coarse channelization should be carried out to  $O(100)$  kHz, or  $\sim 2500$  channels. Maximum coarse channel width as a function of station size is shown, for  $k = 0.01$ , in Figure 6.8.

$C_{C-coarse}$  and  $\bar{C}_{C-fine}$  are only weak functions of the number of channels generated. Rather, they are dominated by the number of antenna and beam streams, respectively, requiring processing. In any array where the field of view to be processed is significantly less than a full hemisphere of sky,  $N_{TOT} \gg N_S \bar{N}_B$ , or equivalently,  $N_A \gg \bar{N}_B$ . Thus the computational demands of coarse channelization always dominate downstream rechannelization costs.




Figure 6.8: Maximum coarse channel width required prior to beamforming, based on Equation 6.13 with  $k = 0.01$ .

$C_C$ ,  $C_{C-coarse}$ ,  $C_{C-fine}$ , and  $C_{C-hybrid} = C_{C-coarse} + C_{C-fine}$  are plotted in Figure 6.9 for  $RBW_{coarse}$  defined by Equation 6.13 with  $k = 0.01$ . Coarse channelization dominates the hybrid channelization costs by over an order of magnitude for all station configurations.

### Memory Use

The memory required to perform an FFT on a data stream scales in proportion to length of the transform and the number of taps used in a potential polyphase filter frontend. The RAM required to implement a coarse channelization on each of  $N_P N_{TOT}$  antennas is, in bytes,

$$R_{C-coarse} = N_{TOT} N_P N_{C-coarse} N_{taps} n_c, \quad (6.14)$$

where  $n_c$  is the number of bytes used to store a complex-valued word in the transform. Similarly, the memory required to form a fine channelization is

$$\bar{R}_{C-fine} = \bar{N}_B N_S N_P N_C N_{taps} n_c. \quad (6.15)$$

Note that  $\bar{R}_C$  scales with  $N_C$  and *not*  $N_{C-fine}$ .

$R_{C-coarse}$ ,  $R_{C-fine}$ ,  $R_{C-hybrid} = R_{C-coarse} + R_{C-fine}$  are plotted in Figure 6.10 for  $n_c = 4$  bytes, and  $RBW_{coarse}$  defined by Equation 6.10 with  $k = 0.01$ . The RAM required for single pass channelization,  $R_C$ , is also shown.

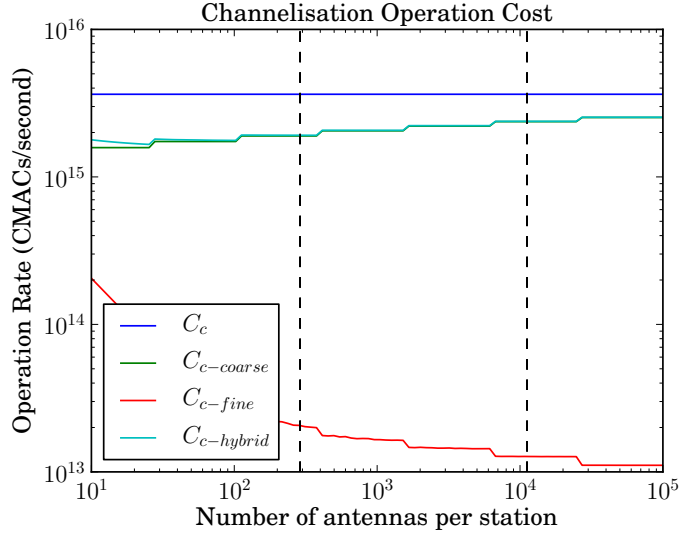


Figure 6.9: Computational costs associated with channelization. The number of frequency channels has been required to be an integer power of 2.

Memory requirements are driven by coarse channelization for  $N_A \gtrsim 200$  antennas/station, with a lower bound set by the second transform stage. SKA-low-2013 requires around 600 times more RAM for single pass channelization compared to two stage.

### 6.3.2 Summary

Figure 6.11 shows the total computational cost of potential SKA-Low-1 configurations utilizing partial, hierarchical and FFT-based beamforming algorithms. Two-stage channelization costs have been included. This result, together with the analysis that preceded it in this section, leads to a number of conclusions. These are summarized below:

1. In the above analysis we considered an SKA-low array with a fixed field-of-view requirement. This opposes the SKA-low-2013 specification, which specifies a single beam across the entire frequency band. We find that a single-beam SKA-low-2013 will have a beamforming and correlation computational cost of fewer than  $10^{15}$  CMACs. This is lower than the computational requirements for channelizing data from the array, suggesting that more beams could be accommodated with relatively little increase in total computation. For 20 square degree coverage at all observing frequencies we find an SKA-low-2013 station requires an average of  $\sim 4$  beams across the observing band. This moderate adjustment would allow equal fields-of-view to be observed at all frequencies.
2. An idealized analysis of an array based on correlation of digitally synthesized beams suggests that, for a dual-polarization array, the number of stations and antennas per




Figure 6.10: Memory required to channelize signals in potential SKA-low configurations.




Figure 6.11: Computational costs associated with three different beamforming architectures with two-stage channelization costs included.

station should be equal for optimal computational efficiency. This optimal point shifts in favour of larger stations when more efficient methods of beamforming are used. However, channelization costs, which are only weakly dependent on station size, reduce the relative importance of beamforming algorithm efficiency at large station sizes. For this reason partial beamforming should be preferred up to station sizes of at least 1000 antennas, since it offers the most flexible beamforming options.

3. Given the flexibility benefits which come with having large stations of multiple, independent fields-of-view, and the shallowness of the optimal computational point, a reasonable target station size is  $\sim 1000$  antennas. This is approximately four times larger (in antenna count) than the size suggested by SKA-Low-2013, but still gives power spectrum coverage at the scales demanded by the EoR science-case.
4. There is relatively little to be gained computationally by performing two-step channelization of antenna and beam signals. However, the memory savings are significant, which potentially make this a sensible processing configuration, particularly for platforms which possess only relatively small quantities of high-bandwidth memory.

## 6.4 An FFT-based correlator

This thesis has explored the ability to obtain correlation matrices using efficient FFTs, rather than computationally costly cross-multiplication. Here we apply the calculations of Section 2.6 to the SKA-low-2013 specifications to evaluate the potential of an FFT-based direct-imaging correlator. Since we have seen that low filling factor arrays can undermine the potential efficiency of FFT-based imaging here we divide the SKA-low array into two conceptually distinct parts. A core<sup>4</sup> ( $\text{SKA}_{\text{core}}$ ), which contains all antennas within a radius  $R_{\text{core}}$ , and an extended array ( $\text{SKA}_{\text{ext}}$ ), which encompasses all stations outside this radius. The distribution of stations used is shown in Figure 6.12. Under this subdivision, the two-part SKA correlator has three components.

1.  $\text{SKA}_{\text{core}} \times \text{SKA}_{\text{core}}$
2.  $\text{SKA}_{\text{core}} \times \text{SKA}_{\text{ext}}$
3.  $\text{SKA}_{\text{ext}} \times \text{SKA}_{\text{ext}}$

---

<sup>4</sup>This definition of the core differs from the meaning in Dewdney (2013), where the core is taken to be the entire central 3 km of the SKA-low array (866) antennas, which excludes only the 45 antennas forming spiral arms out to  $\sim 100$  km.




Figure 6.12: The configuration of the core of SKA-low-2013. Reproduced from Dewdney et al. (2010).

Since SKA-low-1 has a filling factor that falls with radius, we consider using an FFT-based imager only for the first correlation contribution. The other two contributions remain are assumed to be computed by traditional cross-multiplication. Given the SKA-low-2013 specifications, it is possible to compare the computational cost of a correlator comprising the three sub-array components above at different sizes of  $R_{core}$ . The results of such an analysis are shown in Figure 6.13, which shows the cumulative number of stations and filling factor of  $SKA_{core}$  as a function of  $R_{core}$  as well as the costs of the three correlation components and their total. The minimum total computational cost is found to be  $\sim 73\%$  of a purely cross-correlation-based correlator implementation, with  $R_{core} \sim 1100$  m.

As with other low-frequency arrays, FFT-based direct-imaging is unlikely to be used in SKA-low-2013 for reasons of computational efficiency. Though a small computational saving is possible, it is likely outweighed by the complication of having a hybrid correlation system, the additional calibration requirements of direct-imaging systems, as well as the potential efficiency increases that FX correlators can obtain by cross multiplication implementations such as that described in Chapter 3. Even though this simple analysis has only been carried out for SKA-low-2013, we note that, since these specifications likely represent the largest quantities of stations, other designs for SKA-low-1 are likely be even less amenable to direct-imaging correlators.

## 6.5 An FFT-based central beamformer

In addition to correlating signals from the SKA-low-1 stations, the SKA Design Reference Mission (SKA Science Working Group, 2011) states that the array will be used for an extensive survey of galactic pulsars. This survey goal, in contrast to the EoR intensity

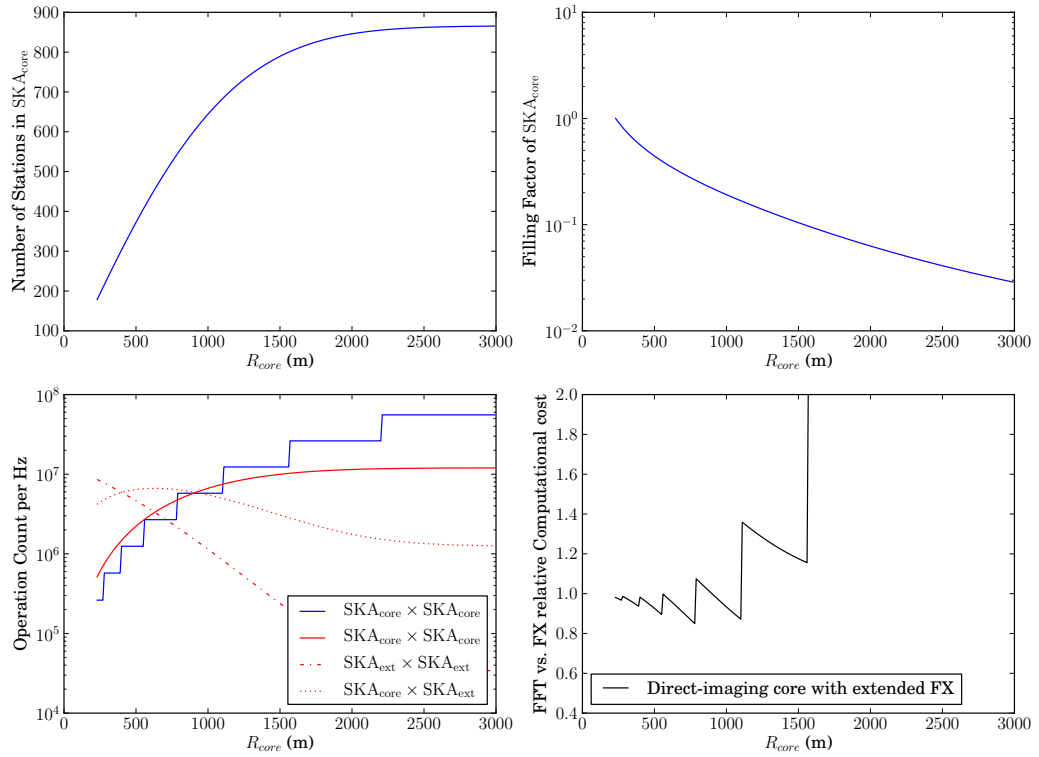


Figure 6.13: Computational cost of a mixed FFT-FX correlator, based on a radially Gaussian distribution of stations estimated from Dewdney (2013) (top left). Number of grid points has been required to be an integer power of 2. Maximum savings of a mixed mode system occur when the central FFT correlator is used to correlate antennas within a radius of  $\sim 1100$  m of the array centre. In this case, the computational load of the complete correlator system is  $\sim 85\%$  of a pure FX implementation.

