

Algorithms for Improving Sensitivity of Radio Interferometric Images

A thesis

Submitted in partial fulfillment of the requirements

Of the degree of

Doctor of Philosophy

by

Srikrishna Sekhar



Indian Institute of Science Education and Research Pune

2018

To Thatha & Patti

CERTIFICATE

Certified that the work incorporated in the thesis titled 'Algorithms for improving sensitivity of radio interferometric images' submitted by Srikrishna Sekhar was carried out by the candidate. The work presented here or any part of it has not been included in any other thesis submitted previously for the award of any degree or diploma from any other University or institution.

Date: 11th June 2018



Ramana Athreya
(Supervisor)

DECLARATION

I declare that this thesis written submission represents my ideas in my own words and where others' ideas have been included, I have adequately cited and referenced the original sources. I also declare that I have adhered to all principles of academic honesty and integrity and have not misrepresented or fabricated or falsified any idea/data/fact/source in my submission. I understand that violation of the above will be cause for disciplinary action by the Institute and can also evoke penal action from the sources which have thus not been properly cited or from whom proper permission has not been taken when needed.

Srikrishna

Date: 11th June 2018

Srikrishna Sekhar
Roll No: 20112016

Acknowledgements

I would like to thank my advisor Dr. Ramana Athreya for his support and guidance throughout my PhD. He offered me the space to explore new ideas, while still providing me with invaluable feedback and direction to my work. I am grateful for his confidence in my abilities, even when my own confidence was lacking. I extend my thanks to my thesis committee members Dr. Ishwara Chandra and Dr. Sanjay Bhatnagar for their comments, critical feedback and constant encouragement throughout my research process.

I would like to thank the Director, IISER Pune and the Department of Physics for the opportunity and the facilities.

Aniket Bhagwat helped immensely with this thesis, particularly with the data analysis for the final chapter. I am very grateful for his help.

Dr. Prasad Subramanian has also been of immense help — in matters both academic and musical — and I remain grateful for his time and generosity.

I would like to thank the staff at GMRT, Khodad for their support, help, and hospitality during my observations. Thanks are also due to Dr. Ruta Kale for arranging the fortnightly meetings at NCRA, which was communal therapy for those of us struggling with our data.

I couldn't have done this without the constant friendship of Amruta, Aparna, and Swetha. I am deeply indebted to you all and am in a constant state of shock that you've put up with me for this long. Count this as a vastly insufficient thank you.

I would also like to express my gratitude to Dr. John Mathew and Dr. Pushkar Sohoni for generously adopting me into the HSS department and bringing some much-needed diversity into my life (not to mention all the music!). The last few years of my life would have looked very different without these experiences, and I remain grateful.

Shweta, Mansi, Chintan, Alka, Sandip, Ron, Vibi, Swapna, Harsha, Al and Polo — Lab-mate, step lab-mates and other wayward vagrants. Thank you for keeping the motivation steady and the spirits high. Jamming and performing with Manasi, Abhijeet, Nachiket, and Trupti gave me some respite from the long hours, replacing it instead with hours of epic music. Life would've been a lot more dull and dreary but for their company and music!

I thank my impossibly patient family for keeping me grounded. Amma, Appa, Anna, Dhivi, Tashi and of late Paru. Thank you for your constant love, support and encouragement. Finally, Abhi — none of this would have been possible without your support and understanding, I am eternally grateful.

It would take another year or two to list out everyone who has helped me through this. Rest assured I remain very grateful for the help and company.

Contents

Acknowledgements	v
Synopsis	xii
1 Introduction	1
1.1 Radio frequency interference	1
1.2 Direction dependent errors	4
1.3 Double double radio galaxies	4
1.4 Structure of this thesis	5
2 Uncorrupting the visibilities - A review	6
2.1 RFI Identification and Removal	6
2.1.1 RFI flagging	7
2.1.2 RFI excision	15
2.1.3 RFI mitigation	23
2.2 Direction dependent calibration errors	26
2.2.1 A-projection	26
2.2.2 Source peeling	29
3 The IPFLAG algorithms	34
3.1 Introduction	34
3.2 Description of the algorithms	35
3.2.1 RFI flagging in the binned UV plane	36
3.2.2 RFI flagging in the baseline time-channel plane	38
3.3 Observations and parameters	49
3.4 Efficacy of the algorithms	51
3.5 Conclusion	64
4 Direction dependent error mitigation	66
4.1 Introduction	66
4.2 Snapshot imaging	67
4.2.1 Flux density variation with time	68

4.2.2	Flux density variation with radial distance	73
4.2.3	Flux density variation with parallactic angle	74
4.2.4	Flux density variation with elevation angle	80
4.2.5	Variable ionosphere	80
4.2.6	Subtracting residual source flux variation	89
4.3	Baseline based defringing	90
4.4	Discussion	96
5	Evolution of spix in restarting AGN	101
5.1	Introduction	101
5.1.1	Quenching of the radio jet	103
5.1.2	Restarting the radio jet	104
5.1.3	Time scales of quiescence	105
5.2	Evolution of the spectral index structure of dying radio lobes	105
5.2.1	Sample of double-double radio galaxies	106
5.2.2	Image Processing	114
5.3	Results and Discussion	118
5.3.1	Hotspot Compactness	118
5.3.2	Spectral index structure	119
5.4	Conclusion	120
6	Summary	127
	Appendix A Spectral leakage	130

List of Figures

2.1	VLA 1.4 GHz beam	28
2.2	Schematic of the Lonsdale scales for ionosphere fluctuations	30
3.1	Flowchart of the IPFLAG algorithm	37
3.2	Comparison of the binned UV plane before and after applying GRIDflag	40
3.3	Illustrations of the defringing process in TCflag	45
3.4	RMS noise vs. UV distance after IPFLAG	47
3.5	Comparison of flux density with TGSS after IPFLAG	50
3.6	Fractional flux density change in strong and faint sources after IPFLAG	55
3.7	RMS noise as a function of radial distance after IPFLAG	56
3.8	RMS noise as a function of total flux after IPFLAG	57
3.9	RMS histogram across the entire image after IPFLAG	58
3.10	Source counts before and after IPFLAG	59
3.11	Reduction of artefacts after IPFLAG	61
4.1	Variation of flux density with time in the J1453+3308 field	69
4.2	Variation of flux density with time in the J1158+2621 field	70
4.3	Variation of flux density with time in the VIRMOSC field	71
4.4	Variation of flux density with time in the TAUBOO field	72
4.5	Variation of flux density with radial distance	73
4.6	Locus of a source on an elliptical beam	75
4.7	Model change in flux density with parallactic angle	75
4.8	Variation of flux density with parallactic angle in the J1453+3308 field	76
4.9	Variation of flux density with parallactic angle in the J1158+2621 field	77
4.10	Variation of flux density with parallactic angle in the VIRMOSC field	78
4.11	Variation of flux density with parallactic angle in the TAUBOO field	79

4.12	Variation of flux density with antenna elevation angle in the J1453+3308 field	81
4.13	Variation of flux density with antenna elevation angle in the J1158+2621 field	82
4.14	Variation of flux density with antenna elevation angle in the VIR-MOSC field	83
4.15	Variation of flux density with antenna elevation angle in the TAUBOO field	84
4.16	Histogram of source position jitter in the J1453+3308 field	85
4.17	Histogram of source position jitter in the J1158+2621 field	86
4.18	Histogram of source position jitter in the VIRMOSC field	87
4.19	Histogram of source position jitter in the TAUBOO field	88
4.20	Residual and model baseline fringe in the TC plane-1	93
4.21	Residual and model baseline fringe in the TC plane-2	94
4.22	Residual and model baseline fringe in the TC plane-3	95
4.23	Single baseline-channel residual fringe and fitted model— Good fit	97
4.24	Single baseline-channel residual fringe and fitted model— Poor fit	98
4.25	Improvements to source structure after DDE mitigation - 1	99
4.26	Improvements to source structure after DDE mitigation - 2	100
5.1	Image of 3C452 at 150 MHz	107
5.2	Image of J0041+3224 at 150 MHz	109
5.3	Image of J0116-4722 at 150 MHz	110
5.4	Image of J1158+2621 at 150 MHz	111
5.5	Image of 3C293 at 150 MHz	112
5.6	Image of J1453+3308 at 150 MHz	113
5.7	Image of J1548-3216 at 150 MHz	114
5.8	Image of J1835+6204 at 150 MHz	115
5.9	Image of J2345-0449 at 150 MHz	116
5.10	Hotspot FWHM at 150 and 325 MHz as a function of spectral age.	121
5.11	Spectral index map of J0041+3224 and J0116-4722	122
5.11	Spectral index map of J1158+2621 and J1453+3308	123
5.11	Spectral index map of J1548-3216 and J1835+6204	124
5.12	Radial spectral index profiles of J0041+3224	125
5.13	Radial spectral index profiles of J0116-4722	125
5.14	Radial spectral index profiles of J1158+2621	125
5.15	Radial spectral index profiles of J1453+3308	126
5.16	Radial spectral index profiles of J1548-3216	126
5.17	Radial spectral index profiles of J1835+6204	126

A.1	Plot of the spectral leakage induced by a Fast Fourier Transform . .	131
-----	--	-----

List of Tables

3.1	The effect of IPFLAG on various image parameters	53
3.2	Percentage change in the median flux density in the three different regimes.	60
5.1	List of source positions, sizes, and redshifts for the sample of DDRGs	108
5.2	List of DDRGs and their reported spectral ages	117

Synopsis

This thesis presents new techniques to flag radio frequency interference (RFI) and mitigate systematic errors in calibration. They were tested on several 150 MHz data sets obtained using the Giant Metrewave Radio Telescope (GMRT), Pune, India. The 150 MHz band is important for studying a variety of astrophysical phenomena, but is under-utilised on account of the difficulties in producing high quality images due to an active RFI environment and direction-dependent calibration errors. The four new algorithms presented here were developed for RFI flagging and mitigation of direction-dependent errors in GMRT 150 MHz data, but may be implemented for any imaging interferometer at any frequency.

The RFI environment at the GMRT 150 MHz band is very active, strong, dynamic and diverse with very different time and frequency characteristics. The issue of persistent and/or broadband RFI has already been addressed to some extent by the algorithm RfiX. Therefore, we have focused on developing two algorithms to identify and flag RFI that is localised in time and frequency, namely GRIDflag (which operates on the entire binned UV plane) and TCflag (which operates on the time-channel plane of an individual baseline). They have been implemented as part of a single RFI flagging pipeline called IPFLAG.

During the course of a synthesis observation the sampling of the visibility is irregular and often sparse. However due to the limited field of view of the instrument, the visibilities can be binned and averaged in the UV plane without loss of information. GRIDflag exploits the redundancy within a UV-bin to identify and flag RFI corrupted visibilities. Different baselines can contribute visibilities to a particular UV-bin, but at different times. The visibilities in a given UV-bin represent the same (or very similar) sky brightness but are sampled under different RFI environments. The algorithm works by “locally” identifying the RFI threshold in each bin and flagging the corrupted visibilities. The use of “local” thresholds combines higher efficacy of RFI detection with conservation of UV coverage. One of the key advantages of this approach is that the procedure typically retains at least some visibilities in most UV bin even while flagging the fainter RFI in the data.

TCflag works on the residual (*i.e.*, source subtracted) baseline time-channel

(TC) plane which often shows fringe structure due to various reasons, such as improper subtraction of source flux, baseline gain fluctuations etc. In the context of RFI flagging these fringes will inflate the estimated RFI threshold, and hence reduce the sensitivity of the RFI detection. We eliminated these fringes in the TC plane by clipping them in the corresponding two dimensional Fourier space. This works because the RFI and residual fringes have different signatures in the Fourier domain - the sinusoidal fringes in the TC plane will appear as compact peaks in the Fourier plane, and the compact peaks in the TC plane (*i.e.*, the RFI) will conversely be spread over several bins in Fourier space. The RFI thresholds are estimated in the absence of the fringes, and the flags are applied back to the original data. The application of these two algorithms have resulted in a 20-50% improvement in image noise, with an accompanying increase in the number of detected sources. This work has been accepted for publication in the *Astronomical Journal*.

Direction dependent (calibration) errors (DDEs) can arise due to a non-isoplanatic ionosphere, an azimuthally asymmetric antenna primary beam, antenna pointing errors etc. These effects result in the variation of source flux and position as seen by the telescope during an observation, and the magnitude of these variations is a function of time and direction. These observed variations cause errors during the imaging and self-calibration processes as they violate their implicit assumption of an invariant sky brightness distribution *i.e.*, the imaging and calibration algorithms subtract a time-averaged “constant” sky from every visibility while the visibilities contain information of an apparently varying sky distribution.

Several groups around the world have been trying to model and eliminate these effects. We succeeded in reducing the errors due to these effects by targeting the consequences rather than the primary cause. Essentially, uncorrected DDEs reduce image depth by spreading dirty-beam like structures throughout the image from the improperly subtracted sources. Therefore, after going through the standard calibration, imaging and self-calibration procedures we carried out ‘snapshot imaging’ of just the brightest sources separately for each 5-10 minutes scan of the data. This effectively relaxes the constraint of an invariant sky over the entire observation, and imposes time invariance to just 5-10 minute intervals. Applying this process of snapshot imaging to the residual visibilities ‘corrects’ for the over- and under- subtraction of strong sources arising from the factors mentioned previously. This procedure greatly reduces the sidelobes of these improperly subtracted strong sources, and improves the image noise as well as the detectability of faint sources. The conceptual and procedural simplicity, with very little increase in computation, comes at the cost of no further improvement in dynamic range of the strong sources themselves. On the other hand, the strongest sources are

usually not of much interest in many studies.

While the snapshot imaging method eliminates the time and positional variability of the strong sources averaged across all baselines, each baseline will still exhibit fringes corresponding to the improperly subtracted strong sources. This can happen if antenna primary beams are either not circular or differ from antenna to antenna. We reduced these baseline-based errors, by fitting and subtracting the predicted model fringes for the brightest 3-5 sources in the sky, baseline-by-baseline and over short (5-20 minute) time intervals. A combination of the two methods results in a reduction in image RMS noise by up to 40% in the vicinity of the strong sources, as well as an overall improvement in systematic patterns across the field.

We applied the above techniques to image a sample of ‘double-double radio galaxies’ (DDRGs) at 150, 325 and 610 MHz using the GMRT. These objects have two distinct sets of radio lobes, where the outer (older) lobes are remnants from a earlier epoch of activity and are no longer being energised. We wished to investigate if the pattern of spectral index steepening away from the hotspot in FRII type radio galaxies was retained in the dead outer lobes of DDRGs. We created high resolution spectral index maps between the 150-325 MHz and 325-610 MHz images to measure the pattern of variation of spectral index, if any. We had hoped that the persistence of such a spectral signature of the ancient location of energisation would validate our attempt for a similar search in galaxy cluster radio halos. However the results indicate that there is no consistent spectral index pattern across our sample of DDRGs which rules out the possibility of using it as a tracer for energisation. It seems that the lobe homogenises within timescales of about 25 Myr.

Chapter 1

Introduction

Low frequency radio observations ($\nu < 1$ GHz) are important for studying a wide variety of astrophysical phenomenon. Typically, observations using wavelengths that are ≥ 1 m are plagued by strong radio frequency interference (RFI) that can be significantly stronger than the cosmic signal. In addition to this, there are a variety of other corrupting effects such as the time-dependent ionosphere, antenna beam asymmetries, etc., that can severely limit the sensitivity of the images at these frequencies. We are in the age of the “next-generation” radio telescope, the Square Kilometre Array (SKA), and its precursors: LOFAR (Röttgering 2003), MeerKAT (Jonas 2009), ASKAP (DeBoer *et al.* 2009), MWA (Tingay *et al.* 2013), and uGMRT (Gupta *et al.* 2017). It is vital to understand how best to correct for or eliminate the sources of error associated with these phenomena in order to achieve the very high sensitivity that these instruments are being designed for.

The Giant Metrewave Radio Telescope (GMRT; Swarup *et al.* 1991) is the only fully steerable large dish radio telescope operating at 150 MHz. But the band is highly under-utilised due to the presence of a strong RFI environment and ionospheric effects. Diffuse emission is particularly susceptible to distortions of source structure due to systematic errors caused by RFI and the ionosphere. Accurately mapping these objects requires the elimination (or mitigation) of these systematic errors that affect the imaging and calibration processes. Adding to the corpus of existing algorithms and tools to address these issues is the subject matter of this thesis. While the algorithms were developed for GMRT data they can also be used for imaging data from any other radio interferometer.

1.1 Radio frequency interference

Radio frequency interference (RFI) is most generally defined as any signal received by the interferometer that is not cosmic in origin. These signals can originate from

a variety of different physical sources. While the potential sources of RFI can be many, we classify the correlated RFI into four different categories on the basis of their behaviour in the time and frequency domains.

Persistent RFI – Persistence refers to the temporal character of the RFI. Typically, if the RFI affects several minutes or more (up to several hours) it can be called persistent. Under certain conditions persistent RFI can be removed from the data while salvaging the original visibilities, which is more desirable than simply rejecting the data.

Sporadic RFI – Sporadic RFI only affects a few time samples at a stretch. However, it can cumulatively affect large fractions of the data and severely hamper the sensitivity of the observation if not appropriately identified and rejected. Its unpredictability in time and frequency makes it difficult to salvage the true visibilities under the RFI.

Broadband RFI – This refers to RFI that occupies a large fraction of (or the entire) observing band. Typically sources such as fluctuating power lines, sparking wires, generators etc., contribute this kind of RFI.

Narrowband RFI – RFI that occupies only one or a few frequency channels is called narrowband. This kind of RFI originates from sources such as mobile telephone towers, satellite communication, FM signals etc.

RFI can occur in any combination of time and frequency characteristics, which makes observing in the presence of an active RFI environment challenging. If this RFI is not mitigated and/or flagged with sufficient sensitivity the final image noise can be several times larger than the theoretical noise limit. RFI can also result in non-zero values of the so-called “closure quantities” which are central to the process of self-calibration (Pearson & Readhead 1984). Non-zero values (for the closure phases) and non-unity values (for closure amplitudes) of these quantities lead to systematic errors in both the self-calibration and imaging processes. Therefore, accurate reconstruction of the sky from the visibilities requires the elimination of RFI.

Different sources of RFI leave different signatures in the data, depending on their spatial and temporal scales of variability. Persistent RFI appears as temporal oscillations (fringes) in the baseline, with the inverse of the fringe stop frequency, whereas sporadic RFI are ‘hotspots’ in the data, affecting only a small region in the time-channel plane of a baseline. Both manner of RFI cause poorer calibration solutions, as well as significant systematic patterns in the image plane.

Each kind of RFI usually requires a different strategy due to their different temporal and frequency characteristics. There has been much effort dedicated toward identifying the best strategies for a particular telescope/RFI environment leading to several different techniques to identify and flag/excise/mitigate the

RFI. A distinction is made between RFI *mitigation*, *excision*, and *flagging*. By *mitigation* we mean a solution that will prevent the RFI from corrupting the visibilities, *i.e.*, remove the RFI from the signal chain in some manner prior to correlating the visibilities. *Excision* involves recovering the original visibilities even in the presence of RFI - this is usually performed in software. Finally *flagging* involves rejecting RFI corrupted visibilities from further processing.

Essentially, all algorithms which flag RFI attempt to estimate the “true” background level (*i.e.* in the absence of RFI) and a threshold level above which the data is deemed to be affected by RFI. The background level may be calculated using a median filter (e.g. Bhat *et al.* 2005; Middelberg 2006). Spectral structure in the visibilities, either due to source structure or due to instrumental bandpass, can be taken care of by piecewise processing across the frequency axis (Winkel *et al.* 2007). More sophisticated algorithms operating at multiple scales have been used in the software suite AOflogger to identify the proper level of the background for determining the threshold for RFI (Offringa *et al.* 2010, 2012).

