

ALMA MATER STUDIORUM  
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN  
**Astronomia**

CICLO XXII

FIS/05 ASTRONOMIA E ASTROFISICA

**Very Long Baseline  
Interferometry  
in Italy**

**Wide-field VLBI imaging and astrometry and prospects for  
an Italian VLBI network including  
the Sardinian Radio Telescope**

*Presentato da*

**John Stanley Morgan**

*Coordinatore Dottorato*  
Lauro Moscardini

*Relatore*  
Daniele Dallacasa

---

*Co-relatori*

Mauro Nanni

&

Franco Mantovani

Esame finale anno 2010



## Abstract

In this thesis the use of widefield imaging techniques and VLBI observations with a limited number of antennas are explored.

I present techniques to efficiently and accurately image extremely large UV datasets. Very large VLBI datasets must be reduced into multiple, smaller datasets if today's imaging algorithms are to be used to image them. I present a procedure for accurately shifting the phase centre of a visibility dataset. This procedure has been thoroughly tested and found to be almost two orders of magnitude more accurate than existing techniques. Errors have been found at the level of one part in 1.1 million. These are unlikely to be measurable except in the very largest UV datasets.

Results of a four-station VLBI observation of a field containing multiple sources are presented. A 13 gigapixel image was constructed to search for sources across the entire primary beam of the array by generating over 700 smaller UV datasets. The source 1320+299A was detected and its astrometric position with respect to the calibrator J1329+3154 is presented. Various techniques for phase calibration and imaging across this field are explored including using the detected source as an in-beam calibrator and peeling of distant confusing sources from VLBI visibility datasets.

A range of issues pertaining to wide-field VLBI have been explored including; parameterising the wide-field performance of VLBI arrays; estimating the sensitivity across the primary beam both for homogeneous and heterogeneous arrays; applying techniques such as mosaicing and primary beam correction to VLBI observations; quantifying the effects of time-average and bandwidth smearing; and calibration and imaging of wide-field VLBI datasets.

The performance of a computer cluster at the Istituto di Radioastronomia in Bologna has been characterised with regard to its ability to correlate using the DiFX software correlator. Using existing software it was possible to characterise the network speed particularly for MPI applications. The capabilities of the DiFX software correlator, running on this cluster, were measured for a range of observation parameters and were shown to be commensurate with the generic performance parameters measured.

The feasibility of an Italian VLBI array has been explored, with discussion of the infrastructure required, the performance of such an array, possible collaborations, and science which could be achieved.

Results from a 22 GHz calibrator survey are also presented. 21 out of 33 sources were detected on a single baseline between two Italian antennas (Medicina to Noto).

The results and discussions presented in this thesis suggest that wide-field VLBI is a technique whose time has finally come. Prospects for exciting new science are discussed in the final chapter.



## Acknowledgements

This research was supported by the EU Framework 6 Marie Curie Early Stage Training programme under contract number MEST-CT-2005-19669 'ESTRELA'. I would also like to thank the Max-Planck-Institut für Radioastronomie in Bonn for their support; both in hosting me for four months at the start of 2009, and also for offering me a PhD studentship immediately after my ESTRELA fellowship came to an end. The vast majority of this thesis was written during this time and having the use of the institute facilities during writing up was invaluable.

Discussions with many people have influenced the direction I have taken in my doctoral research. In roughly chronological order I would like to acknowledge: Steven Tingay, who introduced me to the DiFX software correlator in late 2007 and whose introduction to wide-field VLBI, scribbled on one sheet of A4, went on to form the basis of this thesis. Franco Mantovani, my mentor in radio astronomy, who has been unerring in his support from when we first started working together in early 2008 right up until the submission of this thesis. Walter Brisken who has patiently responded to innumerable enquiries from me over the last two years and who, on the three occasions we have met in person, has always left me full of ideas to explore. In particular he deserves credit for a great deal of help in the implementation and mathematical description of the UV-shifting technique, and for the idea of expanding the baseline vectors to a higher-order characterisation of the delay across the primary beam. Walter Alef and Helge Rottmann who immediately took an interest in my work on DiFX before anyone could have guessed that it would lead to anything useful and who have been extremely supportive ever since. I would also like to thank Adam Deller and the rest of the 'DiFX community' for their friendly help and advice.

At IRA I would like to thank Mauro Nanni, my supervisor, for his constant good humour, for regular conversational Italian practice and for giving me total freedom in finding a PhD project and pursuing it; also Daniele Dallacasa, my university supervisor, and Hans de Ruiter, PI for the Bologna ESTRELA node, for help in negotiating Italian bureaucracy.

Thank you to my family for tolerating my insistence on living so far away; I wish I could bring them with me wherever I go. Thank you also to my friends all over the world, especially those I've made in Italy and Germany, who have left me with so many happy memories from the last three years, and who have made the low points tolerable.

John Morgan

Bonn, 12 March 2010



# Contents

<b>1</b>	<b>Radio Astronomy</b>	<b>1</b>
1.1	Radio Astronomy . . . . .	1
1.1.1	Radio Sources . . . . .	1
1.1.2	Propagation from the object to the Earth's atmosphere . . . . .	2
1.1.3	The Atmosphere . . . . .	2
1.2	The Antenna . . . . .	2
1.2.1	Parabolic antennas . . . . .	3
1.2.2	Phased Arrays . . . . .	5
1.2.3	Receivers . . . . .	7
1.2.4	Down-conversion and sampling . . . . .	8
1.3	Radio Interferometry . . . . .	9
1.3.1	Resolution of an interferometer . . . . .	9
1.3.2	Sensitivity of an interferometer . . . . .	10
1.3.3	Field of view of an interferometer . . . . .	11
1.3.4	Correlation . . . . .	13
1.3.5	The delay model . . . . .	14
1.3.6	Correlator Output . . . . .	14
1.4	Very Long Baseline Interferometry . . . . .	16
1.4.1	Timing . . . . .	16
1.4.2	Delay Modelling . . . . .	17
1.4.3	Fringe Rotation . . . . .	18
1.4.4	Transport of data to the correlator . . . . .	18
1.5	Radio Interferometers . . . . .	19
1.5.1	Connected Element Arrays . . . . .	19
1.5.2	VLBI Arrays . . . . .	20
1.5.3	SKA Pathfinders . . . . .	21
<b>2</b>	<b>Observing &amp; Data Reduction</b>	<b>23</b>
2.1	Flagging (data editing) . . . . .	23
2.2	Calibration . . . . .	24
2.2.1	Closure quantities . . . . .	24
2.3	Phase (or delay) calibration . . . . .	24
2.3.1	Clock correction . . . . .	25
2.3.2	Calibrating for the effects of the atmosphere . . . . .	25
2.3.3	Correcting instrumental errors . . . . .	25
2.3.4	Parallactic angle correction . . . . .	26
2.4	Phase calibration using astronomical sources . . . . .	26
2.4.1	Fringe fitting . . . . .	26
2.4.2	Using a phase calibrator . . . . .	26
2.4.3	Residual errors after phase referencing . . . . .	28
2.5	Amplitude calibration . . . . .	28
2.5.1	Autocorrelation Correction . . . . .	28
2.5.2	Gain Calibration . . . . .	28
2.6	Bandpass calibration . . . . .	29
2.7	Determining information from the visibility function . . . . .	29
2.8	Imaging . . . . .	30

---

2.8.1	Generating the dirty map . . . . .	30
2.8.2	Deconvolution . . . . .	31
2.9	Self-calibration . . . . .	31
2.10	Subtraction of confusing sources . . . . .	32
2.11	Interferometry with a limited number of antennas . . . . .	32
2.11.1	Two Antennas . . . . .	32
2.11.2	Three Antennas . . . . .	33
2.11.3	Four Antennas . . . . .	33
<b>3</b>	<b>VLBI in Italy</b> . . . . .	<b>35</b>
3.1	Radio astronomy in Italy . . . . .	35
3.1.1	Croce del Nord . . . . .	36
3.1.2	Medicina and Noto . . . . .	36
3.1.3	Sardinia Radio Telescope (SRT) . . . . .	36
3.1.4	Matera . . . . .	36
3.2	Infrastructure Required . . . . .	36
3.2.1	Recorded VLBI . . . . .	39
3.2.2	e-VLBI . . . . .	40
3.2.3	Correlation . . . . .	40
3.2.4	Supplementary Data . . . . .	40
3.3	Array Parameters . . . . .	40
3.3.1	Possible collaborations . . . . .	40
3.3.2	Wider collaborations . . . . .	40
3.3.3	Sensitivity . . . . .	41
3.3.4	UV coverage . . . . .	41
3.4	Science with an Italian VLBI array . . . . .	42
3.4.1	Maser Astrometry . . . . .	42
3.4.2	Bandwidth Synthesis Astrometry and Geodesy . . . . .	42
3.4.3	Geodesy . . . . .	43
3.4.4	Spacecraft tracking . . . . .	43
3.4.5	Testing e-VLBI technologies . . . . .	43
3.4.6	Antenna Tests . . . . .	43
3.5	Conclusions . . . . .	43
<b>4</b>	<b>A correlator at IRA</b> . . . . .	<b>45</b>
4.1	Different Correlator Architectures . . . . .	45
4.1.1	ASIC Architecture . . . . .	45
4.1.2	FPGAs, GPUs and other multi-purpose hardware . . . . .	46
4.1.3	Software correlators on specialised multi-purpose computers . . . . .	46
4.1.4	Software correlation on commodity PCs . . . . .	46
4.2	The DiFX software correlator . . . . .	47
4.2.1	Data flow . . . . .	48
4.3	The Cluster at IRA . . . . .	48
4.4	Cluster parameters . . . . .	48
4.4.1	CPU . . . . .	48
4.4.2	Network . . . . .	48
4.4.3	Reading the input data . . . . .	49
4.5	Characterising the cluster . . . . .	49
4.5.1	Network . . . . .	49
4.5.2	Disks . . . . .	50
4.6	Benchmarking the cluster for DiFX . . . . .	52



4.6.1	Optimum number of threads . . . . .	52
4.6.2	Assessing the performance of the cluster with different numbers of stations . . . . .	53
4.6.3	Benchmarking the cluster for disk-based correlation . . . . .	53
4.6.4	Maximum performance for a 6-station experiment . . . . .	53
4.7	Conclusions . . . . .	53
<b>5</b>	<b>Wide-field VLBI</b> . . . . .	<b>57</b>
5.1	VLBI sensitivity . . . . .	58
5.1.1	Recorded VLBI . . . . .	58
5.1.2	e-VLBI . . . . .	59
5.1.3	Other limitations on bandwidth . . . . .	59
5.2	Wide-field VLBI and the primary beam . . . . .	60
5.2.1	Sensitivity over the entire primary beam . . . . .	60
5.2.2	Combining pointings . . . . .	60
5.2.3	Mosaicing with multiple pointings . . . . .	61
5.2.4	Phase response across the primary beam . . . . .	62
5.2.5	Correcting for the primary beam . . . . .	63
5.3	Planning a wide-field VLBI observation . . . . .	63
5.4	Time-average and Bandwidth Smearing . . . . .	64
5.4.1	Limitations imposed by Correlator Hardware . . . . .	65
5.5	Calibration and imaging of wide-field datasets . . . . .	66
5.5.1	Calibration . . . . .	66
5.5.2	Imaging . . . . .	66
5.5.3	Non-coplanar Effects . . . . .	66
5.5.4	Parallel Imaging Algorithms . . . . .	66
5.6	Multiple UV datasets . . . . .	67
5.7	UV Shifting . . . . .	67
5.7.1	Residual delays and residual delay rates . . . . .	67
5.7.2	Transforming visibility datasets . . . . .	67
5.7.3	UV shifting to first order . . . . .	68
5.8	Accurate UV shifting . . . . .	69
5.9	Averaging and Imaging . . . . .	70
5.10	Conclusions . . . . .	70
<b>6</b>	<b>Testing the UV-shifting algorithm</b> . . . . .	<b>71</b>
6.1	A method to test UV-shifting algorithms . . . . .	71
6.1.1	Preparing test datasets . . . . .	71
6.1.2	Evaluating the performance of the UV-shifting algorithm . . . . .	72
6.1.3	Correlation parameters . . . . .	72
6.2	The correlator delays . . . . .	72
6.3	Results . . . . .	74
6.3.1	Shift 1 . . . . .	77
6.3.2	Shift 2 . . . . .	77
6.3.3	Shift 3 . . . . .	81
6.4	Variation of the delay over the wide field . . . . .	81
6.5	Conclusions . . . . .	81

<b>7</b>	<b>VLBI Observations of the 1320+299 Complex</b>	<b>85</b>
7.1	Motivation . . . . .	85
7.2	1320+299 . . . . .	85
7.2.1	The four-station array . . . . .	87
7.3	Observation . . . . .	87
7.3.1	The observation . . . . .	87
7.3.2	Observation technique . . . . .	88
7.4	Modifications to the correlator . . . . .	88
7.4.1	LO offsets . . . . .	88
7.5	Initial Calibration and imaging . . . . .	88
7.5.1	Initial flagging and calibration . . . . .	88
7.5.2	Imaging . . . . .	89
7.6	Initial imaging of the three components . . . . .	89
7.6.1	Imaging of the phase reference calibrator J1329+3154 . . . . .	89
7.6.2	Phase-referenced imaging of the three components . . . . .	90
7.6.3	The astrometric position of the A-component . . . . .	90
7.6.4	Fringe fitting on the A-component . . . . .	92
7.6.5	Further self-calibration of 1320+299A . . . . .	93
7.6.6	Correcting for the primary beam . . . . .	95
7.7	Applying A solutions to other datasets . . . . .	96
7.7.1	Impact of the A component on the B and C datasets . . . . .	96
7.7.2	Peeling 1320+299A . . . . .	97
7.7.3	1320+299A as an in-beam calibrator . . . . .	98
7.7.4	Conclusions . . . . .	99
7.8	Tiling the primary beam . . . . .	99
7.8.1	Area to map . . . . .	99
7.8.2	Generating the UV data . . . . .	99
7.8.3	Subtracting A . . . . .	101
7.8.4	Results . . . . .	102
7.9	Conclusion . . . . .	103
<b>8</b>	<b>A single-baseline VLBI survey of 22 GHz calibrators</b>	<b>105</b>
8.1	Motivation . . . . .	105
8.2	Baseline Parameters . . . . .	105
8.3	Observation Strategy . . . . .	105
8.4	Sources . . . . .	106
8.5	The Observation . . . . .	106
8.6	Correcting for Clock shifting . . . . .	106
8.7	Checking correlated data integrity . . . . .	108
8.7.1	Maser Jumps . . . . .	108
8.7.2	Impact of re-syncing of the data . . . . .	108
8.8	Results . . . . .	109
8.8.1	Effect of the troposphere . . . . .	109
8.9	Conclusions . . . . .	111
<b>9</b>	<b>Conclusions</b>	<b>113</b>
9.1	Science with an Italian VLBI network . . . . .	113
9.2	Wide-field VLBI . . . . .	113
9.2.1	Tiling the entire primary beam . . . . .	113
9.2.2	Surveys of compact objects . . . . .	114
9.2.3	Challenges . . . . .	115

---

9.3	The Delay across the primary beam . . . . .	115
<b>A</b>	<b>Software</b>	<b>117</b>
A.1	Correcting LO offsets in DiFX . . . . .	117
A.2	Post-correlation UV shifting . . . . .	118
A.2.1	Limitations of current software . . . . .	118
A.2.2	Where to Implement . . . . .	118
A.2.3	Implementation . . . . .	118
A.2.4	Time-averaging . . . . .	119
A.2.5	Uses . . . . .	119
A.3	Running batch jobs on the cluster . . . . .	119
<b>B</b>	<b>The IRA-DiFX software correlation system</b>	<b>121</b>
B.1	The pydifx framework . . . . .	121
B.1.1	wrapper scripts . . . . .	121
B.1.2	vex2calc . . . . .	121
B.2	Other Scripts . . . . .	122
B.2.1	log2clock . . . . .	122
B.2.2	log2tsys . . . . .	122
B.3	Mathematical functions . . . . .	122
<b>C</b>	<b>RFIX</b>	<b>123</b>
C.1	Philosophy . . . . .	123
C.2	Implementation . . . . .	123
C.3	Using RFIX . . . . .	123
C.3.1	Data flagging . . . . .	124
C.3.2	Data display . . . . .	124
C.4	Possible future work . . . . .	124



# Conventions

In this document:

- ‘CPU core’ refers to a single core of a multi-core CPU, or a single-core CPU.
- ‘Compute process’ refers to a number-crunching process of the DiFX software correlator (‘core’ in the DiFX instructions).
- ‘Antenna’ refers to a single station of an interferometer network.
- ‘UV data’ and ‘UV dataset’ refer to the correlated data (either the visibilities or the dataset including calibration tables etc.).
- ‘UV coordinates’ and ‘UVW coordinates’ refer to the baseline vectors.
- ‘Scan’ refers to a single pointing of the antennas.
- ‘Sub-band’ refers to what AIPS calls an IF: a contiguous set of spectral channels which have all passed through the same station electronics.
- HPBW and FWHM usually have the same meaning in the context of antennas (unless something other than Power is being measured). HPBW is used unless a function other than the power function or functions in the abstract sense are being referred to: FWHM is used for these cases.

## Properties of VLBI Antennas mentioned in this thesis

Antennas are often referred to by their two letter DB Codes. These are given below along with some general information about the antennas. There is no DB Code yet listed for the Sardinia Radio Telescope so the abbreviation SRT is used instead.

DB Code	Antenna	$d$ (m)	Latitude	Longitude	Network(s)
Ef	Effelsberg	100			EVN
Mc	Medicina	32			EVN
Ma(Mr)	Matera	25			Geodesy
Nt	Noto	32			EVN
Pa	Parkes	64			LBA
SRT	Sardinia	64			EVN
Wz	Wettzell	25			Geodesy
Ts	Tsukuba	32			Geodesy

## Abbreviations used in this thesis

AIPS	Astronomy Imaging Processing System
AOC	(NRAO) Array Operations Centre
API	Application Programming Interface
ASI	Agenzia Spaziale Italiana (Italian Space Agency)
ATCA	Australian Telescope Compact Array
BBC	Base-Band Converter
DBBC	Digital Base-Band Converter
DBE	Digital Back-End
Decl.	Declination
EOPs	Earth Orientation Parameters
erf	Standard Error Function
FWHM	Full-Width Half Maximum
FPGA	Field-Programmable Gate Array
GMRT	(Indian) Giant Metre-wave Radio Telescope
GNSS	Global Navigation Satellite System
GPU	Graphical Processing Unit
HWHM	Half-Width Half Maximum
HPBW	Half-power Beam Width
IRA	Istituto di RadioAstronomia ((Italian) Institute of Radio Astronomy)
INAF	Istituto Nazionale di AstroFisica (National Institute of Astrophysics)
JIVE	Joint Institute for VLBI in Europe
LBA	(Australian) Long Baseline Array
LCP	Left-hand Circular Polarisation
LO	Local Oscillator
Maser	Microwave Amplification through Stimulated Emission of Radiation
PSF	Point Spread Function
RA	Right Ascension
RAID	Redundant Array of Inexpensive Disks
RCP	Right-hand Circular Polarisation
RMS	Root Mean Square(d)
RFI	Radio-frequency Interference
SEFD	System Equivalent Flux Density
SI	le Système International d'unités
SNR	Signal to Noise Ratio
SRT	Sardinia Radio Telescope
TAI	Temps Atomique International (International Atomic Time)
UT1	Universal Time 1
UTC	Coordinated Universal Time
VLA	Very Large Array
VLBA	(US) Very Long Baseline Array
VLBI	Very Long Baseline Interferometry
WSRT	Westerbork Synthesis Radio Telescope
WIDAR	Wide-band Digital Interferometric ARchitecture
VERA	VLBI Exploration of Radio Astrometry (Japanese Array)

# Chapter 1

## Introduction to Radio Astronomy

In this chapter I will discuss the hardware and techniques used for radio astronomy observations. I will begin by describing the radio sources and their propagation to the receiver. I will then describe radio antennas and their properties. Much of what is contained in these first sections will also be applicable to radio observations other than those that use interferometry technique, but have a strong bias towards the information relevant for radio interferometry.

I will then introduce radio interferometry and associated concepts such as correlation. I will then introduce Very Long Baseline Interferometry (VLBI): the specific type of radio interferometry used for observations presented in this thesis and the variety for which the techniques presented in this thesis are most relevant.

I will end by describing the properties of various radio interferometers in use and in development today.

### 1.1 Radio Astronomy

#### 1.1.1 Radio Sources

Radio astronomy studies the longest wavelength electromagnetic radiation which can be detected. At least for Earth-based instruments this corresponds very roughly to a frequency between  $\sim 10$  MHz and  $\sim 1$  THz (i.e. a wavelength of between 3 m and 0.3 mm)<sup>1</sup>. Objects which emit radio waves include:

- The Sun and other stars
- Jupiter (and other Jupiter-class planets)
- Supernova Remnants
- Quasars and Galaxies
- Micro-quasars
- Pulsars
- Spacecraft
- Molecular clouds

There are many different mechanisms which can generate radio emission. These include thermal emission (black-body and free-free), synchrotron emission, artificial transmissions, atomic (or molecular) emission lines and masers.

---

<sup>1</sup> There are many decades of band below 10 MHz but observing them from below the ionosphere is very difficult. Observations have been made below 10 MHz in space (Kaiser 1990) and indeed from the ground in the early days of radio astronomy (Kellermann 2004)

### 1.1.2 Propagation from the object to the Earth's atmosphere

Many physical effects can have an influence on the propagating radiation, the most important one being absorption. This and other subtle effects may be relevant for some observations; for example pulsar dispersion and scintillation. Another important influence on radiation is redshift (due to movement of the source with respect to the Earth and Expansion of the Universe).

It should be noted that radio observations are often made precisely to understand these intermediate processes. Examples include Redshift surveys, Faraday rotation studies and measuring the Sunyaev-Zel'dovich effect on the Cosmic Microwave Background.

### 1.1.3 The Atmosphere

Except for rare occasions when radio observations are made from space the atmosphere will be the one of the main contributors to unmodelled effects which must be calibrated for.

#### The Ionosphere

The ionosphere extends from 60 to 2000 km above the Earth's surface (Fomalont and Perley 1999). Its principal effect is to refract the incoming radiation. Faraday rotation also rotates the plane of polarisation of the incoming radiation. The excess path  $L_i$  is given by the equation

$$L_i \propto -\nu^{-2} N_e \quad (1.1)$$

where  $\nu$  is the frequency and  $N_e$  is electron column density. A typical value of  $L_i$  is  $\sim 2$  m at 2 GHz.

The (negative) delay that the negative excess path causes incurs a phase shift and since the phase shift  $\phi$  for a given delay is proportional to frequency:

$$\phi = \frac{2\pi c}{\nu L_i}. \quad (1.2)$$

It is important to note that if this effect were entirely static it would have absolutely no effect on astronomical observations. However this effect is inherently unstable, not only varying with time on a variety of cycles, but also from place to place.

#### The Troposphere

At frequencies above 2 GHz where the effect of the ionosphere is less significant, the effects of the troposphere begin to dominate. The troposphere is modelled as having a dry component (6 km thick) and a wet component (2 km thick). The dry component is fixed at 2.3 metre delay at the zenith. The wet component adds a delay of between 0 and 0.3 m.

## 1.2 The Antenna

In this section I introduce different kinds of antennas and their properties. Parabolic antennas have been used for all of the observations presented in this thesis and have been used for almost all radio astronomy from the late 1970s until the present day. There is one significant exception: when connected arrays such as the VLA



or the WSRT are used as an element in a VLBI array they are “phased up” to work as phased arrays. In this case the individual elements are still parabolic antennas, but the response of the entire array is that of a phased array.

Other kinds of antenna such as horns (e.g. the Holmdel Horn Antenna), cylindrical telescopes (e.g. the Croce del Nord, see sect. 3.1.1) and Krauss-type (such as the Nançay Decimetric Radio Telescope) have been used in the past and some remain in use.

More recently phased-array designs have returned to fashion. Examples include LOFAR and the MWA (see sect. 1.5)

### 1.2.1 Parabolic antennas

Parabolic antennas use a parabolic reflector (or mirror) to focus all of the received radiation from the target area on the sky onto a small area. One advantage of this approach is that the receiver may be cooled. It also means that the receiver is only sensitive to a small patch of sky. This may be an advantage (elimination of strong signals outside the region of interest) or a disadvantage (limited field of view). Just as with reflecting telescopes at other wavelengths, various optical configurations are possible, and often several focal planes are used within a single instrument (the SRT will use three, see Olmi and Grueff 2006).

#### Resolution of a parabolic antenna

Following Napier (1999), the voltage pattern ( $F$ ) of a parabolic antenna is given by the Fourier transform of the aperture function. The power function  $A = F^2$ . Since the Fourier transform of a top-hat function is the sinc function

$$\text{sinc } \theta = \frac{\sin \theta}{\theta}, \quad (1.3)$$

The radial power function  $A$  for a perfect rotationally symmetric, uniformly illuminated parabola (i.e. top-hat aperture function) at a radial offset  $\theta$  from the pointing centre is given by

$$A(\theta) = A_{max} \cdot \text{sinc}^2 \frac{\theta}{\sigma} \quad (1.4)$$

where

$$\sigma = 2 \cdot \ln(2) \cdot \text{FWHM}. \quad (1.5)$$

The FWHM of the power function (herein referred to as the HPBW or the primary beam) is given by

$$\sin(\text{FWHM}) = 1.22 \cdot \frac{\lambda}{d}. \quad (1.6)$$

where  $d$  is the diameter of the parabola and  $\lambda$  is the wavelength of the incident radiation. The width of the HPBW of the primary beam for radio frequencies according to this equation is shown for several antennas in fig. 1.1.

#### Aperture function of the reflector

Many antennas, however, do not have an aperture function which is top-hat shaped. This is because it is desirable to suppress the sidelobes to reduce increasing the system temperature with spurious radiation. The optics of the telescope




Figure 1.1: The HPBW of the main beam for several antennas in arcsecs (left y axis) and degrees (right y axis) as a function of frequency.

will also have an influence on whether the feeds can or should “illuminate” the dish uniformly. The many considerations that informed the design of the optics of the SRT are set out in sect. 3 of Olmi and Grueff (2006). For the VLA dishes uniform illumination is sought (see Napier 1999, sect. 2). For most purposes (correcting for the main beam for sources away from the pointing centre for example) we assume a radial *Gaussian* power function.<sup>2</sup>

$$A(\theta) = A_{max} \cdot \exp\left(\frac{-\theta^2}{2\sigma^2}\right) \quad (1.7)$$

where

$$\sigma = \frac{\text{FWHM}}{2\sqrt{2 \cdot \ln(2)}} \quad (1.8)$$

with FWHM given by equation 1.6.

Figure 1.2 compares the primary beam approximated by equations 1.4 and 1.7 for 100 m telescopes and 25 m telescopes. It is clear that the Gaussian approximation is not valid much beyond the HPBW as the null in the  $\text{sinc}^2$  function is not accounted for. For this reason the Gaussian is often truncated at approximately double the HPBW. Empirical values for the VLA suggest that the main error may be the effective diameter of the dish rather than the shape of the function (see fig. 1.3). In reality the sidelobes are often not radially symmetric due to aperture blockage: see for example fig. 3 of Reich et al. (1978) for Ef and fig. 3-11 of Napier (1999) for the VLA antennas. At least for antennas with an Alt-Azimuth mount the primary beam orientation with respect to an astronomical source will change with hour-angle. For all antennas the primary beam will change to some degree

<sup>2</sup>Even for arrays where uniform illumination is sought the Gaussian approximation is often used (e.g. the VLA Holdaway 1999, sect. 1.2)

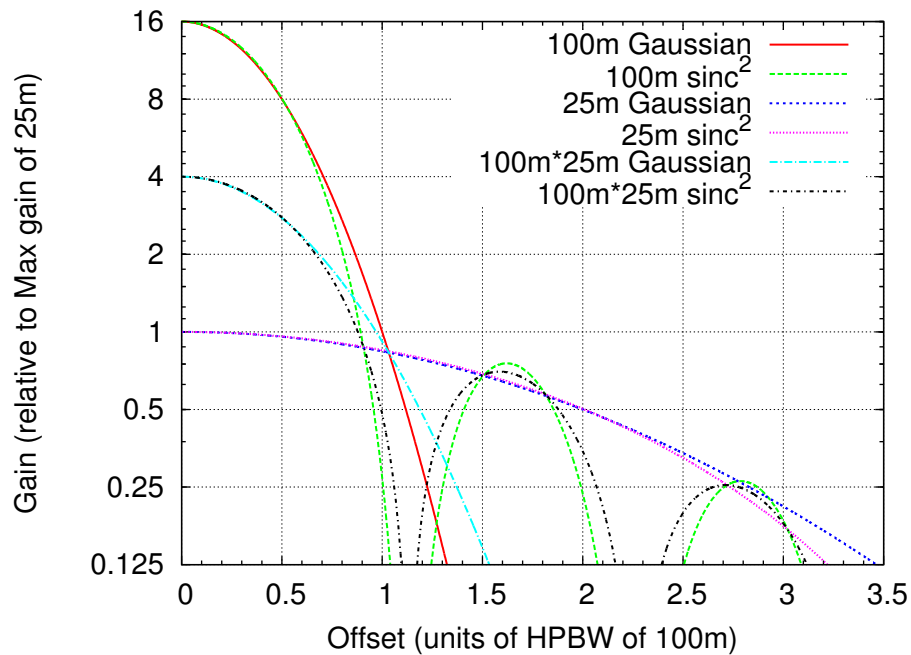


Figure 1.2: Approximations of the sensitivity pattern of a 100 m telescope and a 25 m telescope approximated first by a Gaussian then by a  $\text{sinc}^2$  function, both of FWHM given by equation 1.6. It is assumed that both have the same peak sensitivity.

with declination due to deformation. Finally, for many larger dishes (for example Ef and Pa), only the central portion of the dish is designed to be sufficiently accurate to be used as high frequency. For these dishes particularly (but for other dishes too) there is unlikely to be a simple linear relationship between primary beam size and frequency.

### 1.2.2 Phased Arrays

Phased arrays consist of an array of antennas whose individual elements are all sensitive to the same patch of sky. An area on the sky (a beam) is then picked out by adding a delay to the signal from each element and adding them together in such a way that signals originating from the area of interest will be added coherently— a process known as ‘beam-forming’. This is similar to the delay model of a correlator (see sect. 1.3.4). Provided there is sufficient processing power, many beams can be formed simultaneously. For telescopes such as LOFAR where the individual elements are sensitive to the entire sky, only limited data bandwidth (and computing power to model the ionosphere and correlate the baseband data) prevent imaging the entire sky at full sensitivity and resolution.

#### Resolution of a phased array

Just as with parabolic antennas, the sensitivity of a phased array is given by the square of the Fourier transform of the voltage pattern of the entire array. The resolution is therefore approximately that of a reflector with the dimensions of the array. It is identical to the instantaneous synthesised beam of the array were it being used as an interferometer (see sect. 1.3.1). For example the sensitivity

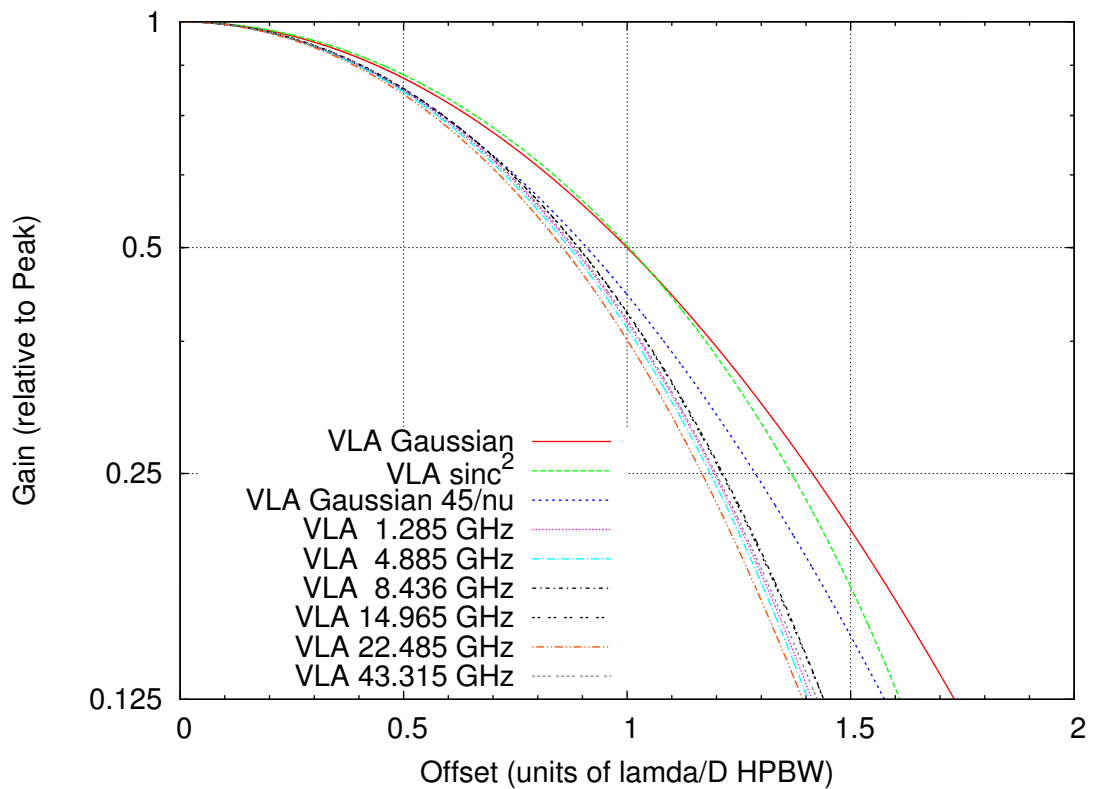


Figure 1.3: The main beam of the VLA determined empirically and modelled using polynomial functions (see the documentation for the AIPS task PBCOR). These have been scaled to frequency so that they can be directly compared with the Gaussian and  $\text{sinc}^2$  approximations and with each other. The  $45/\nu$  approximation is from Ulvestad et al. (2009). Note that these are still radially symmetric distributions and continue to drop off with increasing offset (i.e. they do not attempt to model the nulls and sidelobes). From this figure it can clearly be seen that the shape of the function beyond the HPBW follows the  $\text{sinc}^2$  approximation more closely than the Gaussian approximation.

pattern of the WSRT in phased-array is a narrow fan beam which changes orientation with respect to the source at different hour angles. The sensitivity pattern of the VLA is a “snowflake” with sensitivity away from the central lobe strongly concentrated in six radial lines.

This narrow beam increases the resolution (in only one direction for the WSRT). This means that when a wide field of view is desired for VLBI either greater bandwidth must be used or sensitivity must be sacrificed: the elements of the connected element array are either correlated individually or, as is more usual, only a single element is used.

Another important point is that the orientation of the entire array will change with respect to the source as the earth rotates, causing a time-dependence in the beam pattern.

### 1.2.3 Receivers

The incident radiation is then detected by some kind of receiver. The many different forms of receiver are not directly relevant to this thesis, however I will make a brief reference to a type of feed which will likely have great importance in the future for interferometry.

#### Multi-feeds and Imaging Arrays

Radio telescopes are typically single pixel instruments: that is, only a single detector is placed in the focal plane<sup>3</sup>. Adding a second feed at an offset in the focal plane greater than the width of the primary beam doubles the amount of sky visible to the antenna for all experiments (and a survey can proceed at twice the speed).

Such a feed is particularly useful at high frequencies where the high resolution reduces the solid-angle on the sky to which the antenna is sensitive. The Medicina 22 GHz feed has 6 feeds arranged in a circle around a central feed — an arrangement which allows 7 parallel traces across the sky (Verma et al. 2008).

This concept can be extended to a two-dimensional imaging array (as commonly employed at optical wavelengths) Examples include Apertif (Verheijen et al. 2008) and ASKAP (Johnston et al. 2008). The disadvantage of this approach is that the receiver may be too bulky to be cooled and so the system temperature (see sect. 2.5.2) is greatly increased. However if the receivers can be made cheaply enough then there is the possibility of doing away with the reflector altogether and building a phased array of these receivers.

For VLBI, the amount of processing required to image a single primary beam is already huge and the value of multi-feeds is questionable for the near future. However one system particularly suited to interferometry is the system employed by VERA, where the second feed may be moved within the focal plane to allow constant monitoring of a suitable phase calibrator.

Another possible VLBI application of a multi-feed would be correlating several multi-feeds for larger antennas where the field of view is reduced. However, since VLBI is bandwidth limited, this would come with a sensitivity penalty.

---

<sup>3</sup>Although recording multiple frequencies and polarisations does result in more than one data-stream.

### 1.2.4 Down-conversion and sampling

This section describes the next step of the process for interferometry, although much will be applicable to single-dish observations as well.

Once the signal has been detected the polarisations are split at the feed, phase calibration tones are added, a low-noise amplifier is used to increase the signal which is then down-mixed to an intermediate frequency (IF)  $\sim 500$  MHz for transmission off the antenna. The signal is then re-amplified and sent to several BBCs each of which will extract a single sub-band. The bandwidth of each sub-band is  $\sim 10$  MHz. 8 or 16 MHz is typical although much less may be sufficient for some spectral line observations. At this point the signal is ready to be digitised<sup>4</sup>.

This is done by sampling the voltage of the antenna repeatedly to give a time-series of voltages herein referred to as a datastream. In most cases, and certainly for VLBI experiments, the maximum sensitivity is limited by the bandwidth. There are two choices to be made: the precision of each sample and the sampling rate.