mapping requirements, mandates a large instantaneous FoV and high time-resolution data. Such data are obtained by forming tied-array beams. These are beams formed by coherent weighting and addition of signals from each station. In essence, this is two-stage hierarchical beamforming of the entire array, with each station playing the role of a single beamformed tile.

In the ideal case, a tied-array beamformer would be capable of completely tiling the field of view of an SKA-low-1 station with tied-array beams. In complete analogy with the arguments for the number of beams required by direct-imaging systems of Chapter 2, the number of beams required for complete coverage of the primary beam is

$$B = \frac{N}{f}. \quad (6.16)$$

Following this fact D'Addario (2013) states that the total computational cost of beamforming is  $\sim \frac{N^2}{f}$  and dwarfs the computational cost of the SKA-low-1 correlator, leading to the conclusion that the number of beams must necessarily be reduced for the system to be practical. Whilst there are other concerns, in particular regarding output data rate, which might motivate the same conclusion, it is certainly not the case that beamforming operations alone need motivate beam reduction. Indeed Figure 6.13, though presented in the context of interferometry, shows that beamforming across the entire primary beam can be accomplished at a lower cost than correlation up to a core radius of  $\sim 1500$  m.

### 6.5.1 The Fractional Fourier Transform

A potential complication of FFT beamforming is that if antenna grid positions are not allowed to vary with frequency (which is only possible in the case of interpolation-gridded antennas) the pointing centres of the formed beams are frequency dependent. This problem can be alleviated using the fractional Fourier Transform (FrFT) to form beams. The FrFT is a generalization of the Fourier transform. In the case of discretely sampled data, the  $N$  point FrFT takes the form

$$F_k(\alpha) = \sum_{x=0}^{N-1} f_x e^{-\frac{2\pi i}{N} kx\alpha}. \quad (6.17)$$

The FrFT gains its name from the coefficients of the transform  $e^{-\frac{2\pi i}{N}\alpha}$ , which are the  $\frac{1}{\alpha}$  roots of unity. In the case  $\alpha = 1$  the FrFT takes the form of a standard discrete Fourier transform.

Crucially, the FrFT can be computed in  $O(N \log(N))$  time, in which case it may be referred to as the fractional Fast Fourier transform (FrFFT). The deconstruction of Equation




Figure 6.14: A schematic illustration of the FrFFT process. Based on Figure 3, Bluestein (1970).

6.17 into a form amenable to calculation using FFTs is obtained by using the identity  $2kx = k^2 + x^2 - (x - k)^2$  to recast Equation 6.17 as (Bluestein, 1970)

$$F_k(\alpha) = e^{-\frac{i\pi\alpha k^2}{N}} \sum_{x=0}^{N-1} f_x e^{-\frac{i\pi\alpha x^2}{N}} e^{\frac{i\pi\alpha(k-x)^2}{N}} \quad (6.18)$$

$$= e^{-\frac{i\pi\alpha k^2}{N}} \sum_{x=0}^{N-1} y_x z_{k-x}, \quad (6.19)$$

where the sequences  $y$  and  $z$  are defined by

$$\begin{aligned} y_i &= f_i e^{-\frac{i\pi\alpha i^2}{N}} \\ z_i &= e^{\frac{i\pi\alpha i^2}{N}}. \end{aligned}$$

In this form, the summation in this Equation 6.18 is revealed to be a convolution. As such, it can be computed via a Fast Convolution algorithm, by utilizing the convolution theorem. The total FrFFT algorithm can thus be performed via the process depicted schematically in Figure 6.14. Given that the multiplicative factors and the Fourier transform of  $z$  can be precomputed, the computational cost of the FrFFT is dominated by the (zero-padded) Fourier transform of  $y$ , and the inverse Fourier Transform required to obtain  $y * z$ . The total cost is therefore approximately twice that of a simple zero-padded FFT.

With the choice  $\alpha = \frac{\lambda_{min}}{\lambda}$ , where  $\lambda_{min}$  is the wavelength of the highest observational frequency, and  $\lambda$  is the wavelength of the channel being beamformed, the FrFFT can be used to generate synthesized beams whose pointing centers do not depend on frequency. In this case the transform can be considered to be the exact frequency-domain equivalent of a set of true-time-delay filters. Figure 6.15 shows the 7 beams nearest to zenith obtained from a 64-element  $\frac{\lambda}{2}$ -spaced linear array by both FFT and FrFFT methods. Beams are shown for the case of observations at wavelengths of  $\lambda_{min}$ ,  $\frac{5\lambda_{min}}{4}$  and  $\frac{7\lambda_{min}}{4}$ . We see that the FrFFT beamformer removes the frequency dependence of beam pointing centres.

We review the use of the FrFFT for beamforming here because it does not appear in radio-astronomy literature, presumably because FFT-based beamformers in general have




Figure 6.15: The 7 beams closest to zenith for a 64-element  $\frac{\lambda_{min}}{2}$ -spaced regular array. Beams have been obtained at three different observing frequencies by use of a Fractional FFT (top) and a standard FFT (bottom)

		Bitwidth (bits)
ADC sampling	$n_s$	8
FFT processing	$n_c$	32 (18 real + 18 imag)
Beams	$n_b$	8 (4 real + 4 imag)
Correlator Output	$n_{corr}$	64 (32 real + 32 imag)

Table 6.2: Assumed bitwidths used in data rate calculations.

not gained much traction in astrophysical applications. However, we note that the use of the FrFFT for beamforming and other physical applications is well-known in other fields (see for example, Curtis, 1998; Bailey & Swarztrauber, 1991).

## 6.6 Data Transport

Of equal importance to the processing requirements of SKA-low-1 are the required data transport rates, both within stations and to a central processing location. In this section we summarize the relationships between station size and data rates assuming, as before, that the top-level array specifications including the total number of antennas are given by SKA-low-2013. The bitwidths of data products at various stages of processing are shown in Table 6.2, and are used in the discussion that follows.

### 6.6.1 Internal Station Data Rates

Assuming digitization of antenna signals with  $n_s$  bit samplers, the input data rate to each SKA-low station is given by

$$D_{S-in} = 2B \times N_A N_P n_s. \quad (6.20)$$

The output rate is determined by the number of beams required and is, for each station,

$$D_{S-out} = B \times \bar{N}_B N_P n_b. \quad (6.21)$$

The correlator input data rate is simply given by  $\bar{D}_{C_i} = N_S D_{S-out}$ . Since  $\bar{N}_B \propto \frac{1}{N_S}$ ,  $\bar{D}_{C_i}$  is independent of station configuration. For the values given in Table 6.2  $\bar{D}_{C_i} = 2.14$  TB/s.


The correlator output rate is

$$\bar{D}_{C-out} = \frac{1}{2\tau_{dump}} \bar{N}_B N_C N_P^2 N_S (N_S + 1) n_{corr}, \quad (6.22)$$

where  $\tau_{dump}$  is the correlator dump rate. A maximum of  $\tau_{dump}$  is given by (Alexander et al., 2009)

$$\tau_{dump} \sim 1200 \frac{D}{b}, \quad (6.23)$$

where  $b$  is the maximum baseline length. In practice, only the visibilities from long baselines need necessarily be dumped at this rate, whilst baselines from pairs of core stations can be averaged for longer periods. Correlator output data rate is shown in Figure 6.16 for the worst (dumping all correlations at the rate required by the longest baselines), best (dumping at the rate required by the core baselines) and average (dumping long and short baselines at different rates) cases. For SKA-low-2013, assuming a baseline-dependent dump rate, the correlator output data rate is  $\sim 7$  TB/s. This is several times larger than the input data rate, and will present a significant challenge to the visibility post-processing pipeline. Output data rate alone may be enough to motivate an increase in SKA-low-1 station sizes. We note that the nominal station size of  $\sim 1000$  antennas recommended in the previous section would result in a factor of four reduction in correlator output data rate compared to the SKA-low-2013 specifications. We further note that by considering the number of beams required in the two-part SKA correlator of Section 6.4, the data output rate associated with such a correlator is some 1–2 orders of magnitude lower than a pure FX implementation. This is a far more compelling reason to use an FFT-based correlator than any computational arguments which can be made.



(a) Fixed field of view with  $\Omega_A = 20$  square degrees. Number of beams has been forced to be integer. (b) SKA-low-2013 with one beam per-channel at all frequencies.

Figure 6.16: Correlator output data rates. The shaded region shows dump rates appropriate for baselines between the longest and core lengths. The solid line shows the output data rate when baselines are split into “core” and “long” subsets, and each accumulated for an appropriate length of time.