RFI excision techniques are generally implemented entirely in software, and use some known property of the corrupting RFI to recover the underlying visibilities. Pen *et al.* (2009a) uses singular value decomposition to separate out the terrestrial and celestial signals from the visibilities, and remove only the terrestrial components. There is also a scheme that projects the RFI as a point source at the North Celestial Pole (Cornwell *et al.* 2004; Golap *et al.* 2005) that depends on the RFI being quasi-constant in time. Athreya (2009) exploits the fact that a stationary source of RFI will acquire the inverse of the fringe stop frequency to fit and subtract RFI per baseline. This scheme requires that the RFI fringe be constant over a good fraction of the fringe period, and that the sampling be fast enough for the fringe to be fit.

Mitigation techniques usually require a modification of the data acquisition pipeline at the telescope. Some methods involve the use of a reference antenna to measure the interfering signal and subtract it from the data to recover the underlying, uncorrupted visibilities (Barnbaum & Bradley 1998; Briggs *et al.* 2000; Hellbourg *et al.* 2014). Each of these methods use the reference antenna in different ways (real-time adaptive cancellation, subspace projection, etc.), and target the excision of broad- and narrow-band RFI from a persistent (typically known) source of RFI such as satellites, cell towers etc. However, these methods do not attempt to excise or flag sporadic RFI. There are also methods that use the reference antennas, or antenna arrays to null the antenna beam in the direction of the known RFI (Kocz *et al.* 2010; Van Der Veen & Boonstra 2004). Fridman & Baan (2001) provides a review of different methods of mitigating interference. Other mitigation techniques have attempted to sample the data stream at ultra high rates (nano-

seconds) to catch and flag the RFI before it corrupts the entire integrated time sample (Buch *et al.* 2016).

A single strategy is rarely effective against all manner of RFI, and typically a combination of one or more of the above methods needs to be employed to effectively identify and remove the RFI from the data. Therefore, we see our efforts as adding to the existing suite of tools. With this in mind we have dedicated the next chapter to a more detailed description of many of the extant RFI tools available.

1.2 Direction dependent errors

Direction dependent errors (DDEs) constitute an entirely independent source of systematic errors in the image. These are processes which occur between the source and the correlator, introducing phase and amplitude errors in the data. These processes include atmospheric effects (diffraction, refraction, absorption, non-isoplanaticity, etc.), primary beam asymmetries, telescope pointing jitter, non-coplanar arrays, wide-field imaging, etc.

There have been various attempts to solve these problems by many groups around the world. We have taken an entirely different approach to this problem by seeking to address the consequences rather than the fundamental causes. As before, we believe that these efforts do not replace each other but are complementary. Therefore, we have described some of these previous efforts in some detail in the next chapter.

1.3 Double double radio galaxies

Double-double radio galaxies (DDRGs) are a class of restarting active galactic nuclei (Osterbrock 1991) that show two distinct pairs of radio lobes. Of these the outer pair of lobes are “dead”, *i.e.* no longer being energised by the radio jet. The inner pair of radio lobes represented the latest episode of nuclear activity. These sources provide an observational test bed to study the evolution of the structure of dead and/or diffuse synchrotron plasma.

The focus of this thesis is primarily on the development of algorithms to improve image sensitivity; the DDRG study was taken up as an illustration of the efficacy of the image improvement effort. However, as a group we are also interested in the study of the diffuse radio halos in galaxy clusters. We expected that the study of the diffuse outer lobes of DDRGs would provide some clues in that direction as well.

1.4 Structure of this thesis

Chapter 2 provides an in-depth review of the existing literature on various RFI flagging, mitigation and excision algorithms as well as algorithms which address DDE issues.

Chapter 3 describes two new algorithms developed to identify and flag intermitted RFI in the binned UV-plane and in the baseline time-frequency plane.

Chapter 4 describes the algorithms for mitigating DDEs associated with a time-dependent ionosphere and an asymmetric antenna primary beam.

Finally, Chapter 5 describes the application of these techniques to study the temporal evolution of the spectral index structure in the outer lobes of double-double radio galaxies.

Chapter 2

Uncorrupting the visibilities - A review

Radio telescope data are generally limited by two major categories of corrupting effects *viz.* radio frequency interference (RFI) and direction dependent calibration errors (DDEs). In general, the experience of the radio astronomy community has been that no single method has proved to be the silver bullet for all situations. Therefore, we felt that an extensive review of existing algorithms should be an important component of this thesis, to provide the context and contrast to the new algorithms being presented in the following chapters.

In the methods described here we have used the same notation as in the original sources (for easier reference), which could result in some internal inconsistencies in this chapter. This is not a major issue since there is no cross-referencing between the methods in this chapter, nor are these other methods used in our algorithms described in subsequent chapters.

2.1 RFI Identification and Removal

This section covers algorithms that are used to identify and remove RFI from the visibility data. For the purposes of this discussion, we classify the methods depending upon whether the interfering signal is *mitigated*, *excised*, or *flagged*. RFI mitigation refers to techniques that prevent the RFI from corrupting the visibilities, excision refers to strategies to recover the underlying visibilities from RFI corrupted data, and flagging refers to strategies that identify corrupted data and discard them from further processing.

2.1.1 RFI flagging

Ideally, one would want to identify and separate the RFI from the cosmic signal in the antenna voltage before correlation (*i.e.* mitigation). This way, one would only have to deal with RFI in the antenna before the correlation spreads it to all the baselines of the antenna. If mitigation were not possible one would want to separate the RFI from the cosmic signal in the visibility and remove the former while retaining the latter (*i.e.* excision). However mitigation and excision are not always possible because (a) the RFI affects only a few visibilities at a time, without any underlying structure, or (b) the spatial and temporal signatures of the RFI are not known and are not easily identifiable, or (c) the data needs to be obtained using specialised hardware which are tailored to the specific RFI environment at a telescope. The predominant method of managing RFI during image analysis is to simply identify and exclude RFI-corrupted visibilities (*i.e.* flagging).

Algorithms to flag RFI from the visibility data are almost always entirely software-based, and are typically applied after the correlation. Almost all the methods described in this section operate on the visibility time-channel plane of the baselines. The primary difference between them is in the method of identifying the background in order to determine the RFI threshold. Flagging works best when the RFI is sporadic in time and/or frequency, affecting only a few contiguous integrations at any given time.

Bhat *et al.* (2005) presented a method to identify RFI through fast sampling of the incoming signal (128 μ s) and simultaneous observations with two telescopes (Arecibo and Greenbank). The data were smoothed with a two dimensional median filter to identify the background level. The RFI was identified after normalising the raw data by the median filtered data. They iterated through this process of median filter, normalisation and flagging several times. This resulted in a more sensitive detection of RFI at every subsequent iteration. The median smoothing improved the contrast between the (smoothly varying) cosmic signal and the (locally impulsive) RFI emission. Though this procedure was developed for single-dish pulsar observations the same may be used for interferometers prior to calibration.

Fridman (2008) provides several different methods of estimating the variance of the data and hence the RFI threshold for a robust and efficient identification of sporadic RFI. The difficulty in estimating the RFI threshold lies in determining the ‘true’ variance of the data (*i.e.*, in the absence of RFI) and hence determining the level at which one can flag (RFI) outliers. Cosmic signals received by a radio telescope are noise-like and hence normally distributed. The thermal noise produced by the telescope is also normally distributed, and hence the variance of

the net signal received can be written as

$$\hat{\sigma}_n^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2 \quad (2.1)$$

where x_i are the received data samples. The probability distribution then can be described by the variance $\hat{\sigma}_n^2$ and the mean μ which is nominally zero. In most cases a typical observation does not strictly follow such a distribution. There are often outliers in the data that render the distribution only approximately normal — the central part of the distribution will follow the regular normal distribution but the tails of curve can be significantly heavier. RFI typically causes such an effect on the distribution. In the presence of RFI that can be characterised as short bursts in both the time and frequency plane, the distribution can be written as

$$F(x) = (1 - \epsilon)\mathcal{N}(0, \sigma_{sys}) + \epsilon F_{RFI}(x) \quad (2.2)$$

where $\mathcal{N}(0, \sigma_{sys})$ is a Gaussian distribution with zero mean and a standard deviation determined by the system temperature. ϵ is the fraction of RFI in the distribution and $0 < \epsilon < 1$. A robust statistical measure must be insensitive to outliers, *i.e.*, to F_{RFI} . We discuss here some of the methods presented in Fridman (2008) to obtain robust, stable estimates of the variance.

Variance of the trimmed data - Given a data sample $x_1 \dots x_n$ which are sorted in ascending order, the k smallest and k largest values are removed. k is given by the trimming fraction γ such that $k = \gamma n$ ($0 \leq \gamma \leq 0.5$). The robust variance of such a sample is given by

$$T_1 = \frac{F}{n - 2k} \sum_{n=k}^{n-k} (x_n - \mu_{trim})^2 \quad (2.3)$$

$$\mu_{trim} = \frac{1}{n - 2k} \sum_{n=k}^{n-k} x_n$$

where F is the factor which makes T_1 consistent with a standard Gaussian distribution. μ_{trim} is the mean of the trimmed sample.

Winsorized sample variance - Winsorization also employs the use of trimming to obtain more stable estimates, however it differs from the previous method

in the details of estimating the sample variance. Given that $k = \gamma n$ as before, and $0 \leq \gamma \leq 0.5$ (as before) the data are winsorized by defining

$$W_i = \begin{cases} x_{k+1}, & \text{if } x_i \leq x_{k+1} \\ x_i, & \text{if } x_{k+1} < x_i < x_{n-k} \\ x_{n-k}, & \text{if } x_i \geq x_{n-k} \end{cases}$$

This process sets the data that are rejected by the previous method to the value at the edges. The sample mean is then defined as

$$\hat{\mu}_w = \frac{1}{n} \sum_{i=1}^n W_i \quad (2.4)$$

and the sample variance

$$T_2 = \frac{1}{n-1} \sum_{i=1}^n (W_i - \hat{\mu}_w)^2 \quad (2.5)$$

Like in the previous case, T_2 needs to be multiplied by a ‘consistency factor’ F to obtain the corresponding variance for a standard Gaussian distribution.

Median absolute deviation - For a given sorted data set $x_1 \dots x_n$ the robust variance is given by

$$T_3 = 1.483 \times M \{|x_i - M(x_i)|\} \quad (2.6)$$

where M is the median operator and is given by

$$M = \begin{cases} 0.5 \times (x_m + x_{m+1}), & n = 2m \\ x_{m+1}, & n = 2m + 1 \end{cases}$$

Exponential weighting - This method was first proposed in Shurygin (2000). The mean and variance are expressed as solutions to the following equations -

$$\hat{\mu}_r : \sum_{i=1}^n (x_i - \hat{\mu}_r) e^{-q_i/4} = 0 \quad (2.7)$$

$$T_4 = \sigma_r^2 : \sum_{i=1}^n \left[\frac{(x_i - \hat{\mu}_r)^2}{\sigma_r^2 - 2/3} \right] e^{-q_i/4} = 0 \quad (2.8)$$

where $q_i = (x_i - \hat{\mu}_r)/\hat{\sigma}^2$ and T_4 is the robust variance. Of the four methods discussed here, Fridman (2008) identified the last method as the most effective in terms of identifying outliers while not over-estimating the true variance in the *absence* of outliers. But the median absolute deviation estimator performs the best in terms of identifying outliers *i.e.*, it has a large ‘breakdown point’. The breakdown point determines the fraction of data that have to be outliers before a particular estimator is no longer accurate.

Each of these estimators have their own advantages and drawbacks. The most obvious advantage of the exponential weighting method is that the data does not need to be sorted, and hence relative phase information is preserved. This is important in the context of pulsar observations, where the relative phase information can help determine whether the data being flagged belongs to an astronomical pulse or a terrestrial source of RFI. These methods are only meant to be used in the presence of sporadic RFI - there are more effective ways to deal with persistent, narrowband RFI such as rejection of components from the power spectrum.

Nita *et al.* (2007) proposed the use of higher order statistics to identify non-Gaussianity in the received signal and hence identify and flag RFI. They define a ‘spectral kurtosis’ (SK) estimator that is computed from the power spectral density (PSD) estimates. Typically, PSD estimates are computed by calculating the Fourier transform of the autocorrelation of the input signal. For a given PSD estimate \hat{P}_k , the signal can be characterised by the mean $\mu_k = \langle \hat{P}_k \rangle$ and the standard deviation $\sigma_k = \sqrt{\langle \hat{P}_k^2 \rangle - \langle \hat{P}_k \rangle^2}$. From these two quantities, Nita *et al.* (2007) define a dimensionless quantity to characterise the variability in the system

$$V_k^2 = \frac{\sigma_k^2}{\mu_k^2} \quad (2.9)$$

where V_k is known as the coefficient of dispersion. They compute these quantities by considering a complex input signal -

$$X_k = \sum_{n=0}^{N-1} w_n x_n e^{2\pi i k n / N}, \quad k = 0, \dots, N-1 \quad (2.10)$$

where $\{w_n\}$ are the set of coefficients that are used to reduce the amount of spectral leakage in the PSD spectrum. Additionally, the real and imaginary components of X_k are defined as -

$$A_k = \text{Re}[X_k] \quad (2.11)$$

$$B_k = \text{Im}[X_k] \quad (2.12)$$

The real part of the input signal can then be used to construct an estimator from higher order statistics -

$$K(f_k) = \frac{\langle A_k^4 \rangle}{\langle A_k^2 \rangle^2} \quad (2.13)$$

which for a Gaussian random process is exactly 3. This estimator is effectively Equation 2.9 to within a constant factor. One can also define a ‘spectral kurtosis’ estimator as -

$$SK(f_k) = \frac{\langle |X_k|^4 \rangle - 2\langle |X_k|^2 \rangle^2}{\langle |X_k|^2 \rangle^2} \quad (2.14)$$

and given that $|X_k|^2 \sim P_k$, it follows that

$$V_k^2 = SK(f_k) + 1 \quad (2.15)$$

For a Gaussian time domain signal, $SK(f_k)$ is exactly 0.

The significance of V_k^2 is that it relates directly to a measurable quantity *i.e.*, the power of the signal and also ties together several theoretical concepts (such as the spectral variance, spectral kurtosis and the time-domain kurtosis) and shows that they are effectively equivalent estimators. The spectral kurtosis (SK) estimator is useful in determining the level of non-Gaussianity in a signal, since to calculate the PSD the input time-series data are typically broken up into several adjacent blocks. Therefore calculating the deviation from the expected variance of V_k^2 per time-domain block will indicate the level to which the signal has non-Gaussian components. We note here that the statistical measures discussed above can be applied to any parameter space of the visibilities *i.e.*, in the UV plane, in the time-channel plane, or any higher dimensional parameter space (e.g., the time-frequency — delay space).

Weber *et al.* (1997) have proposed a generalised chi-squared test to determine

whether the input signal has been corrupted by RFI. It requires no additional hardware apart from a quantised correlator, which are employed in most radio interferometers. The statistical test relies on the quadratic difference between the sample mean and the ensemble averaged sample mean over a set of Q measurements of the noise s . They define a vector \mathbf{X} that is formed from the Q sample means w_k

$$\mathbf{X} = [w_1, w_2, \dots, w_Q]^T \quad (2.16)$$

where the superscript T denotes the transpose operation. In the absence of RFI, \mathbf{X} converges to a multi-dimensional Gaussian mean vector \mathbf{X}_0 , with a multi-dimensional covariance matrix \mathbf{R}_0 . Both these quantities depend on the properties of the noise, and in the ideal case the noise is assumed to be Gaussian. In order to detect RFI, they devise a test function $\mathcal{C}(s)$ to compute the error in quadrature between the ideal \mathbf{X}_0 and the measured \mathbf{X} . $\mathcal{C}(s)$ is weighted by the inverse of the covariance matrix \mathbf{R}_0 .

$$\mathcal{C}(s) = (\mathbf{X} - \mathbf{X}_0)^T \mathbf{R}_0^{-1} (\mathbf{X} - \mathbf{X}_0) \quad (2.17)$$

If $\mathcal{C}(s)$ is larger than some threshold value λ then the sample is determined to be corrupted by RFI. The values of \mathbf{X}_0 and \mathbf{R}_0 are determined by connecting a noise generator to the correlator prior to the observation.

Middelberg (2006) describes a semi-automated software package named PIEFLAG to identify and flag RFI in radio data from both single dishes and interferometers. PIEFLAG identifies RFI in corrupted channels by comparing it to a ‘clean’ reference channel which is relatively RFI free. Therefore the data are required to be bandpass calibrated prior to running PIEFLAG; the data must also be visually inspected to provide the reference channel to use. PIEFLAG uses two different algorithms to identify RFI in the time-channel plane of a single baseline, namely amplitude based thresholding and standard deviation (RMS) based thresholding. The amplitude thresholding works by calculating the median visibility amplitude for the reference channel, $x_{b,p}$ and the median of the differences to the median, $y_{b,p}$. This value is calculated independently for each baseline b and pointing p . The difference of each data point to the median $x_{b,p}$ is calculated and if the difference is larger than $kn \times y_{b,p}$ then the visibility is assigned a ‘badness’ value of k . This badness value does not translate directly into flags but is rather a book-keeping method to determine how badly a given visibility is affected by RFI. This method

is most effective when the astronomical signal is comparable to, or less than, the system noise. In the presence of very strong cosmic emission this algorithm is not very effective in determining which data are true outliers due to RFI.

The amplitude based thresholding is coupled with an RMS based method that works as follows - The RMS values for small sections of the reference channel (typically a few minutes) are calculated, and the median RMS value $z_{b,p}$ stored. This median value is taken to be representative of ‘good’ data, and is used as a comparison with the other channels. If the RMS in any section of any other channel exceeds $mz_{b,p}$ where m is typically 3, then that entire section is flagged. This algorithm does not assign badness values since groups of visibilities are considered together, and if their RMS is large enough it is indicative that the visibilities are legitimately corrupted by RFI. This algorithm performs well to identify bad data even in the presence of strong source signal. The rare cases where the RFI increases the amplitude without increasing the RMS will then be missed by this method.

In order to combine the results from the two methods, the data are again analysed on a sliding window basis. The sum of the badness of the visibilities within the window is calculated, and if this value exceeds 1 the entire section is then flagged. The flags can then optionally be extended, by counting the percentage of flagged data within a given window. If this fraction is large enough (e.g., > 0.15) then the entire window is flagged. In this manner, RFI is identified and flagged over all the baselines in a visibility data set. This algorithm is most effective when the RFI environment is not very active, and broadband RFI is absent. In the case that all the channels are affected, even the reference channel will be affected and hence the statistics will not represent uncorrupted data.

Offringa *et al.* (2010, 2012) also uses the visibility time-channel plane of a baseline to identify and flag RFI. The algorithms proposed by them are implemented in the software package AOflogger, which iteratively applies their algorithms to the visibility data in a fully automated manner. The data are analysed post-correlation, and can hence be applied to archival data as well. There is nothing specific to any one interferometer within the algorithms, and hence these can in principle be applied to any data obtained from any radio interferometer. Like the previously described method, this algorithm seeks to identify intermittent RFI. It is also sensitive to persistent, narrowband RFI or intermittent, broadband RFI. AOflogger is therefore sensitive to RFI that is localised in either time, frequency, or both.