The sampling rate depends on the bandwidth of the sub-band. A bandwidth of  $\Delta\nu$  MHz consists of a continuum of frequencies from DC ( $f = 0$ ) to  $F = \Delta\nu$ . In order for the frequency  $F$  to be sampled there must clearly be at least 2 samples in each  $(1/F)$  time period. This imposes a minimum sampling rate of  $2 \times F$ , the Nyquist rate. Finite sampling will not cause limitations provided that the sampled signal is properly filtered: sampling causes aliasing of every frequency  $f$  onto frequency  $(f \bmod F)$ .

Maximum sensitivity per bit is achieved with very low numbers of bits per sample. Even with only a single bit per sample the sensitivity is 64% of that which would be achieved with no quantisation. Typically for VLBI either one or two-bit sampling is used both of which have almost identical sensitivity per bit (Thompson et al. 1986, tables 8.1 and 8.1). Note that motivations other than increased sensitivity may lead to the use of a higher number of bits per sample; for example, ensuring that an astronomical signal can still be recovered after subtraction of RFI. Sampling at higher than the Nyquist rate can recover some of the sensitivity lost due to the small number of bits per sample (Thompson et al. 1986, equation 8.14).

For interferometry it is vital that the sampling rate be tied to a frequency standard. For connected-element arrays a single time standard is broadcast to all antennas. The various local oscillators must also be tied to this time standard. In the case of VLBI, time stamps are added within the datastream itself to provide an absolute reference time.

An additional complication can be fixed at this stage: each of the antennas will have a different velocity with respect to the source compared to the correlation centre of the array. Typically for connected element interferometers this is corrected by adjusting the LO frequency. For VLBI observations this is corrected within the correlator (see sect. 1.4.3).

The use of such a low bit-rate per sample essentially discards the amplitude information (of the voltages not the correlation coefficients). This is restored either by observing a source of known amplitude or by using regular (every few minutes) measurements of the off-source system temperature (see sect. 2.5.2).

---

<sup>4</sup>Digital back-ends are in development in which the signal is digitised at intermediate frequency and the final down-conversion to baseband is done digitally.

## 1.3 Radio Interferometry

Motivated by the lack of resolution of even the largest antennas and the lack of radio-frequency imaging arrays, interferometry allows high resolution and multi-pixel images to be made at radio frequencies. The high resolution is obtained by combining the signals of two distant antennas to form an interferometer whose fringe-spacing (resolution) is given by  $\lambda/D$  where  $\lambda$  is the wavelength and  $D$  is the projected baseline (the distance between the antennas as viewed from the source). Multi-pixel images are synthesised by viewing the source on many baselines, perhaps by using multiple antennas but also taking advantage of the fact that as the Earth turns, so does each baseline's orientation with respect to the source — a technique known as “Earth Rotation Synthesis”. When this technique was first developed (Ryle et al. 1950) it was only possible to measure the so-called “visibility function” and attempt to exclude or admit various models of the sky brightness (see sect. 2.7).

However the increasing power of digital computers, and the FFT algorithm (which allows very efficient calculation of a discrete Fourier transform) allowed direct imaging of the sky brightness (see sect. 2.8).

The images produced by an interferometer have important limitations when compared with those that would be produced by an imaging array covering the focal plane of a filled-aperture instrument the size of the interferometer: such an instrument would have a maximum resolution (there would be an upper limit to the spatial frequencies it could sample) but structures of all sizes up to the size of the image would be visible. Interferometers on the other hand, tend only to measure restricted spatial frequencies; so not only would there be a minimum size of object that would be seen, but also a *maximum* (baselines can come close to zero spacing but rarely reach it).

Some spatial frequencies will be sampled less than others, and the spectrum of spatial frequencies which have been sampled will be different for different angles across the source. The less data there are the less statistical likelihood that the most likely solution for sky brightness is, in fact, the correct one (see sect. 2.9).

Nonetheless interferometry provides the only practical way to image the sky at high resolution at radio frequency. A large part of this thesis relies on the technique of sampling each baseline not only at zero delay (signals combined to sample the sky brightness at the centre of the source) but also at a number of other delays, allowing the sky brightness further from the central “phase centre” to be calculated. These issues will be returned to in chap. 5, but first the basic properties of an interferometer (which don't require a full understanding of how the signals of an interferometer are combined) will be derived (its resolution, sensitivity and field of view). Then the details of combining (correlating) the data from the individual elements will be described and the output of this process will be considered. Finally the limitations of interferometry with small numbers of antennas are considered and the technique of VLBI is introduced.

### 1.3.1 Resolution of an interferometer

As mentioned, the fringe spacing of a single baseline is given by  $\lambda/D$  where  $\lambda$  is the wavelength and  $D$  is the projected baseline length. For the ensemble of baselines measured in an observation (different antennas, baseline orientations and frequencies) it could be thought that the maximum resolution in a particular direction across the source would simply be given by the longest baseline (and

highest frequency) along that axis. However unless we throw away all of the data *except* that which gives us the highest resolution (losing lots of sensitivity in the process) the highest spacial frequencies sampled in the image will inevitably be convolved with lower spacial frequencies, reducing the resolution. We could also choose throw away the data on the longer baselines (especially if it is clear that no flux has been resolved and they contain only noise which would increase the SNR of the data which would clearly decrease the resolution).

In other words, we have a choice in how we weight our data. The *maximum* resolution is given by the longest projected baseline and the minimum by the shortest.

Perhaps a good way to visualise the resolution of an interferometer is considering the “dirty beam”. After Fourier-inversion imaging a “dirty image” is produced which is the measure of sky brightness convolved with the “dirty beam”. This is simply the Fourier transform of the voltage pattern of the entire array.

Usually a resolution close to the maximum is assumed.

### 1.3.2 Sensitivity of an interferometer

First we must note that the power of our source and the sensitivity of our instrument is measured in Flux Density. This is measured in Janskys where

$$1 \text{ Jy} = 1 \times 10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1}. \quad (1.9)$$

A measure of the sensitivity of an antenna is the SEFD (as defined in Wrobel and Walker 1999). Essentially it is the sensitivity of a source that would give the same amount of power as the noise measured by the antenna when off-source. This takes into account the efficiency, collecting area and system noise of the antenna. Tables of SEFD at each frequency band are available for all antennas<sup>5</sup>.

The RMS noise  $\Delta S$  of a particular baseline is given by Wrobel and Walker (1999) equation 9.23 ( $N = 2$  case):

$$\Delta S = \frac{1}{\eta_s} \sqrt{\frac{SEFD_1 \times SEFD_2}{2 \cdot \Delta\nu \cdot \tau_{int}}} \quad (1.10)$$

where  $SEFD_1$  and  $SEFD_2$  are the SEFDs of the two antennas,  $\Delta\nu$  is that total instantaneous bandwidth (including all recorded polarisations),  $\tau_{int}$  is the integration time and  $\eta_s$  is the efficiency loss related to the data recording and the correlator (typically dominated by the sampling loss see sect. 1.2.4).

To find the RMS minimum thermal noise in an image generated by an array  $\Delta I_m$ ,

$$\Delta I_m = \frac{1}{\eta_s} \frac{SEFD_{array}}{\sqrt{2 \cdot \Delta\nu \cdot \tau_{int}}} \quad (1.11)$$

where, for a homogeneous array

$$SEFD_{array} = \frac{SEFD_{antenna}}{\sqrt{N(N-1)/2}} \quad (1.12)$$

which is derived from the more general case

$$SEFD_{array} = \left( \sum_{i=1, j=i+1}^{i=N-1, j=N} (SEFD_i \times SEFD_j)^{-1} \right)^{-1/2} \quad (1.13)$$

---

<sup>5</sup>E.g. for the EVN telescopes see Polatidis (2009).



This is simply the sum over all baselines where each baseline is composed of antennas  $i$  and  $j$ . See the EVN Calculator<sup>6</sup>.

Note that  $\Delta\nu/pol$  (the instantaneous bandwidth per polarisation in Hz) is related to the (digital) recording rate by

$$\Delta\nu/pol = \frac{datarate}{2 \times n_{pol} \times n_{bits}} \quad (1.14)$$

where the *datarate* is in bits per second and the factor of two comes from the Nyquist sampling rate. Note also that  $\eta$  changes depending on  $n_{bits}$  (see Thompson et al. 1986).

### 1.3.3 Field of view of an interferometer

The field of view of an interferometer is limited to the area of sky to which each of the antennas is sensitive. The field of view for each of its elements is given by the resolution of that element acting alone. See sect. 1.2 for the resolution of different types of antenna.

In Strom (2004) it is shown that the beam response  $A_{12}$  of a *baseline* is given by

$$A_{12} = F_1 \times F_2 \quad (1.15)$$

where  $F_1$  and  $F_2$  are the voltage patterns of the two antennas.

#### Size of the primary beam for homogeneous arrays

For the case of a homogeneous array  $F_1 = F_2$  and so the “primary beam envelope” of a homogeneous baseline is identical in shape to the sensitivity pattern of a single dish. This is shown in fig. 1.2 for antennas of 100 m and 25 m assuming both a Gaussian and  $\text{sinc}^2$  primary beam sensitivity pattern.

#### Size of the primary beam for heterogeneous arrays

One rather intractable problem is the size of primary beam we should assume for a heterogeneous array.

Let us consider a heterogeneous array consisting of one 100 m telescope and several 25 m telescopes. Let us first take conservative position that the primary beam of the array is given simply by the HPBW of the largest antenna (i.e. that which has the smallest primary beam). The flaw in this argument is that since the sensitivity at the pointing centre of the 100 m dish is so much higher, its sensitivity at the edge of its primary beam is  $16\times$  higher than the sensitivity of a 25 m dish at the edge of *its* primary beam. Indeed we can see from fig. 1.2 that assuming a Gaussian function, the 100 m dish only becomes less sensitive than the 25 m dish at more than  $2\times$  the supposed edge of its primary beam.

Figure 1.2 also shows the approximate primary beam of a baseline of a 100 m and a 25 m dish given by

$$A_{12} = \sqrt{A_1 \times A_2} \quad (1.16)$$

(c.f. equation 1.15) where  $A_1$  and  $A_2$  are the  $\text{sinc}^2$  and Gaussian power functions are also plotted. Assuming a  $\text{sinc}^2$  beam (since the Gaussian approximation is

<sup>6</sup>available at <http://www.evlbi.org/cgi-bin/EVNcalc>

clearly not valid as we approach the null) this function only reaches the HPBW sensitivity of the 25 m at approximately  $2\times$  the HPBW. The sensitivity at this radius should therefore be at least as strong as a baseline of two 25 m telescopes would be. It is therefore safe to extend the assumed primary beam somewhat beyond that of the largest antennas.

An increase in the sensitive area by a factor of four significantly increases the area over which in-beam calibrators can be sought. However it should be noted that the baseline (be it estimated to be a Gaussian function or a  $\text{sinc}^2$  function) is very steep beyond the HPBW and even a small error could result in a very large drop in sensitivity: again assuming a  $\text{sinc}^2$  beam, a further offset of  $0.05\times$ HPBW beyond  $2\times$ HPBW reduces the sensitivity by a further factor of two. Two possible sources of an error of this magnitude are noted below.

### Pointing Errors

$0.05\times$  the HPBW would be an error in pointing of less than  $4''$  at above 10 GHz (fig. 1.1) for a 100 m telescope. This is comparable with the pointing error of Ef (Bach 2009).

### Frequency Dependence

It should also be noted that size of the primary beam is frequency dependent and for wide-band observations will change significantly across the band. Fortunately the fractional bandwidth ( $\Delta\nu/\nu$ ) is smaller for higher-frequency observations, so pointing errors and frequency dependent primary beams counteract each other to some degree.

However this would cause particular problems for geodesy and astrometry observations carried out using the bandwidth synthesis technique (see sect. 2.3.2.) In this case the maximum field of view would be dictated by the primary beam at the highest frequency and careful correction would have to be made to the primary beam at every frequency.

### True shape of the Primary Beam

Finally and most importantly, it should be reiterated that both the Gaussian and the  $\text{sinc}^2$  functions are only approximations. As noted in sect. 1.2.1 the shape of the primary beam varies with hour-angle, declination and non-linearly with frequency. The only way to be sure of the sensitivity beyond the HPBW is through empirical observation *at* the desired frequency. A procedure for doing this is outlined in Fomalont and Perley (1999) sect. 10.1.

### Applying primary beam correction

Since it is more usual for lower resolution interferometers to image the entire primary beam, correction of the primary beam is well understood for arrays such as the VLA and ATCA. It is usually done as the final step once the image is prepared, by applying a radially symmetric polynomial which has been derived empirically (see Fomalont (1999a) sect. 1.3. See also fig. 1.3 for plots of these polynomials for the VLA).

A method for correcting for the primary beam for heterogeneous arrays is discussed in sect. 5.2.5.

### 1.3.4 Correlation

In the correlator the datastreams (a stream of voltage samples from the antenna) are aligned and compared (multiplied). The results of these multiplications are then accumulated over a certain time period. We will explore the output of a correlator assuming that two stationary antennas, equidistant from a radio source, produce datastreams which are then fed into a correlator.

#### Complex correlation

Consider the monochromatic case, i.e. a single sinusoid (a delta function in frequency space): If the signals in the two datastreams are exactly in phase the accumulation of the correlation will be a large positive number. If the signals in the two datastreams are exactly  $180^\circ$  out of phase the sum of the correlation will be a large negative number. However if the signals in the two datastreams are exactly  $90^\circ$  out of phase the correlation will average to zero — just the same as if there were no signal at all. For this reason it is necessary to do a complex correlation, where the imaginary part is calculated by delaying one of the datastreams by  $90^\circ$ . The amplitude of the complex number produced will be exactly constant for any phase difference between the two incoming monochromatic waves and the phase of the complex number produces will be the same as the phase difference between the incoming datastreams.

#### Multiple lags

However, we cannot assume a periodicity in our signal (other than those imposed by sampling theorem for sampled data and the filters). Let us imagine that there is an unmodelled instrumental delay in one of our datastreams. For noise-like continuum astronomical data, the correlation function drops off steeply as we move away from a zero delay shift between the two signals. It is exactly this property which allows precise astrometry (measurement of the location on the sky the flux is coming from) and geodesy (precise measurement of the baseline). This means that our correlator will only produce a large amplitude if the absolute offset between the two incoming datastreams is very small. We can deduce this delay error by correlating the datastreams with multiple offsets between them. The delay error can be found by looking at the amplitudes of the various lags. The source of the delay error could also be that the source of the radiation is not equidistant from the two interferometer elements (or that our coordinates for the location of the antennas is wrong).

It can also be shown (Romney 1999, equation 4-6) that the Fourier transform of the correlation function is a spectrum of the cross correlation.

Thus the necessity to correlate multiple lags derives from three reasons.

1. To ensure that the peak of the correlation function is found even if there is a delay error.
2. To allow the measurement of the sky brightness away from the point perpendicular to our interferometer.
3. To allow multiple spectral channels from one datastream.

If we assume that the signal has no repetition whatsoever, it will appear in only one lag. If we assume that we have a single sine wave this will appear in all lags, with each lag having a phase shift.

### Multiple time integrations

Now we remove the assumption that the antennas are stationary and equidistant from the source. As the Earth turns the baseline orientation with respect to the source changes. For sources away from the North pole the projected baseline length will change as well. By accounting for the distance from the source to each interferometer (by adding a delay to each datastream) and the movement of the elements with respect to the source (by adjusting the frequency of the datastream or changing the delay with time) we can ensure that for a point-like source at the “phase centre” (dictated by our model for adjusting the delays) the correlation function remains peaked at zero lag.

### XF correlators

What has been described is an XF correlator: first the correlation coefficients are found for each lag to produce a discrete correlation function, then this correlation function is Fourier transformed to produce a spectrum.

### FX correlators

Another possibility is to take a chunk of sampled data from each datastream ( $2 \times n_{channels}$  samples where  $n_{channels}$  is the number of spectral channels we require), Fourier transform it and then multiply the results for the two datastreams together.  $N$  chunks are then accumulated with  $N$  given by

$$N = \frac{2 \times n_{channel} \times \tau_{int}}{f_s} \quad (1.17)$$

where  $\tau_{int}$  is the desired integration time. The results of XF and FX correlators are virtually indistinguishable (see Romney (1999) sect. 6.4 for some minor differences) and the choice of one or the other depends on the ease and efficiency of implementation.

### 1.3.5 The delay model

Whether the antennas are on the surface of the Earth or in space, they will be moving rapidly with respect to the source and each other. In order to take account of this we choose a point (the geocentre for VLBI) and define a plane passing through this point and perpendicular to the point in the sky that we are interested in. This point in the sky is the “phase centre”. We now adjust the delay on each datastream so that it is as if each antenna was located on this plane.

### 1.3.6 Correlator Output

In this section we consider the output of the correlator (UV Data) in two ways: first as simply as a data structure (disregarding the physical meaning of the data) and secondly as a sampling of the UV plane. We assume that both the X and F operations have been done, and that so the output is in frequency space.

### UV data as a computer dataset

While the above description is vitally important when manipulating the data, equally important is understanding the basic structure of the data. We have

$N(N - 1)/2$  baselines where  $N$  is the number of antennas. For each baseline we will have a number of sub-bands. If both polarisations are recorded then we will have two polarisations (or four if the cross-polarisations have been correlated). Each sub-band will also be split into channels (from the Fourier transform in the correlator). Finally we have the time integrations.

So essentially UV Data is 4 (or 5) dimensional with axes of time, baseline<sup>7</sup>, sub-band, (polarisation), and spectral channel. The order above (spectral channel being the fastest-varying index, time being the slowest) is typical for data direct from the correlator.

Other data usually included are the delay model used, and various other data useful for calibration and imaging.

### Coordinates for correlated data

Consider a single datum: one complex number representing the result of multiplying the data from two particular antennas together at a particular frequency at a particular time. The phase representing the relative phase of the two waveforms, and the amplitude representing the correlation coefficient. At this time the baseline (the vector connecting one antenna to the other) will have a particular orientation. The relevant description of this baseline is its orientation with respect to the phase centre scaled by the wavelength of the radiation.

We therefore choose one vector  $w$  as the baseline in the direction of the source  $\vec{s}$ . Note that the vector  $w$  is exactly the same as the delay described in sect. 1.3.5. We need two more orthogonal vectors to describe the baseline in 3D space and we can choose their orientation: we orient  $u$  towards the East and  $v$  towards the North.

We now define a plane perpendicular to  $w$  passing through our source. Consider a phase centre  $\vec{s}'$  located a distance  $\sqrt{l^2 + m^2}$  away from  $\vec{s}$  where  $l$  and  $m$  are the cosines of the angle offset in northerly and southerly direction respectively. We can calculate  $w'$ : the baseline vector for the *same* baseline but for a different phase centre from the original baseline vector  $u$ ,  $v$  and  $w$ :

$$\vec{w}' = w + lu + mv. \quad (1.18)$$

The UV data therefore consists of:

1. The right ascension and declination of the phase centre
2. Baseline vectors for each baseline and time integration
3. The visibility data themselves

The equivalence of  $w$  and the delay  $\tau$  and the formulation of  $u$  and  $v$  both as baseline vector components and derivatives of  $w$  with respect to  $l$  and  $m$  come from solving two different problems: using a delay model at the correlator to stop the fringes, and defining a coordinate system for describing the output of the correlator. These relationships are very important when transforming a UV dataset to have a different phase centre (see sect. 5.7.2).

<sup>7</sup>This could also be indexed by the two antennas which make up each baseline

## 1.4 Very Long Baseline Interferometry

VLBI is a technique whereby antennas which aren't necessarily physically connected, or in contact in any way during the observation, can still act as an interferometer. This is usually done by recording the data, with timestamps attached which is then correlated at some later date. e-VLBI (where the data are not recorded but instead sent directly via a network connection to the correlator) blurs the distinction between VLBI and connected-element interferometry somewhat. However the defining difference is that connected-element arrays distribute the LO signal and the individual antennas are synchronised to  $\sim 1$  ps (even if this is over astonishingly long distances for example e-MERLIN, see McCool et al. 2008). Although this frequency standard may be flawed, provided it does not fluctuate significantly on the timescale of microseconds (the likely delay between the signal arriving on different baselines) datastreams will remain synchronised. For VLBI there is no time-stable communication between the antennas. Each antenna has its own time standard and these must somehow be compared during the observation.

This section deals with the complications introduced by this lack of timing connection between the antennas and the extremely long baselines.

### 1.4.1 Timing

The electronics of the antenna (for example the LO oscillators and the sampler) are driven by some frequency standard. VLBI only became practical once oscillators were developed which were sufficiently stable. These frequency standards are chosen for their stability over this time range of the observation rather than absolute accuracy (Thompson et al. 1986, Ch. 3.2). Usually a Hydrogen Maser is used but Rubidium can also be used where lower accuracy is acceptable.

Throughout the observation the local frequency standard is compared against a fixed time standard which is accessible to all stations. These days GPS time is used universally. The noise on each individual measurement is quite high and ideally should be measured every few minutes with these results averaged on the timescale on which the frequency standard is stable (for Hydrogen masers  $\sim 10$  h) with  $\sim 4$  points to allow a polynomial fit. In practise an offset or an offset and rate is usually sufficient. A more likely problem is that the maser is that it will jump at some point during the observation. This is quite rare (but see chap. 8).

Another potential difficulty is that since the observing frequency is most likely different from the frequency of the Hydrogen maser, the LO frequency has to be derived from it. There is always the possibility that there will be some error on this frequency. This is also rare (but see chap. 7!).

Until the advent of GPS the offset was not known sufficiently accurately to correlate and find the delays within the lags being correlated. With constant measurements of the masers against GPS time it should be possible to avoid test correlations to find the clock parameters, nonetheless these are frequently required if the stations do not provide sufficiently accurate clock parameters. Even where the clock parameters provided are good, calculating a delay and rate for each station using the cross correlations from a strong source will give more accurate results (albeit referenced to one antenna).

### 1.4.2 Delay Modelling

As mentioned previously, all interferometry requires that the difference in arrival time of the radiation from a source at different antennas be accounted for. With very long baselines, not only are the magnitudes of most of the effects bigger, but the higher resolution results in a greater impact of any error. Indeed in many cases (for example the station coordinates) it is the absolute error rather than the relative error which determines the effect of the error on the data (it is precisely this quality which allows the position of VLBI stations to be determined so accurately in geodesy).

However an important point to bear in mind is that the correlator model does not have to be *perfect*. Even for high-precision geodesy and astrometry, it simply has to be *good enough* to find fringes. Other techniques exist in order to correct for any residual errors in post-processing which, in most cases, will be dominated by unmodelled atmospheric effects. In cases such as geodesy where the absolute error of the applied model is extremely important, full accountability (i.e. exact recording of the model applied in the output file of the correlator is essential). Providing this book-keeping is done, the applied model can be subtracted in post-processing and a more refined one applied.

The usual approach to VLBI is to calculate the phase centre with respect to the geocentre, and so the quantity required is the light travel time between the antenna and a plane passing through the geocentre and perpendicular to the phase centre of the source (see sect. 1.3.5). By defining it in this way it sounds like a purely geometric quantity: it is not. What we are really interested in is the difference in arrival time between the two antennas. If the radiation arriving at one antenna has passed closer to the sun than the radiation arriving at the other then the relativistic effects caused by this must be accounted for. Defining the delay as the light travel time between the antenna and the geocentric plane is merely a convenience.

The terms which go into the model used by the correlator are described in detail in Sovers et al. (1998). A few are described briefly below.

#### Geometric Modelling

The main contribution to the delays is the location of the antennas with respect to the source. This is dominated by the diurnal movement of the antennas but includes many other terms.

#### Earth Orientation Parameters

This is essentially difference between the modelled orientation of the Earth and the unmodelled orientation of the Earth determined by regular coordinated VLBI observations. It is parametrised as follows:

1. UT1-UTC (and TAI-UTC)
2. X pole
3. Y pole

TAI ticks in SI seconds and is the most stable time standard that can be derived from atomic clocks. UT1 is tied to the rotation of the Earth with respect to the extragalactic objects which make up the celestial reference frame (sidereal in its

Table 1.1: Magnitude of EOPs and Magnitude of Errors at different times after the observation: Error 0 is the error if the correlation is done at the same time as the observation (e.g. e-VLBI) assuming that data is already available for the previous day (data are usually released at 05:00:00 UTC). Predictions are also available for up to 2 months in advance, the predictions for the day and the following day are also shown. Taken from on MJD 55200.

	UT1-UTC ms	UT1-UTC mas at equator	X Pole/Y Pole mas
Magnitude	0 – 1000	0 – 15	$\pm 500$
Error -15	0.004 – 0.016	0.06 – 0.24	0.02–0.07
Error -10	0.01	0.15	0.09
Error -5	0.03	0.45	0.09
Error -1	0.06	0.85	0.1
Error 0	0.1	1.6	0.5
Error 1	0.2	3.1	0.8
Error 2	0.3	4.5	1.1

derivation but with a day length of approximately one mean solar day). UTC also ticks in SI seconds but leap seconds are added (or subtracted though this has never been done in practice) to keep UT0-UTC less than 1 s. Typical magnitudes (and errors) of these parameters are shown in table 1.1.

Interpolation of the observed results mean that these parameters are known with maximum accuracy approximately two weeks after the observations take place. This means that the error will be greatest for e-VLBI. However, as with all model errors, provided that the correlation is done with sufficient time and frequency resolution, this can be fixed in post-processing.

### 1.4.3 Fringe Rotation

In connected-element interferometry the LO of each antenna is adjusted to take account of the differing station velocities with respect to the source. In VLBI the complexity of the delay model means that for simplicity (and repeatability) it is far preferable to deal with this within the correlator. The delays have already been calculated so no further modelling need be done, but the fact that the delays change with time must be dealt with.

Essentially the input datastreams must be stretched and squeezed — difficult since they are discrete samples. One technique is first to align the datastreams to an accuracy of one sample, and then apply the sub-sample correction as a post-FFT phase shift.

### 1.4.4 Transport of data to the correlator

Historically the principal medium used to transport the data has been tapes. However, around the turn of the last millennium, the price of standard computer hard-disks dropped to the level where their use was an economic proposition (Whitney 1999).

#### Disk Recording

The most commonly used system is known as ‘Mark 5’. The units on which the data are stored are Mark 5 ‘8-packs’ or ‘Modules’: rugged enclosures which contain 8 identical commodity 3.5” hard disks. The units which read and write



the 8-packs are known as 'Mark 5 units' A number of standards are used to record the data.

Competing disk systems include the Japanese K5 standard which nonetheless uses the same 8-pack.

### **e-VLBI**

Another recent development is that of e-VLBI. This uses high-speed data connections to transport the data directly from the antenna to the correlator, obviating the need to record the data at any point. Efforts have been made with the EVN to make this a production feature: the e-EVN routinely observes at  $512 \text{ Mbit s}^{-1}$  and tests have been run as high as  $4 \text{ Gbit s}^{-1}$  on single baselines.

## **1.5 Radio Interferometers**

Since the properties of several instruments will be discussed in the following chapters, it may be useful to describe these arrays briefly.

### **1.5.1 Connected Element Arrays**

#### **Very Large Array**

The VLA consists of 27 antennas (25 m diameter) located in New Mexico. The pads for the antennas are on three straight arms and each arm has a railway track on which the antennas can be moved. Four main configurations are used with various hybrids also possible. The shortest baseline is 39 m, the longest baseline is 36 km.

The VLA is currently being upgraded to become the EVLA. The main object of this upgrade is greatly increased bandwidth.

#### **MERLIN**

MERLIN consists of 6 antennas located across central England. The Lovell 76 m and Cambridge 32 m form the (almost) longest baseline at 217km. In addition 3 antennas identical to the VLA antennas are situated in Knockin, Darnhall and Pickmere and a further ex-radar antenna is located in Defford. When required, the Lovell can be substituted with the Mark II antenna  $25 * 38 \text{ m}$ .

Like the VLA, MERLIN is undergoing a sensitivity upgrade to become e-MERLIN by replacing the radio links with high-speed digital optical fibre links.

#### **Westerbork Synthesis Radio Telescope**

The Westerbork Synthesis Radio Telescope is a co-planar array consisting of 14 equatorially mounted 25 m dishes. It is used as an element in the EVN. Apertif is a project to fit several or all of the antennas with imaging arrays (Verheijen et al. 2008).

#### **ALMA**

The largest radio astronomy project of recent years has been the Atacama Large Millimetre Array (ALMA) formally known as the MMA (Tarengi 2008). This very ambitious project aims to build a connected-element array which can operate at

any frequency from 30-850 GHz where the atmosphere is sufficiently transparent. In order to achieve this it is being built at an altitude of 5000 m on Altiplano de Chajnantor in the Atacama desert in Chile. It will eventually consist of approximately 50 12 m dishes (possibly with additional 7 m dishes forming a compact array).

## 1.5.2 VLBI Arrays

### VLBA

The VLBA consists of 10 antennas very similar (though not identical) to the VLA antennas. The longest baseline is from the US Virgin islands to Hawaii (8610 km). The remaining telescopes are located among the 48 contiguous states with a heavy weighting towards to southern ones to ensure that low elevation sources can be observed.

It is the only VLBI array of its size which was designed as such and is operated almost continuously. The current standard data rate is 512Mbps. However a sensitivity upgrade is in process which will eventually lead to data rates of 4Gbps

### EVN

The EVN originated out of informal VLBI observations using suitably equipped European radio telescopes. The Joint Institute for VLBI in Europe (JIVE) now coordinates an array of more than 20 telescopes<sup>8</sup> most of which are in Europe but also in other countries such as South Africa, Russia and China. Arecibo (Puerto Rico) often participates too.

The EVN observes three sessions of three weeks each a year. In addition there are several e-VLBI sessions a year and triggered target of opportunity sessions are also available. 512Mbps is the standard datarate.

### LBA

The Australian Long Baseline Array (Phillips et al. 2009) observes twice a year along with occasional e-VLBI sessions. It consists of Parkes (64 m), Mopra (22 m), ATCA (6 dishes 22 m) and the geodetic dishes Hobart (26 m) and Ceduna (30 m). The NASA deep-space dishes at Tidbinbilla near Canberra often participate as well. 512Mbps is the standard data rate for Parkes, Mopra and ATCA. The other antennas participate at 256Mbps.

### VERA

VERA is a purpose-built Japanese array of 4 dishes (20 m in diameter) located across islands in Japanese territory. It is used primarily for measuring the trigonometric parallax and proper motion of Galactic masers. It usually works at 22GHz (but may operate at between 2 and 43 GHz). Baseline lengths range from 1200 km to 2300 km.

### VLBI2010 and AuScope

Geodesists are planning their future activities in the framework of VLBI 2010. Parameters for future geodesic dishes are likely to include:

---

<sup>8</sup>The telescopes able to participate vary somewhat with frequency and session

- Very high synthesised bandwidth feeds extending to higher frequencies (with backward compatibility with existing X/S observatories)
- Smaller, more agile 12 m-class dishes, often with more than one co-located at a single site.

AuScope is a network of 3 such antennas spread across Australia. One is co-located with the Hobart geodetic antenna. A further similar dish is located in New Zealand.

### 1.5.3 SKA Pathfinders

#### Low frequency

LOFAR, the MWA and the LWA are among a new generation of low frequency instruments planned. They are all phased arrays and will operate at between 10 MHz and 300 MHz.

#### Other SKA pathfinders

In addition the two proposed sites for the SKA will be home to a pathfinder consisting of a large number of small dishes. In Australia ASKAP will be an array of 36 antennas (12 m) each equipped with a focal plane array. MeerKAT will be located on the South African proposed site. It will consist of 80 dishes (12 m) and be similar in specification to the EVLA.



# Chapter 2

## Radio Interferometry Observing Techniques and Data Reduction

In chap. 1 I detailed all of the process of observing up to and including correlation. In this chapter I will cover everything that happens from the recording of the data onwards, following on from chap. 1. This chapter covers all of the calibration and imaging techniques used for observations presented in this thesis (see chaps. 7 and 8) and is written very much from a VLBI perspective, although much of what follows would also be applicable to connected-element radio interferometry.

In most cases, the correlation need not be carried out by the observer, and so what follows is a complete summary of the data reduction process. The techniques used for calibration (in particular phase calibration) determine which sources must be observed along with the target source and how; they may also determine to some degree the selection of the target. Therefore this chapter also discusses some aspects of planning an observation. For wide-field observations this is also discussed in in sect. 5.3.

### 2.1 Flagging (data editing)

Before calibration can begin, it is necessary to remove any data which is so badly corrupted that including it in any calibration or analysis would do nothing but degrade the final result. It is also used to remove any data for which there is clearly some error with one of the antennas: any times for which the antenna was unable to observe; times in which there were recording problems with the antenna; or times for which the antenna was not on source. Most of this can be done automatically using records of what happened at the station.

A step which currently requires manual intervention is the removal of RFI. Any RFI which is sufficiently strong to cause problems with the data has very high amplitude and there is no hope of recovering the underlying astronomical system<sup>1</sup>, certainly not with the limited number of bits per sample used for VLBI. As datasets increase in size, it will become increasingly difficult to do this manually as the time required will rapidly become infeasible. See appendix C for an attempted solution to this problem.

In addition, it usual to discard the outer edge (approx. 5%) of each sub-band. It may also be necessary to discard the first and last time integration of each scan since these may not represent a full time integration.

---

<sup>1</sup>notwithstanding various techniques for RFI mitigation at the antenna or even during correlation

## 2.2 Calibration

This is the process of estimating the true visibility function of the source from the visibilities that have been measured. The measured visibilities will have all kinds of errors, however it is possible to account for these errors by assuming that the errors in amplitude and phase are entirely antenna-based. That is, each antenna has a complex gain (amplitude and phase) which varies with time.

### 2.2.1 Closure quantities

This assumption allows us to determine the amplitude and phase errors analytically provided that there are sufficient antennas (Thompson et al. 1986, chap.. 11)

#### Closure phase

With two antennas there is one observable (the delay at which the two datastreams of the baseline correlate) for two unknowns (the unmodelled delay at each of the two antennas). With three antennas, however, there are three observables for three unknowns. This means that it is possible to solve for each unmodelled delay using simultaneous equations. Note that the absolute delay is not calculated; it will remain relative to a chosen reference antenna.

#### Closure amplitude

The measured amplitude of each baseline contains the product of the amplitudes of each antenna. Two independent baselines (with no antennas in common) allow the calculation of the amplitude of each antenna.

## 2.3 Phase (or delay) calibration

The fundamental unit which is measured by an interferometer is the difference in arrival time of the wavefront at two different elements. Once all effects are properly calibrated, any non-zero phase difference in the visibilities at different stations will be due to the observable (be that source location, source structure or antenna location). In other words, a point source at the phase centre will show zero phase difference for all elements for a known array geometry. However, even assuming that the correlator delay model has taken care of the geometric considerations (generally a fairly good assumption) there remain certain residual effects which mean that the phase of a point source at the phase centre will *not* be static with frequency and time. These are listed below:

- Clock errors
- Instrumental effects consisting of
  - different delays introduced by the station electronics to each sub-band
  - complex delay across the bandpass
- The atmosphere consisting of
  - The ionosphere
  - The troposphere

Generally, with the exception of the clock parameters, none of the calibration data collected during the observation are used in the correlator delay model. CALC9.1 (Petrov 2008) for instance does attempt to model the troposphere but corrections are based only on the date and time of the observation and the location of the antenna and the source. However, it is, in principal, possible to use (for example) meteorology information from the antennas in the construction of the correlator atmospheric model (the correlator model VTD does this. See Petrov 2009).

### 2.3.1 Clock correction

Clock offsets are due to the fact that for VLBI, each station has its own frequency standard (which also defines the time-stamp of each sample during the observation: hence the term clock). Comparisons of the maser to GPS time taken throughout the observation will be sufficient to observe fringes on a strong source, however this will still leave a large residual delay. By fringe-fitting, delay and rate solutions can be found which are generally used throughout the observation (it is usually sufficient to correct for delay and rate for hydrogen masers which can be determined from a single fringe-fit solution on a strong source). Established arrays such as the EVN and VLBA do such “fringe tests” regularly to determine relative clock offsets.

It is likely that in most cases some residual clock errors remain after correlation. This is not explicitly corrected for most astronomy observations (although it is modelled for geodesy experiments, see below) however the phase calibration techniques which follow in sect. 2.4 will normally correct any unmodelled clock offsets.

### 2.3.2 Calibrating for the effects of the atmosphere

Other than the calibration techniques described in sect. 2.4 and the *a priori* modelling terms in the correlator delay model, the effects of the atmosphere on the phase are not usually explicitly corrected for.

It is worth mentioning the possibility of using publicly available data obtained from multi-frequency GPS observations to model the ionospheric delay. This is described in Erickson (1999) and Walker and Chatterjee (1999).

### Bandwidth Synthesis

This technique uses observations at multiple frequencies to model the contributions of the ionosphere and troposphere to the data. Observations are made with a high instantaneous range of frequency (typically sub-bands of equal size are placed throughout the available bandwidth with a mixture of narrow and wide spacings which reduces the ambiguity in the delay solutions found). Using these techniques a frequency-independent component (due to clock errors and the troposphere) and a frequency-dependent component (due to the ionosphere) can be modelled.

### 2.3.3 Correcting instrumental errors

After the feed, the signals from different sub-bands pass through different electronics. This leads to differing delays for each offset, leading to phase offsets between the sub-bands. This is a particularly significant problem for observations

which use the bandwidth synthesis technique: in this case the offsets between the sub-bands once instrumental effects have been subtracted are an observable which is used to model the delay due to the atmosphere.

Since these delays do not usually vary significantly during an observation, for most astronomy observations it is sufficient to carry out a “manual phase cal”: a strong, unresolved source is fringe-fitted (see sect. 2.4.1) for delay (no rate) for each sub-band individually to provide individual solutions for the delay.

This technique can be improved on by using phase cal tones. These are generated by injecting very narrow pulses at a certain rate (typically 1 MHz) into the signal at the earliest opportunity (usually directly after the feed). The frequency that the phase cal tones are injected with is tied to the time standard of the antenna.