## 6.7 Implementation Platforms and Upgrade cycles

Detailed prescriptions for various realizations of SKA processors exist in the literature of the past several years (D’Addario, 2011; Carlson, 2010; Szomoru, 2011; Ford et al., 2012). These largely focus on either custom ASIC or FPGA technology, although a recent focus has been on the use of Graphical Processing Units (GPUs, see Owens et al. (2008) for a comprehensive introduction). These are massively multi-cored processors which can be used to efficiently accelerate some algorithms when paired with traditional CPU hosts. The correlation operation required in radio astronomy is one such effectively accelerated algorithm, and GPUs have seen wide use in recent telescope backends (Wayth et al., 2009; Greenhill et al., 2012; Clark et al., 2011). In this section, we briefly examine two technology candidates: GPUs and FPGAs for their suitability as beamforming and correlation processors for SKA-low-1. We then consider potential technology upgrade paths, which may impact the cost of SKA-sized processing over a period of several decades.

In the analysis of this section we use the Nvidia K20X GPU, and Xilinx XCVX980T FPGA as representative examples of current generation GPU and FPGA technology. The computational capacity of these systems is summarized in Table 6.3, which describes the arithmetic performance and external bandwidth of the two platforms. In the case of the GPU external bandwidth is limited by the PCI Express 2.0 16-lane interface, which has a maximum data transfer rate of 8 GB/s in each direction. In the case of the FPGA, the external bandwidth is taken to be the sum of the 72 13.1 Gb/s transceivers present in the

	Technology	Arithmetic Performance	External Bandwidth	
	Nvidia K20X	GPU	3.95 TFLOPS	8 GB/s (bi-directional)
	Xilinx XC7VX980T	FPGA	1.08 TMACs/s	117 GB/s (total)

Table 6.3: Capacity of current examples of GPU and FPGA technology. Data extracted from Nvidia (2012), (Xilinx Inc., 2012b). FPGA arithmetic performance is based on a 300MHz clock speed.

XC7VX980T.

## Beamforming

With some basic assumptions, it is possible to calculate the number of processing nodes required for an SKA-low-1 beamformer. We assume the following:

- SKA-low-2013 specifications of a single beam per 289 antenna station.
- Ring-based beamforming. This is not necessarily the optimal configuration, but the requirement for each node to connect to a bus containing all beams at all frequencies maximizes the I/O requirements of the beamformer. In any case, for a single beam, beamformer architecture is unlikely to be critical.
- Nyquist sampling of antenna signals (i.e. 600 MHz sampling clock), and ideal utilization of I/O. This is certainly an underestimate of required bandwidth, but gives a rough lower bound to required input I/O.

Under these assumptions, the data rate into each node in the beamforming ring is given by (from Equations 6.20, 6.21)  $D_{s-in} + D_{s-out}$ . The data rate out is given by  $D_{s-out}$ . The RAM use, required for channelization is given by  $R_c$  (Equation 6.14). The computation rate is given by  $C_c + C_b$  (Equations 6.10, 6.6).

Under these assumptions Figure 6.17 shows how many antenna signals could be accommodated per processing node in a chain arrangement, by plotting each of the quantities above as a fraction of the total resources available on each processing node. In the FPGA case, the total number of antennas is limited by processing and RAM requirements, both due to the channelization of antenna data streams. Again we note that this means that additional beams can be generated with very little cost increase. In the GPU case, the system is overwhelmed by I/O requirements for  $\gtrsim 5$  antennas per node, implying that  $\sim 60$  GPUs per station and around  $\sim 50\,000$  in total would be required for beamforming in the SKA-low-1. Given the extremely low computation rates per node in this configuration it is clear that GPUs represent a completely unfeasible processing platform for beamforming

implementation. Conversely, in a single-beam SKA-low-2013 station, a single Virtex 7 node is likely to be able to process  $\sim 25$  antenna signals, utilizing almost all the computational power of each chip. This implies a total of  $\sim 12$  FPGAs per station, and  $\sim 10\,000$  FPGAs in the total array.

## Correlation

A similar calculation can be performed for the correlator. This time, input bandwidth is given by  $\bar{D}_{C_i} = N_S D_{S-out}$  (Equation 6.21). The output bandwidth is given by  $\bar{D}_{C_o}$  (Equation 6.22). The computation rate is given by  $C_X$  (Equation 6.7). In the case of an FPGA processor, we assume that the 4-bit multiplication rate is 4 times that of the stated MAC rate given in Table 6.3 (see Chapter 3).

Given that the correlation can be parallelized by frequency, in order to calculate the number of nodes required to construct a complete correlator, we consider what fraction of the total correlator bandwidth a total node can process. The fraction of resources used on the two systems as a function of processed bandwidth is shown in Figure 6.18. In the case of correlation for SKA-low-1, both platforms become computation limited, and each have similar capabilities, with each node able to process  $\sim 0.1\%$  of the total correlator input bandwidth, for a total of  $\sim 1000$  processing nodes in the correlator system. We note that, despite Figure 6.6 suggesting that a single-beam SKA-low-2013 array will be dominated by correlator costs, the subsequent addition of channelization requirements, along with the ability of FPGAs to process correlator multiplications four times more efficiently than beam-former computations, results in a correlator with only 10% of the processing nodes of the beamformer.

As processing power increases, FPGA-technology is not likely to be limited by I/O bandwidth for any conceivable system with specifications similar to SKA-low-2013. However, GPU-based correlators are likely to become I/O bound if improvements in PCI-E transfer rates cannot keep up with arithmetic performance improvements. Ford et al. (2012) estimate that arithmetic performance of GPUs is likely to double every two years, whilst I/O bandwidth will double approximately every four years. At this rate a GPU correlator is likely to become I/O bound in around 8 arithmetic doubling cycles, or 16 years. This fact, coupled with the likely ability to reduce the correlator output data rate by post-processing of correlations, suggest GPU technologies may remain competitive (in a performance per-node sense) with FPGA-based alternatives in the coming years.




Figure 6.17: Fraction of arithmetic, memory and bandwidth resources used by a first stage (i.e. interfacing with digitizer) beamformer based on FPGA (top) and GPU (bottom) technology.





Figure 6.18: Fraction of arithmetic and bandwidth resources used by a correlator based on FPGA (top) and GPU (bottom) technology processing up to 0.3% of the total array bandwidth

### 6.7.1 Upgrading Hardware

The simple analysis above suggests that whilst FPGA technologies will almost certainly perform the bulk of station-level processing, GPUs may be a competitive technology in the SKA-low correlator. Whilst many estimates of the relative costs of such systems have been made, most recently by Zarb-Adami et al. (2013), in this section we briefly examine the effects of potential upgrade cycles to these comparisons.

Given an expected operational lifespan of an instrument, and knowing the purchase and running costs of an instrument, it is relatively simple to estimate when hardware should be upgraded. When the cumulative future running costs of the current hardware exceeds, over the remainder of an instrument’s life, the total running and upgrade costs of newer technology, an upgrade is economically efficient.

To calculate an optimal upgrade path, we construct a simple model based on a telescope system built on a given technology platform. Deploying or upgrading a given technology incurs the following costs:

- Purchasing hardware (e.g. FPGAs / GPUs).
- Developing hardware and software (e.g. PCBs, FPGA firmware, GPU code).
- Infrastructure (e.g. for providing power to the system, computing racks). This is assumed to be reusable in further upgrades.

but results in a system with improved performance and reduced power consumption. In the following analysis we posit a simple technology progression where processing capacity doubles every two years whilst power consumption and purchase prices remain the same.

Running the system for each year incurs the running costs:

- Power costs.
- Replacement costs for failed hardware.

Zarb-Adami et al. (2013) provide quantitative values to use for each of these components. We use two example systems from this document; the FPGA-based Uniboard<sup>5</sup>, and the GPU-based “Volta” system described. These are detailed in Table 6.4.

Figure 6.19 shows the cumulative running costs of a GPU and FPGA correlator for SKA-low-2013 based on an operational lifespan of 10 years beginning in 2017<sup>6</sup>. A number

---

<sup>5</sup><http://www.radionet-eu.org/uniboard>

<sup>6</sup>These plots were inspired by similar work presented by Jason Manley. See “From MeerKAT to SKA-1s Dish Array: A Digital Perspective” available as a video presentation at <https://science.nrao.edu/science/meetings/casper2012>

Technology	GPU	FPGA
Compute unit	4 GPUs + Host CPU	1 Uniboard (8 FPGAs)
Hardware Cost / unit	6,700 (1500/GPU) (700/CPU)	10,900 (800/FPGA) (4500 infrastructure)
NRE / build or upgrade cycle	250,000	442,000
NRE / initial build	100,000 (Rack & cabling)	30,000 (Rack & cabling)
	2.5/W (power)	2.5/W (power)
Power (W)	1,250 (250/GPU ) (250/CPU)	439
Modules required (2017 build)	109 (436 GPUs)	70 (560 FPGAs)
Power costs (Euro/W/year)	2.5	
Replacement parts	5% / year	

Table 6.4: The parameters of the GPU and FPGA-based correlator systems, based approximately on numbers given by Zarb-Adami et al. (2013). All costs are in 2013 Euros.

of potential upgrade paths are shown, with the optimal highlighted. In the case of a GPU-correlator, the total cost of ownership is minimized over 10 years by upgrading in 2019 and 2023, reducing the total cost by 36%. Alternatively, a single upgrade in 2021 results in only a marginally increased cost. The costs of an FPGA-based correlator, with higher NRE-to-running-cost ratio, are not so readily improved by upgrades.

Whilst the absolute costs detailed in Figure 6.19 have relatively large uncertainties, this simple analysis brings to light two important points. Firstly, when costing potential systems for the SKA it is important to take into account the different upgrade paths which may be most effectively utilized by different technology platforms. In the case of the GPU and FPGA correlator comparison presented here, in the absence of upgrades the cost of ownership of the GPU correlator is  $\sim 150\%$  higher than that of the FPGA system. When factoring in an optimal upgrade strategy, this GPU premium is reduced to  $\sim 66\%$ . Secondly, the impact of early upgrades on cost should lead to serious consideration of a slow, phased rollout of computing power. It seems unlikely that in the process of bringing SKA-low-1 online a central processing system capable of simultaneously handling the complete bandwidth and field-of-view of the array need be necessary immediately. In the case of processing platforms with low development costs such as GPUs, where quantities of scale in hardware manufacturing are not important, a reasonable strategy might be to deploy a correlator in 2017 which is powerful enough to enable the commissioning process, but need not necessarily meet the requirements of the telescopes science goals, which in any case will not be delivered




Figure 6.19: Cumulative costs of running a GPU (left) and FPGA-based (right) correlator from 2017 to 2027, whilst following different upgrade paths. In the case of a GPU correlator, following the optimal upgrade path results in a saving of 36% on the total cost of ownership of the system compared to a policy of no-upgrades.

for some years later.