They use a sliding two dimensional window over the time-frequency plane of a single baseline in order to estimate a smooth surface that represents the astronomical signal within the window and over the baseline. If the visibilities are given by

$V(\nu, t)$, the smooth surface is given by $\hat{V}(\nu, t)$. The requirement for smoothness of $\hat{V}(\nu, t)$ is because source structure is expected to be smoothly varying while RFI causes sharp edges in both the time and frequency directions. The smooth surface is fit within a given window and is used to divide the raw visibilities to better identify outliers. The sliding window is deemed to perform better than tiled windows (where the window is tiled along the time and frequency axes) since there are no ‘tile edges’ due to the fit in a sliding window. They additionally weight the samples in the sliding window according to the distance from the central pixel before calculating the mean value at the centre that determines the surface. Therefore the smoothed surface is given by -

$$\hat{V}(\nu, t) = \frac{\sum_{i=-N/2}^{N/2} \sum_{j=M/2}^{M/2} W_d(i, j)(W_F \odot V)(\nu_i, t_j)}{\text{weight}} \quad (2.18)$$

where the sliding window is of size $N \times M$ along the frequency and time axes respectively, and the \odot symbol represents element by element multiplication. W_d is the weighting function that weights data according to their distance from the centre of the window and W_F is the function that keeps track of data that have been flagged in a previous iteration. W_F is simply either 0 or 1 depending on whether the visibility has been flagged. The denominator is given by -

$$\text{weight} = \sum_{i=-N/2}^{N/2} \sum_{j=-M/2}^{M/2} W_d(i, j)W_F(\nu + i\Delta\nu, t + j\Delta t) \quad (2.19)$$

The above two equations are effectively convolutions, and can therefore be rewritten as -

$$\hat{V} = [(W_F \odot V) \otimes W_d] \oslash (W_f \otimes W_d) \quad (2.20)$$

where \otimes denotes a convolution, and \odot and \oslash denote an element by element multiplication and division respectively. If W_d is given by a two dimensional Gaussian, then the above equation can be separated in each dimension and calculated independently. Once the background surface has been estimated and divided out from the raw visibilities the RFI thresholds are calculated. The two methods discussed by Offringa *et al.* (2010) are related to the cumulative sum, or CUSUM method (Basseville & Nikiforov 1993). The first method is termed VarThreshold, and uses combinatorial thresholding to identify RFI corrupted samples. Typically

RFI thresholds are calculated by looking at each visibility in relation to the overall mean and standard deviation and if the value is above some cutoff it is flagged. Combinatorial thresholding expands on this idea. If two samples A and B are independently not above some cutoff threshold χ_1 , then they are considered together to check whether the combined sample $A \cup B$ is above some lower threshold χ_2 . If not, they are combined with another neighbour and the threshold is now at an even lower χ_3 and so on. As more samples are connected with each other, the threshold correspondingly is reduced. The decision about whether to sample a particular sample $R(\nu, t)$ along the frequency axis is determined by

$$\text{flag}_{\nu_M}(\nu, t) = |R(\nu + (i - j)\Delta\nu, t)| > \chi_M \quad (2.21)$$

where M is the number of samples in combination and i and j traverse $0 \dots M-1$ in a nested manner. The flagging rule for the time direction is similarly determined, and it is flagged if either rule is satisfied.

The second algorithm proposed is termed SumThreshold and is similar to VarThreshold in considering visibility samples in combination. The difference is that in this case, a sum of samples is used to determine whether the sequence should be flagged. Therefore individual samples that are below the threshold may still be flagged if their sum is large enough. The samples are classified in ascending order of thresholding *i.e.*, the lowest threshold is used to classify samples as RFI first before moving on to higher thresholds. At every iteration, samples classified as RFI will be left out of subsequent sums and replaced by the average threshold level. In this manner, the algorithm avoids over-flagging the data while still retaining sensitivity to outliers.

2.1.2 RFI excision

The process of partitioning the telescope signal into a cosmic component and the RFI, and subtracting the latter component, is known as RFI excision. It is not possible to excise the RFI in every case, however in the specific cases where the functional form of the RFI lends itself to subtraction these methods tends to work better than their counterparts which flag the visibilities.

Briggs *et al.* (2000) proposed the use of an ‘RFI reference signal’ to subtract the RFI from the received signal. This, like the previous method, assumes that the source and/or direction of the incident RFI is known and is persistent in time. This method also assumes the RFI is narrowband in nature. The RFI mitigation is performed post-correlation, and they use the phase and amplitude closure relations to subtract the interfering signal to recover the underlying cosmic

signal. Again, similar to the previous method discussed this requires the use of a reference antenna that will receive the RFI but will not receive the astronomical signal. They define four complex spectra

$$S_1(f) = g_A A_A + g_1 I + N_1 \quad (2.22)$$

$$S_2(f) = g_B A_B + g_2 I + N_2 \quad (2.23)$$

$$S_3(f) = g_3 I + N_3 \quad (2.24)$$

$$S_4(f) = g_4 I + N_4 \quad (2.25)$$

where A_A and A_B are the astronomical signals, N_i is the Gaussian noise in each signal, and I is the interfering signal. Each signal is modulated by the associated complex gain, which are the Fourier transforms of the input response functions. From the above four equations, the power spectra can be computed. Under the assumption of stochasticity of the input signal, and that the gain terms are constant over the time window, the power spectra will be of the following form -

$$\begin{aligned} P_1 &= \langle S_1 S_1^* \rangle \\ &= |g_A|^2 \langle |A_A|^2 \rangle + |g_1|^2 \langle |I|^2 \rangle + \langle |N_1|^2 \rangle \\ P_2 &= \langle S_2 S_2^* \rangle \\ &= |g_B|^2 \langle |A_B|^2 \rangle + |g_2|^2 \langle |I|^2 \rangle + \langle |N_2|^2 \rangle \\ P_3 &= \langle S_3 S_3^* \rangle \\ &= |g_3|^2 \langle |I|^2 \rangle + \langle |N|^2 \rangle \\ P_4 &= \langle S_4 S_4^* \rangle \\ &= |g_4|^2 \langle |I|^2 \rangle + \langle |N|^2 \rangle \end{aligned} \quad (2.26)$$

The terms within the angular brackets are averages over the time window . The above equations are the auto-correlation power spectra and are real valued. The complex cross power spectra for any given combination of the input data channel is of the form

$$\begin{aligned} C_{ij} &= \langle S_i S_j^* \rangle \\ &= g_i g_j^* \langle |I|^2 \rangle + g_i \langle I N_j^* \rangle + g_j^* \langle N_i I^* \rangle + \langle N_i N_j^* \rangle, \text{ for } i \neq j, j > 2 \end{aligned} \quad (2.27)$$

and

$$\begin{aligned}
C_{12} &= \langle S_1 S_2^* \rangle \\
&= g_A g_B^* \langle A_A A_B^* \rangle + g_1 g_2^* \langle |I|^2 \rangle + g_1 \langle I N_2^* \rangle + g_2^* \langle N_1 I^* \rangle + \langle N_i N_j^* \rangle
\end{aligned} \tag{2.28}$$

otherwise.

In order to subtract the terms of the form $|g_i|^2 \langle |I|^2 \rangle$ from Equation 2.26, the corresponding terms can be determined from the closure relations obtained from the cross spectra. The complex cross spectra from 3 data channels can be combined in the following form:

$$\begin{aligned}
|g_1^2| \langle |I|^2 \rangle &= \frac{g_1 g_3^* g_1^* g_4}{g_3^* g_4} \langle |I|^2 \rangle \\
&= \frac{C_{13} C_{14}^*}{C_{34}^*}
\end{aligned} \tag{2.29}$$

The above equation estimates the interferer power from the combination of the complex cross spectra across a triangle of baselines. The phases can be estimated from a similar equation, given by

$$g_1 g_2^* \langle |I|^2 \rangle = \frac{g_1 g_4^* g_2^* g_3}{g_3 g_4^*} \langle |I|^2 \rangle \tag{2.30}$$

$$= \frac{C_{14} C_{23}^*}{C_{34}} \tag{2.31}$$

Once the amplitude and phase of g_1 and g_2 are known, they can be subtracted from the astronomical signal (Equation 2.25). In this manner, by using an external reference antenna and by calculating the cross-power amplitude and phase, Briggs *et al.* (2000) manage to recover the underlying astronomical signal. It must be noted here that their method implicitly assumes that there is a single interferer in every channel, or at least that each interferer contributes identically to each channel. In the case that there are multiple interferers that vary across frequency, the above relations will no longer hold true. They also assume that the interfering signal stays constant within a single integration period, and the noise within that same integration period will be negligible.

The GMRT 150 MHz band has a very active RFI environment, and is particularly plagued by persistent, broadband RFI. If this RFI is not appropriately excised it could result in very large fractions of the data being flagged. Athreya (2009) proposes a post-correlation fringe subtraction technique that uses the behaviour of (spatially) stationary RFI to subtract broadband, persistent RFI from

an interferometer. A radio interferometer usually applies a ‘fringe-stop’ pattern that stops a source at the phase centre from fringing. This is achieved by adding an additional delay to one of the antennas of the baseline and this delay is a function of both time and baseline. The fringe stop frequency is given by:

$$\nu_F = -\omega_E U_\lambda(t) \cos \delta(t) \quad (2.32)$$

where ω_E is the angular velocity of the Earth’s rotation, δ is the declination of the source, and U is the instantaneous spatial frequency component (along the hour axis $HA = 0, \delta = 0$). A source of RFI that is spatially stationary with respect to the interferometer, and has been correlated, will pick up a fringe rate that is exactly the fringe stop frequency. If there are several such sources of RFI they will all add vectorially resulting in a net RFI signal. However this net signal will also fringe at exactly the fringe-stop frequency. Therefore even the presence of multiple sources of RFI does not affect the functional form of the RFI in the data. The effect of several different sources will be to reduce the time-scale over which the amplitude and phase of the signal is approximately constant.

In the presence of RFI, the observed visibilities take the form

$$V_{\text{OBS}} = V_{\text{TRUE}} + A e^{i[2\pi\nu_F(t)t - \Phi]} + N \quad (2.33)$$

where ν_F is the fringe stop frequency in Equation 2.32, and A and Φ are the amplitude and phase of the RFI in the baseline. If the amplitude and phase of the RFI remain constant over a large enough fraction of the fitting window, then the observed data can be fitted for estimates of A and Φ and subsequently V_{TRUE} can be estimated after subtracting the fit. Subtracting the RFI from the fit salvages the original visibilities and does not violate the closure relationships. This technique is applicable to any interferometer that implements fringe-stopping, and can be applied to archival data since the implementation is entirely in software. The primary limitation of this method is that it requires the fringe period to be small enough that a substantial fraction of the fringe is encompassed within the fitting window. The size of the fitting window is determined by the requirement that the fringe amplitude is approximately constant and a substantial fraction of the fringe falls within the window. This condition is not met if the fringe amplitude varies too rapidly within a given window, or if fringe phase varies too slowly within a window. Under these conditions the fringe fitting will fail.

The above method is somewhat related to the method proposed by Golap *et al.*

(2005). They map a source of (stationary) RFI to the celestial pole as a source at the pole will also be stationary with respect to the interferometer and will hence not have an intrinsic fringe rate. This RFI, mapped as a point source at the pole, can be subtracted out from the visibilities. They cast the observed visibilities in the following manner:

$$V_{ij}^{obs} = g_i g_j^* V^{source} + a_i a_j^* k_i k_j^* P \quad (2.34)$$

where V_{ij}^{obs} are the observed visibilities, V^{source} are the ‘true’ visibilities due to the cosmic source, g_i and g_j are the antenna primary lobe gains and a_i and a_j are the sidelobe gains for antennas i and j . k_i and k_j are phase-only propagation terms from the source of RFI to the antenna, and P is the RFI power. The strategy is to use two independent copies of the data set - the first copy is self-calibrated to the astronomical signal after subtracting the best model for the RFI affected visibilities. The second copy is self-calibrated for the RFI source at the pole after subtracting the best model for the celestial sources. This process is repeated iteratively until convergence, or some other stopping criteria. The gain solutions used for the self-calibration are specified as -

$$S = \sum_{ij} w_{ij} \left| V_{ij}^{obs} - g_i g_j^* V_{ij}^{model} - a_i a_j^* k_i k_j^* P \right|^2 \quad (2.35)$$

This method assumes that there is only a single source of RFI, with a quasi-constant amplitude and a consistent phase. In the presence of multiple interferers, although they will all be mapped to the celestial pole their amplitude will be varying and the above equation will not appropriately encapsulate that variation. However if these multiple interferers occupy different spectral channels then each channel can be treated independently with the above algorithm. Finally, the assumption is that the interferer contributes the same power to all the baselines of any given antenna. This is equivalent to saying that the interference signal is required to obey the closure relations. If the power received varies on a baseline basis, the closure relationships are violated and the above algorithm will no longer be effective. The algorithm of Athreya (2009) is not constrained by this requirement that the RFI source contribute the same antenna temperature to all antennas/baselines.

Pen *et al.* (2009b) propose another method to identify and excise broadband RFI at the GMRT. In a manner complementary to Athreya (2009) this technique relies on the inherent differences of the correlator output between a terrestrial

source of RFI and a celestial source in the absence of fringe tracking. Even though a terrestrial source of RFI typically has an intensity that varies with baseline, the delay between the two antennas of the baseline is fixed as a function of time. On the contrary, a celestial source contributes a time-dependent delay to the two antennas of a baseline as a function of time (as the Earth rotates). This dichotomy is exploited to identify terrestrial sources of RFI. They use the process of singular value decomposition (SVD) to excise the broadband RFI while retaining the cosmic signal. They assemble the visibilities into a two dimensional, rectangular matrix $V_i(\nu, t)$ where i is the baseline, ν is the spectral frequency and t is time. They find that broadband, terrestrial sources flicker synchronously over all frequencies and in all the baselines. Therefore the visibilities corresponding to the RFI are factorisable in the form

$$V_{RFI}(\nu, t) = \sum_{\alpha} L_i^{\alpha}(\nu) T_i^{\alpha}(t) \quad (2.36)$$

where $L_i(\nu)$ is called the ‘visibility template’ and $T_i(t)$ is called the ‘temporal template’. Each product $L(\nu)T(t)$ is a singular eigenvector of the matrix. Celestial sources produce a fringe with a unique frequency for each baseline and hence cannot be factored in a similar manner. Therefore under an SVD, celestial sources will produce very small eigenvalues. Such a decomposition can be performed when the eigenvectors are orthogonal. However there are reasons why these eigenvectors may not be perfectly orthogonal. For example, if the broadband RFI is caused by electrical arcing (which will tend to be higher at the peak of the AC waveform), this will be common to many baselines. In general this will not cause a problem if the (same) RFI has very different lags between the antennas of different baselines. The authors could eliminate only the strongest 100 SVD eigenvalues from a matrix with over 10^9 entries. This technique tends to occasionally misidentify celestial sources as RFI if the fringe rate is low enough and therefore adversely affects baselines with a small value of $|u|$. The authors have empirically determined that excising the strongest 100 eigenvalues has an adverse impact only on baselines that have $|u| < 10\lambda$. Increasing the number of eigenvalues excised increases the baseline length which is affected, while decreasing the number of eigenvalues used leaves a substantial amount of RFI in the visibilities.

Leshem *et al.* (2000) proposed a spatial filtering technique to project out the interfering signal from the cosmic signal. They use the measurement equation

$$\mathbf{R}(t) = \mathbf{\Gamma}(t)\mathbf{A}(t)\mathbf{B}\mathbf{A}^{\mathbf{H}}(t)\mathbf{\Gamma}(t) + \mathbf{A}_s(t)\mathbf{R}_s(t)\mathbf{A}_s^{\mathbf{H}}(t) + \sigma^2\mathbf{I} \quad (2.37)$$

where $\mathbf{\Gamma}$ are (diagonal) matrices describing the slowly varying antenna gains, $\mathbf{A}(t)$ is the array response matrix of the sources in the sky. \mathbf{B} is a diagonal matrix, with each entry containing the brightness of each source in the sky. \mathbf{A}_s is the array response to RFI, and \mathbf{R}_s is the correlation matrix of the RFI. $\sigma^2\mathbf{I}$ is the noise covariance matrix, assuming white noise. If the measured covariance matrix can be written as a sum of the covariance of the astronomical visibilities \mathbf{R}_v (the first term in Equation 2.37), white noise, and a single RFI of power σ_s^2 -

$$\mathbf{R} = \mathbf{R}_v + \sigma_s^2 \mathbf{a}\mathbf{a}^H + \sigma^2 \mathbf{I} \quad (2.38)$$

then a projection matrix \mathbf{P} can be defined such that

$$\mathbf{P} = \mathbf{I} - \mathbf{a}(\mathbf{a}^H\mathbf{a})^{-1}\mathbf{a}^H \quad (2.39)$$

and this projects out the interferer \mathbf{a} , since

$$\mathbf{P}\mathbf{a} = 0 \quad (2.40)$$

They describe a method to estimate the projection matrix \mathbf{P} from the eigenvalue decomposition of the measured covariance matrix \mathbf{R} , and therefore the knowledge of the interferers \mathbf{a} is not required. Since \mathbf{a} can potentially vary on very short timescales, the projection operation will have to be done on similar timescales to effectively excise the RFI from the visibilities.

Kocz *et al.* (2010) develop a method to project out the interfering signal using spatial filtering (in a similar manner to Leshem *et al.* 2000) optionally with a reference antenna. They state that the reference antenna itself does not provide any additional improvement to the RFI excision, unless the gain of the reference antenna is significantly higher than the primary antenna in the direction of the interferer.

The above algorithms are intended to identify and subtract unknown sources of RFI that have unknown temporal and frequency characteristics. Ellingson *et al.* (2001) describe a method to identify and remove the narrow band interference from the GLONASS (Hofmann-Wellenhof *et al.* 2007) navigation satellite. The GLONASS signal consists of two signals, a ‘‘coarse/acquisition’’ (C/A) signal that is effectively a sinusoidal signal which undergoes a phase shift of 0° or 180° every $1.96 \mu s$ which is known as the ‘chip rate’. Each group of 511 chips is a single logical unit, and is unchanging for any given GLONASS satellite. The second component

of the GLONASS signal is similar to the first, except that the chip rate is 10 times higher. Since the transmitted power for both components is approximately the same, the power spectral density of the second signal is roughly ten times less than for the first. Therefore, mitigating just the C/A signal can result in a significant improvement in sensitivity. This signal can be written as -

$$s_t(t, \omega_c) = c(t - \eta)b(t - \eta)e^{i\omega_c t + \phi} \quad (2.41)$$

where $c(t)$ represents the group of 511 chips, known as the ‘pseudorandom noise’ (PN) code. $b(t)$ represents the data bits, ω_c is the central transmission frequency and ϕ is the arbitrary phase introduced during modulation. $c(t)$ can only be +1 or -1, and changes at the chip rate of 511 kHz. η is the unknown time offset from the beginning of any given PN code, since the start times are not known prior to observation. The corresponding signal received at the telescope can be written as -

$$s_r(t) = G(t)P(t)s_t(t - \tau, \omega_c + \omega_d) \quad (2.42)$$

where τ is the propagation time, and ω_d is the frequency change due to the Doppler effect. $P(t)$ contains all the propagation losses etc., and $G(t)$ subsumes the antenna pattern response, the receiver response and the feed response. In addition to the above time-dependent variation, due to the nature of a radio telescope there will be several frequency dependent effects as well. Due to the finite bandwidth of the telescope the sharp phase jumps in the signal will be blunted. There may also be multipath scattering effects that cause a frequency dependent phase. However, the net effect of all these phenomena can be described as a single frequency dependent (but almost time independent) bandpass term $H(\omega)$. This $H(\omega)$ is slowly varying, and is not crucial to determining $s_t(t)$. However while subtracting the fitted signal from the data, the effect of $H(\omega)$ must be taken into account. Given the unknown propagation time τ , Equation 2.41 should be re-written in terms of $\mu = \tau + \eta$.