These pulses form a “picket fence” in frequency space. Careful measurements of the phases of these phase cal. tones after the baseband data have been digitised (with a sampling rate also determined by the antenna time standard) will allow exact subtraction of the delay due to the instrumentation (Walker 1999a)<sup>2</sup>

The only instrumental delay which now remains unaccounted for is the cable from the maser to the point where the pulses are injected. This will be time variable to some degree due to changes in the length of the cable caused by temperature variations and stretching as the telescope moves. For the best possible correction of instrumental delays, the cable length is measured throughout the observation and this is added to the delay.

### 2.3.4 Parallaxic angle correction

Circular polarisation feeds are generally used for VLBI. For altazimuth-mounted telescopes this leads to a diurnal turning of phase which must be corrected.

## 2.4 Phase calibration using astronomical sources

### 2.4.1 Fringe fitting

For stronger sources it is now possible to fringe fit these sources: that is, a linear fit is made to the variation of phase with respect to frequency and time, correcting for residual errors in the delay and rate respectively. Unfortunately, making this fit also removes any residual delay and rate which is not due to unmodelled delay but due to source position. This means that the astrometric position of the source is lost. In addition, if the integration time required to detect the source is large compared to the time over which the residual delay and rate can be considered to be roughly constant (the so-called coherence time which is frequency dependent but of the order of minutes) it will not be possible to detect the source.

### 2.4.2 Using a phase calibrator

The solution is to observe both the target source, and a phase calibrator: a known source which is bright and compact enough to be detected well within the coherence time on all baselines (or at least one baseline to each antenna). The phase calibrator can then be used to determine the phase solutions and these can then

---

<sup>2</sup>Notwithstanding any  $2\pi$  ambiguities in the system. These are resolved by using the “manual phase cal.” technique to first order and using the phase cal tones to improve the solutions: both with better instantaneous accuracy and taking into account their variation with time.



be applied to the target. This has two advantages. First the astrometric precision is preserved presuming that a high-precision position is known for the phase calibrator. Secondly, even if the target is too weak to be fringe-fitted within the coherence time, after application of the phase solutions from the phase calibrator it is possible that the source may be detected since the smaller residual delay allows greater time-averaging of the data.

There are two distinct ways in which a phase referencing observation can be carried out. The choice is dependent upon whether the nearest suitable phase calibrator is within the primary beam.

### Choosing a Suitable phase calibrator

By selecting the baseline of the most sensitive to least sensitive element and setting  $\tau_{int}$  to the coherence time at the specified frequency, equation 1.10 can be used to work out the minimum source sensitivity that can be detected with fringe fitting (i.e. the minimum flux of a potential phase calibrator) <sup>3</sup>.

### In-beam phase calibration

If the separation of the target source from the nearest available phase calibrator is sufficiently small that the calibrator is detectable by all antennas, even when they are pointing at the target source, it will be possible to spend the entire observation pointing the antennas at the target source. If a source consists of several co-located components, one of which is bright enough to serve as a phase calibrator, then only a single correlation need be carried out. In the case of a larger separation of the target and calibrator it may be necessary to correlate twice: once on the target source and once on the phase calibrator.

This technique has two advantages. Firstly all of the available observing time can be spent observing the target. Secondly since the observations are simultaneous, even time variability of the phase on short timescales is accounted for. For this reason, if there is a choice of targets all of which would achieve the desired science goal, any target for which in-beam phase referencing is possible should be chosen (all other things being equal).

The chances of a lucky coincidence between calibrator and target are better at lower frequencies since the primary beam is larger. Indeed at lower frequencies it is likely that several calibrators exist somewhere in the primary beam (Garrett et al. 2005a). Prior VLBI observations of the field of interest are extremely useful as they allow possible calibrators to be found.

The increasing sensitivity of VLBI arrays means that it will become more common that in-beam calibration can be used. It is also one possible application of the wide-field techniques presented in this thesis. However it is important to note that unless the phase calibrator has been studied *a priori* and already has an accurate position, the positional accuracy will be lost, unless either astrometric techniques are used or the observation is referenced to a high-quality known source.

---

<sup>3</sup>Sources above decl.  $-30^\circ$  which may be useful for phase calibration can be found in the list of VLBA calibrators available at <http://www.vlba.nrao.edu/astro/calib/>

### Nodding phase calibration

In cases where in-beam phase calibration is not possible, nodding phase calibration can be used. Observations of the target are interspersed with observations of the phase calibrator(s). The delay solutions determined for the calibrator scans can then be interpolated to provide solutions for the target scans.

The frequency with which it is necessary to switch between target and calibrator is frequency dependent and details are available in Wrobel et al. (2000). However, as a rule of thumb the time spent observing the target will be approximately half the total observing time.

The density of phase calibrators on the sky is now sufficiently high (at least in the northern hemisphere) that it should be possible to apply this technique to any source on the sky.

### 2.4.3 Residual errors after phase referencing

Proper use of this technique should allow the removal of all atmospheric, instrumental and clock errors to first order. Second order errors will remain however, and for cm-wavelength observations they are likely to be dominated by differences in the atmospheric delay between the target and the calibrator (Fomalont and Perley 1999). Residual errors due to the tropospheric delay differing between calibrator and target can be corrected for by observing multiple phase calibrators at different elevations (Mioduszewski and Kogan 2009).

## 2.5 Amplitude calibration

The output of the correlation process is a set of correlation coefficients: zero when there is zero correlation and one when the input baseband datastreams match exactly. For a bright source approximately one sample in  $1 \times 10^{-3}$  to  $1 \times 10^{-4}$  correlates (Sovers et al. 1998). Some method must now be used to relate these correlation coefficients to the flux density of each baseline.

### 2.5.1 Autocorrelation Correction

First the autocorrelations can be used to correct for any errors in the levels used in sampling. N.b. this correction is harmful for data correlated with some hardware correlators which do not calculate all of the cross-terms for multi-bit data (Mark Kettenis, Jive, Priv. Comm).

### 2.5.2 Gain Calibration

For connected-element interferometry it is usual to observe a flux calibrator. This technique is problematic for VLBI since all sources which are detected by VLBI are variable.

Instead, single dish observations are used to determine the gain of the antenna (and the “gain curves”: the variation of the gain with elevation). System temperature measurements are taken at each antenna (typically at every sub-band) every few minutes throughout the observation by measuring the increase in system temperature when noise from an artificially generated source is added to the sky signal.

These gain and system temperature values  $T_s$  are then used to calculate the flux density  $S_c$  from each correlation coefficient  $\rho$  using the following formula:

$$S_c = b \rho \sqrt{\frac{T_{s1} T_{s2}}{K_1 K_2}} \quad (2.1)$$

where  $K$  is the gain of the antenna in  $\text{Jy K}^{-1}$ .  $b$  accounts for digitisation loss and any correlator-specific scaling factor (Walker 1999b).

## 2.6 Bandpass calibration

The characteristics of the receiver and the filtering applied to the signal mean that it is necessary to calibrate the gain across each sub-band. The nature of this gain is complex (i.e. it affects both amplitude and phase) however the main effect is a sharp reduction in gain towards the edge of each sub-band. This is done by using the amplitude and phase of a strong source (which has no spectral structure).

One subtlety for VLBI is the fact that the delay model has shifted the reference frame of the cross correlations to the centre of the Earth. This means that as the rate of the applied delay model changes the bandpass shape (due to the bandpass filter) and it will shift within the band. Post-processing software corrects for this subtlety, but the outer edges of the band may have to be discarded.

## 2.7 Determining information from the visibility function

Once calibration is completed it is possible to infer information about the source by looking directly at the visibility function of the source.

Firstly the visibility amplitudes as a function of UV distance (projected baseline length) can be examined. Using this technique it may be possible to discriminate between different rotationally symmetric models for the source (see Doeleman et al. (2008) and Pearson (1999) fig. 16-2).

It is also possible to determine basic source structure by examining the function of phase with respect to frequency and time. Examples of parameters which can be determined in this way include: distance of the source from the phase centre (a point source distant from the phase centre will have a linear slope of phase with respect to frequency and time); single source or double source morphologies (double sources have a visibility function which oscillates in amplitude); and point-like or extended source (Pearson 1999, fig. 16-1).

### Solving for the antenna positions (Geodesy)

If the object of the observation is *solely* to determine either the position of the source or the position of the antenna, then some of the preceding calibration will be unnecessary. The object will simply be to make a best fit to a model with many free parameters, including a residual delay function for each antenna, the location on each antenna and the location of each source, using plausible assumptions about how the delay of each antenna changes with time and how the delays due to the atmosphere depend on source location in the sky at each antenna, and how the delay changes with frequency.

It should be possible to unpick the phase behaviour due to errors in the three *a priori* coordinates needed to define the location of each antenna, and the two coordinates needed to define each source position, since source location errors cause

different phase behaviours depending on whether the error is in right ascension or declination (the situation is similar for errors in antenna location).

For best results many sources must be observed over as wide an area as possible over a full 24-hour session (or as long as possible). More details are given in Fomalont (1999b).

## 2.8 Imaging

An interferometry observation consists of sampling the visibility function over as much of the UV plane as possible. Once appropriate calibration has been carried out, and provided certain conditions are met, it is possible to Fourier transform the visibility function to form a function of sky brightness over a certain solid angle of the sky.

The intricacies of how this is carried out are rather complicated. For a full treatment see Briggs et al. (1999), Cornwell et al. (1999) and Thompson et al. (1986) chap. 11. In this section I will restrict the discussion to a basic outline and to the parameters which can be varied during the imaging process, and how these parameters affect the performance.

### 2.8.1 Generating the dirty map

. This is the first stage in imaging, and may be sufficient for some purposes (such as searching for detected sources).

#### Preparing the UV data to be inverted

Assuming that a discrete Fourier transform will be used (specifically an FFT) the image to be generated is discrete (pixelated). The most obvious parameters that must be set are then the dimensions of the image (in number of pixels) and the dimension of each pixel (the angle on the sky which they represent in arcseconds). These parameters in turn determine how the UV data are to be gridded (the UV data are averaged into a matrix with each point representing the value of the visibility function at a certain point on the UV plane).

Choosing a very high resolution for the resulting image will not lead to any degradation of the resulting image (though it will lead to higher computation time), however if the size of each pixel is large compared with the resolution of the instrument, not all the UV data will lie within the bounds of the input UV grid causing data from longer baselines to be discarded.

#### Weighting the Data

As discussed in sect. 1.3.1, UV data from baselines of different lengths contributes to different spacial frequencies in the image. By changing the weight of different samples of the visibility function it is possible to optimise the image for sensitivity, for structure on different spacial scales, or to attempt a compromise.

#### The dirty map

The result of this gridding and inversion is a map of the sky convolved with the “dirty beam” the Fourier transform of the distribution of the measurements on

the UV plane (which are also dependent to some degree on how the UV data are weighted and gridded).

### 2.8.2 Deconvolution

Various techniques exist to remove the effects of the dirty beam, leaving a map of the estimated true sky brightness. The most widely used technique (and the one used for the imaging in chap. 7) is some variation of the CLEAN algorithm. First the dirty beam is calculated, then the brightest point in the image (or a selected portion) is located and the dirty beam is subtracted. This process is repeated on the data that are left (after subtraction) until the residuals are noise-like (no significant peaks remain). The data have now been separated into “clean components” (consisting of a series of coordinates where flux was subtracted and the amount of flux subtracted from the image at that point) and the residuals (the flux that was left in the image after the subtraction process was completed).

The clean components are then converted into a map by taking each clean component and representing it as the amount of flux subtracted, convolved with the “restoring beam”, typically an elliptical Gaussian similar in dimensions to the main lobe of the dirty beam.

#### Parameters

Many parameters can be varied here: the subset of points on the map which are considered when choosing where to place the dirty beam; the loop “gain” (the fraction of the peak flux which is subtracted at each iteration); the number of iterations; and the size and shape of the restoring beam.

## 2.9 Self-calibration

Self-calibration involves using the target source itself as a calibrator. Source structure must be accounted for when a source is used for phase calibration and the fact that the source is being observed in the first place usually implies that the source structure is (at least to some extent) unknown. Nonetheless, in all cases it is possible to make at least a few assumptions about the pattern of sky brightness. At the very least it can be assumed that the sky brightness is zero or positive across the whole sky.

Usually Self-calibration is done as an iterative procedure: first some *a priori* model of the source structure is used as a first guess. If the structure of the source is entirely unknown, then a point source can be used with a flux determined from the zero-spacing flux density of the UV dataset<sup>4</sup>.

Next the UV data are corrected by using the model of the target source and the closure quantities defined in sect. 2.2: the UV data are divided by the model, transforming it into that of a pseudo-point source (i.e. the structure of the model is now accounted for). The complex antenna gains are then allowed to be free parameters. Solutions are found by constraining them using the closure quantities defined in sect. 2.11. These solutions are then applied to the data which is then imaged. After deconvolution (for example using the CLEAN algorithm) a new

<sup>4</sup>For calibrated data UV data flux density is inversely proportional to baseline length. The zero-spacing flux density can be determined by extrapolating this relation back to zero.

(hopefully) better model of the source emerges. This new model can then be used as the input model for another round of self-calibration.

Several different parameters can be used to refine this process. Firstly (at least for the first few rounds) it may be better to allow only the phases (not the amplitudes) to vary. Secondly, the exact method used to solve for the phase (and amplitude) may be varied: namely the way residuals further away from the fit are weighted. Finally, during the imaging process there are many parameters which may be changed (as described in sect. 2.8). These parameters will likely be changed as the self-cal process is iterated and the model (hopefully) converges towards a meaningful one. In general, phase-only self-calibration is used initially and amplitude calibration is used later.

## 2.10 Subtraction of confusing sources

Often the sidelobes of nearby sources increase the background noise in a field. Where the “confusing source” is close enough to the target field that an image can be made encompassing it, the usual cleaning technique can be used to remove the sidelobes. However it is possible that even a distant source, much too far away to be imaged could affect a field in this way.

In this case an attempt can be made to subtract the source directly from the visibility function: a process known as “peeling”. This is achieved by taking a simple model of the source (perhaps the clean components from a deconvolution), doing a direct Fourier transform of this model, and subtracting it from the UV data.

There is also a more nuanced procedure for doing the same thing: First the UV data are self-calibrated on the confusing source (using the model of the confusing source). Next the model is subtracted from this dataset. Finally the self-calibration solutions are reversed, shifting the UV data back to the original field. This may be done multiple times for multiple confusing sources. This procedure was used for the VLA Low-frequency Sky Survey (Cohen et al. 2007, see also the AIPS procedure PEELR)

## 2.11 Interferometry with a limited number of antennas

Since much discussion in this thesis relates to VLBI with a small number of antennas, the limitations of this are briefly described.

### 2.11.1 Two Antennas

With a single baseline it is possible to determine the amount of flux concentrated on scales corresponding to the fringe spacing of the baseline. With Earth rotation synthesis this can be done for almost all angles across the source (The angles will be somewhat restricted if the source is too close to the equator or at low elevation). Simple astrometry can be done, but this will be limited both by phase ambiguity and by large uncertainties in calibrating out the atmosphere, especially at low frequencies. Nonetheless the positions determined are likely to have a higher accuracy than most other techniques. Real-world applications include the daily determination of UT1-UTC. This is done with a single long East-West Baseline (e.g. Wettzell-Tsukuba). Bandwidth synthesis is used both to correct for the ionosphere and avoid phase ambiguities (see sect. 2.3.2). Images can be made

(see Moscadelli and Porceddu 2006), however, since only one spacial frequency per angle is measured (for a single frequency), the cleaned image will simply show the unresolved flux in the shape of the restoring beam. This configuration is very susceptible to instrument errors, since the loss of a single antenna means that the observation is wasted.

### 2.11.2 Three Antennas

Three Antennas allow improved calibration through the calculation of phase closure. Although conventional imaging is not possible the UV data can be fitted to models for those baselines on which the source is detected. This was commonly done in the early days of interferometry (Ryle et al. 1950) and VLBI. It continues to be used in cases where only a small number of telescopes are suitable for carrying out an observation: mm-VLBI, for example (Doeleman et al. 2008).

### 2.11.3 Four Antennas

Provided that the source is detected on at least four baselines, four antennas is sufficient to do imaging through the calculation of the closure amplitude (see sect. 2.2). Atmospheric and instrumental effects can be calibrated for with sufficient accuracy. However, the spacial frequencies that can be resolved may be limited and the self-calibration solutions obtained (see sect. 2.9) may not be unique.





# Chapter 3

## VLBI in Italy

In this chapter I will discuss the motivation for an Italian VLBI network, the logistics involved, and the feasibility of such a project. An Italian VLBI network would be intermediate in resolution between MERLIN and larger VLBI arrays and could potentially provide unique capabilities in the following ways:-

- Different operating frequencies allowing the study of objects which can't be studied by other arrays <sup>1</sup>
- Undertaking projects which don't require a larger array and/or would not be likely to be allocated time.
  - Monitoring of sources
  - Single baseline surveys
- Having exceptionally fast turnaround from proposal to observation and correlation
- High single-baseline sensitivity
- Real-time capabilities

I will begin by summarising the radio astronomy infrastructure which exists in Italy and that which would be required to sustain a VLBI network. Then I will discuss the performance of the Italian array and how this may be enhanced by the *ad hoc* collaboration of other VLBI antennas around the world. I will then discuss the science that could be done with the Italian antennas which has been proposed in the literature; and also research that is done by scientists in Italy which could be supported by such a network. I will end by proposing some other experiments and techniques which could be carried out on (or at least tested using) an Italian network.

### 3.1 Radio astronomy in Italy

There are four radio telescopes in Italy, three of which are specifically designed for VLBI (although they are also used for single dish work). In addition, another radio telescope is under construction in Sardinia.

---

<sup>1</sup>Noto currently has the widest frequency coverage for VLBI of any antenna in the EVN. The SRT will have continuous coverage from 0.3 to 150 GHz

### 3.1.1 Croce del Nord

The Northern Cross (Croce del Nord) is composed of a single cylinder for the East-West arm and 64 cylinders for the North-South arm. It has the second-largest collecting area of any telescope in the Northern Hemisphere. It was inaugurated in 1967 as a transit instrument, mechanically pointed only in elevation and designed for observations at 408 MHz. The North-South cylinders are phased up using an electro-mechanical system. In recent years, experiments have been done adjusting the phases of the individual receiving elements along a single cylinder, allowing beam-forming in an East-West direction. Renewed interest in low-frequency radio astronomy, and in particular LOFAR has raised the possibility of using the Croce del Nord as an element in such an array on an *ad hoc* basis. Its main weakness is the fact that it is only able to observe a single vertical polarisation (Y). Nonetheless various tests are being made using the Croce del Nord as a SKA testbed.

### 3.1.2 Medicina and Noto

The Medicina and Noto antennas were constructed in the 1980s as VLBI antennas. Since then they have been important participants in VLBI observations (principally within the EVN network) and in addition they are used for some single dish work. Both are 32 m in diameter and they are almost identical. The main difference is that Noto has an active surface which corrects for structural deformation induced by movement of the antenna and permits observations at higher frequencies. Noto has receivers for all the EVN bands: 92, 49, 30, 21, 18, 13, 6, 5, 3.6, 1.3 and 0.7 cm. Medicina has receivers for the EVN bands from 21 cm to 1.3 cm. Medicina is also equipped with a high-speed fibre connection which allows it to participate in EVN e-VLBI sessions — the only Italian antenna to do so. Both are being used extensively as test-beds for the SRT receivers and other technologies.

### 3.1.3 Sardinia Radio Telescope (SRT)

The Sardinia Radio Telescope is currently under construction approximately 30 km North of Cagliari at an altitude of 700 m in one of the driest areas within the Italian territory. It will have a main reflector 64 m in diameter and will use three different foci to allow space for enough receivers for continuous coverage from 300 MHz to 115 GHz. Like Noto it will also have an active surface allowing high antenna efficiency even at high frequency.

### 3.1.4 Matera

The Matera 20 m Antenna is owned by ASI. It is equipped with VLBI recording equipment, but only for the S/X dual-band observations currently used for geodesy (see sect. 2.3.2). There are no plans to upgrade this antenna nor to equip it with a fibre.

## 3.2 Infrastructure Required

There are three (and there will soon be four) antennas equipped for VLBI observations within Italy. This raises the possibility of the Italian antennas working together independent of an external network. Although VLBI observations have already been carried out, a certain amount of infrastructure would improve the



Figure 3.1: Map of Italy illustrating the locations of Italian radio telescopes (underlying map — credit: Eric Gaba (Sting - fr:Sting) and NordNordWest)

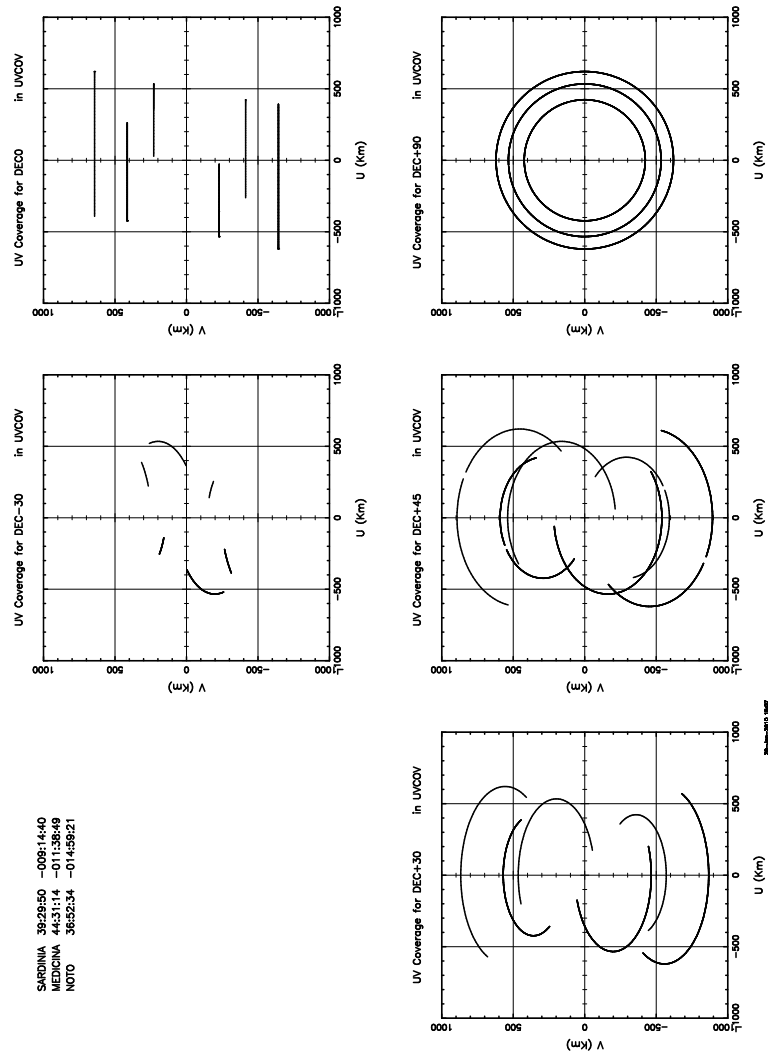


Figure 3.2: UV coverage for a three-station Italian array (SRTMcNt) for sources at  $-30$ ,  $0$ ,  $30$ ,  $45$  and  $90^\circ$  declination respectively. Each track is for when the source is at more than  $15^\circ$  elevation for both stations.

Table 3.1: Locations of the Italian antenna in Cartesian coordinates (m); and latitude (dms), longitude (dms) and altitude above sea level (m). Cartesian coordinates for Mc, Nt and Ma are taken from the SCHED catalog and should be very accurate. Latitude, longitude and altitude are approximate. The position of the SRT is very approximate.

Antenna	X (m)	Y (m)	Z (m)	Latitude (N)	Longitude(E)	Altitude (m)
SRT	4864925	791819	4035511	+39°29'50"	+09°14'40"	700
Mc	4461370	919597	4449559	+44°31'13"	+11°38'49"	67
Nt	4934563	1321201	3806484	+36°52'34"	+14°59'21"	143
Ma	4641939	1393003	4133326	+40°38'58"	+16°42'14"	543

Table 3.2: Baseline length in km between VLBI stations in Italy. Split into two components: Z (component in direction of the axis of rotation of the Earth) and X and Y combined (co-planar component)

Antenna	Component	Mc	Nt	Ma
SRT	<b>All</b>	<b>592</b>	<b>581</b>	<b>649</b>
	XY	423	534	641
	Z	414	229	98
Mc	<b>All</b>	–	<b>894</b>	<b>597</b>
	XY	–	621	507
	Z	–	643	316
Nt	<b>All</b>	–	–	<b>445</b>
	XY	–	–	312
	Z	–	–	327

practicality of regular observations a great deal. This is summarised below, assuming that correlation would take place at the institute headquarters in Bologna.

### 3.2.1 Recorded VLBI

All the Italian antennas have the recording equipment necessary for VLBI work. This is also guaranteed for the SRT. For the observations undertaken with the Italian antennas in this thesis (see chaps. 7 and 8) the data were recorded at the stations onto standard Mark5 8-packs which were then shipped to Medicina. Here the data were copied onto standard computer disks using the Mark5 unit at the antenna. These disks were then transported to Bologna for correlation.

This process would now be somewhat simpler thanks to the high-speed fibre connection between Bologna and Medicina. It would even be possible to 'play back' one station on the mark5 unit in Medicina and correlate it in Bologna as if it were an e-VLBI observation (data from other stations would be required to be on disk in Bologna or playing back from disk at another institute however).

Another possibility would be to purchase some Mark5 playback units for installation in Bologna (at a unit cost of approximately €20 000). Alternatively, an *ad hoc* system of recording could be developed: for example, recording onto RAIDs or standard disks at the station, and then shipping them either to Bologna or transporting them from the antenna to a nearby cooperating institution which has a high-speed network capabilities. It may be possible to use systems already developed to some degree, for example the LBA data recorder (Phillips et al. 2009).

### 3.2.2 e-VLBI

Medicina and the Bologna headquarters are already equipped with high-speed fibre connections to each other and to other European antennas and institutes. Since the DiFX correlator installed in Bologna is being routinely used for VLBI with the LBA in Australia (Phillips et al. 2007; Hancock et al. 2009), it would be a relatively simple task to carry out real-time VLBI with another antenna, outside of Italy.

As well as reducing substantially the amount of hardware required for correlation, e-VLBI also provides quick turnaround time which is advantageous for many potential applications of the Italian array (see sect. 3.4). Tests have been made with e-VLBI at up to 4Gbps. In principal the network connection to Bologna should be able to carry this traffic from two antennas. The load on the correlator would be within the capabilities of the Bologna cluster provided that the output resolution required was not excessive.

### 3.2.3 Correlation

The clusters in Bologna are easily capable of correlating three or four station experiments, even with high bandwidth (see sect. 4.6).

It should be noted, however, that experience in correlation is required for successful operation and large VLBI correlator operations such as the VLBA AOC and JIVE are complex operations requiring a large technical and administrative staff.

### 3.2.4 Supplementary Data

As well as carrying out correlation, correlation centres also have to handle various other data and ensure data integrity. Some of these data are essential to the correlation (EOPs, clock information etc.) others are simply checked and packaged up for the astronomer to use (e.g. system temperature throughout the experiment and weather information).

Scripts have been developed for handling this, and the necessary information for correlator operators has been summarised in a manual (see Morgan (2008) and appendix B).

## 3.3 Array Parameters

The different types of VLBI, and the possibility of partnering with other VLBI antennas on an *ad hoc* basis, present a huge amount of parameter space. An attempt has been made to summarise it here.

### 3.3.1 Possible collaborations

The capabilities of the array could be greatly improved through *ad hoc* collaborations with other VLBI antennas. Just a single longer baseline to East Asia, North America or just to Onsala would allow high-quality astrometry and geodesy.

### 3.3.2 Wider collaborations

25–30 m class antennas are generally under-utilised since they are not used for so much single-dish work as the larger antennas. This raises the possibility of

Frequency (GHz)	SEFD (Jy)			
	SRT	Mc	Nt	Ma
0.3	76	–	76	–
1.0	–	–	1025	–
1.5	29	490	820	–
2	–	400	770	1036
5	30	170	260	–
7	31	840	1100	–
9	23	–	–	–
13	42	320	770	2737
17	72	–	–	–
23	124	700	800	–
32	54	–	–	–
43	98	–	900	–
86	369	–	–	–
100	460	–	–	–

Table 3.3: SEFD of the SRT and other Italian dishes at various bands. Projected Values for the SRT are taken from Cenacchi (2004). Some liberty has been taken with matching up the bands for which a projected values is available for the SRT with the standard EVN bands. If a band is unlisted for the SRT this does not mean that it won't be available (the SRT has continuous coverage over all these bands and beyond) but that a projected value was not available.

an international small-antenna VLBI array specialising in projects (such as wide-field VLBI) where small antennas are sufficient or even preferred. If IRA were to become a proven correlation centre it could take on some or all of the correlation responsibilities.

### 3.3.3 Sensitivity

The estimated sensitivity of the SRT at different frequencies is calculated and tabulated in Cenacchi (2004) table 16. This is reproduced along with values for the other Italian dishes in table 3.3.

### 3.3.4 UV coverage

Designing an interferometer to maximise UV coverage is a complicated task (for an overview see Holdaway and Helfer 1999). With such a small number of antennas, complex analysis of the UV plane would be meaningless. However, we can make a few observations. Most VLBI networks are longer in an east-west direction largely due to outside constraints (There is a lack of Radio telescopes in South America and Africa). On the other hand, the Italian VLBI array is longer in the N-S direction and there is a concentration of dishes in the South. This reduces the effect of foreshortening for observation of lower declination sources and ensures that the most of the antennas will have a longer track on such sources. At the same time the location between the 36th and 45th parallels ensures that sources near the poles are also observable at high elevation.

### Spatial frequencies

Perhaps the main weakness of this array is the fact that the baselines are all of very similar length. The VLBA has baselines ranging from 236 to 8611 km (with a roughly logarithmic distribution of baseline length) whereas the Italian baselines vary by less than a factor of 2. UV coverage for a variety of source declinations is shown in fig. 3.2. Baseline projections vary from  $\sim 200$  km to  $\sim 800$  km. This means that only a short range of spatial frequencies will be sampled. This, along with the small number of stations, will make it difficult to make out a great deal of source structure. (for imaging with the VERA array see Manabe et al. 2004). Nonetheless it would still be a capable instrument for pinpointing individual sources and seeing structure in sources composed of multiple point-sources (see sect. 3.4.1 below).

### Low declination

Observations of sources down to zero declination would be easy with the full 8-hour track above  $20^\circ$  elevation at Mc. It is difficult to say what would be a minimum declination that could be observed. The VLBA can observe down to  $-45^\circ$  declination albeit with greatly decreasing accuracy and sensitivity (see Petrov et al. 2005) despite having 8 of its 10 antennas located above  $30^\circ$  latitude.

Note that the problem of similar baseline lengths is exacerbated for low-declination sources since the longer Mc-No baseline will be foreshortened.

## 3.4 Science with an Italian VLBI array

### 3.4.1 Maser Astrometry

Science that can be done by the SRT acting with the other Italian dishes is outlined in chap. 5 of Brand et al. (2005). This report also discusses the role that the SRT can play in VLBI with the EVN and other VLBI networks (this is also discussed in Mantovani 2006). One of the suggestions is using SRT, Mc and Nt as the VLBI interferometer to determine high-accuracy positions of 6.7 GHz Masers.

These are described as small spots of emission spread along a line or arc 0.1 to  $1''$  in length. The SRT report identifies one problem as being the poor *a priori* knowledge of the position of the masers (as little as a few arcminutes) due to the inherent low resolution of single-dish observations used for detection and monitoring of masers. A preliminary proof-of-concept observation, observed using Medicina and Noto and correlated in Bonn has been carried out and is described in Moscadelli and Porceddu (2006).

Use of the DiFX software correlator is particularly appropriate in this case as there are no limits on time or spectral resolution. This means that fringe searches for sources far from the phase centre would be made easier.

### 3.4.2 Bandwidth Synthesis Astrometry and Geodesy

Both Medicina and Noto have X/S feeds and have participated in geodesy experiments. Matera is also equipped for observations of this kind. Multi-frequency synthesis errors for the X/S system are typically  $\sim 40$  ps  $\approx 1$  cm. For the Mc-Nt baseline this corresponds to a positional uncertainty of 2.8 mas. The collaboration of a more distant station would increase this significantly.



### 3.4.3 Geodesy

Geodesy usually favours longer baselines on different continental shelves. Medicina is on the Eurasian shelf while Noto is on the African shelf. This means that the relative motions of these two plates can be measured using the Italian telescopes alone. The area covered by the Italian baselines is geologically interesting and it would be possible to monitor the baseline lengths for the Italian telescopes during seismic events in Italy. If e-VLBI capabilities were available, this could even be done in real time (Tingay et al. 2009).

### 3.4.4 Spacecraft tracking

A small VLBI network is ideally suited to spacecraft tracking. A fast turnaround is desirable.

In addition to simple orbit determination there is a great range of novel techniques to be explored. Avruch et al. (2006) gives a concise overview of standard techniques for orbit determination, as well as the enhanced techniques that VLBI offers for extremely high accuracy. In addition research is being done in Italy to explore the feasibility of observing satellites with geodetic systems (Tornatore and Haas 2009). This involves observing GNSS (a generic term for satellites related to various global satellite navigation systems) through the same optics which is used to observe the astronomical objects of the ICRF. This allows better tying of the reference frame defined by GNSS satellites to the ICRF.

### 3.4.5 Testing e-VLBI technologies

The huge growth in potential for e-VLBI likely to be seen in the coming years is discussed further in sect. 5.1 and in Spencer et al. (2008). To summarise: DBEs will allow extremely high data rates out of the antenna as long-distance network rates continue to outstrip disk speeds. A single baseline would occupy a niche between e-MERLIN and the EVN in terms of data rate per station) and would allow IRA to be at the forefront of development of new e-VLBI techniques at a very exciting time. Even with a single e-VLBI baseline (connections of two antennas to the correlation centre in Bologna) there would be potential for Italy to play a leading role in the development of high-bandwidth VLBI techniques.

### 3.4.6 Antenna Tests

Finally, having the ability to test VLBI modes without sending the data outside the institute will no doubt be useful for testing new back-ends and the antennas themselves. This may be particularly useful during SRT commissioning. It would allow, for example, holography experiments to characterise the shape of the parabola to be correlated, possibly even with an on-site installation of DiFX.

## 3.5 Conclusions

It is clear that while the Medicina, and Noto antennas are of great value within the EVN and the Sardinia Radio Telescope will be an important addition (Mantovani 2006) an Italian array is feasible and adds great value to these antennas. The value of a fibre to Noto would be huge: it would allow experiments with extremely-high

bandwidth VLBI to be made between Medicina and Noto, with correlation taking place at Bologna Headquarters.

# Chapter 4

## A correlator at the Istituto di Radioastronomia

In this chapter I will give an overview of different correlator architectures. I will then introduce the DiFX software correlator (Deller et al. 2007) and describe the installation which runs at the IRA computing centre. I will describe how computing hardware of the centre was benchmarked for running correlations, and present the results of this analysis.

### 4.1 Different Correlator Architectures

In the earliest days of interferometry the antenna voltages were not digitised and correlation was achieved using analogue electronic circuitry. Digital correlators are now universal at radio wavelengths and until recently have been grouped into hardware correlators and software correlators. Most correlators over the last 40 years have fallen into the former category with a few notable exceptions (one of the very first VLBI correlators was a multipurpose computer, see Alef 2004).

Now, however, there are many other digital architectures well-suited to correlation and a list of modern correlator architectures is something of a continuum from traditional ASIC-based correlators to pure software correlators, implemented on commodity PCs.

One peculiarity that all of these architectures must deal with, is that compared to many high-performance computing tasks, the number of operations done on each incoming bit is very low. Therefore correlation is often limited by the speed at which the data can be transported onto the processing unit.

#### 4.1.1 ASIC Architecture

Traditional hardware correlators consist of ASICs: Application-Specific Integrated Circuits consist of semiconducting chips hard-wired to perform the digital operations necessary to achieve a specific task. These are distinguished for their extremely high non-recurring engineering costs, but extremely good performance: the VLA correlator was switched off in January 2010 after being used for 30 years — an exceptional lifespan for high-performance digital hardware. ASICs are most economic when made in large quantities — and therefore make most economic sense when many correlators (or one particularly large one) are being built at once: all of the ASICs required for all the correlators (and required spares) can then be manufactured in a single run. Hardware correlators based on an ASIC architecture continue to be commissioned and built today: for example, those built for the EVLA and e-MERLIN using the WIDAR architecture.

ASIC-based correlators are generally clocked machines; that is the speed at which they run is tied to the speed of the incoming data and it is not possible to

read the data at a slower speed to allow further processing. This puts very sharp constraints on parameters such as the number of spectral channels per sub-band for these correlators (For example the JIVE correlator: see Campbell 2004).

#### 4.1.2 FPGAs, GPUs and other multi-purpose hardware

Some correlators use specialist yet more multi-purpose hardware such as FPGAs and GPUs. FPGAs can carry out anything that ASICs can do with the difference that while each individual chip is expensive they are reconfigurable and further chips can be bought as and when needed.

The market for high-resolution video and gaming has produced very high-performance processors optimised for vector calculations. This is achieved through very high parallelisation and threading and dedicated memory closely connected to the processing unit. These include the graphics processing units used in graphics cards and the cell processor used in the PlayStation 3. Both of these may be used for correlation, however the because of the large amount of data involved in correlation, bandwidth problems can occur. As the use of these unorthodox architectures for scientific applications increases, efforts are being made to make their use straightforward through the use of APIs to standard programming languages<sup>1</sup>. This reduces the expertise required somewhat however this remains a specialised area of computer programming.

GPUs are used for the prototype MWA correlator (see Wayth et al. 2009) Some of these architectures are compared for LOFAR correlation in van Nieuwpoort and Romein (2009).