## 6.8 Conclusions

This Chapter has explored some of the technical requirements which result from the currently proposed SKA-specifications. We find that for SKA-low-2013 specification (expanded to enable, on average,  $\sim 4$  beams) beamforming should be performed in a partial mechanism, since the costs savings of FFT-based, and hierarchical beamforming systems are not likely to be significant in the SKA unless station sizes are made significantly larger than 1000 elements. This station size may well prove a reasonable target for SKA-low-1, as required field of view could be obtained using flexible independent beams. Stations consisting of 1000 elements, with a physical diameter of  $\sim 70$  m are twice that argued for by the EoR working group. However, a 70 m maximum size limit is a more reasonable requirement based on interferometric arguments. Expanding station size to  $\sim 1000$  elements also results in a significant improvement in correlator output data rate, with a reduction of  $\sim 4$  times. This brings the correlator output to a more manageable (though still undesirably high) rate of  $\sim 0.8$  TB/s.

In a basic analysis of the number of processing nodes required both for beamforming and correlation we have seen that an SKA-low-2013 array which may initially appear to be dominated by correlation processing is actually likely to require significantly more processing at station level. This is due to the significant channelization costs required by each antenna

signal path, as well as the relative inefficiency of implementing beamforming operations rather than the low-bitwidth multiplies required by a correlator. Given the channelization dominance, which is less significant in the case of larger stations with more beams, a practical cost-reducing solution may be to reduce the number of digitally processed data streams. Other than by reducing the total number of antennas in the array, this could also be achieved by performing analogue beamforming on small groups of antennas.

Direct imaging is unlikely to bring any large improvement in computation efficiency if used in SKA-low-1 correlation applications, having only a small potential computational benefit when compared to standard FX processing. However, FFT-based correlation can potentially reduce the SKA-low-1 output data rate by at least an order of magnitude. Regardless of any potential use in the SKA correlator, FFT-based tied-array beamforming appears to be the only sensible option if many synthesized beams are required for time-domain science applications. The central 1500 m of SKA-low-2013 stations could be beamformed via FFT for approximately the same computational cost as the correlator.

The current SKA-low-1 project implementation plan shows FPGAs to be the cheapest technology available for the correlator. Whilst this may indeed be the case, the relative cost of GPU-based processing solutions is significantly reduced by considering plausible upgrade scenarios, in which case they may be more competitive than they appear. In any case, it is no doubt safe to assume that GPUs will play a significant part in the SKA, even if they are not employed in the correlator. If this is the case, the same arguments for upgrades are likely to be valid, and should always be taken into account when constructing instrument implementation costings.

## Chapter 7

# Conclusions

This thesis has investigated the wide variety of methods that can be used to process data from distributed multi-element radio telescopes. In it, we have explored both the theoretical algorithms used in such processing systems, and their real-world implementations and limitations. We have focused on the use of direct-imaging telescopes as alternatives to interferometers based on FX correlation. This thesis documents the design and deployment of a 32-antenna digital backend which has been deployed on the BEST-2 array in Medicina, Italy. I have used this deployment to directly compare FX and direct-imaging methods. The BEST-2 backend has also facilitated the development of a GPU-based transient search machine. Early results from all these systems have been presented.

Alongside the work presented on direct-imaging systems, this thesis has also outlined significant optimizations to low-bitwidth correlators. The previous chapter has also made computational arguments which inform the design decisions soon to be made in the SKA low-frequency array. In this final chapter, I draw together the conclusions of the work presented in this thesis.

### 7.1 Direct-Imaging as an FX Correlator Alternative

The use of interferometers in radio astronomy has enabled rich and varied science. From exquisite measurements of the CMB to high-resolution, widefield surveys, interferometry has allowed radio astronomy to be competitive and complementary to other astronomy disciplines. In particular, as direct power-spectrum probes, telescopes formed of many two-element interferometers have proved phenomenally successful.

Early interferometers employed coherent addition of signals, whilst modern interferometers use powerful digital multiplication circuits to form correlation matrices. These correlation matrices can be used to construct images of the radio sky. In this case, arrays are treated as if they are composed of large numbers of two-element interferometers. An

alternative processing method is to treat an array as a single device used to coherently add signals. This is the paradigm adopted in direct-imaging systems. Signal summations with different weightings can be used to form many simultaneous beams covering the sky. Whilst direct imaging has appeared in literature as a new way of obtaining an image of the sky, it is arguably the most natural way to combine signals from antenna array elements, being completely analogous to the summation of signals in “traditional” mirror and lens-based telescopes. Despite this fact, visibility-based imaging has generally been used in preference, owing to the many pleasing practical properties of correlator-based systems.

Firstly, correlators automatically exclude autocorrelation data in measurements. This significantly improves sensitivity and allows analogue schemes for mitigating autocorrelation noise to be avoided. Secondly, gain errors in antenna arrays naturally factor into per-antenna contributions, making them far easier to calculate and deal with on a per-visibility, rather than per-pixel, basis. Thirdly, the number of antennas (and two-element interferometers) in an array may be far fewer than the number of pixels in the image such an array produces. In this case it is often more efficient to compute and average visibilities, rather than beam powers in real-time.

For direct imaging to be a realistic alternative to FX correlation-based imaging, it is not enough to simply note that FFT-beamformers have an  $O(N \log N)$  computational cost versus the  $O(N^2)$  of visibility-based imaging. Each of the practical advantages of FX correlation listed above must be to be matched by any proposed direct-imaging implementation.

**Autocorrelations:** True gridded arrays can be used in direct-imaging systems without irreparably contaminating images with autocorrelation contributions. This is not the case in interpolation-gridded arrays. Autocorrelations must be explicitly measured and removed from data, or short baselines must be correlated independently via a non-imaging system.

**Gain Errors:** Modern direct-imaging literature notes that real-time calibration systems are necessary to form FX-equivalent images with direct-imaging systems. It is usually explicitly stated that an FX correlator is necessary to obtain coefficients for such a calibration. In this thesis I have attempted to emphasize, both using simple mathematical arguments and using the BEST-2 direct-imaging system, that real-time calibration is a relatively minor drawback to direct-imaging systems. We have seen that only crude calibration is necessary to maintain almost ideal signal-to-noise levels, and non-real-time image-calibration will almost certainly be necessary to deal with direction-dependent artifacts.

**Computation Rates:** We have seen that the large fields of view, moderate resolutions, and high filling factors required for low-frequency science goals are potentially well suited to direct imaging telescopes. However, in the case of most current telescopes, there is no compelling computational reason to prefer a direct-imaging system over an FX correlator. This is largely because of the filling-factors seen in many low-frequency arrays. Though EoR and BAO science cases demand highly filled arrays, in existing telescopes filling factors are still significantly less than 1. Furthermore, each operation in a direct imaging system is likely to be significantly more costly than the FX equivalent. In contrast, we have seen that data output rates from current large, low-frequency arrays can be significantly reduced by direct image processing. It is possible that concerns about output data-rate may be more effective in driving the adoption of direct imaging systems than FX correlation costs.

## 7.2 Future Developments

While direct-imaging systems may not offer significant benefits when compared with currently deployed FX correlators, we need not be disheartened. HI intensity mapping goals are being aggressively pursued, and will almost certainly result in the construction of larger- $N$ , higher filling-factor arrays than have been seen to-date. It is these telescopes where direct-imaging deployments can offer significant benefits.

This thesis has substantially addressed concerns about real-time calibration. Future work should focus on non-real-time calibration of direct-images. In the BEST-2 deployment of Chapter 5, full correlation matrices were used for image-calibration. In future, the question of how much correlation data is needed for robust calibration should be answered.

In BEST-2 we constructed a scalable system that is readily compatible with the flexible “packetized-correlator” framework developed by the CASPER collaboration. An evolution of the BEST-2 direct-imaging system should aim for larger-field, higher-resolution mapping. This will allow exploration of such calibration challenges as direction-dependent calibration as well as delivering intensity mapping science requirements.

We have seen in the BEST-2 deployment that synthetic beams formed by direct-imaging methods can enable time-domain science to be conducted with little extra frontend processing cost. In many beamforming applications where moderate numbers of beams are required, the computationally cheap beamforming afforded by spatial FFT-based processing methods may well motivate their adoption. Tracking of multiple independent beams is potentially a challenging task, but can be accommodated via time-dependent antenna interpolation grids. In any case, there are many science goals which do not require tracking beams, such as blind transient event searches. In this case, FFT-based beamforming applications can

be ideal. In particular such methods represent the only effective way to achieve large-field, high-sensitivity coverage of the sky for SKA time-domain science goals.

### 7.3 FX Correlator Developments

Whilst total operation rates required by signal processing systems are important, we also need to consider the ease of implementing such systems on physical hardware. When considering processing alternatives to FX correlation one must take into account the fact that low-bitwidth multiplication required by the correlation algorithm can often be implemented extremely efficiently. In particular, one such optimized correlator was described and implemented in Chapter 3. Here, we found great opportunities for performance increases by considering the correlation operation as a whole, rather than optimizing each individual multiplication, as has been attempted in some previous works. By utilizing  $18 \times 18$ -bit multipliers to implement four unsigned multiplies simultaneously, the computational cost of each correlation operation is essentially quartered when compared with each operation in a competing direct imaging system.

The correlator presented in Chapter 3 also represents a major first step towards extending the success of the CASPER FPGA libraries to those in the radio astronomy community who would prefer not to use the Simulink-based CASPER toolflow. The new correlation module, available in Verilog, is extremely portable and flexible. Through a simple MATLAB wrapper, it can also be incorporated into existing Simulink-based CASPER designs. I hope that in the future this development mechanism—low-level description of processing modules, with high-level wrappers—will be more widely adopted within the CASPER community.

### 7.4 The Square Kilometre Array

In Chapter 6 of this thesis, some computational arguments were applied to the low-frequency SKA Phase 1 (SKA-low-1) array. This array has not yet been fully specified, with input from the scientific community being actively welcomed in determining its future. A number of aspects of the reference design of SKA-low-1 were found to be sub-optimal. We find that a single-beam SKA-low-1 is dominated by signal channelization costs, in which case more beams could be processed with little computational penalty. Further, given the shallowness of the optimal computation point with respect to station sizes and quantity, we suggest  $\sim 1000$  antennas as a reasonable target station size. This is approximately four times larger than that currently specified, but, if operated with four times as many beams, could deliver the same science, with more flexibility and lower data output rate.