To effectively subtract the GLONASS signal from the observed data, μ and ω need to be accurately determined. In order to do this, a function J is defined as the correlation between $s_r(t)$ and the signal received at the telescope $x(t)$.

$$J(\mu, \omega_d) = \frac{1}{T_J} \int_t^{t+T_J} x^*(t)s_r(t - \mu, \omega_d)dt \quad (2.43)$$

where T_J is the correlator length. The correct values of μ and ω_d are those that maximise the function $J(\mu, \omega_d)$. If μ and ω_d are completely unknown the maximisation problem constitutes a two-dimensional brute-force search known as ‘acquisition’. This is essentially the same problem solved by GPS and GLONASS receivers when they are first switched on. Once the initial value for these parameters is determined, keeping track of them in time is fairly straightforward since they both change slowly. The parameters can be adjusted by sensing the drift in the correlation peak and making the appropriate changes. Once μ and ω_d are known, the GLONASS signal can be appropriately subtracted as a function of time from the astronomical signal. The authors show that the GLONASS signal is effectively removed from the data without affecting the underlying astronomical signal. The so-called ‘toxicity test’ to check whether this procedure is ‘toxic’ to the cosmic signal involved observing real data of an OH maser in the presence of the GLONASS signal. Post subtraction of the RFI, the maser signal remained unchanged. To further confirm that the astronomical signal was indeed untouched, two simulated lines (occupying only a single spectral channel) were added to the data and the subtraction of RFI was done subsequently. The authors find a > 20 dB reduction in the RFI while leaving the simulated signal mostly untouched. This technique is shown to be very effective in removing signals with a known signature and should be extensible to other satellites such as the GPS satellites. Apart from these satellites, there are also several modes of terrestrial communication that use a similar modulation technique to transmit signals which are also tractable using this cancellation strategy.

2.1.3 RFI mitigation

RFI mitigation strategies usually rely on specialised hardware (such as additional antennas, or a correlator with a high sampling rate etc.) to identify and remove the RFI corrupted signal prior to correlation and hence prevent the RFI from affecting the visibilities.

Buch *et al.* (2016) have implemented an algorithm that uses high speed sampling (Nyquist rate for 100s of MHz) at the uGMRT (Gupta *et al.* 2017). They use an appropriately scaled median absolute deviation (MAD) estimator to identify outliers in the time series of each antenna prior to correlating their signals to mitigate broadband RFI. This impulse in the time series data will result in broad-band RFI when the same is Fourier-transformed to obtain the bandpass. A typical GMRT ‘standard’ observation has integration time between 2s and 16s (Lal 2013). Flagging the RFI at the level of Nyquist samples will not result in any gaps in the data, even as the output will be RFI-free. The real-time algorithm

proposed by Buch *et al.* (2016) replaces the flagged samples by either a zero or with a value drawn from an equivalent noise distribution. They have not indicated which of the two options is to be preferred.

Barnbaum & Bradley (1998) proposed the use of an additional low gain antenna with a small collecting area that will be the ‘reference’ channel to detect RFI. The reference signal is then adjusted for differences in gain between the primary and reference antennas and subtracted from the primary signal. This results in the subtraction of RFI, and the recovery of the original signal prior to correlation of the visibilities. They defined the input signal to the radio telescope as

$$v(n) = s(n) + i_p(n) \quad (2.44)$$

where $v(n)$ is the input signal, and n indexes the time series of acquisition. $s(n)$ is the astronomical signal, and $i_p(n)$ is the input RFI that is seen by the primary antenna. The reference antenna will see the same interference with a different gain, $i_r(n)$. They then define the minimisation function $\epsilon(n)$, which is a difference between the input signal $v(n)$ and the scaled value of the interference, $Ai_r(n)$.

$$\epsilon(n) = v(n) - Ai_r(n) \quad (2.45)$$

The algorithm minimises $\epsilon(n)$ and $\epsilon(n - 1)$ in a least squares sense, and uses the result to update the coefficient A of the reference RFI signal. In this manner, the reference antenna adaptively modifies the coefficients to track the change in the amplitude of the received RFI as well as the difference in gains between the primary and reference antennas. This allows for a more robust subtraction of the RFI from the primary antenna and hence a more accurate recovery of the original astronomical signal of interest.

Hellbourg *et al.* (2014) provide a method that employs a reference antenna that tracks an interfering source, and uses a subspace tracking method known as the ‘power method’ to account for changes in the amplitude of the interferer. In a situation where there are M antennas in the array, and N interferers impinging on this array (with the condition $N < M$), They define a $M \times 1$ column vector called the ‘spatial signature vector’ (SSV) where each component in the vector corresponds to the signal received at one antenna in the reference array. They define the covariance vector \mathbf{R}

$$\begin{aligned}\mathbf{R} &= \langle x(t)X^H(t) \rangle \\ &= \mathbf{A}_r \mathbf{R}_r \mathbf{A}_r^H + \mathbf{R}_s\end{aligned}\tag{2.46}$$

where \mathbf{R}_r is the diagonal RFI covariance matrix, containing the RFI powers in the main diagonal and \mathbf{R}_s is the full $M \times M$ containing both the cosmic source as well as system noise. $\mathbf{A}_r = [a_{r1} \dots a_{rN}]$ is an $M \times N_r$ matrix, and each a_r is an RFI SSV. In calculating the overall instantaneous \mathbf{R} , the contribution of the cosmic source can be neglected, since at any instant their contribution is far below the noise. In order to continuously correct for the RFI, given that the RFI is not constant in time, they use an iterative procedure termed the ‘power method’ to calculate the dominant eigenvector in \mathbf{R} and subtract it from the covariance matrix. The assumption here is that the RFI is significantly stronger than the cosmic emission at any given instant.

The power method is a fairly straightforward method to determine the dominant eigenvector in a matrix. Given a random vector \mathbf{u}_0 , at every iteration i

$$\mathbf{u}_{i+1} = \mathbf{R}\mathbf{u}_i\tag{2.47}$$

$$\mathbf{u}_{i+1} = \mathbf{u}_{i+1}/\|\mathbf{u}_{i+1}\|\tag{2.48}$$

until $\|\mathbf{u}_{i+1} - \mathbf{u}_i\| < \epsilon$, where ϵ is some small real-valued number.

Once this method has converged, the dominant eigenvector can be projected out of the matrix recovering the cosmic signal. The paper also describes a method to calculate the N^{th} most dominant eigenvector, to track and excise several sources of RFI but goes on to demonstrate that excising the most dominant source of RFI improves the input signal significantly. Therefore that particular method has not been included here. To demonstrate the efficacy of the algorithm, they successfully excised data corrupted with RFI from the Galileo satellite (Hein *et al.* 2002). Finally, Mitchell *et al.* (2002) proposes a method similar to the above algorithms, using two reference antennas rather than one to more accurately track the changes between the interfering and cosmic signals. They continuously update the weights of the reference voltage to accurately subtract out the interfering signal from the primary antenna.

This section has briefly looked at several RFI mitigation, excision and flagging algorithms that have been implemented at various telescopes and in various software packages. Some algorithms are specialised (Ellingson *et al.* 2001, e.g.) and their applicability outside their intended scope is limited. There are other

techniques that are developed in a telescope agnostic manner (Athreya 2009; Middelberg 2006; Offringa *et al.* 2010) that have a more general applicability. However in general we find that no single strategy is capable of effectively removing all kinds of RFI. The various techniques that are developed usually work well in their intended regime, but may not be as effective when used at a different telescope, with a different RFI regime. In context of the GMRT, this is of consequence since there are only a handful of algorithms developed to excise and/or flag RFI at low frequencies (Athreya 2009; Buch *et al.* 2016; Pen *et al.* 2009b).

2.2 Direction dependent calibration errors

There are several existing algorithms that attempt to calibrate for Direction Dependent Errors (DDEs) like an azimuthally asymmetric antenna primary beam (Bhatnagar *et al.* 2008), a non-isoplanatic ionosphere (Intema *et al.* 2009), or a more general algorithm that solves for arbitrary amplitude and phase errors (Noordam & Smirnov 2010; Van Weeren *et al.* 2016). In this section, we discuss the methods which have been implemented in software to better contextualise the techniques presented in Chapter 4.

2.2.1 A-projection

A-projection (Bhatnagar *et al.* 2008) is an algorithm that corrects for the azimuthal asymmetry of the antenna primary beam, as well as a time-dependent pointing offset. Its measurement equation, following the Hamaker-Bregman-Sault notation (Hamaker *et al.* 1996), is

$$V_{ij}^{Obs} = M_{ij} \int M_{ij}^{Sky}(\mathbf{s}) I(s) e^{2\pi i \mathbf{s} \cdot \mathbf{b}_{ij}} d\mathbf{s} \quad (2.49)$$

where V_{ij}^{Obs} is the observed visibility vector (full polarisation), \mathbf{s} is the direction on the sky, \mathbf{b}_{ij} is the baseline length in units of wavelength, and I is the image (Stokes) vector. M_{ij} and $M_{ij}^{Sky}(\mathbf{s})$ are matrices, representing the direction independent and direction dependent gains respectively. The above equation is effectively the same as Equation 4.1, except the 2x2 Jones matrices of Equation 4.1 have been replaced by 4x4 Mueller matrices. The diagonal terms of the Mueller matrix correspond to the polarisation products, and the off-diagonal terms are the polarisation leakages. A Mueller matrix is formed from the outer product of two Jones matrices. A Jones matrix describing two orthogonal polarisations p and q can be written as

$$\mathbf{J}_i^{Sky}(\mathbf{s}) = \begin{bmatrix} \mathbf{J}_i^p & -\mathbf{J}_i^{pq} \\ \mathbf{J}_i^{qp} & \mathbf{J}_i^q \end{bmatrix} \quad (2.50)$$

where pq represents the polarisation leakage from p to q and qp is the reverse. \mathbf{J}_i is the response of antenna i to the sky brightness distribution. From this, the sky Mueller matrix is simply the outer product of the Jones matrices for two baselines.

$$\mathbf{M}_{ij}^{Sky}(\mathbf{s}) = \mathbf{J}_i^{Sky} \otimes \mathbf{J}_j^{Sky*}(\mathbf{s}) \quad (2.51)$$

where \otimes indicates the outer product.

All the terms within the integral in Equation 2.49 involve a Fourier transform from the visibility plane to the image plane and consequently cannot be solved for outside the imaging process. Traditional calibration algorithms, such as self-calibration (Pearson & Readhead 1984) perform the calibration step outside and separately from the imaging step. These solutions are encompassed by \mathbf{M}_{ij} in the above equation. Corrections for direction dependent effects cannot be solved independently from the imaging step, and have to be incorporated into the imaging algorithms.

The various details of antenna construction and geometry usually result in an azimuthally asymmetric antenna pattern; indeed, a prime focus or Cassegrain antenna is almost guaranteed to be asymmetric to some degree. For an altitude-azimuth (alt-az) mounted antenna, this voltage pattern rotates on the sky as the source is tracked; therefore off-axis sources will appear to have a time-variable flux in all Stokes parameters. Therefore \mathbf{M}^{Sky} is both time-varying and a function of baseline ij . Figure 2.1 shows the Stokes RR primary beam (first diagonal term of \mathbf{M}_{ij}^{Sky}) of the VLA at 1.4 GHz. Clearly, the main lobe shows some degree of asymmetry, with a much larger asymmetry in the first and second sidelobes.

In addition to the beam shape asymmetry, every antenna can have a time-dependent pointing error, or pointing jitter. For most telescopes (particularly at low frequencies) the pointing error is only a few percent of the primary beam size. However the net effect is to introduce a time-dependent linear phase on the source structure. The magnitude of the phase is different for every antenna, and hence for every baseline. In order to correct for both the asymmetric primary beam as well as the pointing error, they define an operator \mathbf{E}_{ij}^P , which is the autocorrelation of the antenna illumination pattern for a given polarisation product P . In the presence of small antenna pointing errors (here small is taken to mean small compared to the half-power width of the primary beam) the phase due to the pointing error is




Figure 2.1: The VLA 1.4 GHz Stokes RR beam, reproduced from Bhatnagar *et al.* (2008).

subsumed into E_{ij}^P , but can be separated from the terms that are common for all the antennas.

$$E_{ij}^P = E_{ij}^{PO} f(\phi_i - \phi_j) e^{i(\phi_i + \phi_j)} \quad (2.52)$$

where E_{ij}^{PO} is the antenna illumination pattern for an ideal telescope, $f(\phi_i - \phi_j)$ is the decorrelation suffered by the signal due to the pointing offset, and ϕ_i and ϕ_j are the phases acquired due to a pointing offsets for antennas i and j . Finally, E_{ij} is related to the measurement matrix M_{ij}^{Sky} as

$$E_{ij} = FT[M_{ij}^{Sky}] \quad (2.53)$$

where FT is the element-by-element Fourier transform operator. The operator E_{ij} is incorporated into the iterative deconvolution algorithm typically used, which is outlined as

1. Initialise the algorithm with either an *a priori* model, or zero.
2. “Major cycle” - Calculate the residual visibilities $V^{Obs} - V^M$ in the forward direction, employing \mathbf{E}_{ij} . Compute the residual image using \mathbf{E}_{ij}^\dagger as an interpolation and gridding operator for re-sampling the visibilities on to a regular grid.
3. “Minor cycle” - Update the model image by using some operator
4. Go to 2 until convergence is achieved.
5. Smooth the final images with the given resolution and add back the image residuals.

The “forward” step in the deconvolution algorithm refers to obtaining the model visibilities given a model image, and the “backward” step involves propagating the residual visibilities to the image plane. The forward step is given by

$$V^M = \mathbf{E}_{ij}^P \mathbf{F} I^M \quad (2.54)$$

where \mathbf{F} is the Fourier transform operator. The backward step is given by

$$I^R = \det(\overline{\mathbf{F}^\dagger [\mathbf{E}_{ij}^P]})^{-1} \mathbf{F}^\dagger V^R \quad (2.55)$$

The use of a gridding operator \mathbf{E}_{ij}^P that is ‘aware’ of the beam asymmetry allows for the errors due to an azimuthally asymmetric primary beam and a time-dependent pointing error to be accounted for and corrected. This method requires a detailed two dimensional model of the antenna primary beam that is not readily available for many telescopes.

2.2.2 Source peeling

The other major source of error (particularly at lower frequencies) is the ionosphere. The ionosphere contains partially ionised gas, and the electron density in the ionosphere varies as a function of time and position. Depending on the spatial scale of homogeneity of the ionosphere and the geometry and field of view (FoV) of an array there can be 4 different regimes through which the ionosphere can affect an observation (Intema *et al.* 2009; Lonsdale 2005) (see Figure 2.2). Case 1 and 2 can be corrected by the standard self-calibration procedure. For cases 3 and 4, the gradients within the ionospheric patches will contribute different phase delays

across the FoV of the instrument and hence a single calibration solution for the entire sky will no longer be valid. The most general case is regime 4, where there can be large variations of ionospheric phase across the FoV.




Figure 2.2: The four Lonsdale regimes of ionospheric phase calibration, depending on the relative spatial scales of ionospheric fluctuations, the array geometry, and the antenna field of view. Image reproduced from Intema *et al.* (2009).

The first proposed solution to this problem was by Noordam (2004), using a technique known as *source peeling*. The principle of peeling is fairly straightforward - Given a bright off-axis source, self-calibration solutions are found in the direction of the source. The contribution of this calibrated source is then subtracted out from the rest of the visibilities, and the procedure repeated for the next brightest source until the desired sensitivity is reached. In practice this is usually achieved in one of two ways:

1. The CLEAN components of the entire field are subtracted from the visibilities, and only the components of the strong source are added back. Self-

calibration is then performed using only the components corresponding to the bright source, and the calibrated source is then removed from the visibilities.

2. The visibilities are rotated toward the bright source, and self-calibration is performed using only the components of this source. The calibrated source is then subtracted from the visibilities, and the visibilities are rotated back to the original phase centre.

This process is quite successful in reducing the errors associated with a non-isoplanatic ionosphere. The largest constraint with the source peeling algorithm is that the peeling solutions are only valid for the handful of directions for which it is determined, leaving the rest of the field largely uncorrected. There will however be an improvement throughout the field from peeling only the brightest sources since the errors associated with those will reduce systematic patterns throughout the field.

Field based calibration (Cotton *et al.* 2004) was developed to calibrate the ionosphere of the VLA 74 MHz all sky survey VLSS (Cohen *et al.* 2007). Broadly, the procedure uses the source peeling algorithm to peel the brightest source within a given polyhedron facet (Cornwell & Perley 1992; Perley 1989), and applies the measured phase solution over all the sources within that facet. This method is applicable in Lonsdale regime 3 (see Figure 2.2), where the gradient of ionospheric phase across the entire array can be approximated to a linear term since the line of sight through the ionosphere for each of the antennas is very similar. Defining the constancy of ionospheric phase to the imaging facet is not necessarily physically appropriate, and is a drawback of the approach.

The SPAM algorithm (Intema *et al.* 2009) uses the peeling solutions to fit for an ionospheric phase screen that encompasses the entire sky. SPAM assumes that the ionosphere is a thin (two-dimensional) layer of electrons at a fixed height above the surface of the Earth. This in turn means that the ionosphere effectively acts as a phase screen, which modifies the incoming phase both as a function of time and direction. The total phase rotation passing through a patch of the ionosphere is given by

$$\phi_{ion}(\mathbf{p}, \xi) = \frac{\phi_{ion}(\mathbf{p})}{\cos(\xi)} \quad (2.56)$$

where \mathbf{p} is the location of the ‘pierce point’, *i.e.*, the location on the phase screen that the incoming wave is passing through, and ξ is the zenith angle. The value of ϕ_{ion} is directly measured through the peeling scheme discussed above. Source

peeling samples the phase screen over a non-uniform set of points, which have to be interpolated to obtain the ionospheric model for the FoV. The interpolation is performed using a physically motivated set of basis vectors known as Karhunen-Loeve (KL) vectors (Van Der Tol *et al.* 2007) which describe the physical scales of variability of the ionosphere more accurately than the more commonly used Zernike polynomials.

The process of fitting the observed phases to a larger phase screen model allows for the phase solutions to be generalised to other regions of the map without assuming that the same phase solutions can apply over a large fraction of the sky. Since the solutions here can vary on a baseline to baseline basis, similar to that discussed in Section 2.2.1, the calibration can only be applied during the imaging process. At the moment SPAM is implemented only in AIPS (Greisen 2003) and applies these corrections while imaging with the polyhedron faceting algorithm (Cornwell & Perley 1992; Perley 1989). SPAM allows for the ionospheric phase to vary within a facet, by taking care of it during the imaging process. SPAM is more robust than a straightforward peeling algorithm because it captures more accurately the varying ionospheric phase across the FoV.

Facet calibration is a more recent development (Van Weeren *et al.* 2016) which attempts to carry out a direction dependent calibration for the entire field of view at the LOFAR. The principle is very similar to both Field-based calibration as well as SPAM. We do not discuss all the details of the algorithm here, as many of the calibration steps are very particular to the LOFAR and are not applicable to other interferometers (such as clock synchronisation between the remote stations and the core). The direction dependent calibration portion of the algorithm is discussed below.

The assumptions made by this algorithm are that (a) The only errors in the direction-independent calibration arise from ionospheric phase and antenna primary beam, (b) the beam is slowly varying with time, (c) the direction dependent errors vary smoothly with direction and (d) the ionospheric phase varies as ν^{-1} . In order to perform a direction dependent phase calibration, first the directions need to be identified. This process is similar to that employed for peeling. The brightest sources in the sky are calibrated first, and then subsequently the fainter sources - this is so that the errors from the brightest sources do not contaminate the calibration solutions for the fainter sources. Given that the assumption is that the DDEs vary smoothly across the sky, the map is broken up into facets which are determined by Voronoi tessellation (e.g., Okabe 1992). Voronoi tessellation is perhaps more appropriate than polyhedron faceting since the distribution of sources across the field is rarely uniform. The tessellation is performed by centring the Voronoi facets on the brightest sources identified, and the phase calibration solutions are

determined simultaneously for all the sources within the facet. The calibration is performed after subtracting the direction-independent calibrated sources from the visibilities, and adding back only the sources within a particular facet. Once a particular facet is calibrated, the sources are subtracted from the visibilities and the next facet is added. In this manner, the entire sky is iteratively calibrated resulting in an overall directionally dependent calibration solution. Initially a few rounds of phase-only self-calibration are performed using a small solution interval to capture the potentially rapid variability of the ionosphere. Following that, an amplitude calibration is performed independently for each polarisation with a larger solution interval. This is because the amplitude errors are expected to be caused by the primary beam which varies on much longer timescales than the ionosphere.