#### 4.1.3 Software correlators on specialised multi-purpose computers

Moving further towards the software end of the correlator spectrum, multi-purpose (yet still specialist) computing architectures can be used for correlation. A prime example is the LOFAR correlator which uses an IBM blue gene supercomputer for correlation (also discussed in van Nieuwpoort and Romein 2009).

#### 4.1.4 Software correlation on commodity PCs

Finally, normal commodity PCs can be used for correlation using the CPU for correlation. Until recently this was only feasible for very low bit-rate observations. However, as noted by Brisken and Romney (2008) (see fig. 5) at least for VLBI, recording rates have increased slowly over the last 30 years compared with Moore's law, and software correlation is now feasible for VLBI.

The main advantage of this kind of correlation is flexibility: by designing the correlator to work on multi-purpose hardware it will run on any cluster designed for scientific data processing. As computer specifications increase the capabilities of the correlator will increase as well.

In addition, the specialist knowledge required to work on the code is much reduced compared with other architectures. This means that many people working on VLBI software and hardware will have the skills required to dive in and equip a software correlator with new features. A version can also be tailored for a specific project which requires a particular observing mode. This flexibility is particularly valuable for VLBI where networks are often ad-hoc and a great number of different antenna back-ends and recording systems are in use.

---

<sup>1</sup>see for example <http://www.khronos.org/openc1/>

There is also flexibility which stems from the higher-level programming. Since the operating systems layer handles access to the CPU and other hardware, software correlators are not clocked machines and can therefore proceed with jobs which require more processing more slowly.

## 4.2 The DiFX software correlator

Since the IRA computing centre is already equipped with a cluster suitable for running a software correlator, this was a sensible option to explore. In late 2006 an early release version of DiFX (Deller et al. 2007) was installed and the first fringe tests with small test correlations began. Since then a number of astronomical observations have been correlated at IRA including all of the observations presented in this thesis.

In the intervening period DiFX has progressed a great deal. It has become the standard correlator for the LBA and AuScope and will very shortly replace the VLBA hardware correlator. Usability has vastly improved and for a standard correlation the necessary skills required to prepare and run the operation can probably be learned in a day.

DiFX is described more fully elsewhere (Deller 2009; Deller et al. 2007). A brief summary follows.

### Implementation

DiFX is written in C++ and relies heavily on the Intel library IPP (Integrated Performance Primitives) for vector operations. It would be reasonably straightforward to transfer DiFX to a different vector library if this were required.

In order to allow distributed execution DiFX makes use of the message passing interface (MPI). This is a standard which has many implementations including MPICH and OpenMPI. It provides an API which allows separate process to relate to, and communicate with, each other. Any implementation of MPI will provide compiler (typically for C, C++ and FORTRAN) which generates specially linked executables; and also an application for launching the processes. These may all be launched on a single machine or across multiple machines (nodes) on a cluster. In the latter case a separate interface (such as ssh) must enable communication between the nodes. The MPI launcher then establishes links between all of the processes. Each process is “aware” of how many processes there are, and is given an index which identifies the process.

### Types of process

There are three different kinds of process in DiFX. The first process is designated the FXManager. This process coordinates the sending of the data from the datastreams to the cores and writes the final output to disk. One process per incoming datastream<sup>2</sup> reads the baseband data from the module (if correlating from 8-pack), file (if correlating from disk) or network socket (if eVLBI is being used). The rest of the processes are assigned to be compute processes (also known as ‘cores’).

---

<sup>2</sup>currently there is always one datastream per antenna

### 4.2.1 Data flow

Large chunks of data are read from disk (or network socket or module). Typically these will be read in chunks of ( $\sim 10$ MB). Rather smaller chunks (a round fraction of the disk read — approximately a quarter) are then passed from each of the datastream nodes to the compute processes exactly as they were on disk.

Each compute node now has a small packet containing all of the baseband data from all of the datastreams for a certain interval of time. This is referred to as a ‘sub-integration’. Within the compute process one FFT block of samples is unpacked at a time for each antenna, Fourier transformed, cross multiplied and added to the accumulation. Typically a single compute process will run on each machine, however the compute processes are multi-threaded. The data received by each core is further divided in time and each thread works on a separate time chunk of data.

After correlation (with the results for all threads averaged together), the results are passed to the FXManager which averages the required number sub-integrations together and writes the results to disk.

## 4.3 The Cluster at IRA

The computer cluster that was used for these tests consists of 12 nodes. Each node of the IRA cluster has two quad-core processors, resulting in a total of 8 CPU cores per node: 96 in total. The nodes of the IRA are connected by two different networks; a 1 Gbit Ethernet network and a 1 Gbit Ethernet network. Storage of approximately 10 TB is available on a networked file system and high-speed connections are possible to Medicina and other institutes in Italy and around the world.

## 4.4 Cluster parameters

There are many considerations which can limit the speed of correlation. The main ones are outlined below.

### 4.4.1 CPU

There is an absolute limit to the amount of correlation that a particular CPU can do. Small increases in performance may be possible through optimisation of the code, but this is outside the scope of this thesis. The aim in optimising the cluster should be 100% CPU utilisation during correlation.

### 4.4.2 Network

The distributed nature of DiFX makes it necessary for large amounts of data to be passed between cluster nodes. This network traffic will be dominated by the packets (a few megabytes in size) which are passed from the datastreams to the cores. The data passing from the cores back to the FXManager should be substantially smaller.

Complications arise when the baseband data is also stored on a networked file system. This additional traffic can be avoided to some extent by placing the datastream processes on the machine which is physically serving the disk.

### 4.4.3 Reading the input data

In most cases this will be the limiting factor for correlation speed; just as it is the limiting factor for the sensitivity of VLBI. For eVLBI it will be speed of the network. For correlation from mark5 units, the correlation will be limited by the speed at which these can be read. If the data is to be stored on disks within the cluster, then this is somewhat complicated by the sheer size of the data. Storing all of the data, even for a single station, on a single commodity disk is not usually feasible. This means that technologies such as RAID must be used to stripe the data across several disks. This can lead to much higher read rates than would be possible for a single disk however this may come with a penalty for the latency.

In addition, there is the fact that a single RAID will probably contain the data for several stations. This means that the device is constantly seeking between different files to deliver the data for all stations. Finally hard disks can read data at different speeds, depending on whether the data is stored near the centre of the platter or near the edge.

### Writing the output file

The correlator could also be limited by the speed at which it can write out the output file to disk. Unless there is some problem with the file system this is highly unlikely to be a limiting factor for normal correlations. However, for wide-field VLBI the larger dataset outputs may approach the volume of the input baseband data! In these cases the speed of the output file system becomes critical. Writing the output file to the same disks as are being used to store the input data could also be problematic, since read-speed may be lost switching from reading the input data to writing the output data.

## 4.5 Characterising the cluster

Before tests are carried out using DiFX itself, the first aim will be to attempt to characterise the performance of the network and the disks more generally.

### 4.5.1 Network

The most simple test of the speed of the network is to measure the maximum throughput of data from one machine to another (for example the utility `tcp` will copy from the memory of one machine to the memory of another). This will generally provide an upper limit of the speed of the network.

A steady stream of data passing from one node to another does not really characterise the network usage of DiFX. A better idea of the network speed is found by measuring the speed at which packets of a few megabytes can be passed around.

This can be done using the utility `NetPIPE`. `NetPIPE` will pass increasingly large sizes of packets, starting at one byte. The speed at which a packet of one byte can be passed gives a measure of the latency of the network.

DiFX, in common with almost all distributed computing applications, uses MPI to pass data and commands from one node to another. There will be overhead associated with these commands. The most accurate way of assessing the speed of the network for MPI is to use the MPI version of `NetPIPE`.

Network	TTCP		NetPIPE		
	Throughput Mbit s <sup>-1</sup>	Latency $\mu$ s	TCP Throughput Mbit s <sup>-1</sup>	TCP Latency $\mu$ s	TCP Throughput Mbit s <sup>-1</sup>
1GBe	893	50.9	896	44.8	1770
10GBe	8510	34.1	8190	33.8	5330
Infiniband				3.5	11 500

Table 4.1: Speed of the network for three different networks: the 1 and 10 Gbit s<sup>-1</sup> networks on the cluster and an infiniband network on a cluster at the MPIfR in Bonn. The maximum throughput was taken as the highest throughput rate of any packet size (usually the largest packet size). The latency was the time taken for the quickest packet (usually the smallest packet size: 1 byte). Note that for the MPI results the packet does a round trip and all of the times are halved. Since the packet is already being sent back before it has all arrived the throughput may appear to be up to double the true network speed.

Three different network architectures were compared: the 1 GB Ethernet network on the cluster in Bologna, the 10 Gbit Ethernet network in Bologna and the infiniband interconnect on a cluster at the MPIfR in Bonn. The infiniband switch has a maximum theoretical throughput of only 8 Gbit s<sup>-1</sup> FIXME check. However it is known for its excellent latency.

The first test was a simple test using `ttcp`. This was carried out on both Ethernet networks in Bologna. 10 Gbit of data were transferred from machine to another as UDP packets. This test showed that the speed of the 10 Gbit network was almost 10 $\times$  faster than the the 1 Gbit network for direct throughput. Indeed, since the CPU usage was almost 100%, it is possible that this could have been a limiting factor.

Next NetPIPE was used to send increasing large TCP packets from one machine to another (fig. 4.1). The results of this comparison were also encouraging: the 10 Gbit network remained 10 times faster than the 1 Gbit network. This also allowed measurement of the latency of the connections for the first time: the 10 Gbit network was marginally better (see table 4.1). Unfortunately the results were not quite as good when the same program was used with OpenMPI. Here it was also possible to compare with the infiniband architecture used on the cluster in Bonn. The difference between the 1 Gbit and 10 Gbit is considerably less here. The performance of the infiniband network exceeds that of the 10 Gbit network for all packet sizes.

In particular the variation in performance for each test appears to be very high for the 10GBe network compared with the other two.

## 4.5.2 Disks

A number of *ad hoc* and more formal tools exist to test the speed of computer disks. The simplest is `dd`, a low-level tool which will read or write chunks of data to or from disk.

`dd` was used to create a file 11GB in size. The disk cache was flushed and then the same file was read. This test was performed on one of the cluster nodes, and also on the computer which hosts the file system for both the IRA cluster and the Bonn cluster. For all tests apart from that on the MPIfR disk server, the CPU usage of `dd` was very low. Even in the latter case it was only 50%. The results were rather peculiar but seemed consistent. The nodes of the IRA cluster have



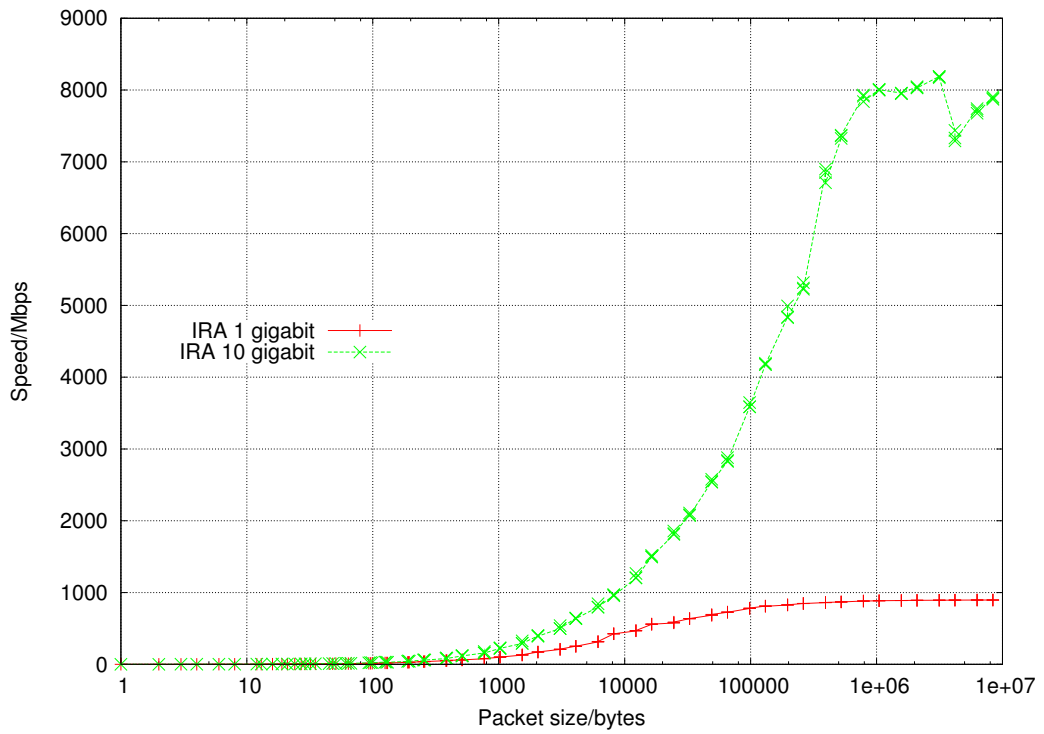


Figure 4.1: Speed of sending different-sized packets from one node to another over the two networks available on the IRA cluster using NetPIPE in TCP mode.




Figure 4.2: Speed of sending different-sized packets from one node to another over the two networks available on the IRA cluster. The Bonn cluster was also tested. The test used NetPIPE compiled for OpenMPI. *Note* that the speeds shown are calculated for “round-trip time”. This means that they actually show approximately double the speed of the interconnect.

Network	host	Read speed MB/s	Write Speed MB/s
IRA	node	209	50.1
	fileserver	167	52.8
MPIfR	node	83.3	97.7
	fileserver	593	567

Table 4.2: Comparison of disk read and write speeds for two different clusters.

a much faster read speed than write speed. This is good for correlation where the baseband data is located on this file system. It seems rather strange, however that the write speed is somewhat higher from the node than from the disk server, however it is a difference of only 20%.

For the MPIfR cluster the speed of `dd` when run on the RAID server itself is extremely high. For the usual case of reading the RAID from the node, the speed was slightly less than that of the IRA cluster.

For DiFX use correlating from baseband data from the RAID it would likely be a very big advantage to run the datastream nodes on the fileserver machine.

## 4.6 Benchmarking the cluster for DiFX

Now that the properties of the cluster have been determined it is possible to test the correlator and see what performance the cluster is capable of. For this an EVN test correlation will be used. Baseband data for 6 stations is available however in the first instance it will not be used. Instead the program `vlbifake` (Phillips and Deller 2009) will be used to generate simulated datastreams. This eliminates the bottleneck of reading from disk, and the performance that the cluster is capable of were it limited by CPU and network alone can be measured.

### 4.6.1 Optimum number of threads

The first step was to determine the optimum number of threads to be used for each compute process. This was done by replicating fig. 2 of Phillips and Deller (2009): a 6-station observation was correlated multiple times, varying both the number of compute nodes and the number of threads. The results of these correlations are shown in fig. 4.3. The setup of the cluster was optimised for consistent results *not* for best performance: one node was assigned to run the FXManager, another 6 were assigned to be datastreams which left up to four processing nodes (one machine was unavailable).

In addition to the compute threads of each compute processes, there is an additional thread which handles sending and receiving of data. For this reason it is generally found that running with one thread less than the number of CPU cores is optimum. Phillips and Deller (2009) found that somewhat less than this was even better. These results show that there is very little difference in performance between 7 and 8 threads per node. If anything 8 is preferable.

### 4.6.2 Assessing the performance of the cluster with different numbers of stations

Next the performance of the cluster with different numbers of compute nodes was determined for different numbers of stations. Each compute node ran with 7 processing threads for all of these runs.

As before, the setup was optimised for consistency and not for performance, however for correlations of less than six stations, it was possible to test performance with more than four compute nodes. The results are summarised in fig. 4.4. The results show a remarkably clear linear relationship processing speed with number of cores for lower data rates. This shows that when limited only by CPU, DiFX scales extremely well.

At higher data rates the results become more unreliable: during the correlations the speed of the correlation varied considerably for these runs. This is very likely not due to DiFX but to the 10 Gbit Ethernet network at IRA. As was noted in fig. 4.5.1, at least for MPI the network performance is very jittery.

### 4.6.3 Benchmarking the cluster for disk-based correlation

Finally the disk bottleneck was restored and the speed of the cluster was tested correlating with different numbers of stations. The results are shown in fig. 4.5 and appear to be more or less consistent with the read speed per node of  $1670 \text{ Mbit s}^{-1}$  found in the disk benchmarking (see table 4.2). The total disk read speed was highest for the four-station experiment with a total disk-read speed of  $1475 \text{ Mbit s}^{-1}$ . There is some room for improvement (perhaps by optimising the disk read size) but only for a performance gain of some 10 % to 30 %. Further gains may be possible by placing the data from each station on a different RAID, each served by a different computer.

### 4.6.4 Maximum performance for a 6-station experiment

In addition to the results shown in fig. 4.4, correlations of vlbifake data were also carried out with several datastream nodes running on a single PC. This caused a noticeable, though not critical degradation in performance. However it was minimal, provided that the total datarate from that node was below  $1 \text{ Gbit s}^{-1}$ . Running with the 6 datastreams processes shared between two nodes (allowing at least one processor core for every datastream process and vlbifake process the throughput per datastream was  $518 \text{ Mbit s}^{-1}$ ). Therefore the cluster is *just* capable of correlating a  $512 \text{ Mbit s}^{-1}$  6-station observation in real time. Further improvement could be seen by moving the datastream processes onto the file servers, freeing even more nodes to act as compute nodes.

## 4.7 Conclusions

The results obtained are excellent for such a modest cluster. They mean that correlating data from all of the Italian stations could easily be done in real time at today's recording rates.

The results also suggest that DiFX scales extremely well to the available computational power, with fig. 4.4 showing a very linear relationship as long as the network is not close to saturation. For the  $10 \text{ Gbit s}^{-1}$  network this was not a problem up to and above  $1 \text{ Gbit s}^{-1}$  and the network results for the Bonn infiniband

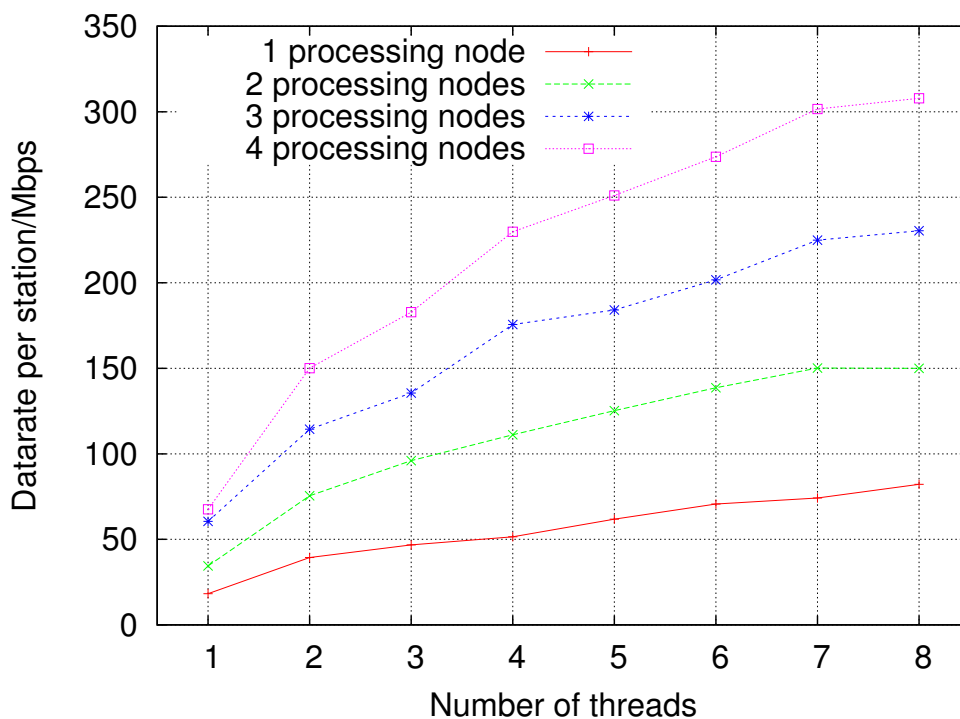


Figure 4.3: comparing the data rate per station for a 6-station observation correlating with different numbers of threads and compute nodes. There was very little difference in performance between 7 and 8 threads per node

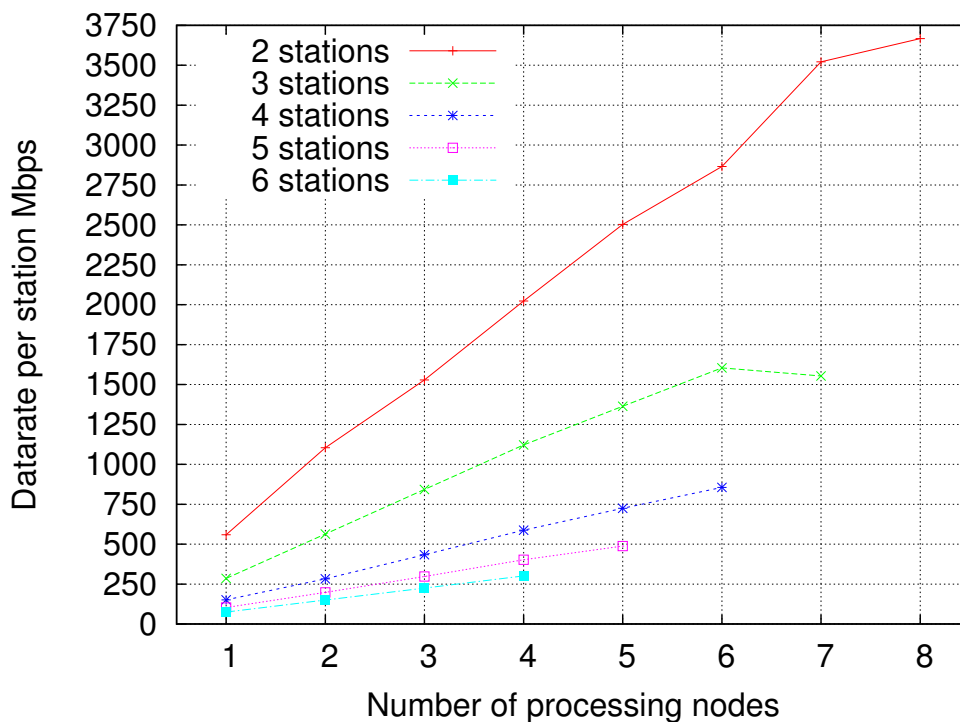


Figure 4.4: Comparing the data rate per station correlating different numbers of station with different numbers of compute nodes.

network show that it is possible to improve on this situation. Nonetheless if the bandwidth per antenna approaches or even exceeds  $10\text{Gbits}^{-1}$  it will likely be necessary to multiplex in frequency bands.

As the number of stations increases further, the baseline-based operations (which scale as approx.  $N^2$  for large  $N$ ) begin to dominate and the amount of operations required to correlate increases quadratically. Deller et al. (2007) show that this effect is noticeable beyond around 10 stations for DiFX. This is not likely to be a problem for the number of stations that would be correlated at IRA.




Figure 4.5: Comparing the data rate per station correlating different numbers of stations from baseband data stored on disk. Correlation was done with all of the baseband data on a single file system.

# Chapter 5

## Wide-field VLBI

The key feature that sets VLBI apart from normal radio interferometry is the longer baselines and resulting higher resolution. One similarity which is often overlooked is that the antennas themselves are broadly similar to those used for shorter-baseline interferometers. As discussed in sect. 1.3.3 it is the antennas themselves which dictate the maximum field of view. This means that while VLBI dramatically increases the resolution, compared with connected-element radio interferometry there is no corresponding decrease in the field of view. This allows very large images to be made.

A useful way to estimate the size of images that can be produced by an interferometer is the ratio of the resolution to the HPBW of the field of view. Both of these values are difficult to calculate exactly and more details are given in sections 1.3.1 and 1.2.1 respectively. However, to get a rough figure it is sufficient to assume that:-

1. For an interferometer the fringe spacing or angular resolution  $\theta$  is approximately  $\lambda/D$  where  $\lambda$  is the wavelength and  $D$  is the baseline length (sect.1.3.1).
2. The primary beam can be approximated by a symmetric Gaussian with a FWHM  $\Theta$  given by  $1.22 \cdot \lambda/d$  (sect. 1.2.1).

For heterogeneous arrays we conservatively assume that that  $d$  is the diameter of the *largest* antenna (whether this is too conservative is discussed in sect. 1.3.3). A broad-stroke measure of the size of image that can be made is given by  $\Theta/\theta (\approx D/d)$ , that is the width of the primary beam divided by the fringe-spacing of the entire array: the number of resolution units across the middle of an image of the primary beam. This is shown for various arrays in table 5.1.

This table shows that even for larger connected-element arrays such as MERLIN the images that can be produced are large, but for VLBI they are staggeringly large: 344 440 resolution units across the primary beam for the VLBA. If this were imaged with one pixel per resolution element (a bare minimum: see sect. 2.8) this would yield tens of gigapixels of image.

However there are three caveats to this:

1. This technique is only useful if there are several sources sufficiently bright and compact for VLBI detection within the primary beam;
2. It is necessary to correlate the data with sufficient resolution in time and frequency to avoid bandwidth and time-average smearing;
3. There must be sufficient computing resources and appropriate techniques to handle the resulting large datasets.

Table 5.1: Comparison of field of view in resolution units for different arrays: diameter of largest antenna  $d$ ; length of longest baseline  $D$ ; the approximate number of resolution units across the primary beam  $D/d \approx \Theta/\theta$ ; approximate channel width as a proportion of sky frequency  $\delta$  and integration time  $t_{int}$  required to image the entire primary beam keeping smearing losses below 5% (see equations 5.8 and 5.9).

Array	$d$ m	$D$ km	$D/d$	$\Delta\nu/v_{sky}$ (kHz/GHz)	$t_{int}$ (s)
GMRT	45	25	556	0	0
ALMA <sup>a</sup>	12	14	1116	350	8.2
VLA <sup>b</sup>	25	32	1280	350	8.2
MERLIN <sup>c</sup>	32	217	8680	66	1.5
EfMaMcNo <sup>d</sup>	100	1334	10334	34	0.78
SRTMcNo(Ma) <sup>e</sup>	64	894	13968	34	0.78
EVN	100	10180	101800	4.4	0.10
VLBA	25	8611	344440	1.3	0.031

<sup>a</sup>Largest configuration

<sup>b</sup>A configuration (the largest).

<sup>c</sup>Without the Lovell 76m.

<sup>d</sup>Array described in sect. 7.3.1.

<sup>e</sup>Array described in chap. 3.

In this chapter I will first set out the challenges faced in properly taking advantage of this exceptional capability. In sect. 5.1 I will look at the sensitivity which will be achievable with VLBI arrays in the near future. I will then address the question of how the sensitivity changes over the wide field, both for homogeneous and heterogeneous arrays (sect. 5.2). I will look at how multiple pointings can be used to cover areas greater than the HPBW of the main lobe of primary beam of the array and how the differing sensitivity to different points on the sky can be minimised and accounted for. I will briefly discuss the planning and scheduling of wide-field VLBI observations (sect. 5.3). I will quantify time-average and bandwidth smearing and look at several ways of solving this problem (sect. 5.4) and also the implications for correlation. I will look at wide-field imaging and identify the need for a technique which will split the data up into multiple smaller datasets which can be handled by today's calibration and imaging algorithms (sect. 5.5). I will then set out how this can be achieved (sect. 5.6). Verification of the accuracy of this technique is carried out in chap. 6, and this technique is used for the observation described in chap. 7.

## 5.1 VLBI sensitivity

VLBI is bandwidth limited: the sensitivity is limited by the amount of data that can be transported to the correlator. Table 5.2 compares VLBI array sensitivity to connected element arrays.

### 5.1.1 Recorded VLBI

After several decades in which the limitations of magnetic recording media meant little increase in bandwidth<sup>1</sup>, significant developments are being seen which al-

<sup>1</sup>228 Mbits<sup>-1</sup> was achieved with the Mark III in the early 1980s see Alef (2004).



Table 5.2: Array dimensions and bandwidth of current or near-future VLBI arrays against two generations of connected-element arrays. For VLBI we assume two bit sampling i.e. bandwidth=bit-rate/4

Array	Total Collecting Area (m <sup>2</sup> )	Bandwidth/Antenna (MHz)
VLA	13253	50
MERLIN	7304	32
EVLA	13253	16 000
e-MERLIN	7304	2 000
VLBA	4908	1 024 <sup>a</sup>
e-EVN	12230	256 <sup>b</sup>
EVN	18554	512 <sup>c</sup>

<sup>a</sup>near future

<sup>b</sup>Connections to some telescopes already at 10Gbps= 2500MHz of bandwidth

<sup>c</sup>Near future see van Langevelde (2008)

low large increases in sensitivity. Now that recording to disk (rather than tape) is universal (see sect. 1.4.4), the data volume which can practically be recorded at the stations is currently entirely dependent on the size of 3.5" disks which are available for an affordable price. 2-4 Gbits<sup>-1</sup> seems to be the limit which is practical in the near-term (Brisken and Romney 2008) and this is the target for the VLBA sensitivity upgrade<sup>2</sup>.

### 5.1.2 e-VLBI

Since fibre connections to many telescopes in the EVN already have a bandwidth of 10 Gbits<sup>-1</sup> and the DBE will allow output of 10 Gbits<sup>-1</sup> van Langevelde (2008) it seems likely that this data bandwidth will be realised in the coming years. In the foreseeable future fibre connections could provide speeds of up to 100 Gbits<sup>-1</sup> (Spencer et al. 2008). At this point having a correlator which can correlate in real time becomes more challenging.

### 5.1.3 Other limitations on bandwidth

Even in the near future the maximum bandwidth of receivers will be reached at least for lower frequencies. This problem is exacerbated by the fact that the frequency ranges observed have to be chosen avoiding very strong RFI at certain frequencies. At higher frequencies a certain bandwidth  $\Delta\nu$  represents a smaller fractional bandwidth  $\Delta\nu/\nu$  and in addition RFI is less of a problem. In this case more than 2 bits per sample can be used which will aid RFI mitigation. As a possible model, e-MERLIN uses 2 GHz bandwidth for 2 polarisations and 3 bits per sample for higher frequency and 0.5 GHz bandwidth for 2 polarisations and 8 bits per sample in L-Band (Spencer et al. 2008). With current VLBI systems, recording of more than 2 bits per sample is not possible however.

<sup>2</sup>Indeed, in this case it really is price that is the limiting factor. Multiple mark 5 units at each station would allow a significant increase in recording rate but with a corresponding increase in media costs.

## 5.2 Wide-field VLBI and the primary beam

Understanding the area of the sky to which the antenna is sensitive is critical for wide-field VLBI since it is this, and *only* this, which limits the field of view once smearing has been eliminated. Since this has already been discussed in some depth in sect. 1.2.1, here we will discuss two issues more specific to wide-field VLBI: the sensitivity of VLBI networks in which the antennas are heterogeneous in size; applying a correction for the primary beam; and how it may be possible to conduct VLBI surveys using multiple pointings.

### 5.2.1 Sensitivity over the entire primary beam

In sect. 1.3.2 we calculated a measure of the sensitivity of an interferometer  $\Delta I_m$ : the RMS noise in an image resulting from an observation with a particular array. This was calculated in equation 1.11, however this equation is only valid at the centre of the beam. As we move away from the centre of the beam the sensitivity is reduced.

Assuming that all of the antennas have the same pointing centre, the sensitivity of the array for any offset from the pointing centre can be derived by using  $A$  (the power function of the telescope over the primary beam) defined by equation 1.7 to scale the SEFD of each antenna (letting  $A_{max} = 1$ ). To calculate the sensitivity at some offset from the phase centre, we need to divide the  $SEFD_i$  of each antenna  $i$  by the power function of that antenna  $A_i(\vec{r})$  where  $\vec{r}$  is the offset from the pointing centre. Thus equation 1.11 becomes

$$\Delta I_m = \frac{1}{\eta_s \cdot \sqrt{2} \cdot \Delta\nu \cdot \tau_{int}} \times \left( \sum_{i=1, j=i+1}^{i=N-1, j=N} \frac{A_i(\vec{r}) \times A_j(\vec{r})}{SEFD_i \times SEFD_j} \right)^{-1/2}. \quad (5.1)$$

This is intuitive: consider the SEFD to be a flux equivalent to a source of measured flux  $X$  at the centre of the beam. If the same source were observed at the HPBW it would have a measured flux of  $X/2$ , and the SEFD would be equivalent to a source of strength  $2X$  measured at the centre of the beam.

$\Delta I_m$  has been calculated for 7 VLBA antennas and 7 EVN antennas assuming a radial Gaussian primary beam given by equation 1.7 and is shown in fig. 5.1.


We see that though the peak sensitivity of the EVN is very high compared with VLBA, the VLBA is sensitive to a much wider area. It would therefore be possible to cover a much wider area with a single pointing of the VLBA that with the EVN<sup>3</sup>.

### 5.2.2 Combining pointings

We may wish to combine data from different pointing centres, possibly overlapping. In this case we cannot assume radial symmetry, but we can still work out the power distribution for a point on the sky with RA  $\alpha$  and Decl.  $\delta$  where antenna  $i$  has a pointing centre at RA  $\alpha_0(i)$  and Decl.  $\delta_0(i)$ . Assuming the usual symmetric Gaussian:

$$A(i) = \exp \left( -\frac{(\alpha - \alpha_0(i))^2 + (\delta - \delta_0(i))^2}{2\sigma^2} \right) \quad (5.2)$$

<sup>3</sup>This assumes the same bandwidth for both arrays. It should be noted that at some frequencies the EVN now has more than 10 antennas of size 32m or less making the coverage of the EVN similar to the VLBA but with a more pronounced high sensitivity at the pointing centre



(a) Single 4-hour pointing with 7 VLBA Antennas. Central peak sensitivity approx  $50\mu\text{Jy}/\text{beam RMS}$   
 (b) Single 4-hour pointing with 7 European Antennas. Central peak sensitivity approx  $25\mu\text{Jy}/\text{beam RMS}$

Figure 5.1: Figures showing the theoretical weak-source limit over the entire primary beam for images made with a conventional single pointing of 7 VLBA antennas compared with 7 EVN antennas (Ef, Jb2, Tr, Mc, Nt, Wb (single dish), On) Arbitrary cut-off at the HPBW edge of the VLBA array ( $100\mu\text{Jy}$ ). This is essentially calculated using exactly the same function as used by the EVN Calculator website except that the SEFD of each baseline is scaled by an inverse Gaussian. The axes are scaled to is  $0.5 \times \text{HPBW}$  of Effelsberg (100m). The calculated peak values agree exactly with the EVN Calculator

with  $\sigma$  given by equation 1.8.

For the purpose of calculating  $\Delta I_m$ , we can assume that to combine data from  $M$  pointings with  $N$  baselines, we can simply sum over  $M \times N$  baselines (i.e. each baseline is counted  $M$  times with the appropriate on-source integration time  $\tau_{int}$  for each pointing).


It should be obvious that, at least for Gaussian-like primary beams,  $A_1 \times A_2$ , and therefore the theoretical sensitivity integrated over the entire primary beam, will be maximised when  $\alpha_1 = \alpha_2$  and  $\delta_1 = \delta_2$ : that is, at any one time all of the antennas should have the same pointing centre<sup>4</sup>.

### 5.2.3 Mosaicing with multiple pointings

This is a technique which is often used with multi-element arrays such as ATCA and the VLA. If correctly spaced pointings are made, UV data from multiple pointings can be used for sources between phase centres. Sources further from one phase centre will be closer to another, and with the correct spacing very uniform noise levels can be achieved: So-called Nyquist Sampling of the sky is illustrated in Holdaway (1995). There is no obvious reason why such techniques might not also be applicable to VLBI.

This may be a particularly valuable technique if applied to heterogeneous arrays. By making multiple pointings, the narrow but sensitive primary beam of the larger antennas can observe over a wider area. By making these pointings overlap the data from the smaller antennas will be applicable to all pointings. This raises the question of whether a heterogeneous array with more collecting area can match the survey speed of a homogeneous array of smaller antennas

<sup>4</sup>There may be a few exceptions cases where this is not desirable. For instance if a single large antenna is used to give a single strong baseline for each antenna on an in-beam calibrator, the smaller antennas could all point at the centre of the target field while the larger antenna could point at the calibrator.



(a) Two overlapping pointings with 7 European Antennas totalling four hours of observing time. Central peak approx.  $35 \mu\text{Jy beam}^{-1}$ . Pointing centre spacing is the HPBW of  $E_f$ .

(b) Four overlapping 1-hour pointings with 7 European Antennas. Central peak sensitivity approx  $40 \mu\text{Jy beam}^{-1}$ . Four pointing centres are placed on a square, the diagonal given by the HPBW of  $E_f$ . Pointing centre spacing is the HPBW of  $E_f$ .

Figure 5.2: As figure 5.1 except that the time is shared by multiple pointings. Note that the Gaussian approximations are not truncated and so over-estimate the sensitivity of  $E_f$  further from the pointing centre.

given an equal total collecting area.


Let us consider using the multiple pointings of the EVN sub-array described in fig. 5.1. Figures 5.2 and 5.3 show what happens when we split the four hours of observation time between two, four, eight and 13 overlapping pointings. We are able to cover successively larger areas of the sky with the sensitivity and area covered approaching that of the VLBA.

It should be noted that beyond the shaded area there is a large area covered by the smaller antennas well within the HPBW of their primary beam. This larger additional area surveyed at lower sensitivity could yield potential in-beam calibrators for subsequent studies.

Using this technique would be relatively straightforward where an out-of-beam calibrator is being used: a scan is made on the calibrator followed by one of the target fields and then another scan is made on the calibrator. Where there are multiple target fields these are cycled through in turn alternating with the calibrator. However using an in-beam calibrator would significantly complicate the use of this technique. Firstly some amount of time will be lost moving from one pointing to another. Secondly it is unlikely that the in-beam calibrator would be detectable for all pointings on all baselines. For a smaller number of target fields it would be relatively easy to transfer phase solutions from adjacent scans to fields without in-beam calibrators. For very large numbers of fields, a safer strategy may be to use nodding calibration in any case to ensure good phase calibration for all fields.