With regards to correlator implementation for the SKA, we have seen that direct imaging is unlikely to offer significant computational savings, though it is capable of reducing the SKA correlator output data rate by an order of magnitude. FFT-based beamforming also offers an extremely effective way of generating large numbers of tied-array beams for time-domain science applications.

In adopting a traditional FX correlator design, and considering currently suggested implementation plans, we have been able to make rough quantitative estimates of the number of processing nodes required. Of particular note is the fact that rapid upgrade cycles can significantly alter the relative running costs of an SKA correlator based on different technology platforms. GPU-based platforms are shown to be significantly more competitive when 2-year upgrade cycles are adopted. Given that a full-scale correlator is unlikely to be immediately necessary during the commissioning stages, I believe that GPUs represent a more realistic solution to a first generation SKA-low-1 correlator than most recent studies suggest.

## 7.5 Concluding Remarks

Attempts to make precision measurements of the HI signal at large redshifts are progressing rapidly. These experiments have led to the construction of some of the largest- $N$  interferometers ever used in radio astronomy. This growth trend shows no sign of stopping. After statistical detection of the EoR power-spectrum signature, even larger, more sensitive telescopes will be required to image this time in cosmic history.

Radio astronomy is rich with opportunities to develop new instrumentation. In some cases, incremental advances, taking advantage of the progression of technology is the most sensible course to pursue. However, in others, changes in processing methodologies can bring about great gains. These changes can be straightforward, such as refactoring the correlation algorithm, or as dramatic as replacing FX correlators with direct-imaging systems – the ease with which we can manipulate signals in digital radio telescopes affords great flexibility.

Whilst optimizing existing algorithm implementations is very useful, it seems extremely likely that the next generation of low-frequency telescopes will provide the tipping point where alternatives to traditional interferometers will become viable, and quite possibly essential. The work in this thesis has laid some of the foundations for this progression, which will no-doubt be an extremely exciting time for those working in astronomy – both scientists and engineers alike.

# References

- Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., et al. (2013). Planck 2013 results. I. Overview of products and scientific results. *arXiv preprint arXiv:1303.5062*.
- Alexander, P., Bregman, J. A., & Faulkner, A. J. (2009). SKA data flow and processing. In *Wide Field Astronomy & Technology for the Square Kilometre Array*, volume 1 (pp.16).
- Allen, L. R. & Frater, R. H. (1970). Wideband multiplier correlator. *Electrical Engineers, Proceedings of the Institution of*, 117(8), 1603–1608.
- Armour, W., Karastergiou, A., Giles, M., et al. (2012). A GPU-based Survey for Millisecond Radio Transients Using ARTEMIS. In *Astronomical Data Analysis Software and Systems XXI*, volume 461 (pp.33).
- Armstrong, R., Hickish, J., Zarb Adami, K., & Jones, M. E. (2009). A Digital Broadband Beamforming Architecture for 2-PAD. In *Proceedings of Wide Field Astronomy & Technology for the Square Kilometre Array* (pp. 273–278).
- Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. (1977). The absolute spectrum of CAS A - an accurate flux density scale and a set of secondary calibrators. *Astronomy and Astrophysics*, 61, 99–106.
- Backer, D. C. & PAPER Collaboration (2009). PAPER: The Precision Array to Probe the Epoch of Reionization. In *American Astronomical Society Meeting Abstracts #213*, volume 41 of *Bulletin of the American Astronomical Society* (pp. 481.07).
- Bailey, D. H. & Swartztrauber, P. N. (1991). The Fractional Fourier Transform and Applications. *SIAM Review*, 33(3), 389–404.
- Banday, A. J. & Wolfendale, A. W. (1991). Fluctuations in the galactic synchrotron radiation. I - Implications for searches for fluctuations of cosmological origin. *Mon. Not. R. Astron. Soc.*, 248, 705–714.

- Barsdell, B. R., Bailes, M., Barnes, D. G., & Fluke, C. J. (2012). Accelerating incoherent dedispersion. *Monthly Notices of the Royal Astronomical Society*, 422(1), 379–392.
- Bartolini, M. (2010). private communication.
- Bassett, B. & Hlozek, R. (2010). Baryon acoustic oscillations. In P. Ruiz-Lapuente (Ed.), *Dark Energy: Observational and Theoretical Approaches* (pp. 246).
- Battye, R. A., Browne, I. W. A., Dickinson, C., et al. (2013). H I intensity mapping: a single dish approach. *Mon. Not. R. Astron. Soc.*, 434, 1239–1256.
- Beardsley, A. P., Hazelton, B. J., Morales, M. F., et al. (2013). The EoR sensitivity of the Murchison Widefield Array. *Mon. Not. R. Astron. Soc.*, 429, L5–L9.
- Bellanger, M., Bonnerot, G., & Coudreuse, M. (1976). Digital filtering by polyphase network: Application to sample-rate alteration and filter banks. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 24(2), 109–114.
- Bennett, C. L., Banday, A. J., Górski, K. M., et al. (1996). Four-Year COBE DMR Cosmic Microwave Background Observations: Maps and Basic Results. *The Astrophysical Journal*, 464(1), L1–L4.
- Bennett, C. L., Larson, D., Weiland, J. L., et al. (2012). Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *ArXiv e-prints*, (pp. 177).
- Bluestein, L. (1970). A linear filtering approach to the computation of discrete Fourier transform. *IEEE Transactions on Audio and Electroacoustics*, 18(4), 451–455.
- Bolton, J. & Slee, O. (1953). Galactic Radiation at Radio Frequencies. V. The Sea Interferometer. *Australian Journal of Physics*, 6(4), 420.
- Boonstra, A. J. & Van Der Veen, A. J. (2001). Gain decomposition methods for radio telescope arrays. In *Proceedings of the 11th IEEE Signal Processing Workshop on Statistical Signal Processing.*, number 1 (pp. 365–368).: IEEE.
- Bowman, J. D. & Rogers, A. E. E. (2011). Results From EDGES. In *American Astronomical Society Meeting Abstracts #217*, volume 43 of *Bulletin of the American Astronomical Society* (pp. 107.08).
- Burke, B. F. & Graham-Smith, F. (2010). *An introduction to radio astronomy*. Cambridge University Press.

- Campbell, D. B. (2002). Measurement in Radio Astronomy. In *Single-Dish Radio Astronomy: Techniques and Applications*, volume 278 (pp. 81–90).
- Carlson, B. (2010). The Giant Systolic Array (GSA) Straw-man Proposal for a Multi-Mega Baseline Correlator for the SKA. *SKA Memo Series*, 127.
- Chang, T.-C., Pen, U.-L., Peterson, J., & McDonald, P. (2008). Baryon Acoustic Oscillation Intensity Mapping of Dark Energy. *Physical Review Letters*, 100(9).
- Cheung, A. C., Rank, D. M., Townes, C. H., Thornton, D. D., & Welch, W. J. (1969). Detection of Water in Interstellar Regions by its Microwave Radiation. *Nature*, 221(5181), 626–628.
- Clark, M. A., Greenhill, L. J., & La Plante, P. C. L. (2011). Accelerating Radio Astronomy Cross-Correlation with Graphics Processing Units. (pp.36).
- Cohen, A. S., Rottgering, H. J. A., Jarvis, M. J., Kassim, N. E., & Lazio, T. J. W. (2004). Deep, high-resolution survey at 74MHz (Cohen+, 2004). *VizieR Online Data Catalog*, 215, 417.
- Cooley, J. W. & Tukey, J. W. (1965). An Algorithm for the Machine Calculation of Complex Fourier Series. *Mathematics of Computation*, 19(90), 297–301.
- Cornwell, T. & Fomalont, E. B. (1999). Self-Calibration. In *Synthesis Imaging in Radio Astronomy II*, volume 180.
- Cornwell, T. J., Anantharamaiah, K. R., & Narayan, R. (1989). Propagation of coherence in scattering: an experiment using interplanetary scintillation. *Journal of the Optical Society of America A*, 6(7), 977.
- Crocce, M. & Scoccimarro, R. (2008). Nonlinear evolution of baryon acoustic oscillations. *Physical Review D*, 77(2).
- Curtis, T. E. (1998). Principles of Sonar Beamforming. In *IMA Conference on Parallel Computation*.
- D’Addario, L. (2008). LWA Fine Delay Tracking. *LWA Memo Series*, 143.
- D’Addario, L. R. (2011). Low-Power Correlator Architecture for the Mid-Frequency SKA. *SKA Memo Series*, 133(March).
- D’Addario, L. R. (2013). Array Beamformers for the SKA. *SKA memo series*, 147, 1–4.

- Davies, R. D., Rowson, B., Booth, R. S., et al. (1967). Measurements of OH Emission Sources with an Interferometer of High Resolution. *Nature*, 213(5081), 1109–1110.
- de Souza, L., Bunton, J. D., Campbell-Wilson, D., Cappallo, R. J., & Kincaid, B. (2007). A radio astronomy correlator optimized for the Xilinx Virtex-4 SX FPGA. In *IEEE International Conference on Field Programmable Logic and Applications* (pp. 62–67): IEEE.
- de Vos, M., Gunst, A., & Nijboer, R. (2009). The LOFAR Telescope: System Architecture and Signal Processing. *Proceedings of the IEEE*, 97(8), 1431–1437.
- Dewdney, P. (2013). SKA1 System Baseline Design. (pp. 1–98).
- Dewdney, P., De Vaate, J.-g., Gunst, A., et al. (2010). SKA Phase 1: Preliminary System Description. *SKA Memo Series*, 130.
- Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. (1965). Cosmic Black-Body Radiation. *The Astrophysical Journal*, 142, 414.
- Fan, X., Carilli, C. L., & Keating, B. (2006). Observational Constraints on Cosmic Reionization. *Annual Review of Astronomy and Astrophysics*, 44, 415–462.
- Ford, D., Faulkner, A., & Alexander, P. (2012). A Software Correlator for SKA 1. *SKA Memo Series*, 139.
- Foster, G. S. (2013). *Large-N Correlator Systems for Low Frequency Radio Astronomy*. PhD thesis, University of Oxford.
- Furlanetto, S. R., Peng Oh, S., & Briggs, F. H. (2006). Cosmology at low frequencies: The 21cm transition and the high-redshift Universe. *Physics Reports*, 433(4-6), 181–301.
- Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. (2004). Statistical Probes of Reionization with 21 Centimeter Tomography. *The Astrophysical Journal*, 613(1), 16–22.
- Gould, D. M. & Lyne, A. G. (1998). Multifrequency polarimetry of 300 radio pulsars. *Monthly Notices of the Royal Astronomical Society*, 301(1), 235–260.
- Green, D. A. (2009). A revised Galactic supernova remnant catalogue. *Bulletin of the Astronomical Society of India*, 37, 45–61.
- Greenhill, L. J., Werthimer, D., Taylor, G. B., & Ellingson, S. W. (2012). A broadband 512-element full correlation imaging array at VHF (LEDA). In *International Conference on Electromagnetics in Advanced Applications (ICEAA)* (pp. 1117–1120): IEEE.