Chapter 3

The IPFLAG algorithms

3.1 Introduction

As mentioned in Chapter 1, a single RFI mitigation strategy is rarely successful, as different sources of RFI leave different signatures in the data depending on their spatial, temporal and frequency characteristics. Persistent RFI causes strong fringes to appear on the baseline, and the amplitude and phase of these fringes can vary as a function of time and frequency. Broadband RFI can affect entire baselines, causing fluctuations in the associated antenna gains. Intermittent RFI, which tends to be more localised in time and frequency, look like 'hotspots' in the visibility data and cause large scale ripples in the image plane. It is difficult to recover the underlying visibilities in this case since there is no analytical expression to describe the signature of intermittent RFI. Although there has been some work done to identify and localise sources of RFI (Pen *et al.* 2009a; Sokolowski *et al.* 2015), there exists no general framework which one can use to classify and excise this kind of RFI.

This chapter deals with techniques that were developed to identify and flag intermittent RFI, with a higher sensitivity and better efficacy than has been previously achieved. We use two different parameters spaces to identify and flag RFI, namely the (median-binned) visibility plane and the visibility time-channel plane of the baselines.

The work presented in this chapter has been accepted for publication by the *Astronomical Journal*.

3.2 Description of the algorithms

The measured visibilities in any polarization in the presence of RFI signal can be written as

$$V^o = G_i G_j^* (V_{ij}^{\text{sky}} + V_{ij}^{\text{RFI}} e^{if_{ij}}) + \eta_{ij} \quad (3.1)$$

where ij are the antenna indices, V_{ij}^{sky} are the visibilities due to the cosmic source, $e^{if_{ij}}$ is the fringe-stop function (to stop the fringe of a source at the phase centre), V_{ij}^{RFI} is the effective visibilities of the RFI after accounting for the differences in the antenna gains toward the RFI and the pointing centre and will be proportional to the flux of the source. That is, $V_{RFI} = k S^{RFI} (G_i^s G_j^{s*}) / (G_i G_j^*)$ where G^s are the gains of the sidelobes through which the RFI are detected, and G are the main lobe gains. In order to self-calibrate the data, the quantity $S^{RFI} (G_i^s G_j^{s*})$ has to be constant for all antennas, which will not be true for RFI coming through different sidelobes in different antennas. Finally, η_{ij} is additive (Gaussian) noise in the system. G_i and G_j are the complex antenna gains, which can be affected by strong RFI even if it remains uncorrelated. This equation is true for any polarisation. $e^{if_{ij}}$ will stop the fringe of a source at the phase centre, but will introduce a corresponding fringe on a stationary terrestrial source like RFI. Previous work has used the form of $e^{if_{ij}}$ to excise RFI while recovering the original visibilities (Athreya 2009). V^{RFI} can be orders of magnitude larger than V^{sky} and can vary with time, frequency and baseline. The second term in the above equation causes poorer self-calibration solutions and introduces systematic errors in the image plane. The errors during self-calibration arise because in general correlated RFI will violate the closure relationships (Pearson & Readhead 1984). Since RFI typically contributes a different effective temperature to each antenna, the correlated power of the RFI will vary across baselines of a single antenna. Therefore the RFI amplitude and phase are not always separable into antenna based terms, and therefore violate the closure quantities.

We focus in this chapter on RFI that is localised in time and frequency in the visibility plane. RFI localised in such a manner appear to be ‘hotspots’ in the visibility data, and consequently cause large scale ripples in the image plane. We explore two visibility parameter spaces to identify and flag RFI corrupted visibilities *viz.* the median binned UV plane of the entire interferometer, and the time-channel plane of a single baseline. The two algorithms described here are called GRIDflag and TCflag respectively, and are combined into an integrated flagging package called IPFLAG.

3.2.1 RFI flagging in the binned UV plane

Individual baselines of an aperture synthesis radio interferometers sample elliptical tracks in the visibility plane (Thompson *et al.* 2008). However, the tracks of different baselines are usually irregularly (and sparsely) distributed over the UV plane. Typically during imaging, these visibility points are interpolated onto a regular grid in order to use the Fast Fourier Transform (FFT) algorithm. This ‘gridded’ UV plane is fundamental to almost all imaging algorithms, and the sampled UV cells within the grid define the UV coverage of the observation. The size of each individual UV cell should be small enough to satisfy the Nyquist limit to avoid aliasing across the largest area being imaged. A single baseline can contribute one (or more) visibilities to a single UV cell - either within a short time interval, or over multiple epochs. Other baselines may also contribute visibilities to the same UV cell – at the same time in redundant arrays (Westerbork), or nearly redundant arrays (shorter baselines of the GMRT or VLA), or at different times (longer baselines of GMRT or VLA). The visibility samples within a UV cell measure approximately the same celestial information, but will differ in the RFI environment that they encounter (since they sample the cell at different times, or simultaneously but from physically different RFI environments). We propose to use this dichotomy to identify and flag RFI affected visibilities.

We define a quantity termed the ‘UV-bin’ which is similar in size to the UV cell. The basic assumption of these algorithms is that in the absence of RFI, the differences in values of all the visibilities falling within a single UV-bin is dominated by the system temperature rather than the source structure. Thus one can assume that all the visibilities within a single UV-bin represent the same spatial Fourier component and any difference between them well in excess of the system temperature can be ascribed to RFI. Further, since the visibility function is expected to be locally smooth for any realistic sky intensity distribution, data from adjacent UV-bins can be combined to calculate statistically secure thresholds to identify RFI. Thresholds calculated in this manner will be more resilient to large outliers in any given UV-bin. This is even more appropriate when applied to the residual visibilities (which are obtained after subtracting as many sources as are detected even in the presence of RFI).

These routines are meant to be used in addition to the standard procedures used for radio interferometric data reduction using packages such as CASA/AIPS (Greisen 2003; McMullin *et al.* 2007). These standard ‘recipes’ typically consist of pre-calibration flagging, several rounds of iterative imaging and self-calibration followed by (often manual) residual visibility flagging. These procedures tend to only take care of the strongest RFI in the data. A typical GMRT 150 MHz

observation analysed using this recipe produces images with an RMS noise of 1.5 - 5 mJy/beam. Applying the RfiX algorithm (Athreya 2009) to the raw visibilities eliminates much of the broadband RFI, and routinely results in images with an RMS noise of 0.7 - 1 mJy/beam. Finally, adding the routines described in this chapter results in images with an RMS noise of 0.5 - 0.7 mJy/beam, which are among the most sensitive achieved with the GMRT at 150 MHz.

The recipe for the scheme is as follows:

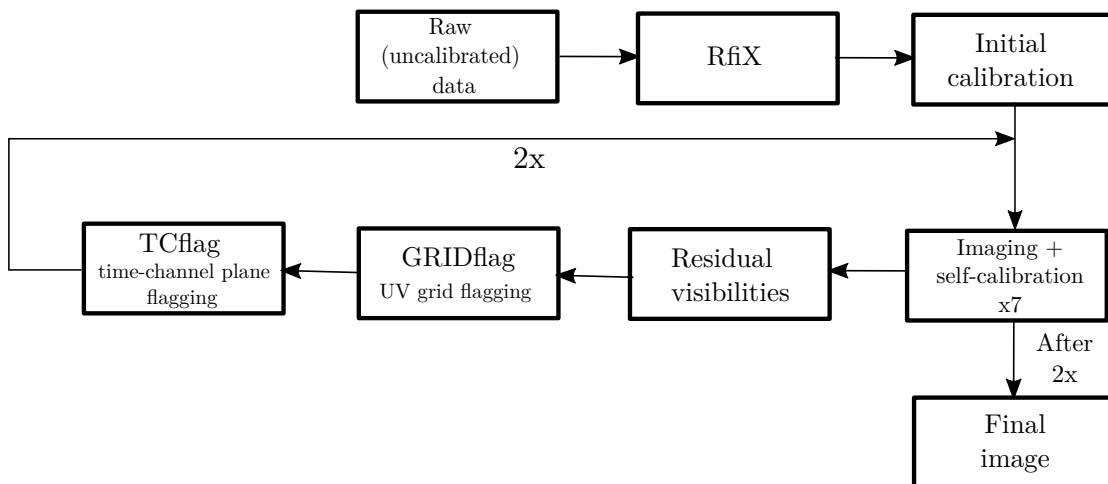


Figure 3.1: Flowchart illustrating the data reduction procedure which includes GRIDflag and TCflag algorithms.

1. Apply the RfiX algorithm to eliminate broadband, persistent RFI and use the standard CASA/AIPS procedures to calibrate, image, and flag the data to obtain the residual visibilities.
2. Allot the visibilities into bins in the UV plane, and calculate the median value and robust standard deviation (with respect to the median) of all the visibilities within each UV-bin. These bins are roughly the same size as the cells used for gridding by imagers.
3. Partition this median binned UV plane into a series of annuli for calculating RFI thresholds as a function of UV distance. Since both source signal and RFI in the residual visibilities tend to decrease as a function of radial distance partitioning the UV plane in such a manner naturally conforms to this change. The choice of width is not critical and is decided by the competing requirements of being sensitive to the changing standard deviation as a function of UV length and having sufficient UV-bins with an annulus.

4. Use the distribution of medians over all the UV-bins within an annulus to exclude highly contaminated bins from the determination of RFI thresholds. The rest of the UV-bins are used to determine a smoothed median background in all populated bins. This background will be a function of UV radius.
5. The smoothed median background and the local standard deviation is used to identify RFI affected data within each bin through any scheme based on standard deviation (*i.e.*, σ) thresholds. The local standard deviation may be derived either from data in the particular UV-bin or from the nearest neighbour bins. The same $n\text{-}\sigma$ threshold may be used over the entire UV plane, since independently calculating the thresholds within each annulus already accounts for the expected variation of standard deviation as a function of UV length.
6. These flags are applied to the original, un-smoothed, and un-binned data, and the entire process of imaging and self-calibration is repeated.

This procedure preserves the UV coverage to a large extent, because (a) the flagging in each cell is processed separately, and in most cases at least a few visibilities survive the process and (b) the combination of using the within-cell degeneracy and analysis within an annulus allows a more finely tuned flagging threshold per cell and per annulus which avoids the large scale rejection of data. Finally, this procedure allows for a smooth variation of standard deviation even with an annulus. The variation of standard deviation may be compensated for by differential weighting of the visibilities, although we have not yet implemented this.

Since only half the visibilities are recorded (the other half are a Hermitian conjugate), all the data have to be appropriately conjugated to lie in the same half plane. The UV-bins will also have to be extended to a few bins into the other half plane to obtain statistics at the edge. This procedure was applied separately and successively to the Stokes RR and LL and stokes V polarisations.

3.2.2 RFI flagging in the baseline time-channel plane

As noted earlier, RFI is most easily identified in the residual visibilities, and one can look for corrupted visibilities in the time-channel plane of individual baselines. However, the (strong source subtracted) residual time-channel plane often has multi-component sinusoidal patterns superposed on the residual source signal. These fringes may be due to either improper subtraction of source flux, or RFI. There are several factors that may result in an incorrect subtraction of the

source flux, such as an azimuthally asymmetric antenna primary beam, a time dependent pointing error, uncorrected gain fluctuations and other baseline-based errors that affect the imaging process. These residual visibility fringes inflate the sigma thresholds above which RFI localised in both time and frequency may be detected. Therefore, we developed a method to eliminate these fringes prior to estimating the cutoff thresholds.

The scheme to identify RFI in the residual time-channel plane is as follows:

1. For each polarization and baseline, take a two dimensional Fourier transform (along the time and frequency axes) of the complex visibilities within a window to obtain the group delay - delay rate (GD-DR) plane. The window has to be large enough to cover a substantial fraction of the fringe, while being smaller than the period over which the amplitude of the fringe changes.
2. In the GD-DR plane, iteratively sigma clip all components above some threshold, thereby eliminating the corresponding fringes in the time-channel plane.
3. Inverse Fourier transform to obtain the fringe-free time-channel plane.
4. Identify RFI corrupted visibilities using any thresholding algorithm (such as sigma clipping) in the fringe-free time-channel plane.
5. Apply the flags identified in the previous step to the original data, and restart the process of imaging and self-calibration.




Figure 3.2: (a) Comparison of the binned UV plane before and after applying GRIDflag for the 3C286 field. Brighter regions indicate visibilities corrupted by RFI.

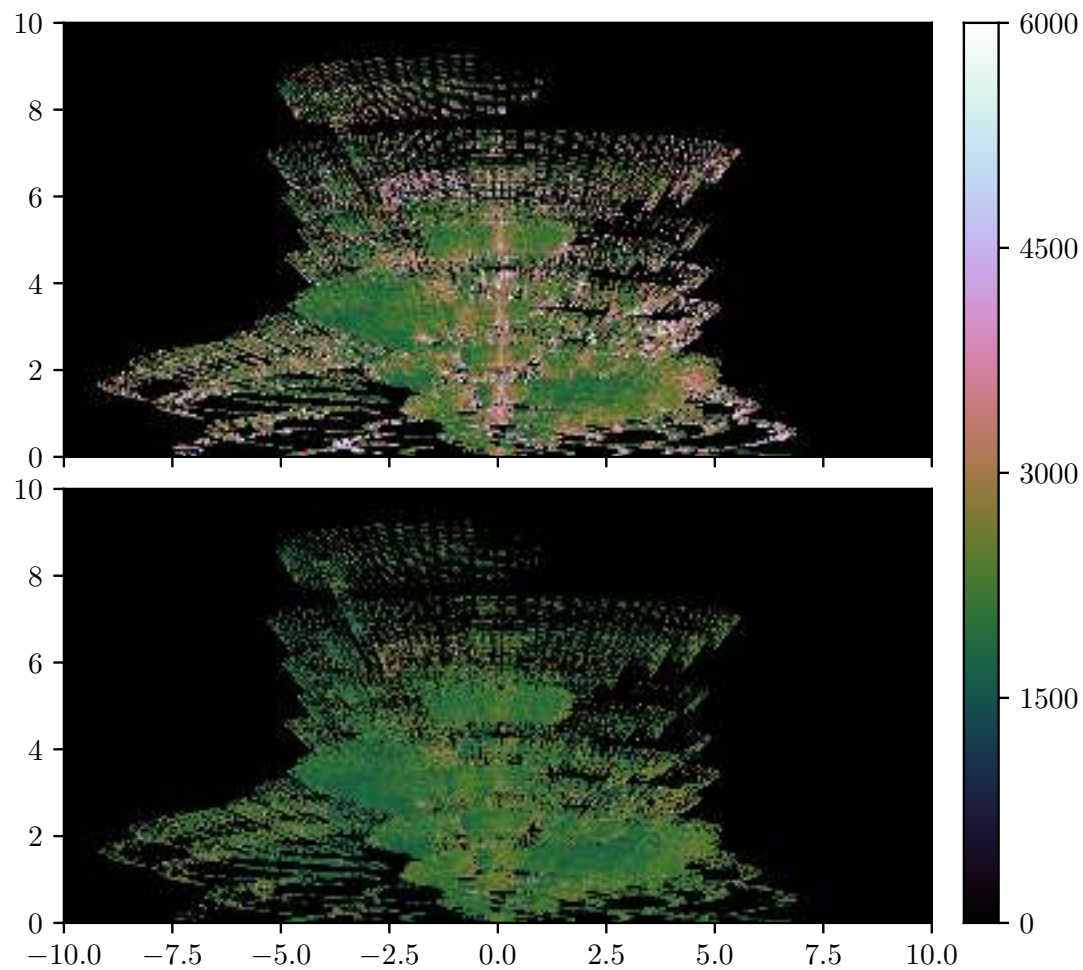


Figure 3.2 continued ... (b) The median binned UV plane before and after GRID-flag for the VIRMOSC field.




Figure 3.2 continued ... (c) The median binned UV plane before and after GRID-flag for the J1453+3308 field.




Figure 3.2 continued ... (d) The median binned UV plane before and after GRID-flag for the J1158+2621 field.




Figure 3.2 continued ... (e) The median binned UV plane before and after GRID-flag for the A2163 field.

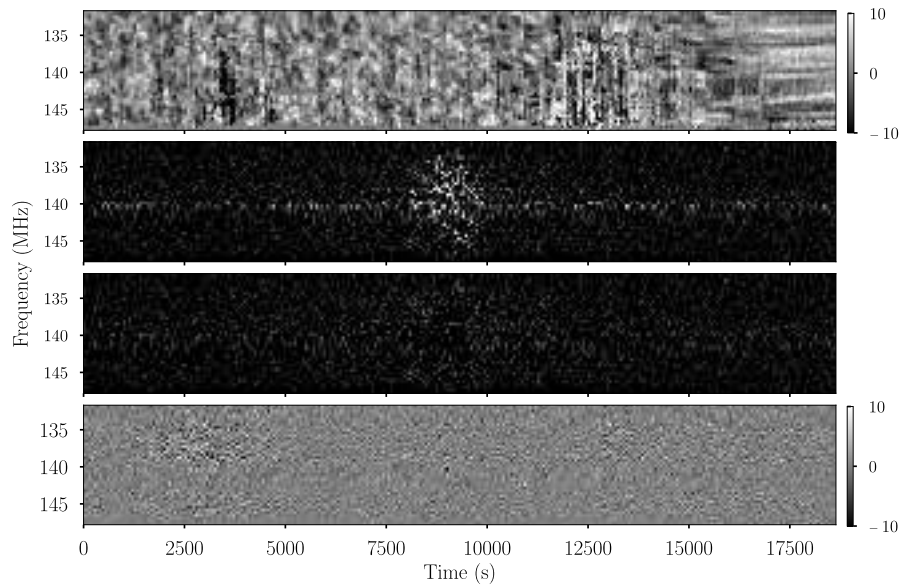





Figure 3.3: (a) Illustration of the defringing process in TCflag. The 4 panels in each figure, from top to bottom, are (i) fringes in the time-channel plane, (ii) its Fourier transform, (iii) fringe peaks clipped in the Fourier plane, and (iv) inverse transformation to obtain the fringe-free time-channel plane.



(b) Figure 3.3
continued ...



(c) Figure 3.3
continued ...



(d) Figure 3.3
continued ...