#### 5.2.4 Phase response across the primary beam

For a perfect parabola the phase response should be flat across the primary beam. Deviations from this are exactly what is measured when the holography technique is used to measure the surface inaccuracies of an antenna (Kesteven et al. 2001, see, for example,). Technical reports on Holography measurements do not usually report this raw data; however, the phase response must be measure over many



(a) Eight overlapping pointing with 7 European Antennas totalling four hours of observing time. Central peak approx.  $45 \mu\text{Jy beam}^{-1}$ . Pointing centre spacing is the HPBW of Ef.

(b) Thirteen overlapping pointings with 7 European Antennas totalling four hours of observing time. Central peak approx.  $50 \mu\text{Jy beam}^{-1}$ .

Figure 5.3: As figure 5.2 but the number of pointings is approaching the full coverage of the VLBA primary beam.

sidelobes ( $\sim 20$ ). Across the central lobe it should be constant (R.C. Walker, NRAO, Priv. Comm.).

### 5.2.5 Correcting for the primary beam

The procedure for correcting the primary beam used for the VLA and ATCA is described in sect. 1.3.3. A similar correction could be done for data from a homogeneous VLBI array (such as the VLBA) however it would be possible to assume that the primary beam does not vary across the image in most cases. For arrays of telescopes that are heterogeneous in size it would seem appropriate to correct the UV data since a different correction will be applied on an antenna by antenna basis (or each baseline can be associated with a certain primary beam following (Strom (2004); see also sect. 1.3.3). In principal a different correction should also be applied at different frequencies. Currently, the fractional bandwidth of VLBI observations is quite low apart from where bandwidth synthesis is used (see sect. 2.3.2). However 1GHz bandwidths are imminent and 2.5GHz bandwidth will arrive in the foreseeable future. Even sharing this bandwidth between two polarisations will mean a large fractional bandwidth for lower frequency observations. Sect. 3.2 of Bhatnagar (2009) states that for the EVLA the time dependence and frequency dependence of the primary beam becomes significant at as little as 10% fractional bandwidth and 10% of the HPBW of the primary beam.

It should also be noted that the primary beam correction is particularly important (and difficult) where the polarisation of the astronomical signal is of interest (this is not addressed in this thesis).

## 5.3 Planning a wide-field VLBI observation

Observing a single-pointing wide-field VLBI observation is little different from a normal VLBI observation (planning these has already been discussed in sect. 2.4), except that since a large tract of the primary beam is likely to be imaged, it is implicitly easy to use in-beam calibration where one is known. Where possible

it is advantageous to perform a low-sensitivity calibrator survey of the region of interest before scheduling the main observation. This can be done in the conventional manner (i.e. observing every potential calibrator for several minutes for two or three pointings) or by using wide-field techniques.

There are some occasions where it may be desirable to supplement an array of smaller dishes with one or two larger dishes: if there is a source close to the centre of the field which is only marginally strong enough to use as an in-beam calibrator, then a single larger antenna could be used to provide stronger detections for this source for at least one baseline to each antenna.

## 5.4 Time-average and Bandwidth Smearing

As discussed in sect. 1.3.6, after correlation we have a UV dataset which samples the visibility function across the UV plane. As the earth rotates the baseline moves through the UV plane; also the length of the baseline in the UV plane is different for every frequency. Hence every visibility datum is an integration of an area<sup>5</sup> on the UV plane whose area,  $A$  is given by

$$A = \frac{\delta\nu}{\nu} \cdot \sqrt{\left(\frac{du}{dt}\right)^2 + \left(\frac{dv}{dt}\right)^2} \cdot t_{int} \quad (5.3)$$

where  $\nu$  is the sky frequency,  $\delta\nu$  is the width of each spectral channel,  $t_{int}$  is the correlator integration time and  $\sqrt{(du/dt)^2 + (dv/dt)^2}$  is the speed the baseline is moving through the UV plane (which is a function of baseline and time). This last term is a function of the declinations and hour-angle of the source, and of the latitude and longitude of both the antennas (see Bridle and Schwab 1999, equations 18-33 and 18-34), however it is maximised for a polar source (and close to maximum for a source close to the zenith). This is because it is entirely a function of the non-radial velocity of the antennas with respect to the source.

To see the effect that this averaging has on our data, we need to understand how the visibility function  $V(u, v, w)$  changes as we move through the UV plane. The clearest way of understanding this is to consider the contribution to  $V$  of point sources different distances from the centre of the field. If the field only contains a single point source at the phase centre, the phases will be static along the time and frequency axes (ignoring un-modelled instrumental and atmospheric effects). A single source that is far from this phase centre will have visibilities that rotate rapidly in phase along the frequency and time axes (see sect. 5.7.1 for more details).

It is clear, therefore, that over-averaging the UV data in time and frequency will lead to a reduction in flux density, and increasing uncertainty in position for sources far from the phase centre. Various formulae to quantify the impact of averaging in frequency and time are referenced here: (Thompson et al. (1986) equations 6.67 and 6.71; Condon et al. (1998) equation 12; Bridle and Schwab (1999) various equations). The smearing depends somewhat on UV coverage and weighting and on the bandpass shape for bandwidth smearing. We adopt Bridle and Schwab (1999) equations 18-29 and 18-43 for their simplicity. They are reproduced below as equations 5.4 and 5.6:

$$R(\delta\nu) = \frac{1}{\sqrt{1 + \beta^2}} \quad (5.4)$$

<sup>5</sup>sector of an annulus of an ellipse

where

$$\beta = \frac{\delta\nu r}{\nu_0 \theta} \quad (5.5)$$

and

$$R(t_{int}) = 1 - \left( 1.22 \cdot 10^{-9} \cdot \left( \frac{r}{\theta} \right)^2 \cdot t_{int}^2 \right) \quad (5.6)$$

where  $r$  is the offset from the phase centre;  $\theta$  is the HPBW of the *synthesised beam*;  $\delta\nu$  is the width of each spectral channel;  $\nu_0$  is the sky frequency at the centre of the channel;  $t_{int}$  is the integration time (accumulation period) of the correlator; and  $R$  is the ratio of smeared flux to unsmeared flux. Note that equation 5.6 is an estimate for the smearing encountered for a 12-hour observation of a polar source (worst case).

Let  $r = \Theta$  (the HPBW of the *primary beam*) and allowing 5% smearing (i.e.  $R(\delta\nu) = R(t_{int}) = 0.95$ ) and since

$$\frac{\Theta}{\theta} = \frac{D}{d} \quad (5.7)$$

we can rearrange to provide these simple rules of thumb.

$$\frac{\delta\nu}{\nu_0} = 0.4 \cdot \frac{d}{D} \quad (5.8)$$

$$t_{int} = 6400 \cdot \frac{d}{D} \quad (5.9)$$

Equation 5.8 can also be formulated as

$$n_{channels} = 2.5 \cdot \frac{\Delta\nu}{\nu_0} \cdot \frac{D}{d} \quad (5.10)$$

where  $\Delta\nu$  is the width of the sub-band. The time and frequency resolution required to limit  $F - F_{att}$  to 5% is shown in table 5.1 for various arrays.

### 5.4.1 Limitations imposed by Correlator Hardware

ASIC-based Hardware correlators (see sect. 4.1) are often clocked machines. That is, the speed at which the data is read onto the processor is set by the speed of the processor. This generally places some constraint on the number of output channels and on the output integration time (for more information see sect. 4.1.1). This puts sharp constraints on the area that can be imaged from a dataset resulting from a single pass of the correlator — indeed this has been a limiting factor in previous wide-field VLBI observations (Garrett et al. 2005b; Lenc et al. 2008).

Taking DiFX (Deller et al. 2007) as an example of a software correlator, no such constraints are placed on the time or frequency resolution (memory usage may become problematic at more than 10000 channels per sub-band; however even for the VLBA, the number of channels required for a 16MHz sub-band at 1.6GHz is 8192 for a smearing loss of slightly more than 5%). In addition it is un-clocked and therefore, even if correlation takes place on a modest cluster the correlation will still proceed, albeit at a slower pace. In fact, as discussed in sect. 4.4, software correlators tend to be bandwidth limited and so correlating with extremely high time and frequency resolution limits correlation speed only if the output rate is high enough to overwhelm the disk to which the output data is being written. Software correlators are therefore a good match to the requirements of wide-field imaging.

## 5.5 Calibration and imaging of wide-field datasets

### 5.5.1 Calibration

Calibration of wide-field datasets provides lots of scope for using sources detected across the field for in-beam calibration. See (Garrett et al. 2004) for more details. There is also a great deal more practical detail on widefield calibration in chap. 7.

### 5.5.2 Imaging

An obvious approach is to image the large dataset directly (Garrett et al. 2001) at least for an initial search for detections (Lenc et al. 2008). Data reduction algorithms such as IMAGR and VTESS in AIPS are able to account properly for time and bandwidth smearing, working on datasets with multiple channels and short integration times. However there is another effect which comes into play for larger images.

### 5.5.3 Non-coplanar Effects

This is a problem specific to Fourier synthesis imaging where a single image is made from one large UV dataset. As discussed in Thompson (1999), the Fourier transform between the visibilities and the sky brightness is only two dimensional if

$$2\pi \cdot w \cdot (\sqrt{1 - l^2 - m^2} - 1) \ll 1 \quad (5.11)$$

where  $w$ ,  $l$  and  $m$  are as defined in sect. 1.3.6.

As can be seen from the equation above, this scales quadratically with distance from the phase centre of the image; this means that as images get bigger (in number of resolution units of diameter) this effect increases. A rule of thumb is that serious errors result if the product of the angular offset in radians times the angular offset in synthesized beams exceeds unity i.e.

$$1.22 \cdot \frac{r^2 \cdot \lambda}{D} > 1 \quad (5.12)$$

Imaging algorithms are in development which will allow this correction to be done efficiently (see for example Cornwell et al. 2003)<sup>6</sup>. This raises the possibility that in the future very large images could be made from very large UV datasets directly.

### 5.5.4 Parallel Imaging Algorithms

Even assuming that non-coplanar effects can be accurately and efficiently corrected, time (and memory requirements) would quickly become astronomical for VLBI images larger than a small fraction of the primary beam with the fastest CPUs currently available.

Some efforts are being made to parallelise these algorithms (see Bhatnagar 2009) however no currently available software suitable for VLBI data reduction supports this. Finally current datasets provide information on how the delay varies across the wide-field image only to first order. It is highly likely that this would also lead to errors in position for sources far from the phase centre for

<sup>6</sup>This report notes: *None of the algorithms in use ... actually make use of the instantaneous planarity of the [VLA].* Therefore all of the algorithms discussed should be applicable to VLBI.



direct imaging, unless a very sophisticated earth model is used for correction at the imaging stage. This is discussed in more detail below and also in chap. 6.

## 5.6 Multiple UV datasets

It is clear then that we are currently unable to generate a wide-field image from a large UV dataset directly. Another approach is to generate multiple UV datasets from a single observation. Provided that the baseband data has been recorded (not necessarily the case for e-VLBI) it is possible to correlate more than once using a different delay model each time. This permits the generation of several datasets phase-centred on different areas of the primary beam. However the correlator time required will soon become prohibitive for a large number of phase centres.

Another possibility is to correlate once with the time and frequency resolution sufficient to avoid significant smearing on any part of the field that is to be imaged. This large single dataset can then be transformed to be phase-centred on a different point. Before writing to disk the data can be averaged in time and frequency. This shifting can be done multiple times, in parallel if necessary, generating multiple datasets. These can then be imaged separately.

## 5.7 UV Shifting

### 5.7.1 Residual delays and residual delay rates

As discussed in sect. 1.3.5, in correlation it is necessary to generate a delay model to align the incoming sampled data as if the antennas were all on a plane perpendicular to the point on the sky being observed. This perpendicular plane defines the 'phase centre'.

The behaviour of the visibility function for sources at the phase centre and far from the phase centre were noted in sect. 5.4: a point source at the phase centre will have visibilities whose phases are static along the time and frequency axes, while a point source that is far from this phase centre will have visibilities that rotate in phase along the frequency and time axes.

This phase rotation along frequency and time axes is due to the residual delay and residual delay rate respectively. Consider an error in the delay: this will result in a phase shift for each datum, however the phase shift caused by the residual delay is frequency dependent, causing rotation of the phase along the time axis. Such a delay is (perhaps counter-intuitively from this viewpoint) known as a non-dispersive delay.

If the delay changes with time, so will the phase caused by this delay, causing rotation of the phase along the time axis. Again the rate of change of the phase caused by this delay will be different for each frequency.

### 5.7.2 Transforming visibility datasets

#### The transform

We want to transform our dataset in such a way that a particular distant source with phase rotating in time and frequency is moved to the phase centre where it will have static phase along both axes. We know that this rotation is caused by the delay model applied. It should therefore be possible to transform a UV dataset simply by:




Figure 5.4: Geometry for a source distant from the phase centre. The UV dataset has already been shifted so that both antennas are lying on a plane perpendicular to the first phase centre. For a source some distance from that phase centre, the wavefront will arrive at the two antennas at different times: i.e. there will be a residual delay (blue line). The baseline vector (orange) is also different for that source.

1. replacing the phase centre coordinates with those of the new phase centre
2. replacing the baseline vectors with those calculated for the new phase centre
3. adding a phase shift to every datum, accounting for the fact that the delay is different at the new phase centre.

### 5.7.3 UV shifting to first order

As discussed in sect. 1.3.6 each datum in a UV dataset contains a set of coordinates which allow us to calculate the delay to first order for points distant from the phase centre (equation 1.18). This simple linear model is implicitly assumed in imaging algorithms when each measurement of the visibility function is mapped onto the UV plane. This means that when a Fourier-inverted map of the sky shows emission a certain distance from the phase centre, the distance from the phase centre to the emission measured on the map assumes a linear delay model. If we construct a purely geometric model (i.e. we disregard all considerations other than the locations of the two antennas with respect to the source) then this model will also be linear (assuming  $l$  and  $m$  are defined as cosine vectors)

#### Improved linear model for the delay across the wide field

Traditionally UVW coordinates have been calculated using a purely geometric model. However a such a model is unlikely even to be the possible linear fit to the more sophisticated correlator model. The UVW coordinates used by the DiFX software correlator are derived from the correlator delays (Walter Brisken, NRAO,

Priv. Comm.):

$$(u, v, w) = c \left( \frac{\partial \tau}{\partial l}, \frac{\partial \tau}{\partial m}, \tau \right). \quad (5.13)$$

That is, for each datum, the delay at a point  $\delta l$  and  $\delta m$  is sampled and a linear fit is made.

### Limitations of a first-order model

Any non-linearity in the derivative of the delay with respect to  $l$  and  $m$  will result in a residual delay error even after shifting to the correct coordinates. Even a cursory glance at the higher-order terms involved in a far more sophisticated correlator delay model shows that these assumptions are likely to break down for larger shifts (For the limitations of a linear model derived empirically see sect. 6.3).

It is important to note that a UV-shift correction is applied *only* to the target source and not to any phase calibrator. This means that external phase calibration will not have any effect in reducing this error and therefore even second-order errors will incur significant positional errors.

### Current UV shifting software

Interferometry data reduction packages commonly have algorithms for performing this shift, for example AIPS has the task 'UVFIX'. UVFIX does attempt to model some of these non-linear effects: the most significant non-geometric term is the effect of differential aberration, a relativistic effect with magnitude  $\sim v_{\oplus}/c$  (a maximum of about  $10^{-4}$ ). However positional accuracy is still lost on longer shifts. This is demonstrated in sect. 7 and has also been noted by others (see Lenc et al. 2008; Middelberg et al. 2008).

There are also other disadvantages to using UVFIX. Converting from raw correlator output to FITS, loading the data into AIPS, converting to multi-source and running UVFIX requires that a potentially huge dataset be written to disk a minimum of four times before any averaging can be done. Moreover, extremely large datasets cannot be handled by AIPS efficiently, if at all<sup>7</sup>.

## 5.8 Accurate UV shifting

By generating two delay models, one for the correlation phase centre and another for the desired phase centre, we can transform the data very accurately and assure a high level of consistency between a shifted dataset and one correlated at that location.

Consider two telescopes located on the Earth at  $r_{1\oplus}^{\vec{}}(t)$  and  $r_{2\oplus}^{\vec{}}(t)$  with respect to a source at phase centre  $\vec{s}$ . During correlation the delay model will shift the datastreams such that it is as if the antennas were located at  $\vec{r}_1^{\vec{}}(t)$  and  $\vec{r}_2^{\vec{}}(t)$  on a plane perpendicular to  $\vec{s}$  and passing through the geocentre. For each time centroid, we want to shift the dataset so that the phase centre is at  $\vec{s}'$  with the new antenna coordinates  $\vec{r}_1^{\vec{}}$  and  $\vec{r}_2^{\vec{}}$  on a plane perpendicular to  $\vec{s}'$ .

This will require each visibility to be shifted in phase to take account of the distance between the two perpendicular planes. Since  $w$  is the component of the

<sup>7</sup>Specifically there is a limitation on  $n_{IF} \times n_{channels} \times n_{polarisations}$  when loading the dataset.

baseline vector in the direction of the source,  $\delta\tau$ , the light travel time between  $\vec{r}$  and  $\vec{r}'$ , is given by

$$\delta\tau = \frac{\delta w}{c} = \frac{(w_1 - w'_1) - (w_2 - w'_2)}{c} \quad (5.14)$$

where  $w$  has units of length. Provided that the baseline vectors are calculated sufficiently accurately they could be used directly to calculate  $\delta w$ . Alternatively we can use the delays  $\tau_1, \tau_2, \tau'_1, \tau'_2$ :

$$\delta\tau = (\tau_1 - \tau'_1) - (\tau_2 - \tau'_2) \quad (5.15)$$

where the  $\tau_n$  is the light travel time between  $r_{n\oplus}^{\vec{r}}$  and  $\vec{r}_n$ . As already mentioned, for the delay models used by DiFX, the baseline vectors are derived from the delays and equations 5.14 and 5.15 are equivalent.

The phase shift  $\phi$  to be applied to each visibility is frequency dependent and is simply

$$\phi = 2\pi(\nu \cdot \delta\tau \bmod 1) \quad (5.16)$$

for each sky frequency  $\nu$ .

## 5.9 Averaging and Imaging

Once the dataset has been shifted and the phases no longer rotate along the time and frequency axes, the dataset can be averaged in time and frequency. The degree of averaging possible will depend on the image desired. If individual sources are being picked out, it will depend on the source structure and the uncertainty in the *a priori* position of the source. If all of the primary beam is to be imaged, the averaging will depend on the spacing of the phase centres. There is clearly some degree of choice between closely-spaced phase centres which can be imaged quickly, and more sparse phase centres which will take longer to image.

In any case, unmodelled uncertainties in the antenna delays typically limit averaging to  $\sim 5$  seconds in time and  $\sim 0.5$  MHz in frequency for most cm wavelength observations.

## 5.10 Conclusions

In this chapter I have addressed a broad range of issues related to wide-field VLBI. I have identified the need to generate multiple UV datasets. I have also set out a way in which the necessary UV shifting and averaging can be carried out accurately and with requiring multiple writes of the un-averaged UV datasets to disk.

In appendix A I will describe the implementation of this algorithm, in chap. 6 I will test the astrometric accuracy of the algorithm and in chap. 7 I will go on to describe an observation which was made to explore the use of this technique.

# Chapter 6

## Testing the UV-shifting algorithm

Before using the UV-shifting algorithm described in sect. 5.8 to get astronomical results, it is necessary to test its accuracy. The most meaningful measure of the accuracy of the algorithm is to compare a dataset correlated on phase centre  $\vec{s}$  and shifted to phase centre  $\vec{s}'$  with a dataset correlated directly with phase centre  $\vec{s}'$ .

This will be tested for 3 shifts, each an order of magnitude larger than the previous. Even the shortest shift tested will be greater than those required for the data analysed in chap. 7 and the largest shift represents the longest shift (in terms of resolution units) which would be carried out for VLBI with parabolic antennas.

In sect. 6.1 I will outline the method of testing. In sect. 6.2 I will look in detail at the delay model applied to this dataset. In sect. 6.3 I will look at the results of this test, both for UVFIX and the CALC-shifting algorithm. In sect. 6.4 I will look at how the delay function varies over the wide field.

### 6.1 A method to test UV-shifting algorithms

#### 6.1.1 Preparing test datasets

First a reference dataset is generated by correlating on the known position of a strong source which is strongly detected on all baselines with minimal residual delay<sup>1</sup>. Next the same baseband data are recorrelated, but with a phase centre offset from the true position of the source. The shifting algorithm can then be used to shift the data to the correct position and the resulting dataset can be compared with the reference dataset.

This procedure will be repeated three times, comparing the CALC shifting algorithm with a reference dataset correlated on the correct phase centre for three different offsets: 1' N×1' W; 10' N×10' W and 100' N×100' W. These offsets produce delays of  $\sim 1, 10$  and  $100 \mu\text{s}$  respectively. Since the frequency of this experiment is 8410 MHz a shift of  $1 \mu\text{s}$  corresponds to 8410 turns of phase. For all datasets a delay error of 10 ps leads to a  $30^\circ$  error in phase. For the first shift the results were also compared with UVFIX both with and without aberration correction.

Clearly this technique will not assess whether any error in phase is introduced by the antenna by observing a source away from the pointing centre. This was addressed in sect. 5.2.1.

---

<sup>1</sup>In principal this is not absolutely necessary but makes the results rather easier to interpret.

### 6.1.2 Evaluating the performance of the UV-shifting algorithm

In order to assess the astrometric precision of the UV transform, the procedure used by Tingay et al. (2009) can be followed. In this case the object of the analysis was the verification of the accuracy of the DiFX software correlator for geodesy. This was assessed by comparing it against a reference correlation performed by a trusted correlator. The procedure is as follows: a single sub-band is extracted, the outer spectral channels are discarded, and the residual differences between the phases and amplitudes are measured and assessed for any systematic error that could degrade the astrometric precision even after calibration and averaging.

This analysis is carried out twice: first the two datasets are vector-averaged along the time axis, the scalar phases of the reference dataset are subtracted from the test dataset, and a least-squares linear fit is made to search for any residual delay error. Then all the spectral channels are vector-averaged and the process is repeated.

### 6.1.3 Correlation parameters

This analysis will be carried out on approximately one minute of data from a scan of the fringe-finder source 3C345 taken from the observation described in chap. 7 on three baselines: Mc-Ma, Mc-Wz and Ma-Wz. Post-correlation the Ef baselines and all but the central 50% of the first sub-band were discarded.

The parameters of the reference dataset (2 s integration time, 128 spectral channels per sub-band) were chosen to match Tingay et al. (2009) and were not sufficient to avoid smearing for the offset correlations. It was therefore necessary to correlate with high spectral frequency and very short integration time and then average down to produce a dataset that would match the reference dataset. The furthest shift requires  $\sim 2000$  channels per 8 MHz sub-band and an integration time of  $\sim 0.04$  s to keep time-average and bandwidth smearing below 10%. Correlating such a large dataset requires care, since high spectral resolution leads to large numbers of samples being used for each FFT; the time range covered by these samples then becomes large compared with the integration time. With 2048 spectral channels per sub-band, choosing an integration time of 0.032 s ensures that an integer number of FFT blocks (125) fit into each integration time<sup>2</sup>. Some error will occur in averaging these time integrations up to a 2 s interval, however this shouldn't have any significant effect on these comparisons.

UV shifting and averaging in frequency was done during the conversion to FITS format, time averaging was done in AIPS using the task SPLIT. After time-averaging the first and last time integrations were discarded. The parameters for each correlation are shown in table 6.1

## 6.2 The correlator delays

In this section I will examine in some detail the delay models which will be applied to the data to effect the shift. They were generated by CALC9.1 (Petrov 2008) using the tool used for generating the delay models used by DiFX: calcif2

---

<sup>2</sup>Early tests where this was not controlled showed loss of fringes for the Ma-Mc baseline after averaging (the baseline with the highest rate) and high scatter on the other baselines. This was likely due to the time centroid not accurately reflecting the data accumulated within the integration time.

Correlation Parameters and magnitude and rate of applied delay									
Dataset	offset arcmin	Channels per 8 MHz	$t_{int}$ s	Delay $\mu$ s			Delay Rate $\text{ps s}^{-1}$		
				Ma-Mc	Ma-Wz	Mc-Wz	Ma-Mc	Ma-Wz	Mc-Wz
Reference	0	128	2	0.0	0.0	0.0	0.0	0.0	0.0
Shift 1	1.414	128	0.4	0.654 183	1.314 06	0.659 875	-29.671	-15.104	14.567
Shift 2	14.14	512	0.1	6.5384	13.1425	6.608 69	-297.69	-152.82	144.872
Shift 3	141.4	2048	0.032	64.5065	131.558	67.051	-3073.3	-1705.6	1367.7

Residual errors on UVFIX datasets				
		Delay Error degrees		$r^2$
Ma-Mc	82	Ma-Wz	174	Ma-Mc
Ma-Mc	82	Ma-Wz	174	Ma-Mc
Ma-Mc	82	Mc-Wz	82	Ma-Wz
Ma-Mc	82	Mc-Wz	82	Ma-Wz
Ma-Mc	82	Ma-Mc	1.0	Ma-Wz
Ma-Mc	82	Ma-Mc	1.0	Ma-Wz
Ma-Mc	82	Ma-Mc	1.0	Ma-Wz

Residual Errors CALC-shifted datasets				
		Delay Error degrees		$r^2$
Shift 1	1.5	3.8	1.5	0.753
Shift 2	15	33	20	0.920
Shift 3	180	360	170	0.368
Shift 1	1.5	3.8	1.5	0.753
Shift 2	15	33	20	0.920
Shift 3	180	360	170	0.368

Table 6.1: Correlation Parameters and delay models used to correct the offset datasets. Residual errors detected after shifting. Delays and rates are also shown in fig. 6.2. UVFIX values are in comparison with the CALC-shifted dataset *not* with the reference dataset. Certain  $r^2$  values have been discarded where phase wrapping made them meaningless

(Briskin 2008)<sup>3</sup>. I will not consider the terms which go into generating the delay model (for an overview of VLBI delay models see Sovers et al. 1998, sect. III). Suffice to say that the main contributions to the delay function are the geometric delay, which for ground-based VLBI has a maximum of  $r_{\oplus}/c = 21.27$  ms. The second largest contribution is the rotation of the earth for which a maximum value is  $\omega_{\oplus} r_{\oplus}/c = 1.6 \mu\text{s s}^{-1}$ . A myriad of other terms contribute at a lower level.

The delays are provided as a series of 6-term polynomials (delay as a function of time): one per antenna per two-minute interval. Only the first half of one two-minute interval was required for this test. The correlator provides the time centroid of each time integration as a the number of seconds since midnight UTC (in double precision). From this difx2fits calculates the offset in seconds from the start of the polynomial interval (as a double-precision float) and the appropriate polynomial is evaluated using the following iterative procedure:

$$y = a_0 x^0 + a_1 x^1 + \dots + a_n x^n. \quad (6.1)$$

$$y_0 = a_0, \quad (6.2)$$

$$y_i = x y_{i-1} + a_i. \quad (6.3)$$

Double precision (64 bit) floating point numbers have a precision of approximately 16 decimal digits (Vickery 2009). Since the largest value that a delay can have is 21.27 ms, a lower bound for the error on a delay difference is  $2 \times 10^{-11}$  microseconds. The true value is likely to be somewhat higher since there will be a cumulative effect on the error as the delays are calculated from the polynomials; however the scheme above minimises this by accumulating the larger values first.

The delays for each antenna with respect to time ( $\tau_{Ma}$ ,  $\tau_{Mc}$  and  $\tau_{Wz}$ ) along with the delay for one antenna at the largest offset ( $\tau'_{Ma}$ ) are shown in fig. 6.1. It can be seen that the first-order delay is by far the most significant term, followed by the rate. The other terms of the polynomial are very small in comparison.

As discussed in sect. 5.8 a phase shift will be applied to each datum due to a delay  $\delta\tau$  given by equation 5.15 which is given again here generalised for a baseline between antenna  $i$  and antenna  $j$ :

$$\delta\tau_{ij} = (\tau_i - \tau'_i) - (\tau_j - \tau'_j). \quad (6.4)$$

Figure fig. 6.2 shows  $\delta\tau$  for each baseline for each of the three shifts. It can be seen from these graphs that the shift delay scales almost (but not quite) linearly with increasing offset, and also that the delay is changing at a sufficient rate that it must be re-calculated for each time integration.

### 6.3 Results

The results for all baselines and shifts are summarised in table 6.1. In addition the results for a single baseline are illustrated in figs. 6.3, 6.4, 6.5 and 6.6.

---

<sup>3</sup>The 1.5.0 version of calcif2 and calserver were used. Modifications to difx2fits to do the shifting were made against difx2fits version 1.5.2 to ensure the highest possible accuracy for the time centroids used to calculate the delays.




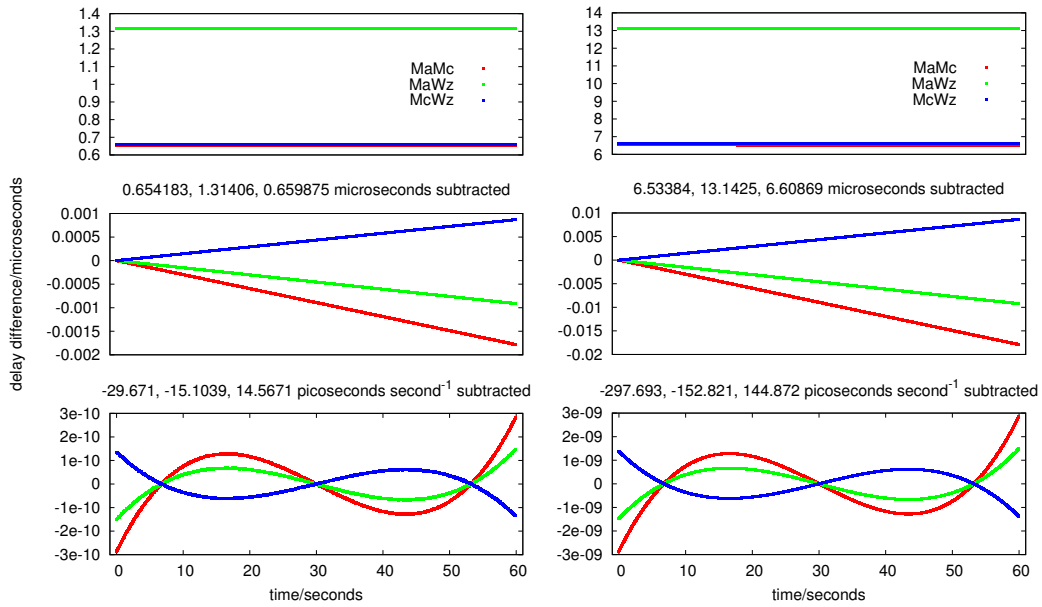
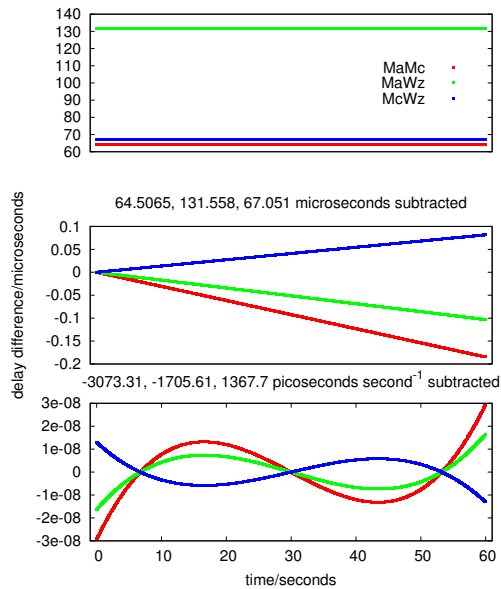


Figure 6.1: Delays  $\tau_{Ma}$ ,  $\tau_{Mc}$ ,  $\tau_{Wz}$  and  $\tau'_{Ma}$  as a function of time.  $\tau'_{Ma}$  is from the largest offset. The first graph shows the actual delays calculated by difx2fits as described in sect. 6.2 and dumped to text file during the FITS conversion with a time resolution of 0.04 seconds. A linear fit was then made to these points. The second graph shows the same data but with the constant offset subtracted. The third shows the data with both the offset and slope subtracted.



(a) Shift 1

(b) Shift 2



(c) Shift 3

Figure 6.2: Delay  $\delta\tau$  on each baseline as a function of time for each of the shifts. These were the actual values used for shifting, dumped to a text file during the FITS conversion process.

### 6.3.1 Shift 1

The first shift corresponds approximately to shifting to the edge of the primary beam of Effelsberg for the observation described in chap. 7. In turn this is equivalent to shifting to the edge of the primary beam for a MERLIN observation without the Lovell 76 m.

#### UVFIX

For UVFIX this shift results in an error on the phase of approximately one quarter of a turn (figs. 6.3). The very tight linear fit to the CALC-shifted data is striking: for the baselines shown in figs. 6.3a both the phase offset and slope with frequency are consistent with a delay error of 27 ps. The  $r^2$  values are very close to 1 for frequency and time for all baselines. This underlines that both of these shifting techniques work on the same principal. It also provides some reassurance that the transformation applied to the data is simply applying a delay shift and not degrading the data in some unexpected way.

The error for all of the baselines is proportional to the delay at the level of just 1 part in 40 000. Nonetheless the delay error this causes would be sufficient to dominate the positional error for an astrometry experiment.

A final note on UVFIX is that the output datasets with aberration correction switched on and off appear to be identical. Inspection of the source code shows that some calculation should be done, but no note is added to the history file making it unclear whether this code is ever accessed<sup>4</sup>.

#### CALC-shifted

For the data shifted using a second calc model an error is just discernible here (fig. 6.4). The slope of the fit is not meaningful but an offset of 1.2 degrees corresponds to a delay of 0.4 picoseconds – 1 part in 1.1 million. For comparison, geodesy experiments using this frequency typically achieve an accuracy of approximately 4 ps.

The magnitude of the offset is comparable with those found by Tingay et al. (2009) when comparing the results from two different correlators; however, the much higher  $r^2$  value shows that in this case the offset is much larger in comparison with the RMS scatter of the points and suggests a systematic error.

### 6.3.2 Shift 2

For this shift the error becomes more significant (fig. 6.5). However the slope of the delay remains small, indicating less than a full turn of error: an error corresponding to a delay which would produce a full turn would also produce a slope across the entire frequency range of  $3^\circ$ , which would likely be discernible in fig. 6.5a.

Indeed the error appears to be approximately  $10\times$  greater than in the previous shift on all baselines, and of the same sign, which strongly suggests a systematic error on the delay at the level of approximately 1 part in 1 million. An error of 4 picoseconds is at the level at which it would have a measurable effect on the highest-accuracy astrometric experiments, assuming that it is systematic across different antennas and sources.

<sup>4</sup>The version of AIPS used for this comparison was 31DEC09. The time stamp on UVFIX.FOR was 2009-09-21.