- Hamaker, J. P. & Bregman, J. D. (1996). Understanding radio polarimetry. III. Interpreting the IAU/IEEE definitions of the Stokes parameters. *Astronomy and Astrophysics Supplement Series*, 117(1), 161–165.
- Hamaker, J. P., Bregman, J. D., & Sault, R. J. (1996). Understanding radio polarimetry. I. Mathematical foundations. *Astronomy and Astrophysics Supplement*, 117, 137–147.
- Harris, C. & Haines, K. (2011). A Mathematical Review of Polyphase Filterbank Implementations for Radio Astronomy. *Publications of the Astronomical Society of Australia*, 28(4), 317.
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. (1982). A 408 MHz All-Sky Continuum Survey. II. The Atlas of Contour Maps. *Astronomy and Astrophysics Supplement Series*, 47, 1.
- Hewitt, J. N. & Murchison Widefield Array Collaboration (2013). The MWA 32-tile Prototype: Deep Integrations and Power Spectra of Foregrounds. In *American Astronomical Society Meeting Abstracts*, volume 221 of *American Astronomical Society Meeting Abstracts* (pp. 108.06).
- Holler, C. M., Jones, M. E., Taylor, A. C., Harris, A. I., & Maas, S. A. (2011). A 2-20 GHz Analog Lag-Correlator for Radio Interferometry. In *IEEE Transactions on Instrumentation and Measurement*, volume 61 (pp. 2253–2261).
- Jacobs, D. C., Parsons, A. R., Aguirre, J. E., et al. (2013). A Flux Scale for Southern Hemisphere 21 cm Epoch of Reionization Experiments. *Astrophys. J.*, 776, 108.
- Jansky, K. (1933). Electrical disturbances apparently of extraterrestrial origin. *Proceedings of the Institute of Radio Engineers*, 20(7).
- Johnson, J. (1928). Thermal Agitation of Electricity in Conductors. *Physical Review*, 32(1), 97–109.
- Johnson, R. C. & Jasik, H. (1984). *Antenna engineering handbook (2nd edition)*, volume -1.
- Jones, R. C. (1941). New calculus for the treatment of optical systems. I. Description and discussion of the calculus. *Journal of the Optical Society of America*, 31(7).
- Landecker, T. L. & Wielebinski, R. (1970). The Galactic Metre Wave Radiation: A two-frequency survey between declinations  $+25^\circ$  and  $-25^\circ$  and the preparation of a map of the whole sky. *Australian Journal of Physics Astrophysical Supplement*, 16, 1.

- 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 Series*, 192(2), 16.
- Liu, A., Tegmark, M., Bowman, J., Hewitt, J., & Zaldarriaga, M. (2009). An improved method for 21-cm foreground removal. *Mon. Not. R. Astron. Soc.*, 398, 401–406.
- Liu, A., Tegmark, M., Morrison, S., Lutmirski, A., & Zaldarriaga, M. (2010). Precision Calibration of Radio Interferometers Using Redundant Baselines. *Monthly Notices of the Royal Astronomical Society*, 408(2), 1029—1050.
- Lonsdale, C. J., Cappallo, R. J., Morales, M. F., et al. (2009). The Murchison Widefield array: design overview. *Proceedings of the IEEE*, 97(8), 1497–1506.
- Lorimer, D. R. & Kramer, M. (2005). *Handbook of pulsar astronomy*, volume 4. Cambridge University Press.
- Lubberhuizen, W. & Kooistra, E. (2007). RSP Firmware Functional Specification. *LOFAR Document*.
- Lutmirski, A., Tegmark, M., Sanchez, N., et al. (2011). Solving the Corner-Turning Problem for Large Interferometers. *Monthly Notices of the Royal Astronomical Society*, 410(3), 2075—2080.
- Magro, A., Hickish, J., & Adami, K. Z. (2013). Multibeam GPU Transient Pipeline for the Medicina BEST-2 Array. (pp.17).
- Magro, A., Karastergiou, A., Salvini, S., et al. (2011). Real-time, fast radio transient searches with GPU de-dispersion. *Monthly Notices of the Royal Astronomical Society*, 417(4), 2642—2650.
- Manley, J., Welz, M., Parsons, A., Ratcliffe, S., & van Rooyen, R. (2012). SPEAD Recommended Practice. *SKA South Africa Document*.
- McMahon, P., Langman, A., Werthimer, D., et al. (2007). Packetized FX Correlator Architectures. *CASPER Memo Series*, 17.
- Mellema, G., Koopmans, L. V. E., Abdalla, F. A., et al. (2012). Reionization and the Cosmic Dawn with the Square Kilometre Array. *Experimental Astronomy*, (pp. 1—84).
- Michelson, A. A. & Pease, F. G. (1921). Measurement of the Diameter of Alpha-Orionis by the Interferometer. *Proceedings of the National Academy of Sciences of the United States of America*, 7(5), 143–146.

- Montebugnoli, S., Bianchi, G., Monari, J., et al. (2009). BEST: Basic Element for SKA Training. In *Proceedings of Wide Field Astronomy & Technology for the Square Kilometre Array*, number November (pp. 331—336).
- Moore, G. E. (1965). Cramming more components onto integrated circuits. *Electronics*, 38(8), 114.
- 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, 1265–1272.
- Morales, M. F., Bowman, J. D., & Hewitt, J. N. (2006). Improving Foreground Subtraction in Statistical Observations of 21 cm Emission from the Epoch of Reionization. *Astrophys. J.*, 648, 767–773.
- Morales, M. F. & Wyithe, J. S. B. (2010). Reionization and Cosmology with 21-cm Fluctuations. *Annual Review of Astronomy & Astrophysics*, 48, 127–171.
- Nagpal, V. (2006). *An FPGA Based Phased Array Processor for the Sub-Millimeter Array*. PhD thesis.
- Netterfield, C. B., Ade, P. A. R., Bock, J. J., et al. (2002). A Measurement by BOOMERANG of Multiple Peaks in the Angular Power Spectrum of the Cosmic Microwave Background. *The Astrophysical Journal*, 571(2), 604.
- Noorishad, P., Wijnholds, S. J., van Ardenne, A., & van der Hulst, T. (2012). Redundancy Calibration of Phased Array Stations. *Astronomy & Astrophysics*, 545.
- Nvidia (2012). Tesla Kepler GPU Accelerators. *Nvidia Tesla K-Series Datasheet*.
- Otobe, E., Nakajima, J., Nishibori, K., et al. (1994). Two-dimensional direct images with a spatial FFT interferometer. *Publications of the Astronomical Society of Japan*, 46, 503–510.
- Owens, J. D., Houston, M., Luebke, D., et al. (2008). GPU computing. *Proceedings of the IEEE*, 96(5), 879–899.
- Parsons, A., Backer, D., Chang, C., et al. (2006). PetaOp/Second FPGA Signal Processing for SETI and Radio Astronomy. In *Fortieth Asilomar Conference on Signals, Systems and Computers* (pp. 2031–2035): IEEE.

- Parsons, A., Backer, D., Siemion, A., et al. (2008). A Scalable Correlator Architecture Based on Modular FPGA Hardware, Reuseable Gateway, and Data Packetization. *Publications of the Astronomical Society of the Pacific*, 120(873), 1207–1221.
- Parsons, A. R., Liu, A., Aguirre, J. E., et al. (2013). New Limits on 21cm EoR From PAPER-32 Consistent with an X-Ray Heated IGM at  $z=7.7$ . *ArXiv e-prints*.
- Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. (2012). A Per-Baseline, Delay-Spectrum Technique for Accessing the 21cm Cosmic Reionization Signature. (pp.30).
- Penzias, A. A. & Wilson, R. W. (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. *The Astrophysical Journal*, 142, 419.
- Perini, F. (2009). Low noise design experience for the SKADS/BEST demonstrator. In *Proceedings of Wide Field Astronomy & Technology for the Square Kilometre Array* (pp. 341–345).
- Perini, F., Bianchi, G., Schiaffino, M., & Monari, J. (2009). BEST receiver experience: general architecture, design and integration. In *Proceedings of Wide Field Astronomy & Technology for the Square Kilometre Array* (pp. 351–354).
- Platania, P., Bensadoun, M., Bersanelli, M., et al. (1998). A Determination of the Spectral Index of Galactic Synchrotron Emission in the 1-10 GHz Range. *Astrophys. J.*, 505, 473–483.
- Prasad, P. & Wijnholds, S. J. (2012). AARTFAAC: Towards a 24x7, All-sky Monitor for LOFAR. *eprint arXiv:1205.3056*, (pp. 10–13).
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). *Numerical recipes in C (2nd ed.): the art of scientific computing*. New York, NY, USA: Cambridge University Press.
- Price, D. (2012). *Radio Astronomy Instrumentation for Redshifted Hydrogen Line Science*. PhD thesis, University of Oxford.
- Qadeer, S., Zafar, M., Khan, A., & Sattar, S. A. (2011). On Fixed Point error analysis of FFT algorithm. *ACEEE International Journal on Information Technology*, 1(3), 3–6.
- Ryle, M. (1952). A New Radio Interferometer and Its Application to the Observation of Weak Radio Stars. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 211(1106), 351–375.