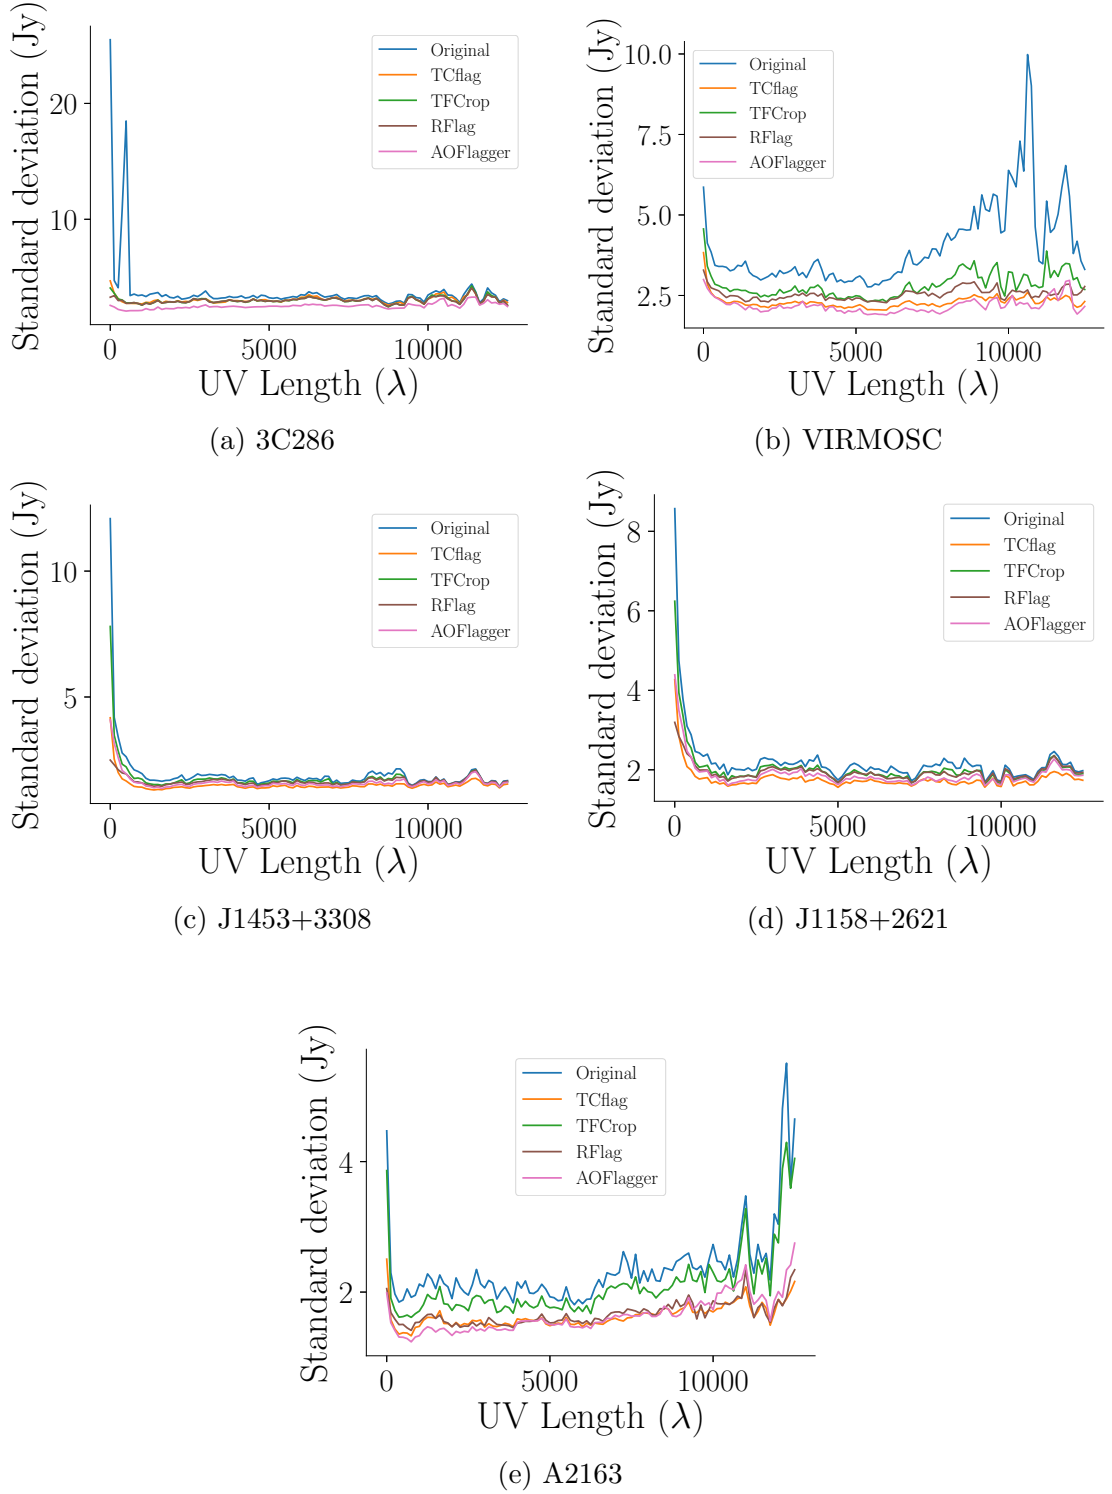


Figure 3.4: A plot of RMS noise in the visibilities as a function of UV distance. TCflag is competitive with the other tested flaggers, and performs marginally better than the rest in three of the 5 cases. In the other two cases, AOFlagger performed better. The relative difference between the median visibility RMS values after flagging is on average $\sim 5\%$. Therefore we can only claim that TCflag is competitive with the other flaggers.

This procedure works because of the difference between the GD-DR plane signatures of source structure and localised RFI. RFI which is localised in time and frequency will be dispersed over several pixels in the GD-DR plane. Conversely, the sinusoidal fringe due to source structure variation will show up as compact peaks, which can then be identified and clipped out. This leaves a fringe-free time-channel plane, with RFI. Figure 3.3 illustrates this process for fringes in different baselines. In each case, the top plots show (from top to bottom) the fringe in the time-channel plane of a baseline, the amplitudes of two dimensional Fourier transform of the time-channel plane, the fringe peaks clipped in the Fourier plane, and the inverse transformed, fringe-free time-channel plane. It is not possible to entirely remove signal from the celestial source using this method due to the issue of spectral leakage (see Appendix A). Briefly, spectral leakage causes a “smearing” of components in the Fourier domain in cases where the input signal does not have an exactly integral number of periods within the window used. This is significant because in the presence of several different frequencies, this ‘smearing’ from each frequency bin will add to the noise floor (in addition to the thermal noise). Therefore signal that contributes weakly to a particular bin may be below the noise floor and hence cannot be accurately subtracted. However, for the purposes of calculating statistics without the bias of the signal it does not matter too much if all the fringes are removed exactly, but rather that the bulk of the fringes are removed so that more sensitive thresholds can be calculated.

Since the fringes in the residual visibilities are primarily due to the incorrect subtraction of source flux, we tried window sizes of 5 - 20 minutes with success. We used a window size of 10 minutes for all the sources, since it also matches the scan breaks in our data. The size of the window does not need to be finely tuned, and the efficacy of this algorithm does not depend critically on this parameter. If the amplitude of the fringe changes significantly within the window, it will smear out the signal in the GD-DR plane. To avoid this, one can also run this procedure using several window sizes in decreasing succession. This method (of Fourier domain clipping) is especially advantageous over others while dealing with fringes due to improperly subtracted sources. The process of the FFT and IFFT do not introduce any additional errors into the system - the numerical error associated with this process is only of the order of ‘double precision’ numbers, *i.e.* $\sim 10^{-12}$. TCflag may be applied to the data at any stage of analysis, since the source fringes are filtered out before flagging. This should not result in any artefacts as only the flags of the original data are being modified.

Figure 3.4 shows a comparison between the time-channel plane flagging algorithm (TCflag) described here, with the standard flagging routines in CASA (TFCrop and RFlag), and with AOflogger. From the plot, all the algorithms per-

form to a similar level on every dataset. In 3 of the 5 data sets TFCrop marginally outperforms the others, whereas in the other two cases AOflogger results in the lowest median RMS value. Therefore we can only claim that TCflag is at least as competitive as the other standard flaggers.

3.3 Observations and parameters

All observations were conducted using the 150 MHz band of the GMRT, with a bandwidth of 16 MHz and a spectral resolution of 62.5 kHz or 125 kHz. The data was recorded with an integration time of 2s. The analysis procedure is outlined in Figure 3.1.

We used the following parameters in IPFLAG while flagging on the residual visibilities:

1. UV-bin size (GRIDflag): 10λ , from the field of view at 150 MHz.
2. Smoothing window for median visibility background (GRIDflag): 5×5 bins.
3. UV-bin annuli width (GRIDflag): 3, 3 and $6.5 \text{ k}\lambda$
4. Fourier transform window size (TCflag): bandwidth \times 10 min.
5. Fourier peak detection threshold (TCflag): $3 \times$ RMS noise
6. RFI threshold (both): $3 \times$ RMS noise

After flagging, we transferred the flags from the data processed with IPFLAG back to the original visibilities and repeated the entire process of imaging, self-calibration, residual flagging, and flag transfer. We then ran one more round of imaging and self-calibration (excluding residual flagging and flag transfer) to obtain the final images. The final images cover $6.25^\circ \times 6.25^\circ$ with a pixel size of $4.5'$. Source extraction was done using the pyBDSF source finder (Mohan & Rafferty 2015), and we used identical parameters across all the images. The source finding ignored extended emission, as well as sources that were below 7σ .

In order to calibrate the flux scale of each of the target fields, standard CASA/AIPS routines were used, employing the Scaife-Heald coefficients (Scaife & Heald 2012) for all the target fields. This process resulted in flux scale discrepancies of up to 15%. This becomes significant while comparing the absolute noise in an image, although it does not affect measurements of the relative change in noise after the application of RFI flagging procedures. We anchored our flux density scale to TGSS-ADR (Intema *et al.* 2017) using the sources common between the two observations. The fractional discrepancy in the scaling of the strong ($> 100 \text{ mJy}$)

and faint (< 100 mJy) sources is shown in Figure 3.5. The reason for such a differential change is unclear, although it may have something to do with the CLEAN bias (Cohen *et al.* 2007; Condon *et al.* 1998) that is known for interferometers. This fractional change ranges from -0.5% to 9% between the sources above and below 100 mJy over all the observed fields.

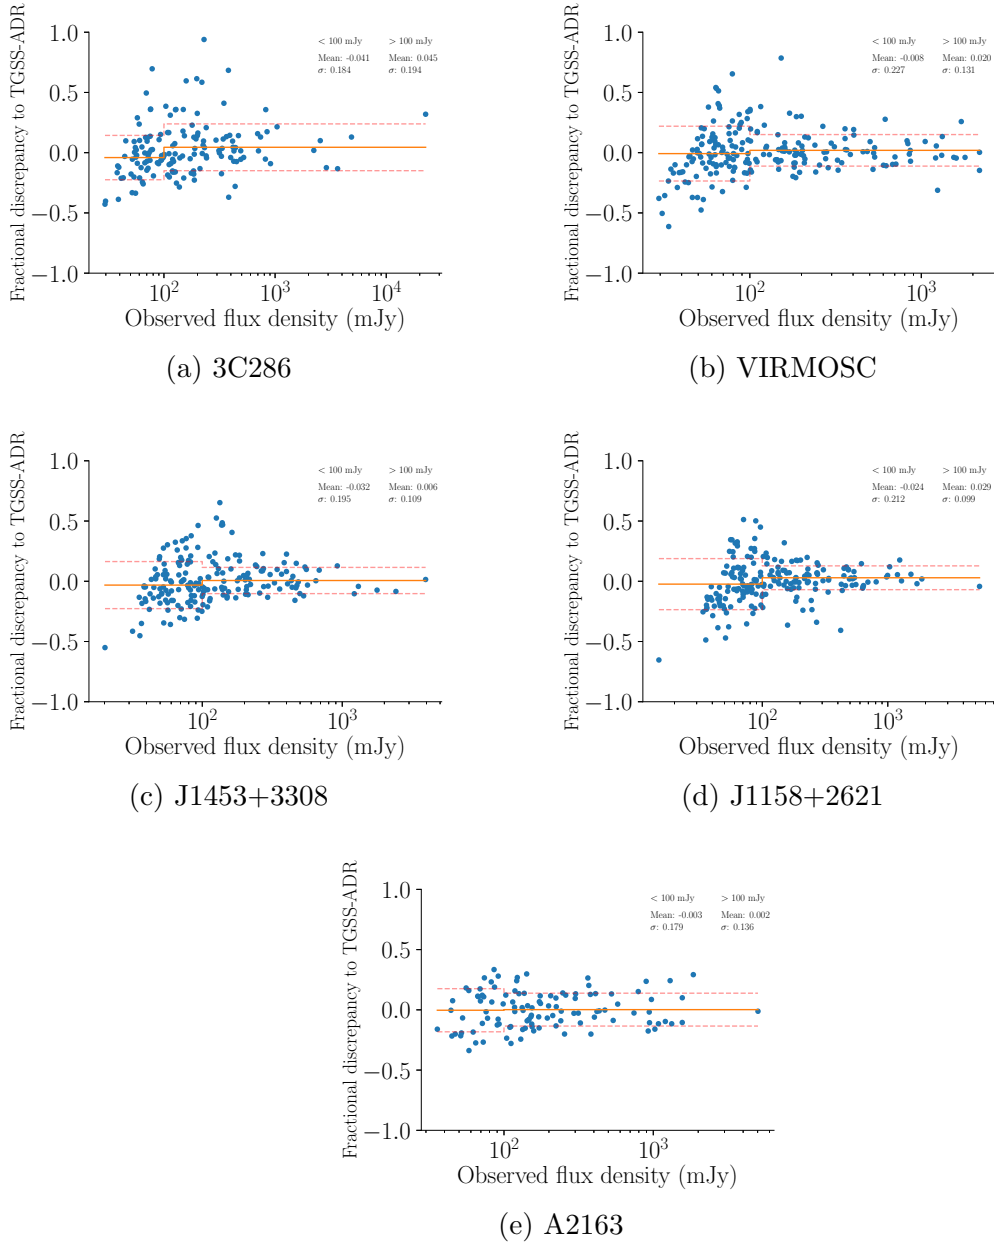


Figure 3.5: The fractional change between the strong (> 100 mJy) sources and weak (< 100 mJy) sources for each observed target field compared with sources in common with TGSS-ADR. We see a similar trend as in Figure 3.6. The values for the mean and standard deviation of the fractional change are indicated on the plot.

3.4 Efficacy of the algorithms

The IPFLAG algorithms were applied to real data from the GMRT at 150 MHz and compared to results without the same. The 150 MHz band of the GMRT is highly under-utilised (due to the presence of very strong RFI) although very important for a variety of astrophysical phenomena. We felt that testing these algorithms against real data offered a more realistic appraisal of their efficacy rather than testing them against simulations with well behaved noise. We carried out the tests on 5 fields, VIRMOSC (GMRT observation code: 14RAA01), J1453+3308 (27_063), J1158+2621 (29_062), A2163 (16_259) and 3C286 (TGSS data, Intema *et al.* 2017).

We briefly summarise each field below:

3C286 – This is a commonly used flux density calibrator source. It is very compact, and has a flux density of 26 Jy at 150 MHz. The field is dominated by point sources with almost no extended emission.

VIRMOSC – A point source dominated field, the strongest of which is ~ 1.7 Jy. There is a single diffuse source of angular size $4'$ and a total flux of 300 mJy.

J1453+3308 and J1158+2621 – Double double radio galaxies, with diffuse outer lobes spanning $4-7'$.

A2163 – A galaxy cluster with a low surface brightness radio halo of angular size $14'$.

Figure 3.2 shows the comparison of the median binned UV plane of tested 150 MHz fields before and after application of the GRIDflag+TCflag routines. The visibilities are shown in the upper half of the UV plane; the lower half has been Hermitian conjugated appropriately. The band of higher intensity seen between $U = [-100\lambda, 100\lambda]$ (e.g., Figure 3.2 c and d) arises due to the application of the RfiX procedure to the data. RfiX can only identify RFI for baselines where the fringe stop phase is substantial within the interval of a few minutes. Therefore it cannot mitigate RFI in the strip between $U = [-100\lambda, 100\lambda]$.

The total data flagged by IPFLAG was between 2.4% - 15.7%, and the corresponding loss in UV-bins was 1.2% - 3.9%. As noted earlier, a larger number of UV-bins survive the process with increased coverage of the target field such as observing for multiple days. Table 3.1 summarizes the dirty beam parameters before and after applying our flagging routines, and in all cases the fractional change is

small. We retained the same restoring beam for all the images. The post imaging comparison plots are shown in Figures 3.6 - 3.10.

These images were compared to those in which these procedures were not applied. Additionally, we also repeated the same procedure, except using the AOFlagger package (Offringa *et al.* 2010, 2012) to flag instead of the IPFLAG routines. AOFlagger operates in the visibility time-channel plane (but not on the UV plane), using various thresholding algorithms (termed ‘SumThreshold’ and ‘VarThreshold’; see Chapter 2) that sets the RFI detection threshold based on the sum of several samples within a sliding window. They also use a morphological operator that is scale invariant to identify and flag RFI. In all of the following plots the results from the original data (*i.e.*, those created without applying AOFlagger/IPFLAG) are plotted in blue, those using AOFlagger in pink, and those using IPFLAG in orange.

Point source fluxes – The plots in Figure 3.6 show the fractional change in flux density (FCF) as a function of original flux density. The flux densities were compared before applying the primary beam gain correction. The FCF was calculated for three flux density regions, *viz.* $S > 100$ mJy, $25\sigma < S < 100$ mJy, and $7\sigma < S < 25\sigma$. We did not consider any sources detected below 7σ . 25σ is a conservative cutoff above which source counts can be considered complete (e.g., Intema *et al.* 2017) and as such defines the lower flux density limit at which the catalog sources are reliable. The FCF values in percentage units are listed in Table 3.2. The FCF ranges from -1% to +6% for the strongest sources. This increase in flux is not unexpected, due to an improvement in calibration after the removal of RFI. The intermediate sources show a larger relative change in their FCF value, compared to the strong sources, by -6% to +1%. The trend continues to the faintest sources detected (at a level between 7σ and 25σ). This reduction was present even in the case of the other flagging algorithms (AOFlagger, TFCrop and RFlag), and source finding algorithms (pyBDSF and aegean; Hancock *et al.* 2012; Mohan & Rafferty 2015). However, we did not see any reduction when an equivalent number of visibilities were randomly flagged in the data sets. We also analysed the image residuals (after subtracting the fitted sources) in the immediate vicinity of the strong point sources, but there is no systematic offset to more negative values which can explain the reduction in flux densities. This effect seems to be independent of the RFI flagging algorithm used. It does not seem to affect the flux density scale since the flux of the brightest sources are unchanged. We suspect that it may have something to do with the CLEAN algorithm we use (that is implemented in CASA). The reduction in noise is larger than the change in flux density of the intermediate sources, and so cannot fully explain the measured

Table 3.1: The effect of IPFLAG on image parameters. The values of the dirty beam major and minor axes in arcsec, median image noise σ_o and σ_f in mJy/beam, and the number of detected sources N_o and N_f are listed for each field. We have also listed the percentage of UV-bins lost and visibility data flagged after IPFLAG.

Source Name	Before IPFLAG				After IPFLAG					
	Major axis	Minor axis	N_o	σ_o	Major axis	Minor axis	N_f	σ_f	Flagged bins %	Flagged points %
3C286	18.4	13.1	374	1.12	18.9	13.0	403	1.01	1.2	2.4
VIRMOSC	23.9	15.9	846	0.55	25.6	15.8	949	0.42	1.7	6.4
J1453+3308	20.2	14.5	862	0.46	18.9	14.1	870	0.38	2.7	15.7
J1158+2621	19.0	15.1	819	0.72	18.9	14.6	844	0.55	3.9	15.5
A2163	28.9	14.3	419	1.30	31.6	15.8	493	1.06	1.5	3.9

fractional change. At present we do not have an explanation for this effect, but it may have something to do with the CLEAN bias (Cohen *et al.* 2007; Condon *et al.* 1998) whose impact is largest for the fainter sources. We recall that a similar reduction in FCF was observed while comparing our fluxes with that of the TGSS-ADR (see Figure 3.5).