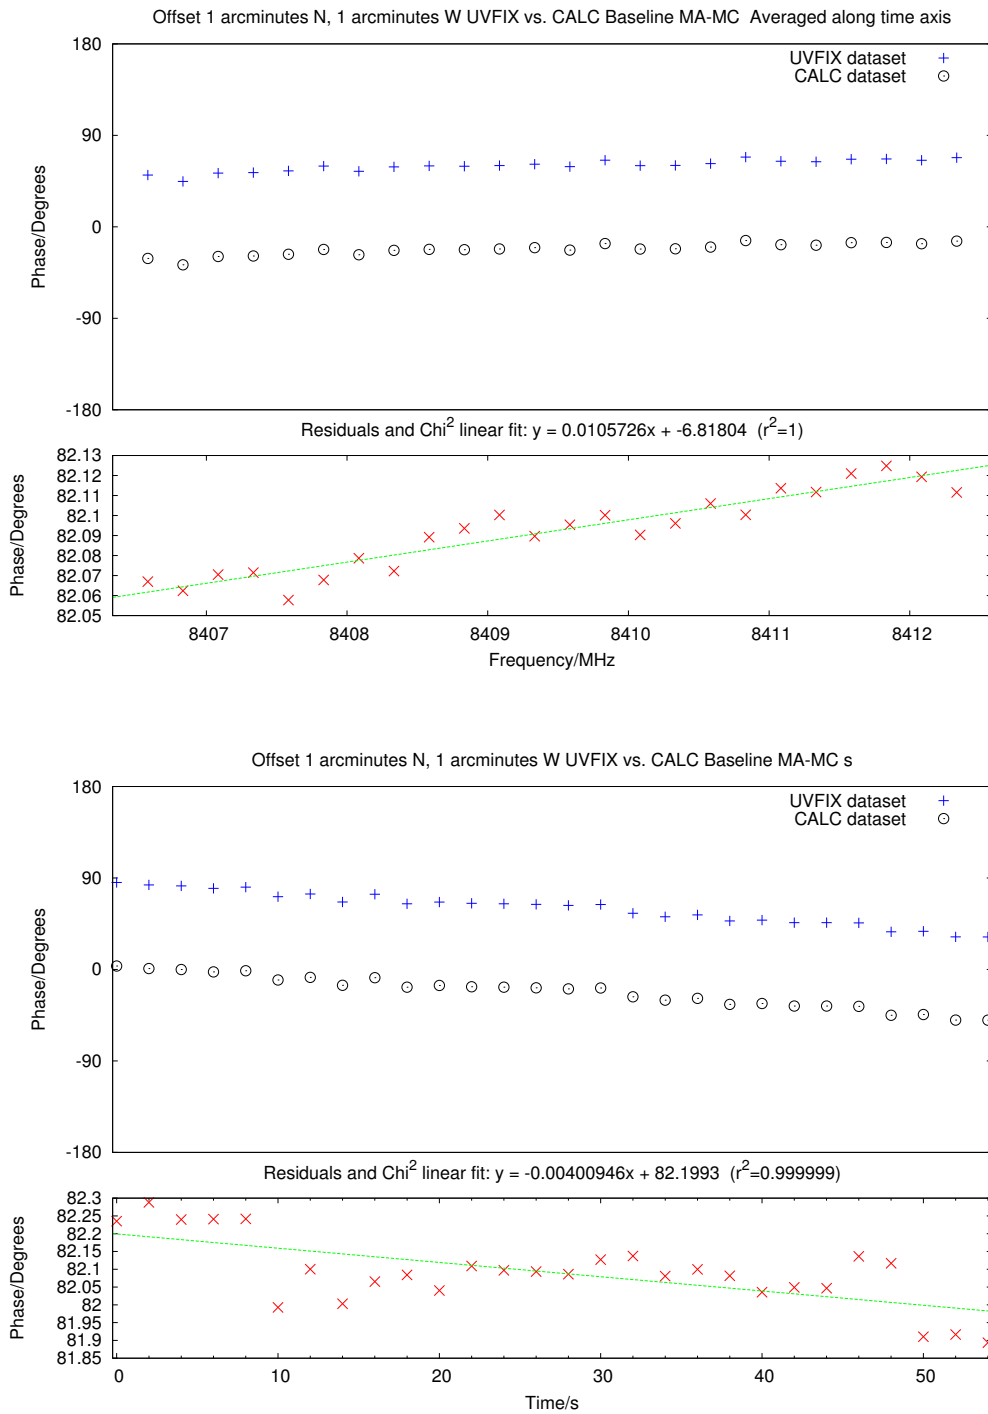


Figure 6.3: Comparison of UVFIX-shifted dataset with CALC-shifted dataset (*not* reference dataset) for a single baseline (Matera-Medicina) with a positional shift of 1 arcmin N and 1 arcmin W. Frequency data has been vector averaged in the time domain, time data has been vector averaged in the frequency domain

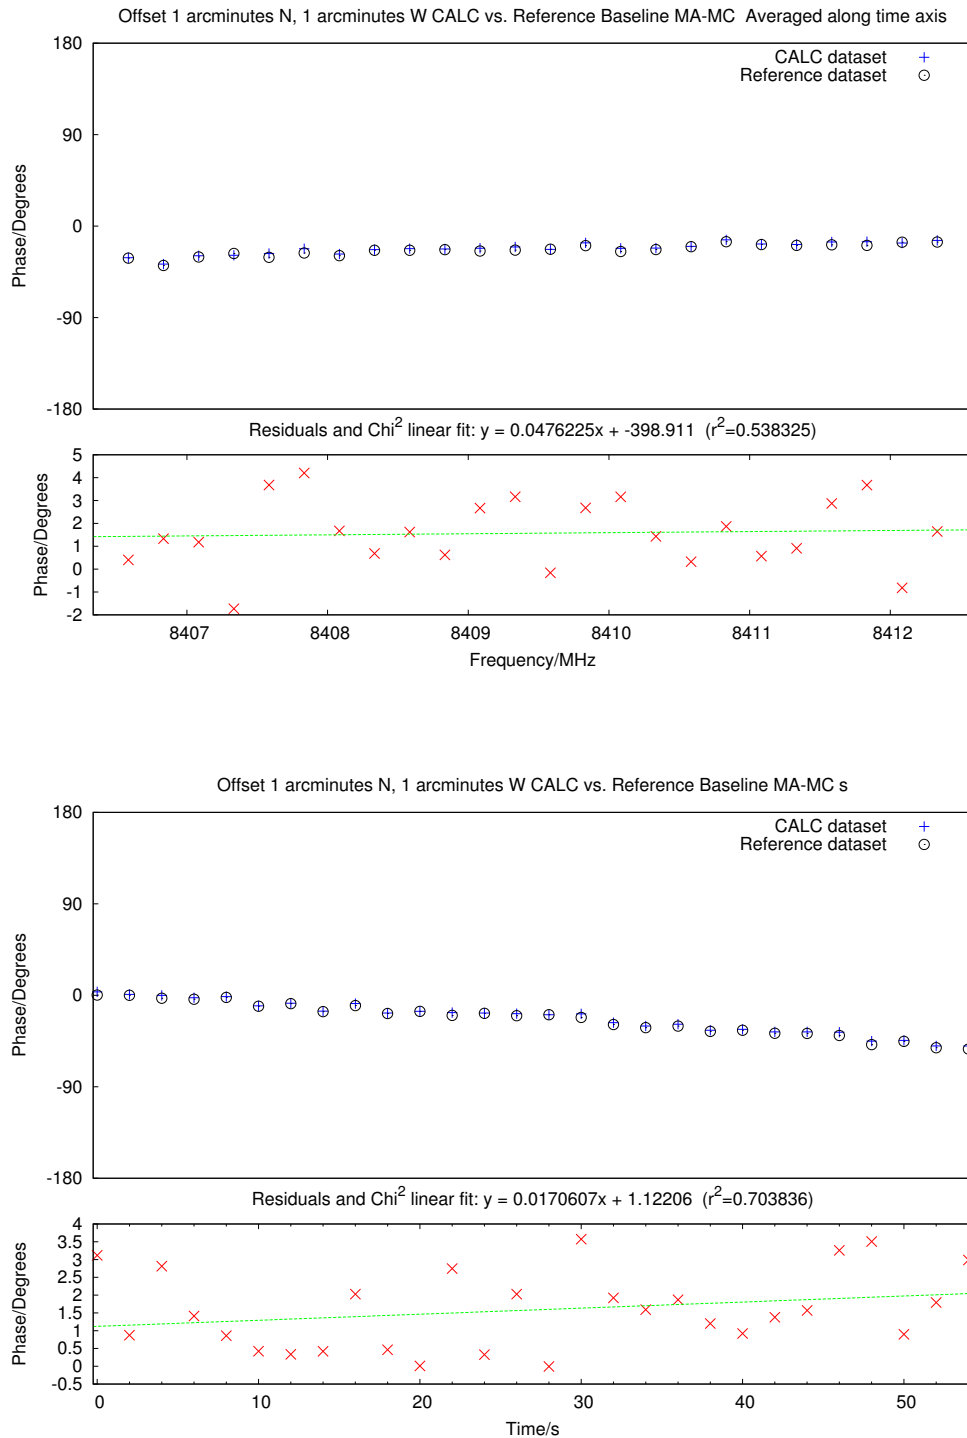


Figure 6.4: Comparison of CALC-shifted dataset with reference dataset for a single baseline (Matera-Medicina) with a positional shift of 1 arcmin N and 1 arcmin W. Frequency data has been vector averaged in the time domain, time data has been vector-averaged in the frequency domain

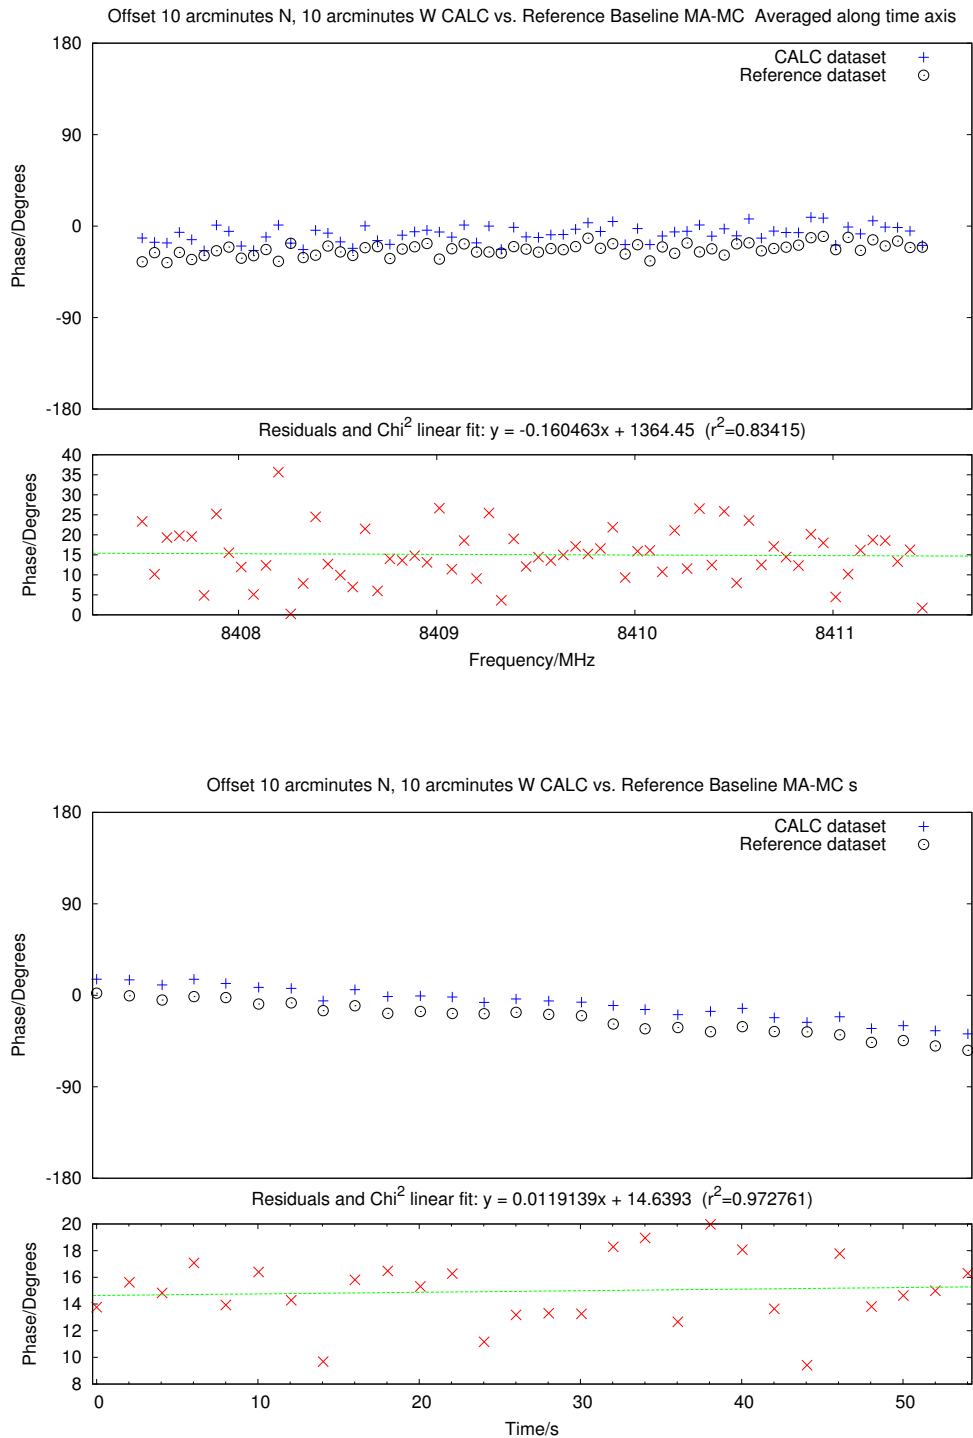


Figure 6.5: Comparison of the CALC-shifted dataset with the reference dataset for a single baseline (Matera-Medicina) with a positional shift of 10 arcmin N and 10 arcmin W. Frequency data has been vector averaged in the time domain, time data has been vector-averaged in the frequency domain

### 6.3.3 Shift 3

The results for this shift remain consistent with an error of one part in 1.1 million.

On the Ma-Wz baseline there is no offset between the reference and the shifted data, however given the results for other offsets and baselines this would be consistent with a single turn of phase. The results for the Ma-Wz cross the  $180^\circ$  boundary and so the Mc-Wz baseline is shown instead (fig. 6.6).

## 6.4 Variation of the delay over the wide field

In addition to the three correlation models (three different offsets and the reference dataset) the delays were also calculated for several other offsets between 0.1 and 100 arcminutes. The results are shown in fig. 6.7. Two fits have been made to this data. The first is a simple linear fit to the origin (0 delay at 0 offset) and the delay at the  $0.3'$  offset. This is very similar to the process used to calculate the U and V coordinates in DiFX (see sect. 5.7.3; calcif2 uses a  $1 \times 10^{-5}$  radian ( $\approx 0.33'$ ) offset). Although the delays are almost linear, the error on the linear fit soon becomes significant. Note the  $\sim 100$  ps offset for 1 arcminute is consistent with the error found for UVFIX at this offset, which implies that any further correction attempted by UVFIX reduces the error by a factor of less than  $\sim 3$ .

A second fit with an extra term was then made using one further point. This extra term gives a function for the delay which is sufficiently accurate for offsets approximately  $10\times$  greater than the simple linear fit. It would be possible extend calcif2 to make this calculation at correlation time. This enhanced function for the delay as a function of  $l$  and  $m$  (perhaps with an additional cross-term) could then be used to apply an accurate shift over a significantly longer shift at any point in post processing.

## 6.5 Conclusions

This method for UV shifting is more accurate than simple shifting using UV coordinates by almost two orders of magnitude, with accurate shifts possible over shifts covering any conceivable error in *a priori* source location.

Crucially, although some astrometric precision is lost for the very long shift, there is little increase in the RMS scatter of the points compared with the original data. The little that there is can probably be attributed to time-average and bandwidth smearing at the level of approximately 10% for these datasets. This means that it will still be possible to determine the source structure with the same accuracy as if the source had been at the phase centre.

The reason for the remaining error is unclear. For global VLBI where the rate of change of delay could be extremely high, a very small error in the time centroid used to calculate the delay could lead to a noticeable error. However the error seen here is not consistent with an error in the time used to calculate the delay (this would result in an error proportional to the rate). It is consistent with an error in the delay, or possibly an error in the frequencies used to calculate the phase shift to apply to each spectral channel.

It should be reiterated that this is a shift into parameter space that would *never* be explored with parabolic antennas: it is likely, though not certain, that this error would be seen for long shifts where such a shift would be sensible (shifting to the edge of the primary beam for the VLBA for example). Tests on

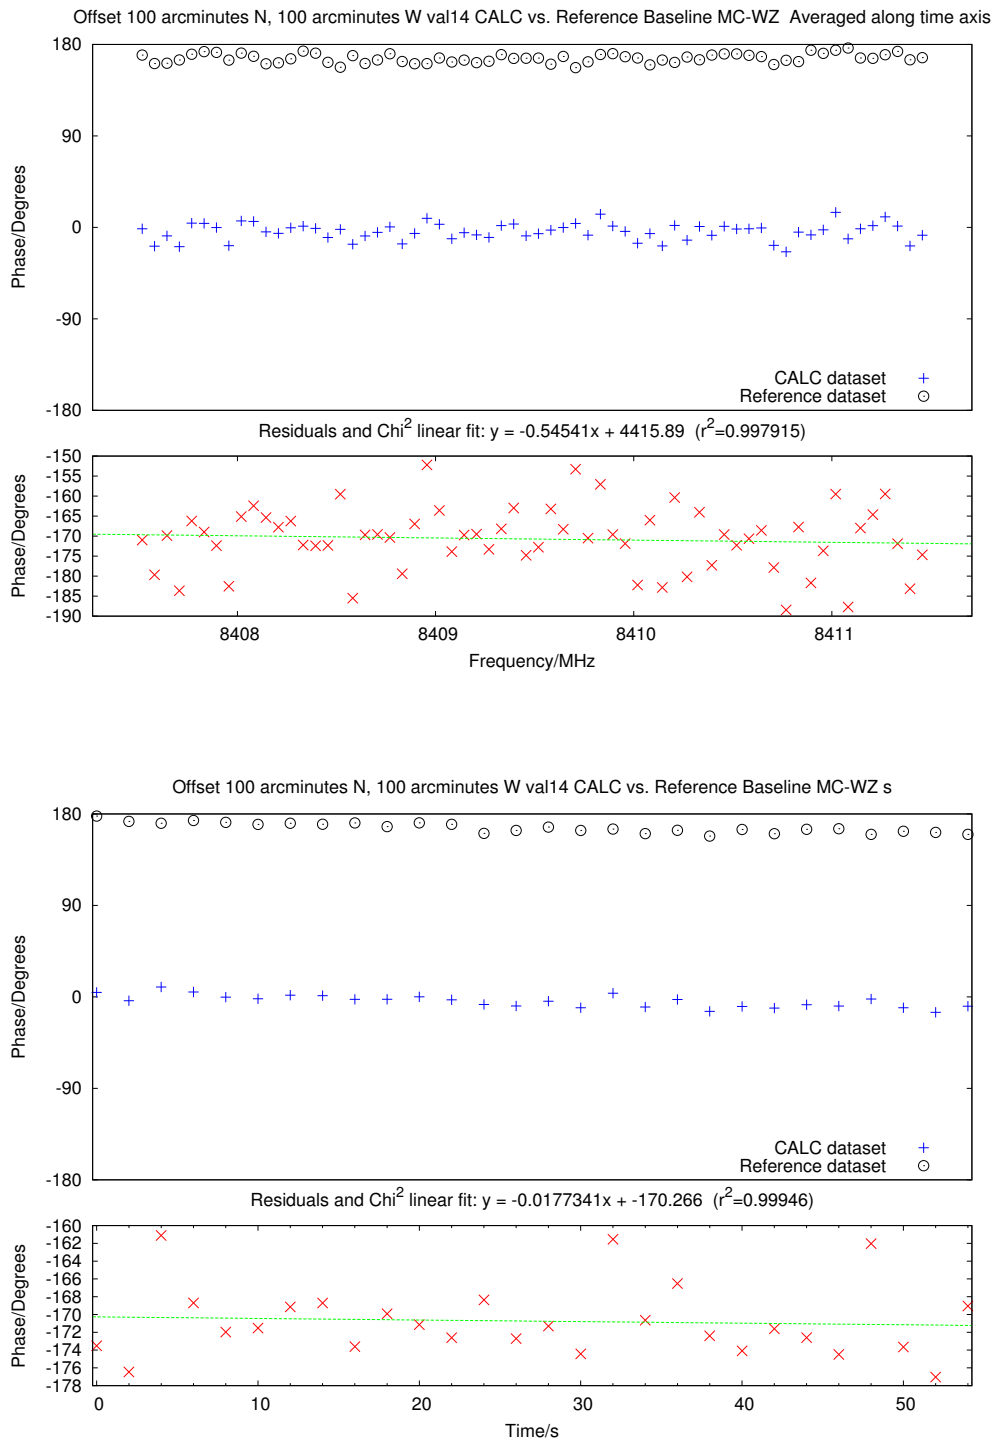


Figure 6.6: Comparison of the CALC-shifted dataset with the reference dataset for a single baseline (Matera-Medicina) with a positional shift of 100 arcmin N and 100 arcmin W. Frequency data has been vector averaged in the time domain, time data has been vector-averaged in the frequency domain






Figure 6.7: Magnitude of the residual delay for different offsets in arcminutes. In the second graph a linear fit has been made to the origin and the  $0.3'$  point and subtracted from the data. The final graph the residuals after a 2nd order polynomial was fitted to the origin, the  $0.3'$  point and the  $3'$  point.

other data would be required to confirm this. In any case, for such large shifts the correlated dataset becomes extremely large and it is likely that correlating to produce multiple phase centres would be necessary anyway.



# Chapter 7

## VLBI Observations of the 1320+299 Complex

### 7.1 Motivation

This four-station *ad hoc* VLBI observation was organised with the aim of testing the performance of the UV shifting method described in chap. 5. A secondary motivation for scheduling the observation was to assess whether wide-field synthesis imaging could be achieved with the four Italian stations (and of course to garner experience in correlation and imaging for this kind of observation). As the SRT was still under construction at the time of the observation and Nt was unavailable, the German antennas Ef and Wz were substituted.

The target field selected was that of the quasar 1320+299 (4C 29.48). This complex shows three discrete components spread over a field approximately 50'' wide, and is therefore a perfect target for wide-field VLBI.

In the following chapter I will first describe the target source, the observation parameters and the initial correlation. I will go on to explain how the UV-shifting algorithm was applied, and how a source search was conducted first over three fields centred on the *a priori* locations of the 3 main components, then over the entire field resulting in an image over 13 gigapixels in size.

### 7.2 1320+299

This source was first observed at 1.4 GHz and 5 GHz using the WSRT (Feretti et al. 1982) and later with the VLA in A and B configurations at 1.4 GHz, 5 GHz and 15 GHz (Cornwell et al. 1986). The source was also observed by VERA at 22 GHz (Petrov et al. 2007) with the component B chosen as the phase centre, however no detection was reported.

VLA and WSRT observations reveal three aligned components: a 20th magnitude quasar is coincident with component A (see fig. 1, Cornwell et al. 1986). Components B and C are both located on the south-west side of component A at distances of approx. 25'' and 50'' respectively.

Component A has a one-sided morphology: a prominent radio core, unresolved by the highest-resolution observations carried out so far (VLA A configuration at 15 GHz). It has peak brightness  $140 \text{ mJy beam}^{-1}$  and an integrated spectral index  $\alpha$  of 0.42 (where  $S \propto \nu^{-\alpha}$ ). A weak outer component is connected to the core region by a jet-like extension.

Component B has an asymmetric structure with a 'jet' which bends towards the north-east, extending to a distance of approximately 3'' from the peak. It has a peak brightness of  $75 \text{ mJy beam}^{-1}$  at 5 GHz and is also unresolved by the VLA in A-configuration. The integrated spectral index of this component is 1.12.




Figure 7.1: The NVSS image of the 1320+299 complex.

The VLA A-configuration 5 GHz image of component C reveals it to have a feature which extends in the north-west direction. It has a peak brightness of  $24 \text{ mJy beam}^{-1}$  and an integrated spectral index of 0.92.

The three components are polarised and the degree of polarisation at maximum brightness ranges between 5 and 19%. No significant depolarization or rotation measure is seen.

As a whole, the structure of 1320+299 looks rather peculiar. There is no conclusive evidence for the three components A, B and C to be physically related. This is confirmed by an image (fig. 7.1) extracted from the NRAO FIRST survey (Becker et al. 1995): this image does not show any bridge of emission connecting the components. On the other hand, according to Cornwell et al. (1986) the probability of an unrelated source with  $\lambda 20 \text{ cm}$  flux density between 250 and 450 mJy lying within  $\sim 25''$  of the central component B is only  $\sim 5 \times 10^{-5}$ . Cornwell et al. do not reach any firm conclusion about the classification of the components of the 1320+299 complex. It is hoped that this observation, the first dedicated VLBI observation of 1320+299 as far as can be ascertained, will allow the true nature of this complex to be determined.




Figure 7.2: Predicted RMS noise in mJy for images across the primary beam of the entire array:  $1 \times 100$  m antenna,  $1 \times 32$  m antenna and  $2 \times 20$  m antenna. Axes are scaled to the HPBW of the 100 m antenna.

### 7.2.1 The four-station array

The 1320+299 complex was observed making use of an *ad hoc* network of antennas, namely Ef, Mc, Ma, and Wz. Since both Ma and Wz are geodetic VLBI dishes, the choice of frequency band was constrained to either the S band ( $\sim 2$  GHz) or X band ( $\sim 8$  GHz). As the X-band primary beam of Effelsberg was still sufficiently large to cover the entire field, this band was chosen to give the highest possible resolution.

The sensitivity and baselines for this observation are broadly similar to what might be expected from an observation using the four Italian antennas sect. 3.1.4. Ef gives rather longer baselines and also a slightly greater sensitivity at this frequency, although for frequencies above 22 GHz, the SRT is projected to have a greater sensitivity than Ef.

Figure 7.2 gives the sensitivity across the primary beam for the observation parameters of our experiment. The central source lies at the centre of the primary beam, the other two sources lie halfway between the centre and the HPBW of Ef. The attenuation due to the primary beam in this area is less than 10%.

## 7.3 Observation

### 7.3.1 The observation

The observing session started on 12th December 2007 at 05:00:00 UT and stopped at 08:00:00 UT. The data were recorded using four base-band converters of the MarkIV data acquisition system and MK5A disk recorder with a sample rate

of 16 megasamples/s and 2 bit sampling for all sub-bands. The four observing frequencies were 8405.49, 8413.49, 8421.49, and 8429.49 MHz. Only the upper side band of 8 MHz was recorded for each channel. We were limited in our observing time and bandwidth by having only 1 Terabyte of storage available at the IRA computing centre at the time.

The MK5A disk packs recorded at each station were sent to the Medicina Radio Observatory, read back on a MK5A unit and loaded onto hard disks. These were then transferred to the computing centre at IRA. The DiFX software correlator installed here was then used to correlate the experiment, and the computer cluster was also used for further data reduction.

### 7.3.2 Observation technique

The observations were carried out using the phase-reference technique: switching between the target source 1320+299 and the calibrator, with a duty cycle of 3 minutes on target and one minute on a point source phase calibrator. The phase reference technique was required to increase our chances of detecting fringes from the individual components A, B, and C and as possible first epoch for future measurements of absolute component motions. The VLBA calibrator J1329+3154 (flux density 0.91 Jy at 8.4 GHz), approximately  $2.5^\circ$  from the 1320+299 complex, was used as reference source.

## 7.4 Modifications to the correlator

When the correlated data were first analysed it was clear that there was a problem with Ef. Careful analysis of the data for Ef showed that while, for a strong source, fringes were visible on very short time integrations they were not visible for integration times longer than a fraction of a second. Looking at how these phases changed with time revealed that they were wrapping twice a second. The diagnosis was therefore that the LO for Effelsberg was 2Hz too high. Modifications were made to DiFX to allow the correction of this offset. These modifications are described in sect. A.1

Although not really necessary for a one-off observation, software was developed to automate processing the log files and preparing the input files for DiFX. This software is described in appendix B.

### 7.4.1 LO offsets

At the time, the DiFX correlator was not equipped to handle the observed LO offset of 2 Hz. This was implemented in the software correlator and this is described in sect. A.1.

## 7.5 Initial Calibration and imaging

A modified version of the EVN ParselTongue pipeline was used to perform initial calibration in an automated fashion using the following procedure.

### 7.5.1 Initial flagging and calibration

The data were loaded using FITLD and an 'NX' table was added using INDXR. The flag table generated by DiFX for the times when the antenna was off-source

was applied, and the outer four channels were flagged using UVFLAG. ACCOR was used to correct the amplitude of the cross correlations by comparison with the autocorrelations. Amplitude calibration was applied using APCAL using system temperature values which had been extracted from the station logs and gain curves taken from another similar EVN experiment. The parallactic angle correction was then applied. The phase calibrator J1329+3154 were fringe-fitted using FRING. Complex bandpass calibration was done using BPASS using the calibrator 3C345. Before and after this calibration a large amount of plots are automatically generated by the pipeline This assists greatly in assessing the quality of the data and diagnosing any problems. In addition the corrections applied are also recorded on plots. The calibration tables produced in this way should be applicable to datasets phase-centred anywhere in the vicinity of J1329+3154.

### 7.5.2 Imaging

Following the initial calibration, the data are spectrally averaged down to a single channel per 8 MHz sub-band using the AIPS task SPLIT. Where calibration will be done at the imaging stage (e.g. self-calibration) it is convenient to have each dataset stored as a multi-source file (this simply affects whether the data are duplicated with calibration applied or whether a calibration table is attached to the data). Therefore all of the sources to be imaged are converted to multi-source files. Three images are then generated: normally- and uniformly-weighted “dirty” maps, and a uniformly-weighted clean map (loop gain 0.05, 100 clean components).

Automated self-calibration is also possible. By default, the sequence of self-calibration starts with phase only, and then alternates between phase and amplitude and phase. After some experimentation this was changed to a series of three phase-only self calibrations, followed by a series of phase and amplitude self-calibrations.

## 7.6 Initial imaging of the three components

As explained in sect. 5.6, in order to image large areas of the primary beam it is necessary to split the area to be imaged into many separate sub-fields. To begin with it was deemed sufficient to explore only three fields, centred on the *a priori* positions of the three components A, B and C. These were taken from Cornwell et al. (1986) and are listed in table 7.1. Three initial shifts were made, generating three UV datasets, each centred on the *a priori* position of one of the three components.

### 7.6.1 Imaging of the phase reference calibrator J1329+3154

If the phase-reference calibrator shows significant structure, then it will be necessary to use a model of this source when using it as a phase-reference source. For this reason an automatic self-calibration of this source was carried out. After 10 rounds of self-calibration (Three phase-only self calibrations followed by seven phase and amplitude self-calibrations; see table 7.2 col. A2 for the convergence of this procedure on the target source) the source appeared to be entirely unresolved, consisting of a Gaussian with a peak flux density of  $662 \text{ mJy beam}^{-1}$ . The RMS flux was  $395 \text{ } \mu\text{Jy beam}^{-1}$ , a dynamic range of 1700. The fact that this source

Component	RA			Decl.		
	h	m	s	°	'	"
A <sub>optical</sub>	13	23	0.873	+29	41	44.970
A	13	23	0.873	+29	41	44.810
B	13	23	2.521	+29	41	32.796
C	13	23	4.138	+29	41	14.875
A <sub>VLBI</sub>	13	23	0.873387(26)	+29	41	44.81409(21)

Table 7.1: *A priori* positions of the components measured from Cornwell et al. (1986) figs. 4a, b and c: VLA in A configuration at 4885 MHz giving an almost circular beam of width  $0.36''$  (equivalent to  $0.02^s$ ). The position of the optical quasar given in Cornwell et al. (1986) and the VLBI position of A determined by this observation is also given. The high accuracy of the VLA position of A compared with the VLBI position suggests a much higher accuracy in the VLA positions of the other two components since they have similarly high SNR in the VLA image. N.B. the errors shown for the VLBI position are the RMS errors on the Gaussian fit to the clean components of the source and do not include any estimate for any systematic errors.

appears unresolved, simplifies any phase referencing to this source since the effects of source structure can be ignored. The flux of J1329+154 on all baselines is shown in fig. 7.4a

In principal, the amplitude and phase solutions derived from this self-calibration could be applied to the target data when using J1329+3154 as a reference source.

## 7.6.2 Phase-referenced imaging of the three components

Each of the fields was then imaged using J1329+3154 as a phase-reference source. To avoid repetition of the initial calibration, further pipeline tasks were developed to copy the calibration tables from the initial calibration, apply them to other datasets and average the data. After this, the standard dirty and clean images were formed with the following imaging parameters: a pixel size of  $0.4 \times 0.4$  mas and  $1024 \times 1024$  pixels, giving a much larger image of each of the three sources than the *a priori* uncertainty in the source position.

Using this method it was possible to detect component A as a  $9\sigma$  point-like source after cleaning (table 7.2 col. A1). However no source was detected at the locations of B or C in the phase-referenced images (table 7.2 cols. B1 and C1). The RMS noise level in the B field was  $700 \mu\text{Jy beam}^{-1}$  and in the C field  $600 \mu\text{Jy beam}^{-1}$  for the normally-weighted dirty images. Although already fairly low, this is somewhat above the theoretic noise limit of  $250 \mu\text{Jy beam}^{-1}$  and even above the RMS noise in the self-calibrated images of J1329+2154.

## 7.6.3 The astrometric position of the A-component

The astrometric position of A with respect to J1329+3154 was determined by re-imaging the phase-referenced A dataset with a pixel size of 0.2 mas per beam. Next the AIPS task JMFIT was used to fit a single elliptical Gaussian to the detected source in the clean map. This yielded the position shown in table 7.1.

Fitting a Gaussian to the uniform dirty map gives a Gaussian with a major axis of  $9.9 \pm 1.4$  mas and a minor axis of  $2.6 \pm 0.40$  mas with a position angle of  $63.8 \pm 3.4$  degrees. The error on the Gaussian fit is  $\sim 0.000026^s = 0.35$  mas in Right Ascension and  $\sim 0.2$  mas in declination. Given the UV coverage (see fig. 7.3) it is unsurprising that the uncertainty in the RA of the source is higher than the uncertainty in declination.



Map	A1	B1	C1	A2	B3	A4	B4	C4	A5	B5
Uniform wt. dirty map	80.2 ±31	1.5±0.8	1.4±1.0	82.1 ±39.6	85.6 ±44	23.3 ±5.1	1.11± 0.91	1.4±1.0	13.4 ±6.2	1.24± 0.93
Natural wt. dirty map	75.2 ±18.5	1.2±0.7	0.8±0.6	77.8 ±22.1	79.9 ±17.6	28.0 ±4.8	0.95± 0.57	0.7±0.6	14.4 ±4.3	1.00± 0.57
Uniform wt. clean map	79.9 ± 9.1	1.4±0.8	1.0±1.4	82.5 ± 6.2	84.6 ± 5.5	21.0 ±3.3	1.08± 0.91	1.4±1.0	13.6 ±5.2	1.27± 0.93
Self-Cal 1 (phase only)	-	-	-	78.1 ± 4.7	-	-	-	-	-	-
Self-Cal 2 (phase only)	-	-	-	78.0 ± 4.7	-	-	-	-	-	-
Self-Cal 3 (amp&phase)	-	-	-	79.6 ± 0.81	-	-	-	-	-	-
Self-Cal 4 (amp&phase)	-	-	-	80.4 ± 0.63	-	-	-	-	-	-
Self-Cal 5 (amp&phase)	-	-	-	80.9 ± 0.54	-	-	-	-	-	-
Self-Cal 6 (amp&phase)	-	-	-	81.3 ± 0.53	-	-	-	-	-	-
Self-Cal 7 (amp&phase)	-	-	-	81.7 ± 0.59	-	-	-	-	-	-
Self-Cal 8 (amp&phase)	-	-	-	82.0 ± 0.63	-	-	-	-	-	-
Self-Cal 9 (amp&phase)	-	-	-	82.0 ± 0.52	-	-	-	-	-	-
Self-Cal 10 (a&p "L1")	-	-	-	82.1 ± 0.51	-	-	-	-	-	-

Table 7.2: RMS noise and peak images for various datasets with different calibration schemes: A1, B1 and C1 are phase referenced against the external calibrator; A2 shows the result of target fringe fitting and self-calibration on A; B3 shows the result of fringe fitting the B dataset on the target source causing A to be shifted to the centre of the dataset; A4, B4 and C4 show these datasets after a model of A has been subtracted; A5 and B5 use A as an in-beam calibrator and A is subtracted from the data. The Self-Calibration schemes applied were: two rounds of phase-only, followed by 6 rounds of amplitude and phase followed by 1 round of amplitude and phase with an "L1" rather than least-squares fit. Values are map peaks in mJy, error is the RMS measured over the entire field (likely to be a slight overestimate for the maps where a detection was present). All values shown are *without* primary beam correction.




Figure 7.3: UV-coverage of the target source

The true error is likely to be dominated by differences in the delay between the phase reference calibrator and the target. A gradient in tropospheric delay of 0.05 ns/degree (a maximum derived from Mioduszewski and Kogan 2009, fig. 4) would result in a maximum error of 6 mas. However the variance of such a large error through the experiment with changing hour angle would have smeared the flux of the source over a noticeably wider area. It is not unreasonable to suppose that the derived position of the source is accurate to at least 1 mas. This is still rather higher than the interferometric resolution of the array (6mas) and obviously any structure on scales of less than this would be unresolved.

#### 7.6.4 Fringe fitting on the A-component

Since the detection of the A-component is strong and can be seen even in the uncalibrated data, there is the possibility of fringe fitting on component A in order to correct for any non-linearity in the change of the delay error which occur during the target scans. This will give improved sensitivity, provided that the

source is strongly detected on at least one baseline to each antenna.

Since the position of A with respect to the external calibrator J1329+3154 was now known with VLBI precision, a new UV dataset was made by shifting to this position (to the nearest milliarcsecond) using the CALC-shifting algorithm. Essentially this means that the delay solutions from phase referencing the external calibrator have been applied to first order. This is a very useful convenience, as fringe-fitting on the target source will now still yield a map centred on the astrometric position of the source and the solutions provided by fringe fitting will not include any source position error and can be applied across the target field.

Fringe-fitting on this new dataset (using Ef as the reference antenna and with a solution interval equal to the 3 minute scan length) produced good solutions for almost all sub-bands, baselines and time intervals (an SNR of at least 9 was found, except for a single scan for which there was no Effelsberg data). This produced a significant improvement in the quality of the image produced as can be seen by comparing columns A1 and A2 of table 7.2: although the RMS noise for the dirty images is significantly higher fringe-fitting on A, the improvement in the dynamic range is over 50% for the cleaned image, suggesting that the added RMS noise is just reduced smearing in the sidelobes, which can then be subtracted more “cleanly” in the CLEAN.

The flux for each baseline for 1320+399 and J1329+3154 are compared in fig. 7.4. Although the SNR is clearly much lower for the Ma and Wz baselines, good fringe-fit solutions were found for all these data and in any case the weight of these data will be much lower.

### 7.6.5 Further self-calibration of 1320+299A


A fringe fit is equivalent to a single round of phase calibration on the target source before averaging in time and frequency. It is also possible to carry out further self-calibration (including amplitude calibration) once the data have been averaged.

The effects of self-calibration on this source were explored extensively in Difmap (Pearson et al. 1994) and it was decided that the most likely model was a point source. Some flux was detected one resolution unit away from the main peak, however it was usually seen symmetrically on both sides of the peak, and it was impossible to pin down to a single location. It appeared to be associated with small phase shifts that were seen only on the Matera baselines.


In the end similar results were obtained by returning to an automated pipeline. Various combinations of different self-cal schemes (phase, phase and amplitude etc.) were experimented with. Little difference was found between them however the best results were found with two iterations of phase-only self cal, followed by 6 iterations of amplitude and phase self-cal the final iteration using “L1” rather than a standard least-squares fit.

### Limitations of the self-calibration technique for this observation

Throughout the automated self-cal imaging process (apart from the final manual clean) the clean box was confined to a small area at the very centre of the image. This was to prevent negative cleaning around the main source: a recognised problem particularly for observations with small numbers of antennas.