- Ryle, M. & Hewish, A. (1960). The synthesis of large radio telescopes. *Monthly Notices of the Royal Astronomical Society*, 120.
- Schawlow, A. & Townes, C. (1958). Infrared and Optical Masers. *Physical Review*, 112(6), 1940–1949.
- Shaver, P. A., Windhorst, R. A., Madau, P., & de Bruyn, A. G. (1999). Can the reionization epoch be detected as a global signature in the cosmic background? *Astron. Astrophys.*, 345, 380–390.
- SKA Science Working Group (2011). The Square Kilometre Array design reference mission: SKA phase 1. *SKA Project Document*.
- Smirnov, O. M. (2011a). Revisiting the radio interferometer measurement equation. I. A full-sky Jones formalism. *Astronomy & Astrophysics*, 527, A106.
- Smirnov, O. M. (2011b). Revisiting the radio interferometer measurement equation. II. Calibration and direction-dependent effects. *Astronomy & Astrophysics*, 527, A107.
- Smirnov, O. M. (2011c). Revisiting the radio interferometer measurement equation. III. Addressing direction-dependent effects in 21 cm WSRT observations of 3C 147. *Astronomy & Astrophysics*, 527, A108.
- Smirnov, O. M. (2011d). Revisiting the radio interferometer measurement equation. IV. A generalized tensor formalism. *Astronomy & Astrophysics*, 531, 159.
- Smith, M. J. S. (1997). *Application-specific integrated circuits*. Addison-Wesley Longman Publishing Co., Inc.
- Spergel, D. N., Bean, R., Dore, O., et al. (2007). Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications for Cosmology. *The Astrophysical Journal Supplement Series*, 170(2), 377–408.
- Switzer, E. R., Masui, K. W., Bandura, K., et al. (2013). Determination of  $z \sim 0.8$  neutral hydrogen fluctuations using the 21 cm intensity mapping autocorrelation. *Mon. Not. R. Astron. Soc.*, 434, L46–L50.
- Szomoru, A. (2011). A UNIBOARD-based Phase 1 SKA Correlator and Beamformer Concept Design. *Context of the SKA Signal Processing Concept Design Review Documents*.
- Taylor, G. B., Carilli, C. L., & Perley, R. A., Eds. (1999). *Synthesis Imaging in Radio Astronomy II*. Astronomical Society of the Pacific.

- Tegmark, M., Eisenstein, D. J., Hu, W., & de Oliveira-Costa, A. (2000). Foregrounds and Forecasts for the Cosmic Microwave Background. *Astrophys. J.*, 530, 133–165.
- Tegmark, M. & Zaldarriaga, M. (2009). Fast Fourier transform telescope. *Physical Review D*, 79(8).
- Tegmark, M. & Zaldarriaga, M. (2010). Omniscope: Large area telescope arrays with only  $N \log N$  computational cost. *Physical Review D*, 82(10).
- the MWA Collaboration (2009). Technical Description of MWA Antenna. *MWA Antenna Critical Design Review Document*.
- Thompson, A. (1998). Quantization Efficiency for Eight or more Sampling Levels. *MMA Memo Series*, 220.
- Thompson, A. R., Moran, J. M., & Jr., G. W. S. (2004). *Interferometry and Synthesis in Radio Astronomy*. WILEY-VCH Verlag, second edition.
- Tingay, S. J., Goeke, R., Bowman, J. D., et al. (2013). The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies. *Proc. Astron. Soc. Aust.*, 30, 7.
- Tingay, S. J., Goeke, R., Hewitt, J. N., et al. (2012). Realisation of a low frequency SKA Precursor: The Murchison Widefield Array. (pp.8).
- Trac, H. & Cen, R. (2007). Radiative Transfer Simulations of Cosmic Reionization. I. Methodology and Initial Results. *The Astrophysical Journal*, 671(1), 1–13.
- Urry, W. L. (2000). A Corner Turn Architecture. *ATA Memo Series*, 14.
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. (2013). LOFAR: The LOw-Frequency ARray. *arXiv preprint arXiv:1305.3550*.
- Wayth, R. B., Greenhill, L. J., & Briggs, F. H. (2009). A GPU-based Real-time Software Correlation System for the Murchison Widefield Array Prototype. *Publications of the Astronomical Society of the Pacific*, 121(882), 857–865.
- Weaver, H., Williams, D. R. W., Dieter, N. H., & Lum, W. T. (1965). Observations of a strong unidentified microwave line and of emission from the OH molecule. *Nature*, 208(5005), 29–31.

- Weinreb, S. (1963). *A Digital Spectral Analysis Technique and its Application to Radio Astronomy*. PhD thesis, Massachusetts Institute of Technology. Research Laboratory of Electronics.
- Werthimer, D., Bowyer, S., Cobb, J., Lebofsky, M., & Lampton, M. (2000). The Serendip IV Arecibo Sky Survey. In *Bioastronomy 99: A New Era in the Search for Life*.
- Williams, J. R. (1968). Fast Beam-Forming Algorithm. *The Journal of the Acoustical Society of America*, 44, 1454.
- Wolszczan, A. & Cordes, J. M. (1987). Interstellar interferometry of the pulsar PSR 1237+25. *The Astrophysical Journal*, 320, L35.
- Wyithe, J. S. B., Loeb, A., & Geil, P. M. (2007). Baryonic acoustic oscillations in 21-cm emission: a probe of dark energy out to high redshifts. *Monthly Notices of the Royal Astronomical Society*, 383(3), 1195–1209.
- Xilinx Inc. (2012a). 7 Series DSP48E1 Slice. User Guide. *Xilinx Datasheets*.
- Xilinx Inc. (2012b). 7 Series FPGAs Overview. Advance Product Specification. *Xilinx Datasheets*.
- Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. (2013). Initial deep LOFAR observations of epoch of reionization windows. I. The north celestial pole. *Astron. Astrophys.*, 550, A136.
- Zarb-Adami, K., Bunton, J., Carlson, B., et al. (2013). SKA1-Low Hybrid PIP. *SKA Documentation*.
- Zernike, F. (1938). The concept of degree of coherence and its application to optical problems. *Physica*, 5(8), 785–795.

## Appendix A

# Convolution Theorem

The convolution theorem states that the Fourier transform of a convolution of two functions is equal to the pointwise product of the Fourier transforms of the two functions themselves. Equivalently, in algebraic form:

$$\mathcal{F}\{f * g\} = \mathcal{F}\{f\} \mathcal{F}\{g\}, \quad (\text{A.1})$$

where an asterisk represents convolution, defined by:

$$f * g = \int f(x)g(z - x)dx \quad (\text{A.2})$$

and the integral takes place over all possible values of  $x$ .

To prove Equation A.1, we start by taking the Fourier transform of both sides of Equation A.2:

$$\begin{aligned} \mathcal{F}\{f * g\} &= \int \int f(x)g(z - x)e^{-2\pi iz\nu} dx dz \\ &= \int \int f(x)g(w)e^{-2\pi i(x+w)\nu} dx dw \\ &= \int f(x)e^{-2\pi ix\nu} dx \int g(w)e^{-2\pi iw\nu} dw \\ &= \mathcal{F}\{f\} \mathcal{F}\{g\}, \end{aligned}$$

where in the second line the variable  $w = z - x$  is introduced to allow the splitting of the two integrals into two distinct Fourier transforms.

## Appendix B

# Two's Complement Number Representation

There are a variety of ways of representing signed numbers using binary counting system. The most widely used of these representations in electronic computation is the two's complement system. In the two's complement system, the bitwise implementations of the fundamental operations – addition, subtraction and multiplication – are independent of the sign of the inputs, and further, do not depend on the operands representing signed, rather than unsigned, numbers.

The two's complement representation of a positive integer is identical to its unsigned representation, whilst the two's complement representation of a negative number is constructed by taking the unsigned representation of its magnitude, inverting all the bits, and adding one. For example:

$$\begin{aligned} -27 &\longrightarrow \text{unsigned representation of } 27 \longrightarrow \text{invert bits} \longrightarrow \text{add one} \\ -27 &\longrightarrow 00011011 \longrightarrow 11100100 \longrightarrow 11100101 \end{aligned}$$

In  $m$ -bit two's complement notation all integers in the range  $[-2^m, 2^m - 1]$  can be represented. In some algorithms, this asymmetry about zero can lead to so-called *D.C.* biases in calculations, and it is common to only discard the value  $-2^m$  by saturating negative values at  $-(2^m - 1)$ .

Two's complement notation is closely related to offset binary notation, in which the most negative value,  $-2^m$ , is represented by all-zeros, and the value 0 is represented by the unsigned binary representation of  $2^m$ . Transformation between two's complement and offset binary representations of numbers is simply accomplished by inversion of the MSB. Table B.1 shows the 8-bit two's complement and offset binary representations of various integers in the range  $[-128, 127]$ .

Decimal value	Two's complement representation	Offset binary representation
-128	10000000	00000000
-127	10000001	00000001
⋮	⋮	⋮
-2	11111110	01111110
-1	11111111	01111111
0	00000000	10000000
1	00000001	10000001
2	00000010	10000010
⋮	⋮	⋮
127	01111111	11111111

Table B.1: Representations of signed integers using the 8-bit two's complement and offset binary schemes.

## Appendix C

# The CASPER X-Engine Architecture

Whilst an overview of the CASPER-designed X-engine is given in Chapter 3, this appendix details the underlying structure of the module’s architecture, and provides a recipe for its implementation. In the diagrams that follow, the X-engine is parameterized as follows:

- $L$  : The number of serially-presented samples to be accumulated in each data window.
- $N$  : The number of antennas presented to the X-engine.
- $A$  : Adder latency within complex multiply operations.
- $M$  : Multiplier latency within complex multiply operations.
- $p$  : The “tap-number”, which runs from 1 to  $\lfloor \frac{N}{2} \rfloor$ .

A top-level schematic of the X-engine design is given in Figure C.1. The only top level inputs are antenna signal inputs (whose order is defined in Figure 3.1) and a synchronization signal, which provides a pulse one clock cycle before the start of a window. After synchronization, the engine processes data in unbroken windows. Where a window of data is input, and further data is not available for processing, an entire window-length gap must be left before another window can be presented. This allows the underlying processing architecture to operate without any per-sample enable signals.

The X-engine is constructed from a series of “taps” as outlined in Section 3.2.1. These are detailed in Figure C.2, which shows the auto-tap stage, and Figure C.3, which shows the structure of a cross tap at stage  $p$  in the processing chain.

Each tap features a complex multiply and accumulate (CMAC) operation, which is depicted schematically in Figure C.4. The CMAC module combines the complex multiplication and accumulation operation with a shift register output. The latencies of this register chain give rise to the output data order detailed in Figure 3.4.