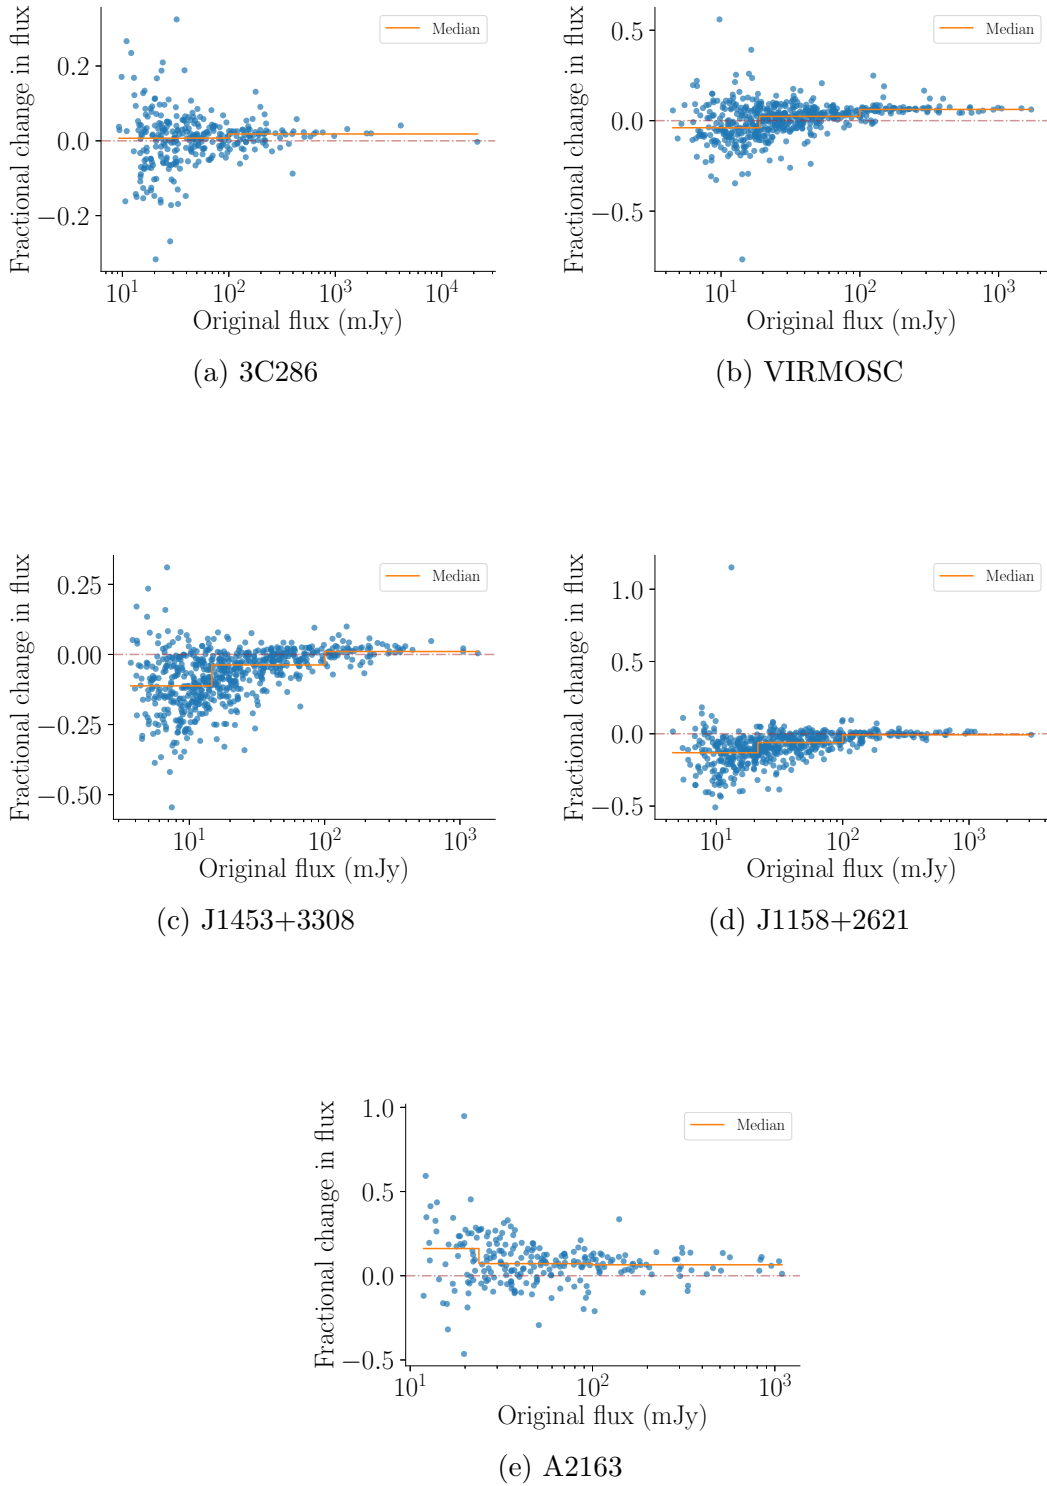


Figure 3.6: Plot of fractional change in flux vs. original flux. The solid line indicates the median value of the fractional change in each of the 3 regimes, *i.e.*, $7\sigma < S < 25\sigma$, $25\sigma < S < 100$ mJy, $S > 100$ mJy. In every case except one (that of A2163) there is a relative decrease of the flux density of the fainter sources compared to the brightest sources.

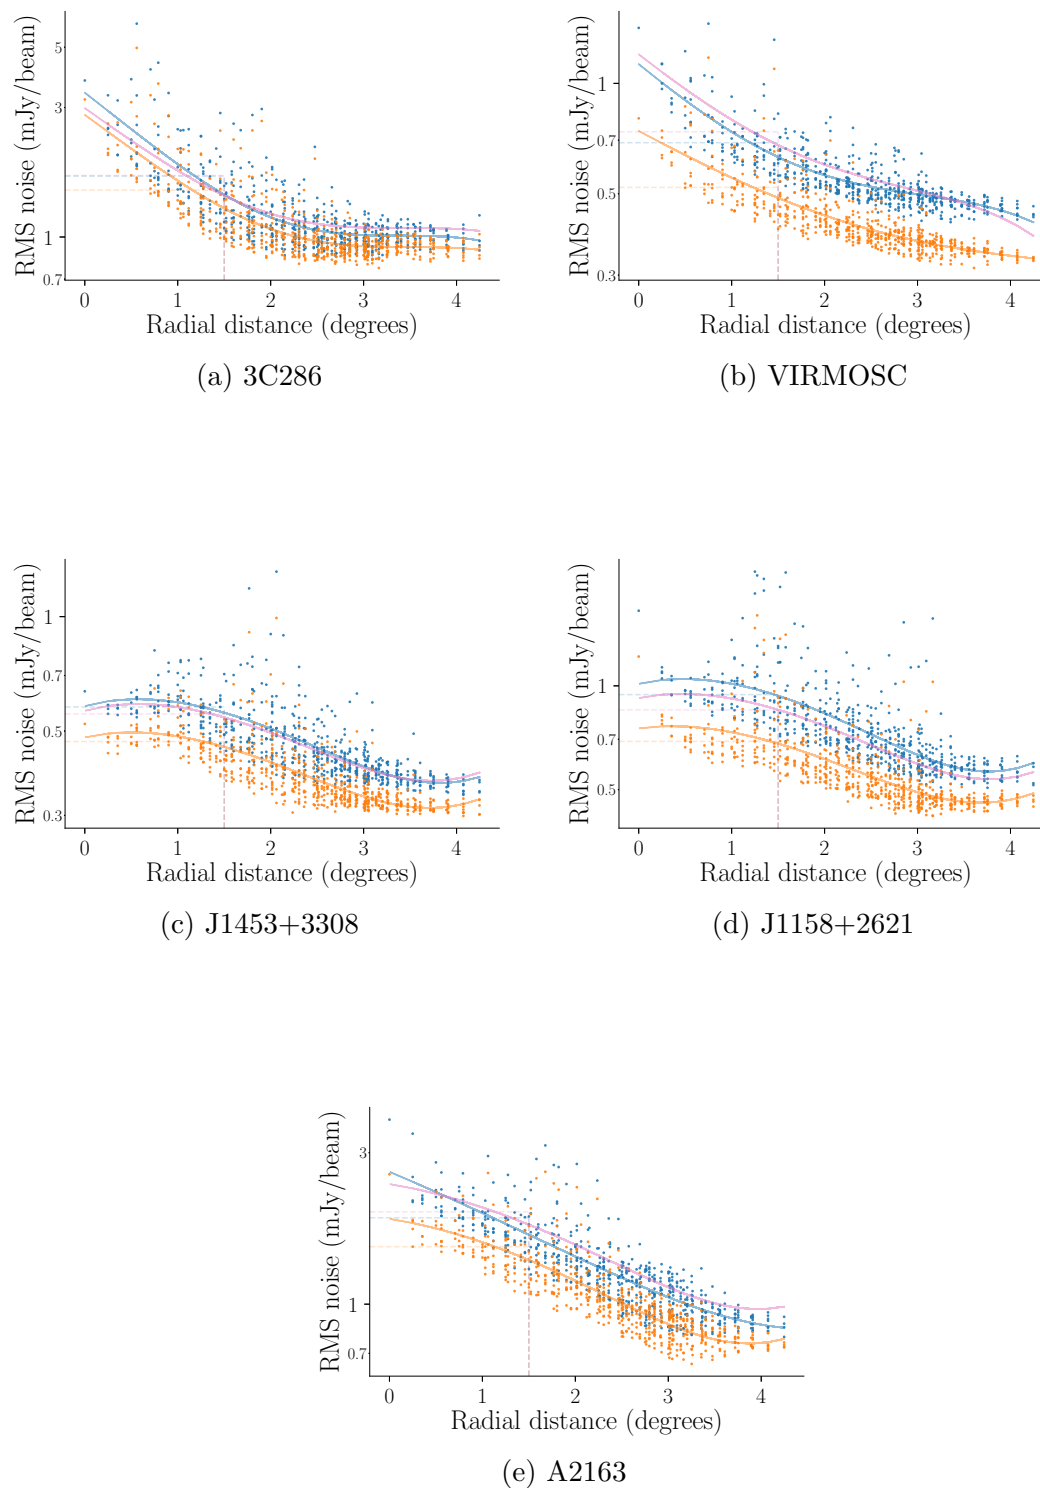


Figure 3.7: Plot of the RMS noise as a function of radial distance from the pointing centre. In this and all subsequent plots, the original data is in blue, IPFLAG data in orange, and AOflagger data in pink (only the trendline). In every case, IPFLAG results in a consistent reduction in RMS across the field.

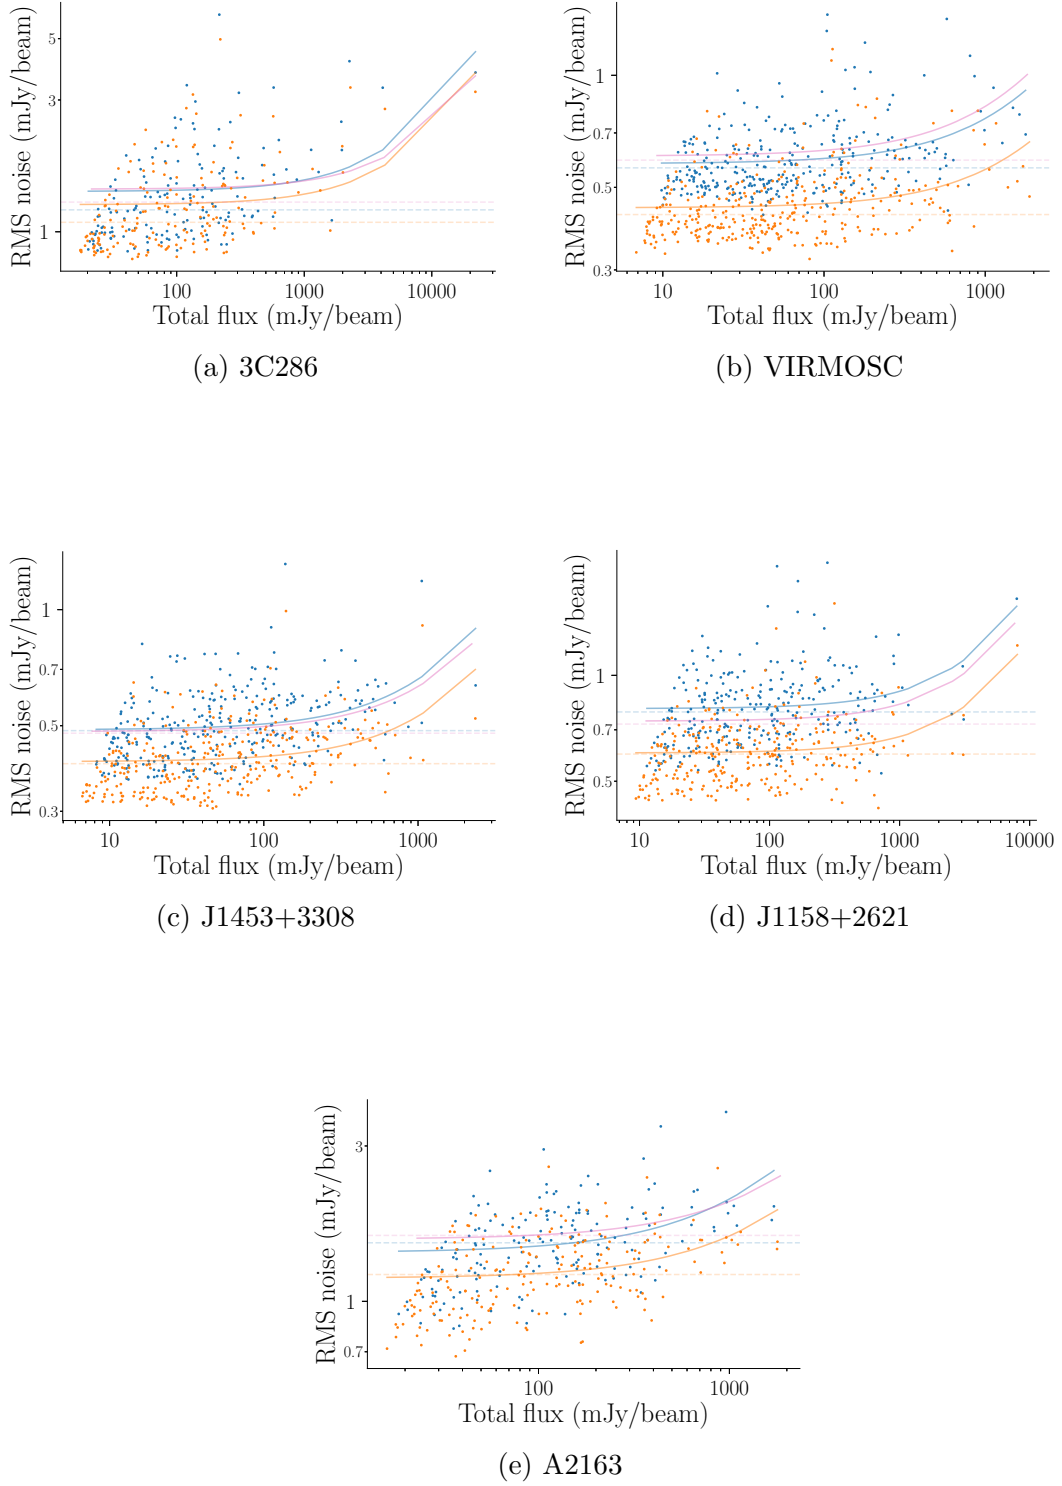


Figure 3.8: Plot of RMS noise as a function of total flux. As with the previous plot, the results obtained after processing with AOFlagger are indicated only with a trendline in the interest of clarity. In general, the RMS noise is expected to go up in the presence of strong sources. We see that IPFLAG again results in a consistent decrease in RMS noise even in the presence of strong point sources

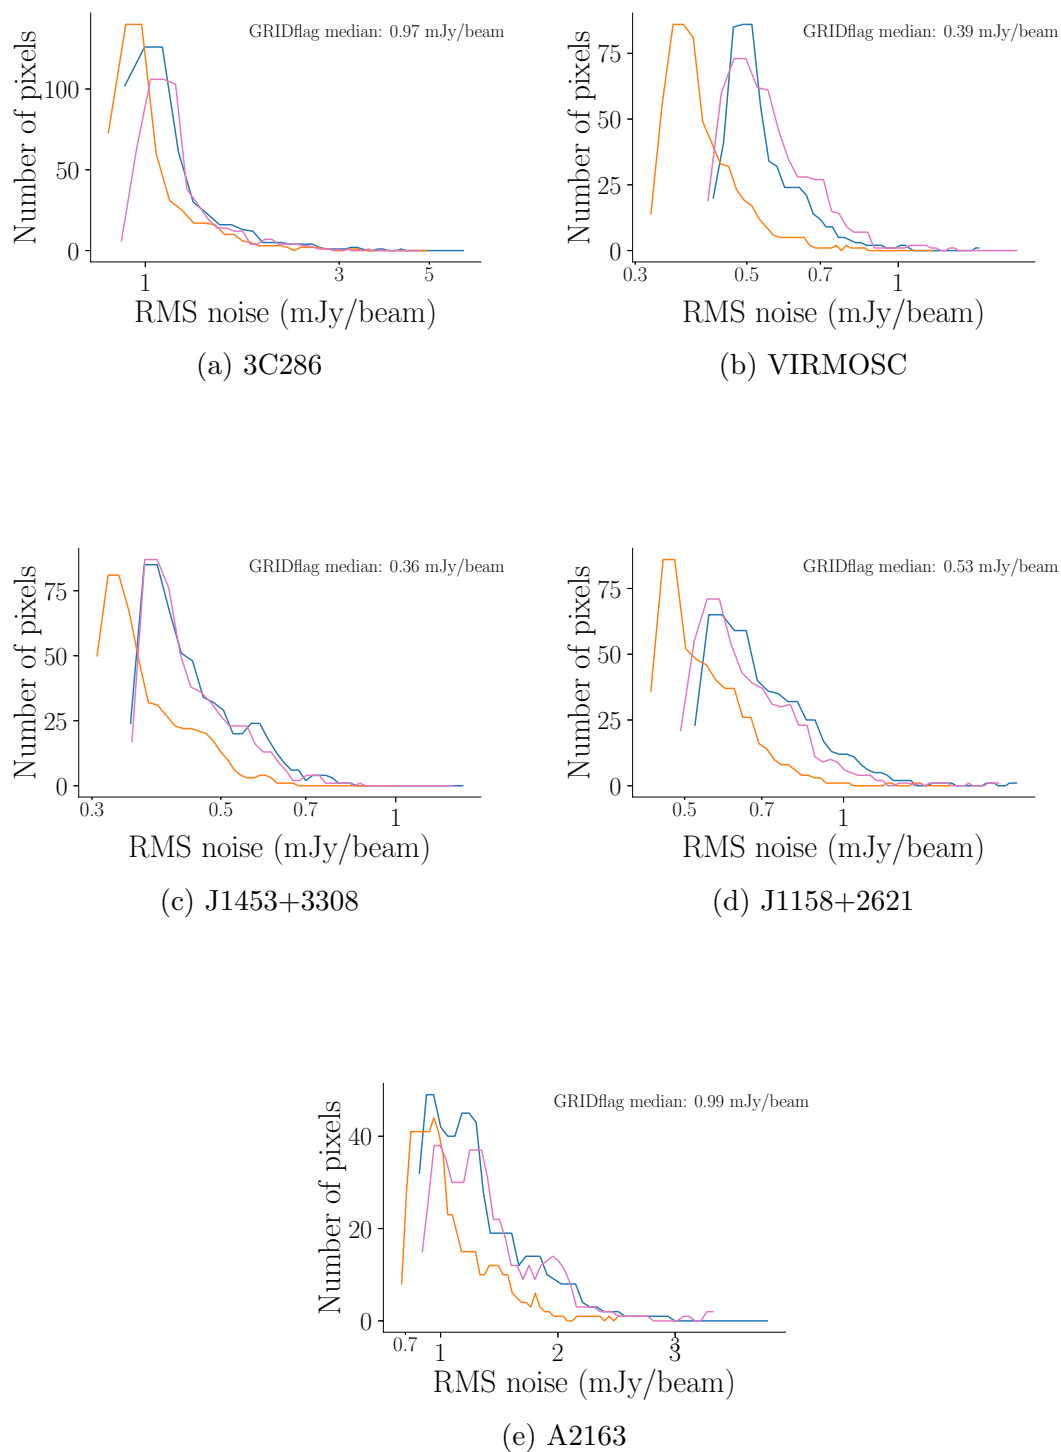


Figure 3.9: Plot of the RMS histogram across the entire field. The peak of the histogram after processing with IPFLAG has shifted toward lower flux densities in all cases.