(a) 1320+299A



(b) J1329+3154

Figure 7.4: Comparison of the flux for the external calibrator J1329+2154 and 1320+299 after self-calibration.




Figure 7.5: The correction applied to UV data from each baseline to correct for the primary beam. The amplitudes of each dataset were *divided* by the corrections shown here to increase the flux on the baselines which would have recorded a lower flux.

Another common with arrays very disparate in size is that the more sensitive baselines tend to over-dominate the final solution. Since various problems had been identified with the data for Wz and Ma, this was deemed to be an advantage. Although this reduces the angular frequency range even further, it permits a dynamic range of 160 (see table 7.2 col. A1). Flagging based on phase closure errors (three antennas) results in the loss of all for that time integration, and so must be done sparingly. Since the data set is small anyway, there is no real advantage is time-averaging the data before imaging. This also ensures that when problem data are flagged it affects the overall dataset as little as possible.

As mentioned above, for the small amount of flux which is not concentrated on the central source, there is no unique solution: it can be placed anywhere one resolution unit or two either north or south of the main peak. This may be true flux which is not constrained by the model, or it may be spurious flux caused by some error which cannot be fixed by self-calibration (a pointing error for example).

### 7.6.6 Correcting for the primary beam

Finally, the effect of correcting the primary beam for component A was tested. At the distance of A from the pointing centre the correction required is only 10% for the Effelsberg baselines and negligible for the others (see fig. 7.5).

Table 7.3 shows the effect of applying this minor correction to all baselines. There appears to be no effect except to increase the overall flux density. The Dynamic range also appears to be slightly higher for the dirty maps when the primary beam correction is applied. This is likely to be due simply to the flux being increased on the stronger baseline, in any case this does not affect the cleaned or self-calibrated maps. The resulting maps were also compared and

Map	No Primary beam correction		Primary beam correction	
	Flux	flux/rms	Flux	flux/rms
Uniform wt. dirty map	82.1 $\pm$ 39.6	2.07	86.5 $\pm$ 32.3	2.68
Natural wt. dirty map	77.8 $\pm$ 22.1	3.51	81.6 $\pm$ 20.2	4.05
Uniform wt. clean map	82.6 $\pm$ 6.2	13.3	86.4 $\pm$ 6.5	13.2
Self-Cal 1 (phase only)	78.1 $\pm$ 4.7	16.5	82.0 $\pm$ 5.0	16.4
Self-Cal 2 (phase only)	78.0 $\pm$ 4.7	16.8	82.0 $\pm$ 4.9	16.8
Self-Cal 3 (amp&phase)	79.6 $\pm$ 0.81	81.1	83.5 $\pm$ 1.0	80.1
Self-Cal 4 (amp&phase)	80.4 $\pm$ 0.63	128	84.4 $\pm$ 0.67	126
Self-Cal 5 (amp&phase)	80.9 $\pm$ 0.54	152	84.8 $\pm$ 0.56	149
Self-Cal 6 (amp&phase)	81.3 $\pm$ 0.53	137	85.1 $\pm$ 0.57	147

Table 7.3: Comparison of the self-calibration convergence both with and without baseline-by-baseline correction applied to the input UV data. Peak flux, and RMS are shown along with the former divided by the latter (dynamic range).

there was no significant difference in morphology. Had A been located closer to the edge of the primary beam, it is possible that this correction would have speeded convergence of the amplitude self-calibration (since what is being applied is, in effect, an amplitude correction for each antenna). Figure 7.6 shows the final image: a point-like source with small secondary peaks (1.5% of the peak flux density) a short distance from the source.

### Manual Cleaning

In addition to the automated self-calibration, the final resulting model was then used as the initial model for a final round of imaging. This was then used to calibrate the original data (amplitude and phase) and the resulting dataset was then carefully cleaned by hand. The clean model produced in this way was then used to produce a map, using a very large restoring beam. This collects all of the integrated flux. The result is shown in fig. 7.7. The flux is  $103 \text{ mJy beam}^{-1}$ . The secondary peak is now much brighter with respect to the main source, and appears only on one side, however with such limited data it is impossible to be sure that it represents a real feature.

## 7.7 Applying A solutions to other datasets

Now that detailed calibration of a single one of these fields has been conducted, the resulting calibration can be applied to each of the others in an automated manner.

A particularly useful mode of operating the automated pipeline here, is to automatically apply the basic calibration, then apply more nuanced calibration manually in AIPS, then restart the pipeline on the modified data and automatically image.

### 7.7.1 Impact of the A component on the B and C datasets

To assess the impact that the A component on the B, it is possible to fringe fit on the target data in B. This self-calibrates on the nearest strong source moving it to the centre of the field. Imaging this dataset (table 7.2 col. B3) shows that it is possible to make a passable map of A (2s integration time and 32 channels per sub-band results in smearing losses of only 8.2% and 0.7% respectively). This




Figure 7.6: Result of 6 rounds of self-calibration on component A (table 7.3)

means that it is likely that at least some of the flux in the B field is the side-lobes of A.

### 7.7.2 Peeling 1320+299A

The AIPS task UVSUB allows the subtraction of a model from a UV dataset.

This was done by running the calibration section of the pipeline (see sect. 7.5.1). Following this the calibration tables were applied manually to the target source in AIPS but no spectral averaging or time averaging was applied. Next UVSUB was used to subtract a model of A. The clean components from the initial phase referenced dataset were used (see table 7.2 col. A1, row 3). Following this subtraction, the spectral channels in each sub-band were averaged together and the imaging section of the pipeline was run on the data (see sect. 7.5.1).

This procedure succeeded in reducing both the peak flux and the RMS noise in the naturally weighted image by about 20%. Moreover, the peak in the image




Figure 7.7: The UV data with 6 rounds of self calibration applied (table 7.3) were Fourier inverted and the resulting dataset was then cleaned carefully by hand and the image shown was generated by restoring the resulting clean model with a very large circular restoring beam ( $7.5 \times 7.5$  mas).

is now further from the centre of the image. This improves our confidence that B has not been detected. It made less difference to the C field (the peak flux was reduced by a small amount and the RMS flux remained unchanged). This is to be expected since the C field is further away.

### 7.7.3 1320+299A as an in-beam calibrator

Since a significant improvement in the quality of the image of the A field was obtained through fringe-fitting on this source, it is possible that a better dataset centred on B may be obtainable by applying this phase calibration, rather than the phase solutions obtained from the calibrator scans.

The B dataset was loaded into AIPS and all calibration tables obtained in the calibration section of the pipeline for dataset A4 (see table 7.2) were copied across and applied to the data. This included the fringe fit on the target source. These calibration tables were applied without averaging and the clean components



from the fringe-fitted dataset (table 7.2 col. A2, row 3) were used to subtract the A source from the data. Finally the channels of each sub-band were averaged together and the imaging section of the pipeline was run

Perhaps surprisingly, given the effect that this had on the detected source, using this calibration had very little effect on the B-field. Although there will be slightly different residual errors at the locations of A and B, A is  $300\times$  closer to B than J1329+3154. However, if we assume that the B field is entirely empty and that we are at the thermal noise limit, there is no reason why improved phase calibration should make any difference.

It is likely, however, that this more nuanced calibration will have a noticeable effect for fields closer to A than B. In order to improve the calibration procedure further then, we must work fields other than Fields B and C.

#### 7.7.4 Conclusions

The difference between the theoretical low noise limit and the maximum peak flux has now reached a level where it is clear that no detection has been made of B or C. It is disappointing that it was not possible to detect two of the three sources, however one of the sources was detected, and this allowed the exploration of wide-field techniques. Moreover, once A has been calibrated and imaged, the solutions for this single dataset can be applied automatically to many other datasets with minimal effort.

## 7.8 Tiling the primary beam

The next step is to devise a strategy to image the entire field by generating multiple UV datasets and imaging each using the automated pipeline.

### 7.8.1 Area to map

Figure 7.8 shows a measure of the sensitivity of each baseline with respect to every other; specifically:

$$\sqrt{\frac{A_i(\vec{r}) \times A_j(\vec{r})}{SEFD_i \times SEFD_j}} \quad (7.1)$$

which is inversely proportional to  $\Delta I_m$  (see equation 5.1). This figure shows that at a radius of approximately  $90''$  from the pointing centre the Ef baselines have a sensitivity comparable to that of the Ma-Wz baseline at the pointing centre. Given that 1320+299A was detected on all baselines, this suggests that it may be possible to detect significant flux from brighter sources out to at least this radius even on Ef baselines. Assuming a pixel size of  $1.4\text{ mas}$ , mapping a circle of this radius would require an image containing 13 gigapixels!

### 7.8.2 Generating the UV data

This was achieved by making a single correlation with the phase centre close to the coordinates of B (the pointing centre). An integration time of  $0.2\text{ s}$  and 512 channels per sub-band was used, resulting in a UV dataset of approximately 5 GB. This resolution keeps time-average and bandwidth smearing to an extremely low level across the entire field. The software described in sect. A.2 was then used to generate multiple datasets.

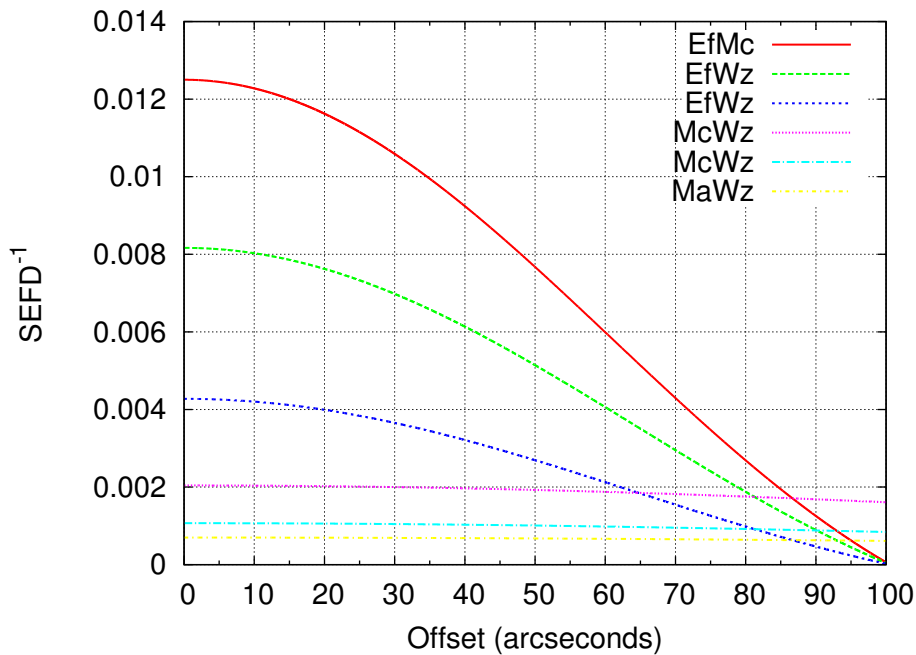


Figure 7.8: A measure of the sensitivity of each baseline taking both the primary beam and the SEFD of each antenna into account (see equation 7.1). The primary beam is modelled using a sinc<sup>2</sup> function as this *may* model the primary beam beyond the HPBW better. In any case it is more conservative.

Shifting in turn to phase centres placed on a grid, and imaging each one allows the imaging of this entire field with each sub-image being produced from a small dataset. It was then necessary to decide the spacing of the phase centres. Spacing the phase centres 5'' apart in declination and 0.5<sup>s</sup> (equivalent to 6.5'') apart in right ascension gives a convenient grid. With images of 4096 × 4096 pixels with each pixel having dimensions of 0.85 × 0.7 mas, the entire area is covered with slight overlap between the images. It can also be seen from the UV coverage diagram (see fig. 7.3) that the resolution is much higher in the north-south direction and so this also gives a reasonable fit to the synthesised beam. Averaging each sub-band to a single channel incurs bandwidth smearing of 4% on this image. Losses due to non-coplanar effects are around ~1%.

The computer that was used to carry out the imaging is capable of generating 8192 × 8192 or even 16384 × 16384 images, so it would have been possible to place the phase centres further apart. However, doubling the separation of the phase centres would have incurred bandwidth smearing of 15% (in addition to non-coplanar effects contributing an error of 5% if not corrected). Avoiding bandwidth smearing would have meant generating images from UV datasets containing more than one spectral channel per sub-band. Table 7.4 shows the results of a brief test of imaging speed with the computer which was used to perform the imaging. It shows little difference in performance up to an image size of 4096 × 4096 pixels. However beyond this the requirement for several spectral channels begins to degrade performance significantly. Smaller sizes can be imaged fast but incur an overhead in the number of UV shifts required.

The vertices for a grid centred on RA 13<sup>h</sup>23<sup>m</sup>2<sup>s</sup>.75 Decl. +29°41'32''.5 with the dimensions described were then calculated for each vertex with an offset of less

Image size	8 MHz Channels		8 MHz Channels		8 MHz Channels		Speed per pixel
	CPU time	Wall time	CPU time	Wall time	CPU time	Wall time	
1024	0.3	0	0.5	0	0.9	1	0.286
2048	0.9	1	1.8	2	3.5	3	0.215
4096	3.7	5	7.2	11	14.3	24	0.298
8192	15.5	30	31.6	63	61.7	142	0.939
16384	69.2	231	133.7	483	260.2	959	1.799

Table 7.4: Time taken to make a dirty map of various sizes with AIPS IMAGR (DO3DIMAGE 0) for datasets consisting of  $4 \times 8$  MHz sub-bands with 1, 2 and 4 spectral channels respectively. The final column shows the time taken in microseconds per pixel for the dataset which is required in order to keep bandwidth smearing below 5%: Namely the first for the 3 smallest images, the second for 8192 pixels per dimension and the 3rd for 16384 pixels per dimension.

No subtraction of A					Subtraction of A				
1.46	2.33	4.72	1.47	1.06	1.37	2.07	4.53	1.19	1.05
1.30	2.24	55.1	6.10	1.18	1.21	1.82	28.9	2.41	1.23
1.11	1.42	58.9	5.30	1.20	1.11	1.35	26.2	3.07	1.28
1.40	1.19	5.12	4.30	1.63	1.45	1.22	2.43	2.62	1.34
1.46	1.40	2.20	1.87	1.28	1.37	1.34	1.74	1.98	1.43

Table 7.5: Peak flux in 25 images surrounding A in mJy (A is located in the central field, close to the border of the two fields which have the highest peak flux) with and without the subtraction of A.

than  $90''$  from the pointing centre. This resulted in a grid covering a circular area containing 782 vertices. A delay model was then calculated for each of these vertices and the large dataset was then shifted to each of the phase centres on the grid in turn. Each was averaged down to a time resolution of 2s and 16 channels per sub band before writing to disk. Each iteration produced a UV dataset in FITS-IDI file format approximately 20MB in size. It took approximately 8 hours to generate the delay models and 24 hours to generate datasets using a single CPU core. These could have been done concurrently.

### 7.8.3 Subtracting A

Before the final imaging of the entire primary beam, the results of subtracting A from neighbouring datasets were assessed. Table 7.5 shows the peak for the 25 fields surrounding A, with and without A subtracted. Table 7.6 shows the RMS pixel values for the 25 fields surrounding A, with and without A subtracted. Certain fields seem to contain sidelobes of A. Although only half of the flux from the A-field is subtracted using this technique, the peaks in neighbouring images

No subtraction of A					Subtraction of A				
7.01	10.5	7.93	5.86	6.39	6.52	8.63	6.88	5.40	6.21
6.20	8.13	30.3	8.05	5.95	6.04	7.41	17.4	6.84	5.76
5.87	6.40	18.4	10.3	5.57	5.72	6.06	25.0	11.9	6.20
6.34	5.87	43.1	12.3	9.85	6.26	6.06	9.36	8.24	10.2
5.61	6.34	6.54	7.27	6.18	5.58	6.32	6.01	11.9	6.57

Table 7.6: RMS flux in 25 images surrounding A in  $1 \times 10^{-4}$  Jy beam $^{-1}$  (A is located in the central field, close to the border of the two fields which have the highest peak flux) with and without the subtraction of A.




Figure 7.9: Each pixel represents a field  $6.5'' \times 5''$ . The flux in  $\text{Jy beam}^{-1}$  is shown for the brightest pixel in each image and the RMS of all of the pixels both with and without primary beam correction applied. In spite of the fact that a simple model of A has been subtracted, this source is clearly visible. Other flux is very likely to be distant sidelobes of A.

are reduced for almost all fields which show spurious high peaks. Significantly, this does not appear to degrade the RMS values of fields which appeared to be clean even before A is subtracted<sup>1</sup>. This, along with the results of this procedure on the B and C fields, suggests that this procedure should be carried out across the wide field.

#### 7.8.4 Results

Finally, the procedure described above was applied to all 782 sub-fields. The results are shown in fig. 7.9. In spite of the subtraction of A, it can clearly be seen as the bright emission slightly to the north-east of the centre of the field. It is also clear that some sidelobes remain. These can be seen most clearly in the RMS flux

<sup>1</sup>This was *not* the case when experimentation was done using the self-calibration procedure described in sect. 2.10.

without PB correction as two bright blobs located symmetrically either side of A (to the north-west and south-east).

Given the symmetry which these two sidelobes have with A, they are almost certainly sidelobes. At twice the spacing from A there is a very bright feature on the edge of the primary beam. The peak fluxes are almost uniform across many fields, suggesting that this is not a true feature.

## **7.9 Conclusion**

It is, perhaps, not surprising that no further sources were found. However this provides a proof-of-concept test of imaging very large fields using the UV shifting technique. No problems were encountered and there is no reason to believe that it could not be applied to much larger fields. The technique is very parallelisable and scalable and the procedure used here could be carried out extremely quickly, even for much larger datasets, if several machines were used for the analysis.



# Chapter 8

## A single-baseline VLBI survey of 22 GHz calibrators

### 8.1 Motivation

The 22 GHz feed has been commissioned for the SRT and is currently being tested on the Mc dish. It consists of 6 feeds arranged in a circular pattern around a central feed (Verma et al. 2008). Currently only the central feed is used for VLBI.

A list of potential calibrators has been compiled for single dish work and will be observed in single-dish configuration. Complementary VLBI observations will give a clear measure of how much flux remains unresolved even on milliarcsecond scales. This will also allow us to measure the performance of the feed for VLBI.

### 8.2 Baseline Parameters

The twin of the Mc dish, the Nt antenna, was used to provide a baseline to Mc. Both antennas are 32m in diameter and the baseline length is 894km. This leads to a resolution at 22.2GHz of approximately 3 mas with a field of view of approximately 100". The 22 GHz feed on Nt was only able to record LCP at the time of the observation. The nominal SEFDs listed for Mc and Nt are 700 Jy for Mc and 800 Jy for Nt, in spite of the fact that Nt has an active surface which should improve the efficiency of the antenna at this frequency. This underlines the improved efficiency of the new feed compared with the older one. Eight contiguous 8 MHz upper sidebands were recorded at frequencies from 22202.49 to 22258.49 MHz inclusive, for a total bandwidth of 64MHz. The data were recorded using eight base-band converters of the MarkIV acquisition system and a MK5A disk recorder with a sample rate of 16 megasamples/sec and 1-bit sampling for each baseband channel.

With a 5 minute fringe-fit interval, this gives an estimated RMS noise level for this baseline of 20 mJy for each single sub-band. This should easily be sufficient to find any compact structure for the bright sources we are looking at.

### 8.3 Observation Strategy

The calibrators were observed in a 24-hour observing session on the Mc-Nt Baseline. The observing session was scheduled to start on 11th February 2009 at 14:00:00 UT and finish 24 hours later.

The schedule was then prepared using the following system:

1. For each 2-hour interval the sources that would be visible were selected

2. Each of these sources was then cycled through in order of declination with a 5-minute scan on each
3. This cycle was repeated until the 2-hour interval was filled

This led to each source being scheduled for observation between 4 and 9 times (depending on declination). This is a similar approach to the scheduling of Geodetic observations, however this was a rather less sophisticated system and simply ensured that all sources would be observed at all available hour angles.

## 8.4 Sources

Table 8.1 lists the sources that were chosen for our survey (Rashmi Verma, IRA, Priv. Comm.). This is a subset of the list of sources compiled as potential calibrators of the feed, including only those which were believed to have compact structure.

## 8.5 The Observation

The MK5A disk packs recorded at Nt were sent to the Medicina Radio Observatory, read back on a MK5A unit and loaded onto hard disks. These were then transferred to IRA for correlation with the DiFX software correlator.

A number of problems were reported for the Medicina antenna. Firstly the observation was curtailed at 10:00:00 UT which meant that some sources around R.A. 19 were only observed at 2 or 3 hour angles. Secondly the Maser jumped by approximately  $1\ \mu\text{s}$  three times during the correlation due to a problem in the ionic pump. Thirdly there was a problem with video converters not having an LO connection for the first 1.5 hours (although this was not seen in the data). During correlation there were sync problems for the Mc datastream suggesting problems with the baseband data recording.

## 8.6 Correcting for Clock shifting

For the interval between each of the clock jumps, it was necessary to make use of the comparisons of the maser to GPS time which were carried out during the observation. A linear least-squares fit was made to the intervals between each of the clock jumps using the `log2clock.py` utility described in appendix B. In each case the rate (slope of the line) was small compared to the scatter of the points and was discarded.

The clock parameters obtained in this way are sufficient to find fringes (in fact, the value of any of the individual readings would have been sufficient, even at this frequency: the standard deviation of the readings was only 14 ns). However the clock parameters determined in this way were not optimal. A test correlation was made on a scan of a strong source near the centre of the time range. Once correlated, these data were loaded into AIPS and a fringe-fit was made not solving for delay rate and setting a delay window of  $2\ \mu\text{s}$ . The data were then correlated in four separate runs corresponding to the four gaps between maser jumps. In the worst case a delay was found of approximately  $0.5\ \mu\text{s}$ . The resulting fringes are shown uncorrected (fig. 8.1) and after correction (fig. 8.2). The delay for Nt



Table 8.1: List of 33 potential 22 GHz calibrators

Source Name	RA			Decl.			Flux (Jy)
	h	m	s	deg	'	"	
3C286	13	31	08.3	+30	30	33	2.32(6)
1803+784	18	00	45.7	+78	28	04	1.32(16)
0153+744	01	57	34.9	+74	42	43	0.32(02)
0212+735	02	17	30.8	+73	49	33	1.56(03)
4C+50.11	03	59	29.7	+50	57	50	2.62(09)
0945+408	09	48	55.3	+40	39	45	1.54(07)
1611+343	16	13	41.0	+34	12	48	1.81(40)
1633+382	16	35	15.5	+38	08	04	1.97(27)
3C345	16	42	58.8	+39	48	37	11.66(39)
0528+134	05	30	56.4	+13	31	55	2.69(22)
0836+710	08	41	24.3	+70	53	42	1.17(05)
1928+738	19	27	48.5	+73	58	02	2.25(14)
0716+714	07	21	53.4	+71	20	36	0.90(03)
3C279	12	56	11.1	-05	47	22	11.54(44)
3C048	01	37	41.3	+33	09	35	1.19(03)
J0650+6001	06	50	31.2	+60	01	45	0.40(04)
3C84	03	19	48.1	+41	30	42	24.71(76)
3C454.3	22	53	57.7	+16	08	54	9.04(16)
3C120	04	33	11.1	+05	21	16	1.70(04)
2134+004	21	36	38.6	+00	41	54	4.72(03)
1642+690	16	42	07.8	+68	56	40	0.88(25)
3C111	04	18	21.3	+38	01	36	2.17(11)
0754+100	07	57	06.6	+09	56	35	1.08(11)
0953+254	09	56	49.9	+25	15	16	1.45(15)
1219+285	12	21	31.7	+28	13	59	0.40(05)
1418+546	14	19	46.6	+54	23	15	1.43(31)
3C274	12	30	49.4	+12	23	28	0.18(0)
3C138	05	21	09.9	+16	38	22	1.310(5)
3C123	04	37	04.4	+29	40	14	3.12(09)
3C147	05	42	36.1	+49	51	07	1.70(02)
3C249.1	11	04	13.7	+76	58	58	0.82(01)
3C309.1	14	59	07.6	+71	40	20	0.00(00)
3C273	12	29	06.7	+02	03	09	1.70(02)

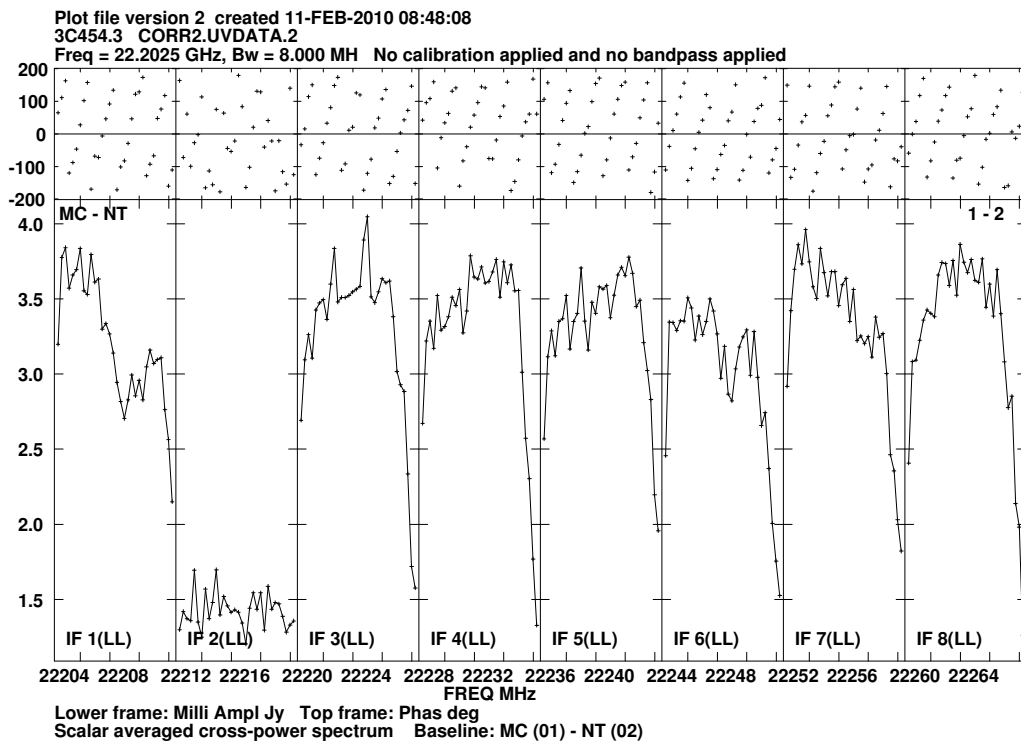


Figure 8.1: Scan correlated using the GPS values obtained from the log file. Even at this moderately high frequency the error on a single GPS reading is not sufficient to lose the fringe, however the residual delay here is clearly unsatisfactory. 4.25 turns of phase across an 8 MHz sub-band implies a delay of approximately  $4.25/8 \approx 0.53 \mu\text{s}$  and indeed this was consistent with the delay calculated by the fringe fit.

was adjusted to take account of this offset and the correlation was run for the whole timerange.

## 8.7 Checking correlated data integrity

Considering these problems it was important to assess the impact of the maser jumps and the sync problems on the data.

### 8.7.1 Maser Jumps

Examination of the GPS logs did not reveal any increasing inaccuracies in the vicinity of the jumps and examination of a scan in which a jump occurred showed good data right up until the jump occurred.

### 8.7.2 Impact of re-syncing of the data

A single scan was chosen and the times of the re-syncing were noted from the correlation log file. A few of the re-syncs were associated with noise in the phase of up to 20 degrees for a few seconds, however the impact on the overall data was judged to be small.




Figure 8.2: The same scan as fig. 8.1 however this time a delay of  $0.57 \mu\text{s}$  has been applied to Medicina in the clock parameters of the correlation. A fringe fit on *these* data results in an SNR of  $\sim 100$  for each sub-band.

## 8.8 Results

Figure 8.2 shows the fringe fit SNR and the maximum and minimum fluxes for all sources. 21 sources out of 33 were detected. Calculating the gain of the two antennas will not be simple since we only have a single baseline. VLBA plots show an unresolved flux of approximately 6 Jy at this baseline length for 3C454.3. Fringe fitting yields an SNR of approximately 100 for each sub-band with a 1 minute fringe fit interval (see fig. 8.2). This is comparable with that predicted by the sensitivity of the baseline.

### 8.8.1 Effect of the troposphere

As would be expected the troposphere appears to have a great effect on the delay errors in the correlator model (see fig. 8.3 c.f. Mioduszewski and Kogan (2009) fig. 4). The opacity of the troposphere will also undoubtedly have an effect on the flux detected, changing both with elevation and throughout the experiment. Examining the system temperatures and flux of individual sources, various effects were seen. For weaker sources it was clear that the system temperature was higher at lower elevations. However many of the sources observed here are extremely bright, and for those sources the opposite effect is seen: the system *increases* at higher elevation. This is probably due to the flux of the source increasing the *measured* system temperature.

Source Name	h	RA	s	deg	Decl.	'	"	$a$ priori Flux (Jy)	VLBA flux	detections	of	Flux (max)	Flux (min)	Fringe fit SNR
3C286	13	31	08.3	+30	30	30	33	2.32(6)		0	5			0
1803+784	18	00	45.7	+78	28	04	04	1.32(16)		4	4	1.4	0.8	30
0153+744	01	57	34.9	+74	42	42	43	0.32(02)		0	4			0
0212+735	02	17	30.8	+73	49	33	33	1.56(03)		7	7	4.0	1.0	39
4C+50.11	03	59	29.7	+50	57	50	50	2.62(09)		7	7	7.0	4.0	100
0945+408	09	48	55.3	+40	39	45	45	1.54(07)		2	5	0.9		22
1611+343	16	13	41.0	+34	12	48	48	1.81(40)		5	5	1.1	1.1	24
1633+382	16	35	15.5	+38	08	04	04	1.97(27)		5	5	1.4	0.2	34
3C345	16	42	58.8	+39	48	37	37	11.66(39)	3.4	5	5	3.4	2.7	73
0528+134	05	30	56.4	+13	31	55	55	2.69(22)	3.5	6	6	1.5	0.5	29
0836+710	08	41	24.3	+70	53	42	42	1.17(05)		5	5	1.1	0.2	26
1928+738	19	27	48.5	+73	58	02	02	2.25(14)		2	2	1.2	1.2	22
0716+714	07	21	53.4	+71	20	36	36	0.90(03)		5	5	2.3	0.5	37
3C279	12	56	11.1	-05	47	22	22	11.54(44)		4	4	8.4	2.2	88
3C48	01	37	41.3	+33	09	35	35	1.19(03)		0	6			0
J0650+6001	06	50	31.2	+60	01	45	45	0.40(04)		0	4			0
3C84	03	19	48.1	+41	30	42	42	24.71(76)	5	6	6	5.8	2.4	88
3C454.3	22	53	57.7	+16	08	54	54	9.04(16)	5.2	4	4	4.2	1.6	57
3C120	04	33	11.1	+05	21	16	16	1.70(04)		3	5	1.5		26
2134+004	21	36	38.6	+00	41	54	54	4.72(03)		2	2	2.0	2.9	30
1642+690	16	42	07.8	+68	56	40	40	0.88(25)		5	5	2.4	3.0	52
3C111	04	18	21.3	+38	01	36	36	2.17(11)		4	4	3.8	2.5	60
0754+100	07	57	06.6	+09	56	35	35	1.08(11)		0	4			0
0953+254	09	56	49.9	+25	15	16	16	1.45(15)		0	6			0
1219+285	12	21	31.7	+28	13	59	59	0.40(05)		1	6	0.8		0
1418+546	14	19	46.6	+54	23	15	15	1.43(31)		0	6			0
3C274	12	30	49.4	+12	23	28	28	0.18(0)		2	4	1.3		21
3C138	05	21	09.9	+16	38	22	22	1.310(5)		0	6			0
3C123	04	37	04.4	+29	40	14	14	3.12(09)		0	6			0
3C147	05	42	36.1	+49	51	07	07	1.70(02)		0	6			0
3C249.1	11	04	13.7	+76	58	58	58	0.82(01)		0	5			0
3C309.1	14	59	07.6	+71	40	20	20	0.00(00)		0	6			0
3C273	12	29	06.7	+02	03	09	09	1.70(02)		4	4	9.5	3.5	78

Table 8.2: Sources observed; coordinates; flux from list compiled by Rashmi Verma; Flux at this frequency and baseline length from VLBA fringe-finder survey. Number of scans in which the calibrator was detected; number of scans which were observed; lowest flux observed; highest flux observed; fringe fit SNR for a 1 minute interval. N.b. 1218+285 was only detected in a couple of sub-bands

## 8.9 Conclusions

In the end, however, there is a very clear line between the sources which showed a clear detection, and those that were not detected. The measured fluxes are broadly in line with the values obtained from the VLBA fringe finder survey<sup>1</sup>.

---

<sup>1</sup>Available here: <http://www.aoc.nrao.edu/~analysts/vlba/ffs.html>

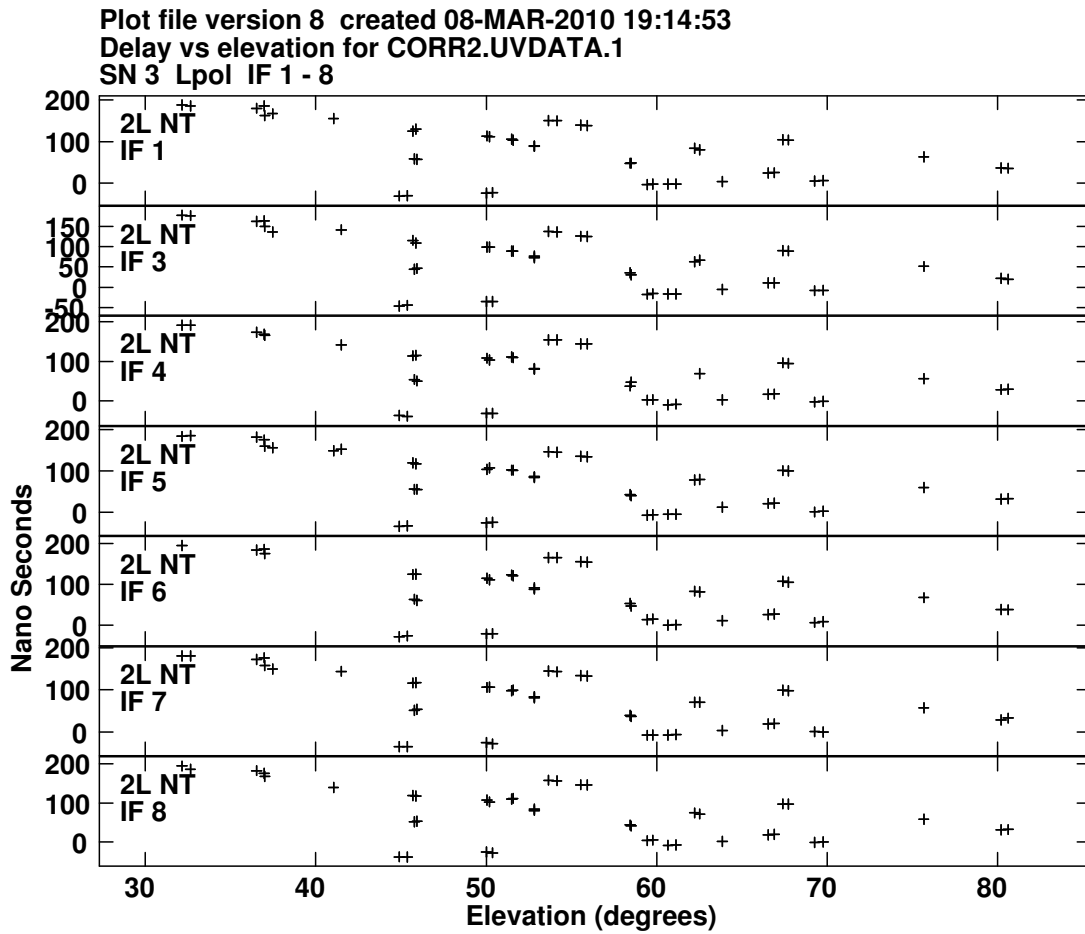


Figure 8.3: Delay solutions for all scans with good fringe fit solutions in the first 8 hours of the observation. A clear dependence on elevation can be seen

# Chapter 9

## Conclusions

This thesis has explored the feasibility of two concepts: VLBI science with a limited number of antennas (specifically those in Italy) and wide-field VLBI. In this final chapter I will summarise the conclusions and briefly discuss possible future work suggested by the conclusions of this thesis.

### 9.1 Science with an Italian VLBI network

Possible science with an Italian VLBI network was discussed in sect. 3.4. Here it is sufficient to add that the results presented in chaps. 7 and 8 underline the potential of a small VLBI network for doing science and for technical studies.

### 9.2 Wide-field VLBI

Prior results (from Garrett et al. (2005b) and Lenc et al. (2008) for example) show the results possible even with the relatively low bandwidths available up until now, and draconian limits on field of view imposed by hardware correlators. With the flexibility of software correlators which are already available and the increasing sensitivity which will become available in the very near future, the way has opened to imaging far larger fields with a greater density of sources.

#### 9.2.1 Tiling the entire primary beam

##### Motivation

In most cases, a VLBI field of interest will have been observed earlier with a connected element array. Since VLBI is bandwidth-limited, it tends to have lower sensitivity compared with connected element arrays, and so in most cases any source that can be detected by VLBI has already been pinpointed. It may seem, therefore, that imaging the entire primary beam would very rarely be worth the trouble. However, there are a few exceptions.

Firstly there is the fact that the Universe does not remain static. If we want to detect new supernova remnants, for example, we may be searching for something that wasn't detectable last time the field was checked. There is also the possibility of searching for transients. Scans of calibrators in particular are a huge potential resource that could be mined if appropriate techniques were developed — each dataset having an in-beam calibration source at the centre of the field!