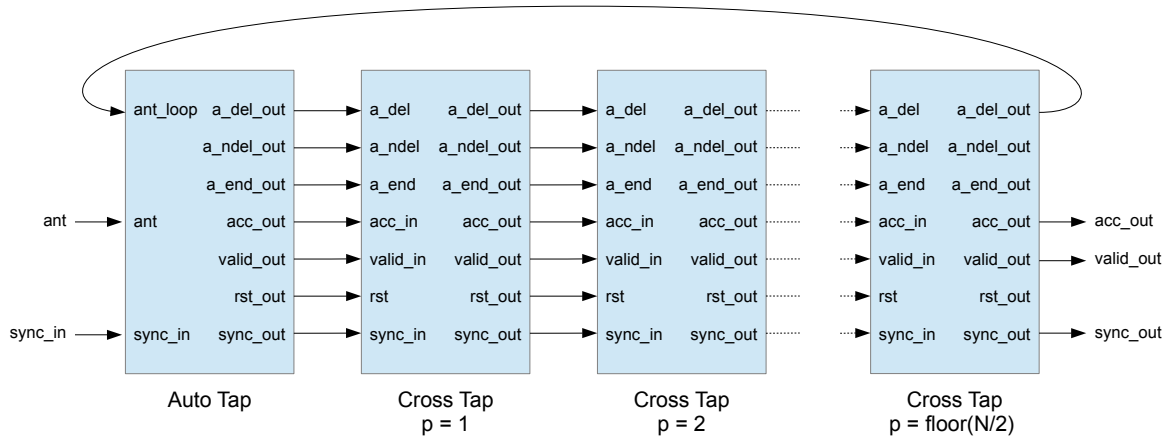


Figure C.1: Top-level schematic of the X-engine circuit.

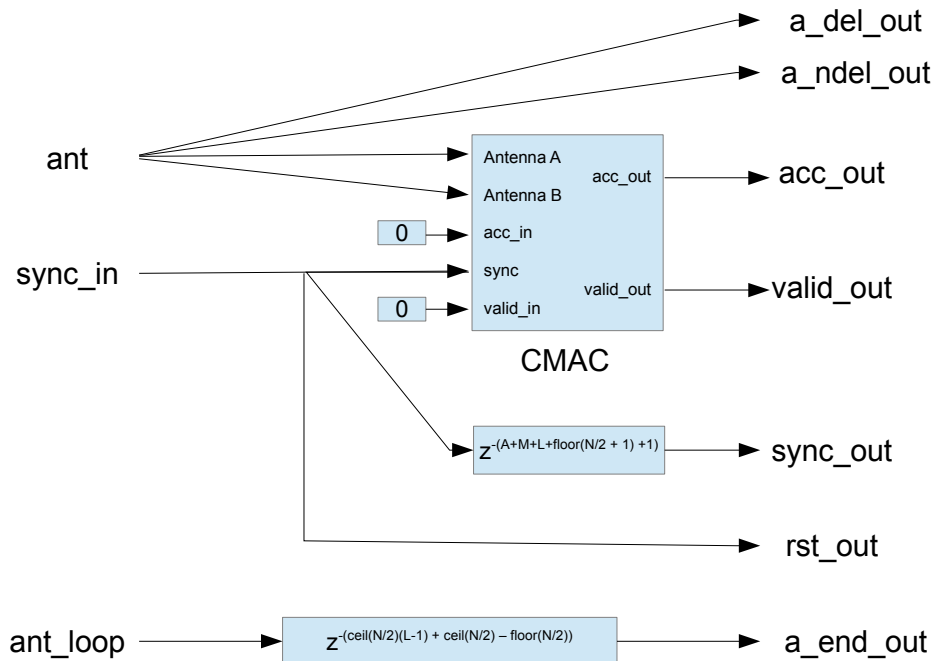


Figure C.2: Schematic of the X-engine “auto-tap”.

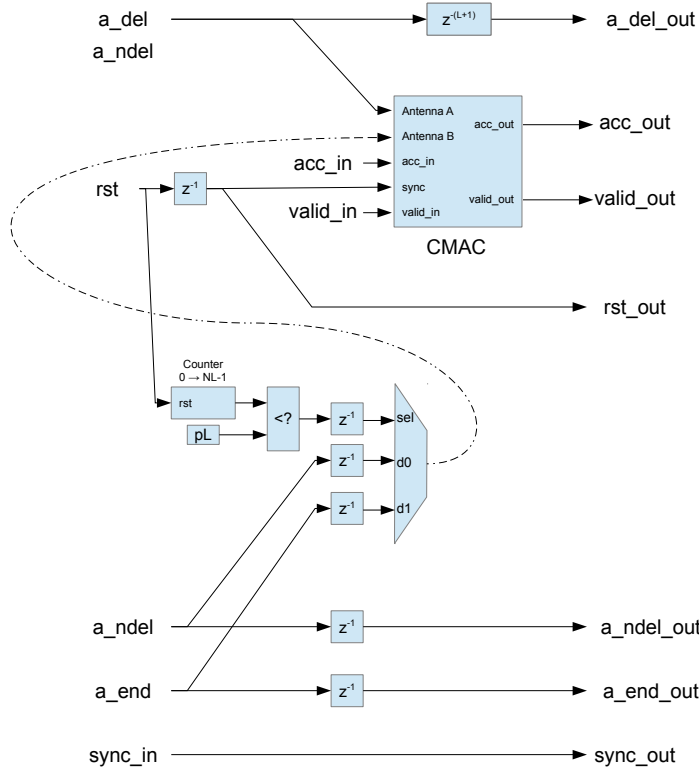


Figure C.3: Schematic of the X-engine “cross-tap” schematic for tap  $p$ .

Where the antenna signal inputs are inherently dual-polarized, as in the standard CASPER correlator (see Figure 3.1), the CMAC module takes the form of four parallel complex multiply and accumulation circuits, one for each polarization product, with common control logic. In the X-engine implementation presented in Chapter 3, the complex multiply operation is replaced by the parallel complex multiply accumulate depicted in Figure 3.10. This is naturally accommodated into the X-engine design, provided  $A$  and  $M$  are correctly set to reflect the latency of the multiply-sum circuit.

The final stage in the X-engine when using unsigned multiplication is to correct the correlation output products, using the circuit depicted in Figure 3.12. The correction process is straightforward, but requires that the correction factors be read out of buffers in the same output order as the X-engine. This order can be generated using the `bl_order_gen` module, whose code is shown at the end of this appendix. In this module, the readout is triggered by arrival of a `sync` pulse from the X-engine, one clock cycle before output data is valid. Following this pulse, the `ant_a` and `ant_b` outputs of the module reflect the antenna pairs present in the X-engine correlation output. Where a double-buffering scheme is used, the `buf_sel` output can be used to select which buffer is currently being read.

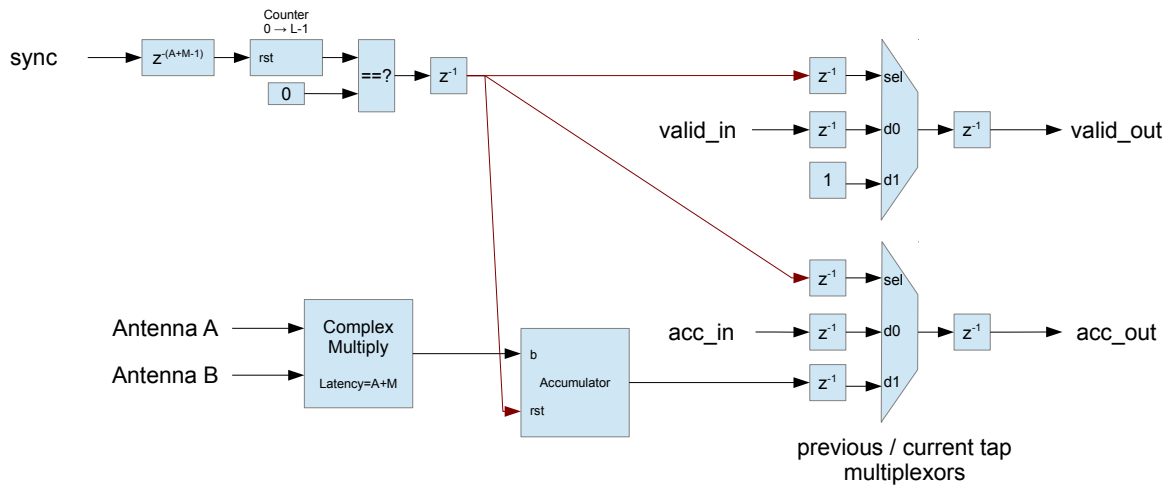


Figure C.4: Schematic depiction of an X-engine Complex multiply-accumulate (CMAC) operation.

```

1 module bl_order_gen(
2     clk,
3     sync,
4     en,
5     ant_a,
6     ant_b,
7     buf_sel
8 );
9
10 // A simple function to calculate ceil(log2(value))
11 function integer log2_func;
12     input integer value;
13     integer loop_cnt;
14     begin
15         value = value - 1;
16         for (loop_cnt=0; value>0; loop_cnt=loop_cnt+1)
17             value = value >> 1;
18         log2_func = loop_cnt;
19     end
20 endfunction
21
22 // The number of antennas in the X-engine
23 parameter N_ANTTS = 16;
24 // The number of bits required to count antennas
25 localparam ANT_BITS = 'log2(N_ANTTS);
26
27 // INPUTS
28 input clk; // The clock
29 input sync; // Connect to the sync output of an X-engine
30 input en; // Connect to the valid signal of an X-engine
31 // OUTPUTS
32 output [ANT_BITS-1:0] ant_a; // First antenna in the correlation pair
33 output [ANT_BITS-1:0] ant_b; // Second antenna in the correlation pair
34 output buf_sel; // Buffer than holds current correction factors
35
36 reg [ANT_BITS-1:0] a=0;
37 reg [ANT_BITS-1:0] b=0;

```

```

38     reg [ANT_BITS-1:0] offset=0;
39     reg buf_sel_reg=0;
40     always @(posedge(clk)) begin
41         if (sync) begin
42             buf_sel_reg <= 1'b0;
43         end else if(a==N_ANTS-1 && b==N_ANTS-1 && en) begin
44             buf_sel_reg <= ~buf_sel_reg;
45         end
46     end
47     always @(posedge(clk)) begin
48         if (sync) begin
49             b <= 0;
50             a <= N_ANTS/2;
51             offset <= N_ANTS/2+1;
52         end else if (en) begin
53             if(a==b) begin
54                 b <= b+1'b1;
55                 a <= offset;
56                 offset <= offset+1'b1;
57             end else begin
58                 a <= a+1'b1;
59             end
60         end
61     end
62
63     assign ant_a = a;
64     assign ant_b = b;
65     assign buf_sel = ant_a <= ant_b ? buf_sel_reg : ~buf_sel_reg;
66 endmodule

```