Secondly, radio telescopes are constantly evolving. Recent developments in VLBI (for example the VLBA upgrade) mean that today's largest VLBI arrays can compete with yesterday's connected-element arrays for instantaneous sensitivity (see table 5.2). Of course e-MERLIN and the EVLA will soon overtake, but it will

be some time before their legacy survey programs have mapped the majority of VLBI targets.

Finally, not every patch of sky will have been observed at every frequency by connected element arrays at sufficient depth, and although most VLBI observations are informed by other radio observations at similar frequencies, there may still be cases where *a priori* knowledge of the field is scarce enough that imaging the largest possible field is worthwhile.

### Techniques

As shown in table 5.1, the 13 gigapixel image presented in chap. 7 was approximately two orders of magnitude smaller than that which would be required to image the entire VLBA primary beam.

Recently a feature has been added to DiFX which allows the correlation of multiple phase centres simultaneously (Deller et al. in prep). A sensible approach would be to correlate multiple phase centres simultaneously to produce  $\sim 100$  datasets slightly smaller than the dataset reduced in chap. 7. Such datasets could then be calibrated and imaged independently, a process which could be carried out on a 10 node computer cluster in a matter of hours.

### Imaging the entire primary beam for every VLBI observation

The ultimate goal would be to image the entire primary beam for every experiment at low enough frequency that the primary beam is sufficiently large (and the population of detectable compact sources is sufficiently high) that several sources can be expected to be found within the primary beam. For gigahertz frequencies, this would allow in-beam calibration for any arbitrary patch of sky (Garrett et al. 2004). This would likely be useful for the primary goal of the observation, but any calibrators found in this way could also be used for subsequent studies and the sources discovered could be interesting in their own right. Also such a system could lead to the detection of transient sources.

In addition to general wide-field techniques required to undertake such work, it would also be necessary to develop techniques to archive the huge amounts of data produced. It is likely that there would be some scope for compressing the data; perhaps by dynamically averaging data in frequency and time depending on how fast the phase varies.

The shifting and averaging technique could also be useful here. A large field could be stored as a single large dataset, and if an archive user is interested in exploring a particular field, a small UV dataset, shifted to the requested coordinates and averaged, could be generated automatically.

### 9.2.2 Surveys of compact objects

Greater VLBI bandwidths lead to greater sensitivity or, if multiple pointings are being made, greater surveying speed. That is, a larger area of sky can be covered to the same sensitivity in the same amount of time. Garrett et al. (2004) fig. 5 shows what the mJy sky would look like to a wide-field VLBI array: there are approximately 100 compact sources brighter than 1 mJy per square degree. The VLBA could reach a sensitivity sufficient to detect a mJy source at the  $10\sigma$  level in approximately 5 minutes at L-band with a bandwidth of  $4\text{ Gbit s}^{-1}$ . Nyquist-sampling the sky it could cover  $\sim$ one degree of sky per hour at this depth. If



finding bright calibrators were the only goal, the sky could be surveyed to a much shallower depth much more quickly.

The trend towards smaller dishes (with larger fields of view) equipped with multi-feed receivers or focal plane arrays will only increase the survey speed (at the cost of an incredible torrent of data).

### 9.2.3 Challenges

Making full use of these opportunities requires the development of new techniques to handle the huge amounts of data required. Software development for VLBI has stalled somewhat in recent years and there is large scope for development of new calibration and imaging techniques.

## 9.3 The Delay across the primary beam

In generating the widefield image presented in chap. 7, one of the most time-consuming steps was generating the delay models for all of the UV shifts required. It would be relatively straightforward to generate these delay models in parallel. However, perhaps there is a better technique: at the end of chap. 6 the possibility was raised of characterising the correlator delay function over the entire widefield. By “sampling” the delay function across the widefield, it would be possible to generate a delay function that is not only a function of time, but of  $l$  and  $m$ .

This model would then be available throughout post-processing. As noted in sects. 5.5.3 and 5.5.4, algorithms are being developed which would allow the efficient imaging of significantly larger fields than are possible with current software. How these algorithms will fare once they are required to work on datasets as large as those analysed in this thesis is unclear, however, their accuracy in generating images spanning more than  $1 \times 10^4$  resolution units would certainly benefit from correcting the delay across the widefield to more than first order.

Once the post-processing software has access to a model of the delay across the widefield, there is no reason why it could not refine it. For connected element arrays at very low frequencies it is already necessary to correct for the changing delay across the image. For the VLA Low-frequency sky survey (Cohen et al. 2007) the time-dependent delay function across the primary beam was calculated by using multiple calibrators which were spread throughout the primary beam.

Gradients in the troposphere can also be measured for VLBI using multiple calibrators (Mioduszewski and Kogan 2009). Such techniques could, perhaps, be used for VLBI, and might be particularly useful for low-frequency work.



# Appendix A

## Software

In this appendix I describe various software projects that were undertaken to allow the correlation of the observations described in chaps. 7 and 8 and the tests described in chap. 4. Other software that has been developed in pursuance of this thesis is described in appendix C.

### A.1 Correcting LO offsets in DiFX

This section describes the modifications made to DiFX to correct for the LO offsets observed in the observations presented in chap. 7.

This could, perhaps, be fixed in post-processing: If the data are correlated with very high time resolution (0.1s would be a minimum to keep averaging losses to an acceptable level) it may be possible to fringe fit the data to remove the turning phase. However it is unlikely that any astrometric measurement would be possible: even though both the target and the phase calibrator would be corrected in the same way, any offset between the two would be flattened by the fringe-fitting algorithm. In addition it is unclear if this would work, since the error caused by the LO offset does not resemble a delay: it is dispersive causing an equal phase rate for all frequencies.

After experimenting with setting an extremely high value for the rate of the clock (i.e. the first derivative of the delay term due to the offset between the frequency standard at Ef and GPS time) we decided that it was necessary to correct for the clock offset within the correlator.

The full internals of DiFX are described elsewhere (Deller 2009), however one of the possible fringe rotation schemes in DiFX is very similar to the one described in sect. 1.4.3; that is, *after* the Fourier transform, the phases are rotated to account for the change in delay.

The correction is very simple:

$$\phi = 2\pi(\delta\nu \cdot t \bmod 1) \tag{A.1}$$

where  $\phi$  is the phase shift applied to every frequency,  $t$  is the time in seconds since the start of the observation and  $\delta\nu$  is the LO offset in Hz<sup>1</sup>. This implicitly applies a constant phase-shift throughout the experiment due to the fact that correction is tied to the start time of the observation. Fortunately absolute phase offsets rarely matter in radio astronomy. Since the correction is done at the level of an individual FFT block it will work even for reasonably large LO offsets. From the Nyquist frequency (sect. 1.2.4) and the FFT size in samples (sect. 1.2.4), we can calculate that provided

$$\frac{\Delta\nu}{n_{channels}} \leq 10 \cdot \delta\nu, \tag{A.2}$$

---

<sup>1</sup>a positive value for a positive offset i.e. observing 1.000 000 001 GHz instead of 1 GHz

amplitude loss due to averaging will be kept to an acceptable level. Thus even with a bandwidth of 16 MHz per sub-band, provided that the number of channels is kept to just 16 it would still be possible to correct for an LO offset of  $\sim 100$  kHz.

## A.2 Post-correlation UV shifting

This section describes the implementation of the algorithm described in sect. 5.8, tested in chap. 6, and used for wide-field imaging in chap. 7.

### A.2.1 Limitations of current software

Current post-processing techniques for shifting of datasets were found to be inadequate for reasons outlined in sect. 5.7.3. To summarise, the following two features are desirable:

1. a minimum number of reads and writes of large datasets
2. preservation of astrometric accuracy

### A.2.2 Where to Implement

To avoid increasing the number of read/writes, the DiFX data path was examined for a point at which the UV shift could be added without incurring another write to disk. After correlation the data are in a format specific to DiFX. They are then converted to FITS-IDI by a program called `difx2fits`. This program also reads the delay model files and so it is a fairly simple task to add a shift at this point.

### A.2.3 Implementation

First `difx2fits` was modified to allow it to take two correlator models as input. The points of implementation which pertain to the accuracy of the algorithm (i.e. the calculation of the delay for each datum) have already been described in sect. 6.2. After the delay difference value  $\delta\tau$  value has been calculated it remains to:

1. calculate the frequency of each spectral channel
2. calculate the phase shift to apply to each spectral channel (from the delay difference  $\delta\tau$  and frequency  $\nu$  for each spectral channel)
3. apply the calculated phase shift

Both the frequencies and the delays were calculated in double precision. In order to calculate the phase shift the modulo of their product must be taken<sup>2</sup> resulting in  $\phi_{turns}$ : a number between 0 and 1 representing the phase shift to be added in number of turns. Once this is done there is no further need for double precision since the visibility data are stored in single precision.  $\phi_{turns}$  must now be applied to the complex numbers representing the visibilities. A phase shift can be applied to a complex number simply by multiplying by a unit vector with the given phase. This unit vector for  $\phi_{turns}$  is given by

$$\phi_{complex} = \cos(2\pi \cdot \phi_{turns}) + i \sin(2\pi \cdot \phi_{turns}) \quad (\text{A.3})$$

---

<sup>2</sup>this can be calculated using the C math function `fmod`

The sin and cos terms are calculated <sup>3</sup> and used to multiply each of the visibility points in turn.

#### A.2.4 Time-averaging

A single-use (for the observation described in chap. 7) version of difx2fits was also written to allow time-averaging in difx2fits. The DiFX output format writes the results for each sub-band separately, and the data is written in baseline (then time) order. For the purposes of time averaging it is therefore necessary to buffer an entire visibility for all baselines. Care must also be taken not to time-average across scan boundaries.

Finally pulsar binning can complicate matters, and difx2fits is also used to assemble the results of multiple correlations into a single fits file. Generalising for all possible use cases was found to be extremely complicated and would probably necessitate a complete rewrite of difx2fits.

For this reason modifications were made with certain restrictions in mind: namely, no pulsar binning and only a single correlator job. To achieve this it was only necessary to buffer a single visibility.

#### A.2.5 Uses

This implementation would allow any model to be swapped into a UV dataset (where integration time and bandwidth limitations allow). There are many possible applications of this. For example it could be used to apply improved EOPs to an eVLBI observation post-correlation (this can also be done in other post-processing software).

### A.3 Running batch jobs on the cluster

As described in chap. 4, for correlator testing, an instance of vlbifake was required to run on several nodes across the cluster. One way of achieving this would be to rewrite vlbifake as an MPI program. However a simpler mode for leveraging OpenMPI to achieve the same effect was found.

OpenMPI allows any application to be run across multiple nodes simply by invoking it using the MPI launcher. These processes then each run in a virtual shell with various environment variables provided (those provided by the host running the process will be present; optionally environment variables can be defined to be passed to each virtual shell, or variables defined on the host launching the MPI job can be distributed.) In addition, several environment variables are defined by OpenMPI on a process by process basis<sup>4</sup>. This includes ones which contains the process number (unique to that process) and the total number of processing nodes.

This allows any script written in a language that can read global variables (bash, python, perl, ruby etc.) to be run in parallel across a cluster. No message passing is possible, however an embarrassingly parallel job can be split into multiple parts and each process can determine which part of the job it should be working on using the environmental variables described.

---

<sup>3</sup>this can be done in single precision using the fused function sincosf

<sup>4</sup>these can be discovered by running 'mpirun env'



# Appendix B

## The IRA-DiFX software correlation system

This is a software package which I developed in the first half of 2008 to provide a complete system for correlating from disk in the Istituto di Radioastronomia. It consists of a library of mathematical functions for astronomy, a simple vex parser and set of scripts. It has largely been superseded by `vex2difx` nonetheless it is described briefly here. A much fuller description can be found in Morgan (2008). An early report on progress with correlation at the Institute of Radio astronomy is given in (Morgan 2009).

### B.1 The `pydifx` framework

All of the tools used for DiFX were made accessible through a python framework. This framework logged all errors and provided a scripting interface for automating correlation jobs. An obvious use of this is to combine this with `ParselTongue` allowing processing from baseband data to imaging in a single script!

#### B.1.1 wrapper scripts

These were written for all of the DiFX tools required for correlation including `mpifxcorr` (the correlator itself); `vex2config.pl` (a perl script for generating the input file for the correlator); `calcif` (for generating the delay model); and `difx2fits` (for converting the correlator output to a FITS-IDI file).

In addition the following software was developed.

#### B.1.2 `vex2calc`

`vex2config.pl` generates a correlator input file from a vex file<sup>1</sup>. However `calcif` requires another file to run. the '`calc`' file contains information on the Earth Orientation Parameters, source positions and antenna positions in order to calculate the delay model. `vex2calc` was written to parse the vex file and produce this information. A vex file parser was written in python to achieve this. This has been tested on a variety of vexfiles describing astronomical and geodetic observations.

---

<sup>1</sup>The vex file is generated when the observation is scheduled and “prescribes a complete description of a VLBI experiment including scheduling, data-taking and correlation”.

## B.2 Other Scripts

### **vex2flag**

This program parses the vex file and produces a flag file giving the times that the antennas were not on-source. This can be in one of two formats: either the format recognised by difx2fits or one that can be used as an input to the AIPS task 'UVFLG'.

### **B.2.1 log2clock**

This will parse a station log file, extract all GPS measurements and make a least-squares fit to the readings to determine a delay and rate for the clock model. The limitations of this technique are mentioned briefly in sect. 8.6. Further problems occur due to inconsistencies (and downright errors) in station log files. Two conventions are possible for the sign of the clock (these should be denoted by gps-fmout or fmout-gps however errors have been observed). Also sometimes microseconds are recorded, sometimes seconds in scientific notation.

log2clock should take care of all of this except the sign convention (in cases where the log file is in error). In addition in some cases it is necessary to provide a time to which to tie the delay and rate.

### **B.2.2 log2tsys**

With some input on the naming convention of the sub-bands, log2tsys will format the system temperatures in the antenna logs. It can output either a file readable by difx2fits or the AIPS task ANTAB.

## B.3 Mathematical functions

Various functions useful for astronomy have been written and are contained within the python astro module. Use of the doctest framework ensures a clear documentation of the exact specification of each of the functions as well as a simple test of the accuracy of the functions on any platform.



# Appendix C

## RFIX: A ParselTongue Framework for automated RFI flagging

As mentioned in sect. 2.1 automatic flagging is urgently needed for radio interferometry data. In addition flagging algorithms are rather lacking in AIPS and GUIs for manually flagging data tend to be rather clunky.

RFIX is an attempt to create a framework in which different flagging algorithms can easily be tested on (old) VLA datasets. This would be particularly well suited to incorporation into existing ParselTongue pipelines e.g. the e-MERLIN pipeline and EVN pipelines.

### C.1 Philosophy

At its heart RFIX uses the definition of UV data provided in sect. 1.3.6. That is, UV data can be described as a multi-dimensional with axes for time, spectral channel, sub-band, antenna, etc. RFI can then be detected by performing statistical operations such as averaging, along these axes.

### C.2 Implementation

The software relies on two other pieces of software: ParselTongue (a python interface to AIPS) and numpy (a numerical library for python). The interface generator swig is also useful. This allows functions to be written in c which can then be called in python. Most statistical functions required will be implemented in numpy, however as a proof of concept, some statistical functions were implemented in C and made callable in python with the use of swig and various wrapper script.

### C.3 Using RFIX

The data for a single scan are read into a numpy array using the ParselTongue "Wizardry" feature. All of the relevant data are stored in an instantiation of the class Visibility. This contains the UV data in a multi-dimensional array object and all of the relevant array boundary information (number of channels, number of times etc.) The times of each of the visibilities is stored as well, along with the frequencies of all the spectral channels.

Once the data are stored in this way, any vector operations can be carried out on them. This has many uses. For example the vector averaging for the plots shown in chap. 6 was done in RFIX; the polynomial fits and subtraction their subtraction from the data were done using numpy, the data were then written to an ascii file and a script was written (automatically using python) to display the data using gnuplot.

### **C.3.1 Data flagging**

Flagging procedures can be designed by the user and then further functions can be used to go through the entire dataset one scan at a time flagging the data according to the user specification. The result is a single ASCII file containing the flags which can then be applied using the AIPS task 'UVFLG'.

### **C.3.2 Data display**

Various functions also exist for displaying data in amplitude and phase against time along with the track through the UV plane of the data.

## **C.4 Possible future work**

A huge handicap of RFX is that it relies on ParseITongue to access the UV data. ParseITongue currently uses RPC which is very slow and read-only. Having the ability to read the data at disk-access speed would be a huge advantage. This project could perhaps be revisited now that development of CASA (which is tied within a python framework) is at an advanced stage.

# Bibliography

- W. Alef, "A Review of VLBI Instrumentation," in *European VLBI Network on New Developments in VLBI Science and Technology*, R. Bachiller, F. Colomer, J.-F. Desmurs, and P. de Vicente, Eds., 2004, pp. 237–244. 45, 58
- I. Avruch, S. V. Pogrebenko, and L. I. Gurvits, "VLBI observations of spacecraft," in *Proceedings of the 8th European VLBI Network Symposium*, 2006. 43
- U. Bach, "A new term for the Effelsberg pointing model," MPIfR, Technical Information — Pointing, 2009. [Online]. Available: [http://www.mpifr-bonn.mpg.de/div/effelsberg/advanced\\_points.html](http://www.mpifr-bonn.mpg.de/div/effelsberg/advanced_points.html) 12
- R. H. Becker, R. L. White, and D. J. Helfand, "The FIRST Survey: Faint Images of the Radio Sky at Twenty Centimeters," *ApJ*, vol. 450, pp. 559–+, Sep. 1995. 86
- S. Bhatnagar, "Plan for Imaging Algorithm Research and Development," NRAO, EVLA Memorandum 139, 2009. [Online]. Available: <http://www.aoc.nrao.edu/evla/memolist.shtml> [Filename evlamemo139.pdf]. 63, 66
- J. Brand, P. Caselli, K.-H. Mack, S. Poppi, A. Possenti, I. Prandoni, and A. Tarchi, "The Sardinia Radio Telescope (SRT) Science and Technical Requirements," Istituto di Radioastronomia, Rapporti Interni 371/05, 2005. [Online]. Available: <http://w1.ira.inaf.it/Library/Rapporti.html> [Filename 371-05.pdf]. 42
- A. H. Bridle and F. R. Schwab, "Bandwidth and Time-Average Smearing," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, and R. A. Perley, Eds., vol. 180, 1999, pp. 371–+. 64
- D. S. Briggs, F. R. Schwab, and R. A. Sramek, "Imaging," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 127–+. 30
- W. Brisken, "A Guide to Software Correlation Using NRAO-DiFX Version 1.1," NRAO, Tech. Rep., 2008. [Online]. Available: <http://www.aoc.nrao.edu/~wbrisken/NRAO-DiFX-1.1/NRAO-DiFX-1.1-UserGuide.pdf>. 74
- W. Brisken and J. Romney, "The Evolution of Hard Drive Prices Over 2000 Days," NRAO, VLBA Sensitivity Upgrade Memorandum 22, 2008. [Online]. Available: <http://www.vlba.nrao.edu/memos/sensi/> [Filename sensimemo22.pdf]. 46, 59
- R. Campbell, "Recent Results from the EVN MkIV Data Processor at JIVE," in *European VLBI Network on New Developments in VLBI Science and Technology*, R. Bachiller, F. Colomer, J.-F. Desmurs, and P. de Vicente, Eds., 2004, pp. 245–248. 46
- E. Cenacchi, "SRT – Il Radiotelescopio della Sardegna," Master's thesis, Università di Bologna, Facoltà di Scienza Matematiche Fisiche e Naturali, 2004, (in Italian). [Online]. Available: [http://www.med.ira.inaf.it/Tesi\\_fatteastro.htm](http://www.med.ira.inaf.it/Tesi_fatteastro.htm) 41
- A. S. Cohen, W. M. Lane, W. D. Cotton, N. E. Kassim, T. J. W. Lazio, R. A. Perley, J. J. Condon, and W. C. Erickson, "The VLA Low-Frequency Sky Survey," *AJ*, vol. 134, pp. 1245–1262, Sep. 2007. 32, 115
- J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick, "The NRAO VLA Sky Survey," *AJ*, vol. 115, pp. 1693–1716, May 1998. 64

## BIBLIOGRAPHY

---

- T. Cornwell, R. Braun, and D. S. Briggs, "Deconvolution," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 151–+. 30
- T. J. Cornwell, D. J. Saikia, P. Shastri, L. Feretti, G. Giovannini, P. Parma, and C. J. Salter, "Extragalactic Sources with Asymmetric Radio Structure II. Further Observations of the Quasar B2 1320+299," *J. Astrophys. Astr.*, vol. 7, pp. 119–129, 1986. 85, 86, 89, 90
- T. Cornwell, K. Golap, and S. Bhatnagar, "W Projection: A New Algorithm for Non-coplanar Baselines," NRAO, EVLA Memorandum 67, 2003. [Online]. Available: [www.aoc.nrao.edu/evla/geninfo/memolist.shtml](http://www.aoc.nrao.edu/evla/geninfo/memolist.shtml) [Filename evlamemo67.pdf]. 66
- A. Deller, "Precision VLBI astrometry: Instrumentation, algorithms and pulsar parallax determination," Ph.D. dissertation, Swinburne University, Faculty of Information and Communication Technology, January 2009. 47, 117
- A. T. Deller, S. J. Tingay, M. Bailes, and C. West, "DiFX: A Software Correlator for Very Long Baseline Interferometry Using Multiprocessor Computing Environments," *PASP*, vol. 119, pp. 318–336, Mar. 2007. 45, 47, 55, 65
- S. S. Doeleman, J. Weintroub, A. E. E. Rogers, R. Plambeck, R. Freund, R. P. J. Tilanus, P. Friberg, L. M. Ziurys, J. M. Moran, B. Corey, K. H. Young, D. L. Smythe, M. Titus, D. P. Marrone, R. J. Cappallo, D. Bock, G. C. Bower, R. Chamberlin, G. R. Davis, T. P. Krichbaum, J. Lamb, H. Maness, A. E. Niell, A. Roy, P. Strittmatter, D. Werthimer, A. R. Whitney, and D. Woody, "Event-horizon-scale structure in the supermassive black hole candidate at the Galactic Centre," *Nature*, vol. 455, pp. 78–80, Sep. 2008. 29, 33
- W. C. Erickson, "Long Wavelength Interferometry," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 601–+. 25
- L. Feretti, G. Giovannini, and P. Parma, "The quasar B2 1320 plus 29," *A&A*, vol. 115, pp. 423–427, Nov. 1982. 85
- E. B. Fomalont, "Astrometry and Geodesy," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 463–+. 30
- E. B. Fomalont, "Image Analysis," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 301–+. 12
- E. B. Fomalont and R. A. Perley, "Calibration and Editing," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 79–+. 2, 12, 28
- M. A. Garrett, T. W. B. Muxlow, S. T. Garrington, W. Alef, A. Alberdi, H. J. van Langevelde, T. Venturi, A. G. Polatidis, K. I. Kellermann, W. A. Baan, A. Kus, P. N. Wilkinson, and A. M. S. Richards, "AGN and starbursts at high redshift: High resolution EVN radio observations of the Hubble Deep Field," *A&A*, vol. 366, pp. L5–L8, Feb. 2001. 66
- M. A. Garrett, J. M. Wrobel, and R. Morganti, "Discovering the microJy Radio VLBI Sky via "Full-beam" Self-calibration," in *European VLBI Network on New Developments in VLBI Science and Technology*, R. Bachiller, F. Colomer, J.-F. Desmurs, & P. de Vicente, Ed., 2004, pp. 35–38. 66, 114
- M. A. Garrett, J. M. Wrobel, and R. Morganti, "Deep VLBI Imaging of Faint Radio Sources in the NOAO Bootes Field," *ApJ*, vol. 619, pp. 105–115, Jan. 2005. 65, 113

- M. A. Garrett, J. M. Wrobel, and R. Morganti, "21st Century VLBI: Deep Wide-Field Surveys," in *Future Directions in High Resolution Astronomy*, ser. Astronomical Society of the Pacific Conference Series, J. Romney & M. Reid, Ed., vol. 340, Dec. 2005, pp. 600–+. 27
- P. J. Hancock, S. J. Tingay, E. M. Sadler, C. Phillips, and A. T. Deller, "e-VLBI observations of GHz-peaked spectrum radio sources in nearby galaxies from the AT20G survey," *MNRAS*, vol. 397, pp. 2030–2036, Aug. 2009. 40
- M. A. Holdaway, "Mosaicing with Interferometric Arrays," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 401–+. 4
- M. A. Holdaway and T. T. Helfer, "Interferometric Array Design," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 537–+. 41
- M. Holdaway, "Comparing MMA and VLA Capabilities in the 36-60 GHz Band," NRAO, ALMA Memorandum 138, 1995. [Online]. Available: <http://www.alma.nrao.edu/memos/index.html> [Filename evlamemo139.pdf]. 61
- S. Johnston, R. Taylor, M. Bailes, N. Bartel, C. Baugh, M. Bietenholz, C. Blake, R. Braun, J. Brown, S. Chatterjee, J. Darling, A. Deller, R. Dodson, P. Edwards, R. Ekers, S. Ellingsen, I. Feain, B. Gaensler, M. Haverkorn, G. Hobbs, A. Hopkins, C. Jackson, C. James, G. Joncas, V. Kaspi, V. Kilborn, B. Koribalski, R. Kothes, T. Landecker, A. Lenc, J. Lovell, J. Macquart, R. Manchester, D. Matthews, N. McClure-Griffiths, R. Norris, U. Pen, C. Phillips, C. Power, R. Protheroe, E. Sadler, B. Schmidt, I. Stairs, L. Staveley-Smith, J. Stil, S. Tingay, A. Tzioumis, M. Walker, J. Wall, and M. Wolleben, "Science with ASKAP. The Australian square-kilometre-array pathfinder," *Experimental Astronomy*, vol. 22, pp. 151–273, Dec. 2008. 7
- M. L. Kaiser, "Reflections on the Radio Astronomy Explorer program of the 1960s and 70s," in *Low Frequency Astrophysics from Space*, ser. Lecture Notes in Physics, Berlin Springer Verlag, N. E. Kassim & K. W. Weiler, Ed., vol. 362, 1990, pp. 3–7. 1
- K. I. Kellermann, "Grote Reber (1911-2002)," *PASP*, vol. 116, pp. 703–711, Aug. 2004. 1
- M. Kesteven, D. Graham, E. Fürst, O. Lochner, and J. Neidhöfer, "Phase-Coherent Holography of the 100 m Telescope," MPIfR, Technical Information — Holography, 2001. [Online]. Available: [http://www.mpifr-bonn.mpg.de/div/effelsberg/advanced\\_points.html](http://www.mpifr-bonn.mpg.de/div/effelsberg/advanced_points.html) 62
- E. Lenc, M. A. Garrett, O. Wucknitz, J. M. Anderson, and S. J. Tingay, "A Deep, High-Resolution Survey of the Low-Frequency Radio Sky," *ApJ*, vol. 673, pp. 78–95, Jan. 2008. 65, 66, 69, 113
- S. Manabe, Y. Tamura, T. Jike, K. Horiai, and Vera Team, "Status and Plan of Geodetic and Astrometric Observations with VERA," in *International VLBI Service for Geodesy and Astrometry 2004 General Meeting Proceedings.*, N. R. Vandenberg & K. D. Baver, Ed., Jun. 2004, pp. 151–+. [Online]. Available: <http://ivscc.gsfc.nasa.gov> 42
- F. Mantovani, "The role of the Sardinia Radio Telescope in Very Long Baseline Interferometry," *Memorie della Societa Astronomica Italiana Supplement*, vol. 10, pp. 98–+, 2006. 42, 43
- R. McCool, M. Bentley, S. Garrington, R. Spencer, R. Davis, and B. Anderson, "Phase Transfer for Radio Astronomy Interferometers, Over Installed Fiber Networks, Using a Round-Trip Correction System," in *Annual Precise Time and Time Interval (PTTI) Meeting*, vol. 40, 2008, pp. 107–+. [Online]. Available: <http://www.pttimeeting.org/archivemeetings/ptti2008.html> 16

## BIBLIOGRAPHY

---

- E. Middelberg, H. Rottmann, W. Alef, S. Tingay, A. Deller, R. Norris, and E. Lenc, "VLBA observations of 96 sources in the Chandra Deep Field South," in *The role of VLBI in the Golden Age for Radio Astronomy*, 2008. 69
- A. Mioduszewski and L. Kogan, "Strategy for Removing Tropospheric and Clock Errors using DELZN," NRAO, AIPS Memorandum 110, 2009. [Online]. Available: <http://www.aips.nrao.edu/aipsmemo.html> [Filename AIPSMEMO111.PS]. 28, 92, 109, 115
- J. S. Morgan, "Radio astronomy data reduction at the Institute of Radio Astronomy," *Memorie della Societa Astronomica Italiana Supplement*, vol. 13, pp. 79–+, 2009. 121
- J. S. Morgan, "The DiFX Software Correlator at IRA," Istituto di Radioastronomia, Rapporti Interni 419/08, 2008. [Online]. Available: <http://w1.ira.inaf.it/Library/Rapporti.html> [Filename 419-08.pdf]. 40, 121
- L. Moscadelli and I. Porceddu, "I-VLBI of molecular masers," *Memorie della Societa Astronomica Italiana Supplement*, vol. 10, pp. 102–+, 2006. 33, 42
- P. J. Napier, "The Primary Antenna Elements," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 37–+. 3, 4
- L. Olmi and G. Grueff, "SRT: design and technical specifications," *Memorie della Societa Astronomica Italiana Supplement*, vol. 10, pp. 19–+, 2006. 3, 4
- T. J. Pearson, "Non-Imaging Data Analysis," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 335–+. 29
- T. J. Pearson, M. C. Shepherd, G. B. Taylor, and S. T. Myers, "Automatic Synthesis Imaging with Difmap," in *Bulletin of the American Astronomical Society*, ser. Bulletin of the American Astronomical Society, vol. 26, Dec. 1994, pp. 1318–+. 93
- L. Petrov. (2008, Sep) Mark-5 VLBI Analysis Software Calc/Solve. [Online]. Available: <http://gemini.gsfc.nasa.gov/solve/> [Accessed 2 Mar 2010]. 25, 72
- L. Petrov. (2009, May) Library VTD for computation of theoretical VLBI time delay. [Online]. Available: <http://astrogeo.org/vtd/> [Accessed 6 Mar 2010]. 25
- L. Petrov, Y. Y. Kovalev, E. Fomalont, and D. Gordon, "The Third VLBA Calibrator Survey: VCS3," *AJ*, vol. 129, pp. 1163–1170, Feb. 2005. 42
- L. Petrov, T. Hirota, M. Honma, K. M. Shibata, T. Jike, and H. Kobayashi, "VERA 22 GHz Fringe Search Survey," *AJ*, vol. 133, pp. 2487–2494, Jun. 2007. 85
- C. Phillips and A. Deller, "Performance testing of the DiFX correlator," in *Proceedings of the 8th International e-VLBI Workshop. 22-26 June 2009. Madrid, Spain Published online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=82.>, p.44*, 2009, pp. 44–+. 52
- C. Phillips, T. Tzioumis, S. Tingay, J. Stevens, J. Lovell, S. Amy, C. West, and R. Dodson, "LBADR: The LBA Data Recorder," in *Proceedings of the 8th International e-VLBI Workshop. 22-26 June 2009. Madrid, Spain Published online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=82.>, p.99*, 2009, pp. 99–+. 20, 39
- C. J. Phillips, A. Deller, S. W. Amy, S. J. Tingay, A. K. Tzioumis, J. E. Reynolds, D. L. Jauncey, J. Stevens, S. P. Ellingsen, J. Dickey, R. P. Fender, V. Tudose, and G. D. Nicolson, "Detection of compact radio emission from Circinus X-1 with the first Southern hemisphere e-VLBI experiment," *MNRAS*, vol. 380, pp. L11–L14, Sep. 2007. 40

- A. Polatidis. (2009, Nov) EVN Status Tables. [Online]. Available: [http://www.evbi.org/user\\_guide/EVNstatus.txt](http://www.evbi.org/user_guide/EVNstatus.txt) [Accessed 24 Jan 2010]. 10
- W. Reich, P. Kalberla, K. Reif, and J. Neidhofer, "Large Dynamic Range Observations with the Effelsberg 100-in Radio Telescope," *A&A*, vol. 69, pp. 165–+, Oct. 1978. 4
- J. D. Romney, "Cross Correlators," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 57–+. 13, 14
- M. Ryle, F. G. Smith, and B. Elsmore, "A preliminary survey of the radio stars in the Northern Hemisphere," *MNRAS*, vol. 110, pp. 508–+, 1950. 9, 33
- O. J. Sovers, J. L. Fanelow, and C. S. Jacobs, "Astrometry and geodesy with radio interferometry: experiments, models, results," *Reviews of Modern Physics*, vol. 70, pp. 1393–1454, Oct. 1998. 17, 28, 74
- R. Spencer, R. Hughes-Jones, and D. Simeonidou, "Towards the Next Generation of e-VLBI," in *The role of VLBI in the Golden Age for Radio Astronomy*, 2008. 43, 59
- R. Strom, "What is the primary beam response of an interferometer with unequal elements?" in *European VLBI Network on New Developments in VLBI Science and Technology*, R. Bachiller, F. Colomer, J.-F. Desmurs, & P. de Vicente, Ed., 2004, pp. 273–274. 11, 63
- M. Tarengi, "The Atacama Large Millimeter/Submillimeter Array: overview and status," *Ap&SS*, vol. 313, pp. 1–7, Jan. 2008. 19
- A. R. Thompson, "Fundamentals of Radio Interferometry," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 11–+. 66
- A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and synthesis in radio astronomy*. New York, Wiley-Interscience, 1986, 554 p., 1986. 8, 11, 16, 24, 30, 64
- S. J. Tingay, W. Alef, D. Graham, and A. T. Deller, "Geodetic VLBI correlation in software," *Journal of Geodesy*, vol. 83, pp. 1061–1069, Jun. 2009. 43, 72, 77
- V. Tornatore and R. Haas, "Considerations on the observation of GNSS-signals with the VLBI2010 system," in *19th European VLBI for Geodesy and Astrometry Working Meeting, held in Bordeaux, France, March 24-25, 2009*, Eds: G. Bourda, P. Charlot, and A. Collioud, Université Bordeaux 1 - CNRS, Observatoire Aquitain des Sciences de l'Univers, Laboratoire d'Astrophysique de Bordeaux, p. 151-155., G. Bourda, P. Charlot, & A. Collioud, Ed., Sep. 2009, pp. 151–155. 43
- J. S. Ulvestad, R. A. Perley, and C. J. Chandler. (2009, Jan) The Very Large Array Observational Status Summary. [Online]. Available: <http://www.vla.nrao.edu/astro/guides/vlas/current/> [Accessed 10 Jan 2010]. 6
- H. J. van Langevelde, "The future of EVN," in *The role of VLBI in the Golden Age for Radio Astronomy*, 2008. 59
- R. V. van Nieuwpoort and J. W. Romein, "Using many-core hardware to correlate radio astronomy signals," in *ICS '09: Proceedings of the 23rd international conference on Supercomputing*. New York, NY, USA: ACM, 2009, pp. 440–449. 46
- M. A. W. Verheijen, T. A. Oosterloo, W. A. van Cappellen, L. Bakker, M. V. Ivashina, and J. M. van der Hulst, "Apertif, a focal plane array for the WSRT," in *The Evolution of Galaxies Through the Neutral Hydrogen Window*, ser. American Institute of Physics Conference Series, R. Minchin & E. Momjian, Ed., vol. 1035, Aug. 2008, pp. 265–271. 7, 19

## BIBLIOGRAPHY

---

- R. Verma, G. Maccaferri, A. Orfei, I. Prandoni, and L. Gregorini, "A new K-band (18-26 GHz) 7-horn multi-feed receiver: Calibration campaign at the Medicina 32 m dish," Istituto di Radioastronomia, Rapporti Interni 419/08, 2008. [Online]. Available: <http://w1.ira.inaf.it/Library/Rapporti.html> [Filename 419-08.pdf]. 7, 105
- C. Vickery. (2009, Jan) IEEE-754 Reference Material. [Online]. Available: <http://babbage.cs.qc.cuny.edu/IEEE-754/References.xhtml> [Accessed 1 Mar 2010]. 74
- R. C. Walker, "Very Long Baseline Interferometry," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, and R. A. Perley, Eds., vol. 180, 1999, pp. 433–+. 29
- R. Walker, "Multi-Tone Per Band Pulse Cal," NRAO, VLBA Scientific Memorandum 8, 1999. [Online]. Available: <http://www.vlba.nrao.edu/memos/sci/> [sci08memo.ps]. 26
- R. Walker and S. Chatterjee, "Ionospheric Corrections Using GPS Based Models," NRAO, VLBA Scientific Memorandum 23, 1999. [Online]. Available: <http://www.vlba.nrao.edu/memos/sci/> [sci23memo.ps]. 25
- R. B. Wayth, L. J. Greenhill, and F. H. Briggs, "A GPU-based Real-time Software Correlation System for the Murchison Widefield Array Prototype," PASP, vol. 121, pp. 857–865, Aug. 2009. 46
- A. Whitney, "Concept for an Affordable High-Data-Rate VLBI Recording and Playback System," MIT Haystack Observatory, Mark 5 Memorandum 001, 1999. [Online]. Available: <http://www.haystack.mit.edu/tech/vlbi/mark5/memo.html> [Filename 001.pdf]. 18
- J. M. Wrobel and R. C. Walker, "Sensitivity," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, & R. A. Perley, Ed., vol. 180, 1999, pp. 171–+. 10
- J. Wrobel, R. Walker, and J. Benson, "Strategies for Phase Referencing with the VLBA," NRAO, VLBA Scientific Memorandum 24, 2000. [Online]. Available: <http://www.vlba.nrao.edu/memos/sci/> [sci24memo.ps]. 28