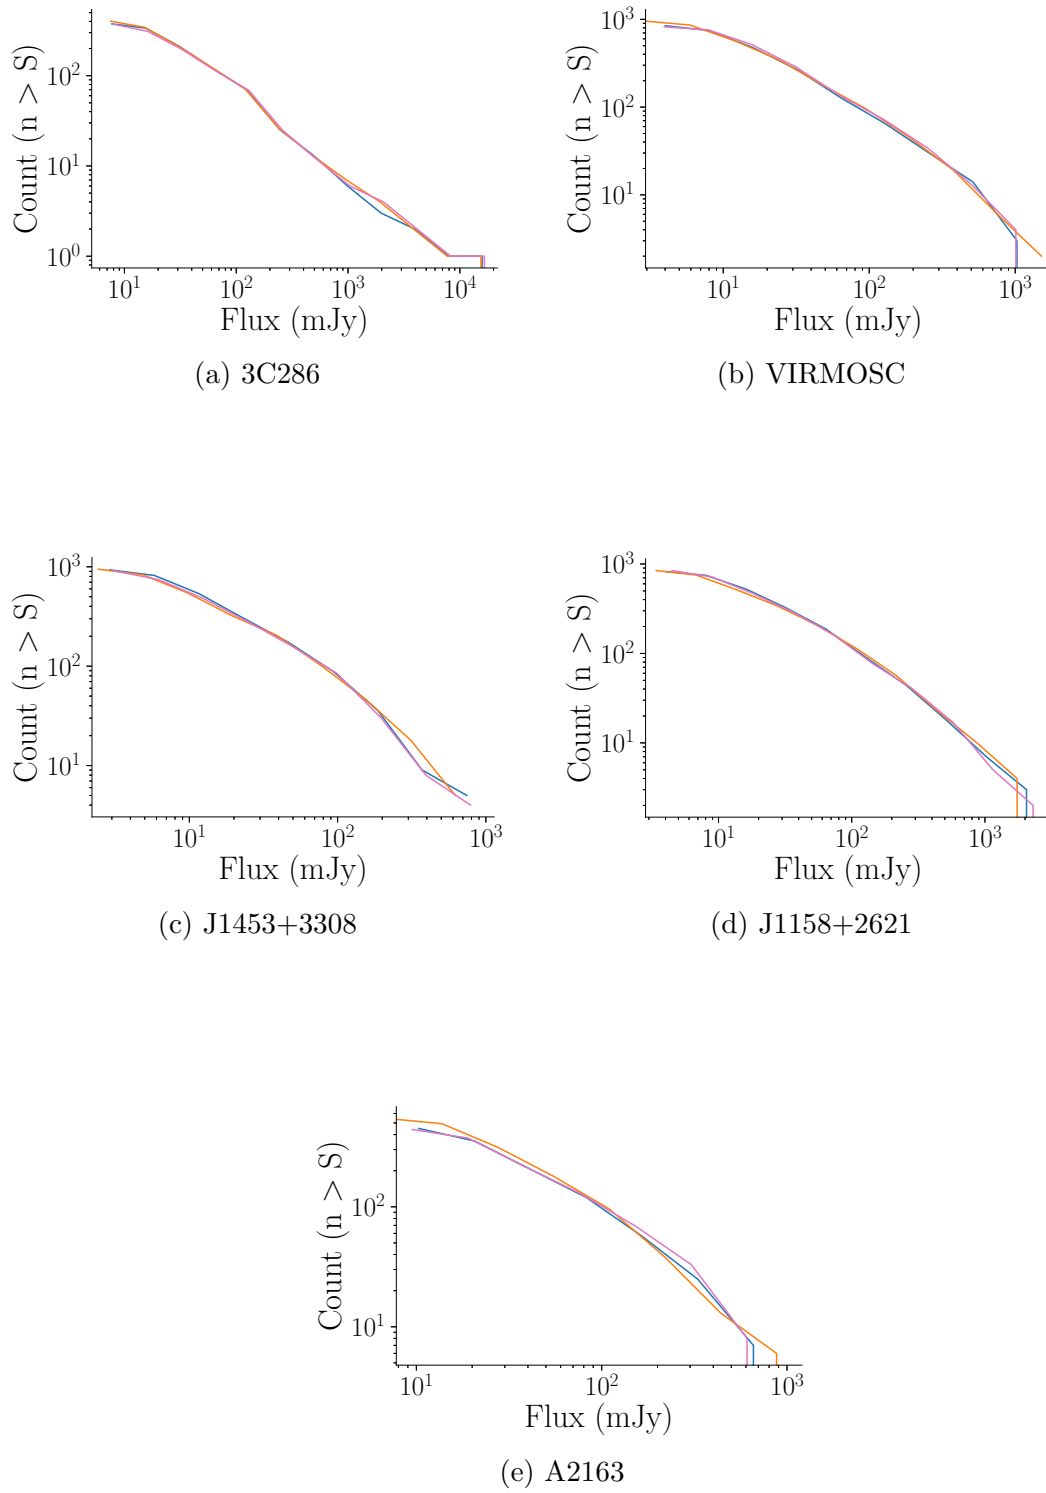


Figure 3.10: Integral source counts from each field for the original data, those processed with AOFlagger and those processed with IPFLAG. In every case IPFLAG results in a larger number of detected sources, down to lower flux densities.

Table 3.2: Percentage change in the median flux density in the three different regimes.

Source name	$7\sigma < S < 25\sigma$	$25\sigma < S < 100 \text{ mJy}$	$S > 100 \text{ mJy}$
J1453+3308	-12	-4	1
J1158+2621	-13	-7	-1
VIRMOSC	-4	2	6
3C286	1	1	2
A2163	15	7	6

Image noise as a function of radial distance – In general, the image noise is known to reduce with radial distance. For this analysis, a 200x200 pixel box was used to calculate the robust RMS noise at 625 locations across the image. The plots in Figure 3.7 show both the scatter and trendline for the un-flagged and IPFLAG data, whereas the AOFlagger data is represented only as a trendline in the interest of clarity. The plots show that IPFLAG has reduced the noise all across the image, indicating that systematics had been present all throughout the field in the original data (and not just at the centre). In some regions of the plot, the AOFlagger trendline indicates a higher RMS in the AOFlagger image than in the original. This is because the self-calibration procedure was applied to heavily flagged data (such as those produced with AOFlagger), and that tends to increase the RMS as a consequence of poor UV coverage and poor source reconstruction.

RMS as a function of source flux – Figure 3.8 shows the RMS noise plotted against the cumulative source flux within each of the 625 boxes mentioned previously. A higher noise level is expected in the presence of strong point sources, due to a higher level of local artefacts.

Noise histogram – Figure 3.9 shows the histogram of the RMS noise across the field (using the same 625 boxes as mentioned above). In every case there is a clear shift in the peak of the histogram toward lower values.

Source counts – A reduction in noise is only validated if a corresponding number of sources are detected at lower flux levels. This last plot, Figure 3.10 shows that IPFLAG does indeed pass this criterion by detecting faint sources to the expected depth. The plot shows the cumulative histogram of the number of sources detected as a function of flux density, *i.e.*, $N(>S)$ above a level of 7σ . This shows that the excess sources come from lower flux densities. Certain portions of the

plot indicate that the curve corresponding to the original data lies above that of the flagged data, which implies that more sources were detected in the original image within that bin. This is mostly an artefact of the plotting package which chose different bin ranges. Additionally, a source might shift into a different bin on account of a change in its flux. The essential point is the extension toward lower flux densities of the IPFLAG'd data.

Artefacts – Figure 3.11 shows the reduction in artefacts in the field after applying IPFLAG.


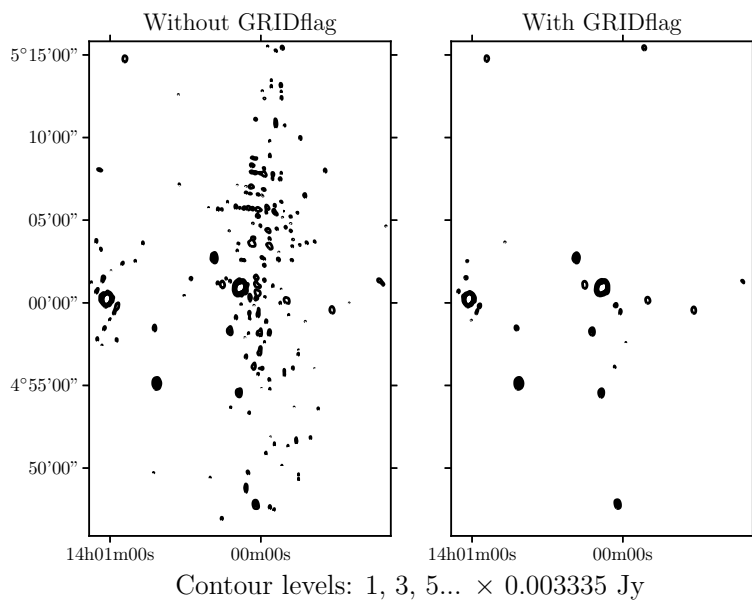
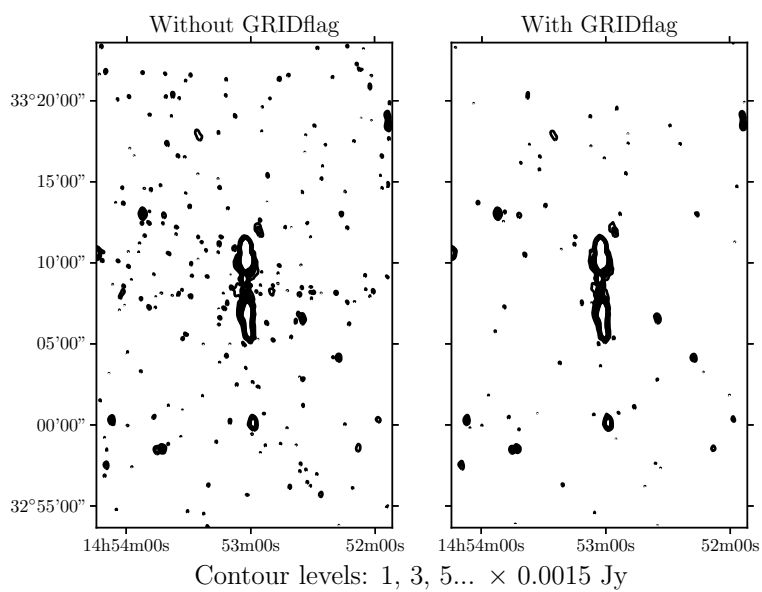


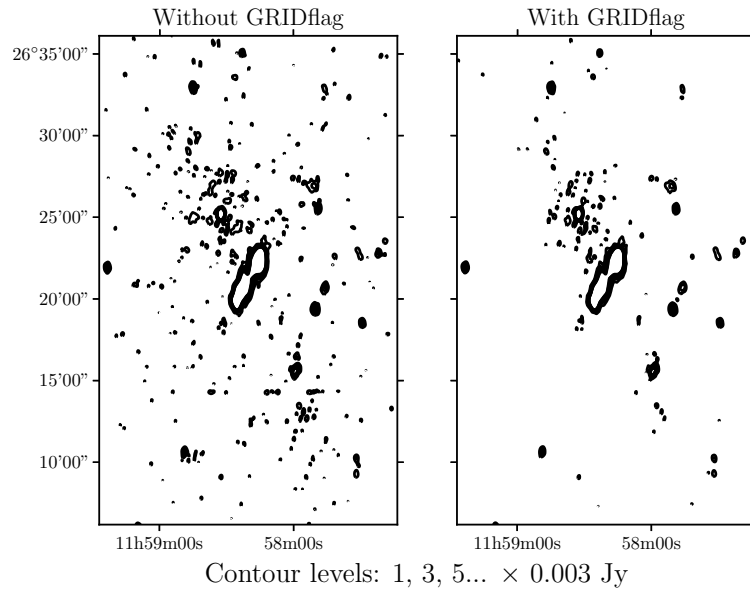
Figure 3.11: (a) A comparison of the noise contours in the inner portion of the 3C286 field. The plot on the left shows the image made from the original data, and the plot on the right shows the image made after the application of the IPFLAG algorithms. The contour levels are indicated below each plot. All the fields show a substantial reduction in artefacts and noise peaks with the application of IPFLAG.



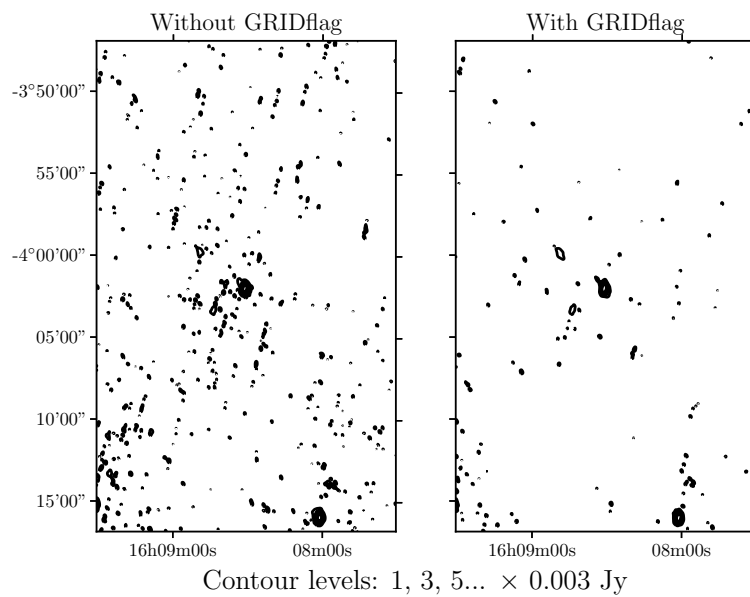
(b) Continued from Figure 3.11. Contours for the VIRMOSC field.



(c) Continued from Figure 3.11. Contours for the J1453+3308 field.



(d) Continued from Figure 3.11. Contours for the J1158+2621 field.



(e) Continued from Figure 3.11. Contours for the A2163 field.

3.5 Conclusion

The IPFLAG algorithms presented here work well across a variety of different data sets. In each case, the image noise is substantially reduced after flagging, thereby increasing the detectability of fainter sources. Although these algorithms were tested on GMRT data, they can be applied to any interferometric data set where the UV plane contains a large number of samples.

We are unaware of any other algorithm that approaches RFI flagging in the manner that GRIDflag does. TCflag is similar to other algorithms such as Rflag, AOFlagger, TFCrop etc., that operate on the visibility time-channel plane of a single baseline while they differ in the manner by which they estimate the background. The combined application of GRIDflag+TCflag, which are components of a single pipeline IPFLAG outperform the other algorithms. AOFlagger was the closest in terms of performance, and we recognise that a more experienced user may be able to obtain better results by optimising its many tuneable parameters to GMRT data.

On the other hand IPFLAG only has six tuneable parameters in all, across both GRIDflag and TCflag. These values did not require much optimisation, and the same values were used across all the fields tested. Of the six, four values are effectively determined by the observation itself *i.e.*, the UV bin size is determined by the field of view imaged, and the three noise thresholds were set to ‘universal’ default values of 3σ . The width of the UV annuli were decided by distributing the number of visibility bins roughly equally across each annulus, however the performance of the algorithms is not strongly tied to the selection of annulus width. While these values may require some tuning for different frequencies and interferometers, they should not require tuning for different fields.

On application of IPFLAG our images have consistently reached an image noise of < 1 mJy/beam with a corresponding increase in source detection; typical GMRT 150 MHz images only reach a noise of 1.5 - 5 mJy/beam, with exceptional effort yielding 0.7 mJy/beam (e.g., Ishwara-Chandra *et al.* 2010). In 3 of the 5 images shown in this chapter, we have reached an image noise of 0.38 - 0.55 mJy/beam which are among the most sensitive maps created from a single synthesis observation at a 150 MHz with the GMRT. These are a factor of ≤ 2 above the theoretical confusion limit at this frequency. Even in the case of 3C286, an exceptionally bright flux density calibration we have reached 1.0 mJy/beam and the peak to RMS ratio is in excess of 21,000 which is unprecedented. All this suggests that IPFLAG works consistently across a variety of data sets, yielding significant improvements in sensitivity.

Since the flagging is performed in such a way as to conserve UV coverage, the

gains in sensitivity are not hindered by a corresponding increase in the sidelobes of the synthesised beam. The efficacy of the algorithm depends on the level of redundancy in the gridded UV plane - The GMRT has good coverage of the short spacings (Swarup *et al.* 1991) at low frequencies, and this coverage is due to get better with the upgraded GMRT (Gupta *et al.* 2017). Further, the algorithm also works at higher frequencies (325 MHz and 610 MHz) but the improvement was not as substantial. We think that this is because the RFI environment is weaker at these frequencies. However, for deep imaging projects such as the MIGHTEE (Jarvis *et al.* 2017) even such a weak RFI environment can significantly hinder the sensitivity. These algorithms are expected to work well for such experiments due to the large redundancy in UV coverage which arises from observing the same field for several nights.

Finally, these algorithms will not work in the presence of persistent, broadband RFI. However, other algorithms exist to mitigate the RFI in these situations (e.g. Athreya 2009). The issues that remain to be addressed primarily arise from a non-isoplanatic ionosphere, the antenna primary beam (asymmetry, pointing jitter etc.) and a wide bandwidth observation. Wide bandwidth observations can observe a varying source structure across the band due to the spectral index structure. In addition to this, the noise characteristics of the receiver and sky can change substantially across the band. Therefore adapting these procedures to the wide bandwidth case should be the next step. Additionally, using IPFLAG in conjunction with recipes like SPAM (Intema *et al.* 2009) should yield further improvements in image quality.

Chapter 4

Direction dependent error mitigation

4.1 Introduction

The sky intensity distribution itself is a constant in most cases and the standard radio astronomy imaging and self-calibration algorithms (e.g. Clark 1980; Cornwell 2008; Högbom 1974; Pearson & Readhead 1984) implicitly assume the same. The Direction Dependent Errors (DDEs) mentioned in Chapters 1 and 2, if not corrected, result in time- and baseline-variability of the sky distribution as seen by the radio interferometer. By baseline variability we mean that the antenna temperature contributed by the source varies from antenna to antenna, and hence to each baseline; *i.e.* the fringe amplitudes of even a point source would vary from baseline to baseline and as a function of time. The implicit assumption of time-invariance by the algorithms result in an improper subtraction of sources, leading to bad residuals in the visibilities and systematic artefacts across the image and consequently poor detectability of faint sources.

Ideally, one would want to calibrate out the DDEs to improve the image quality. Several groups have adopted this “fundamental” approach to tackle the issues of DDEs (e.g. Bhatnagar *et al.* 2008; Cotton *et al.* 2004; Intema *et al.* 2009; Noordam 2004; Van Weeren *et al.* 2016) with varying degrees of success. In this chapter we tackle not the root causes of DDEs but the symptoms, and propose ways of mitigating their consequences.

The strongest sources contribute most of the artefacts arising from the assumption of the constancy of sky intensity. Our procedures *do not* correct for DDEs of any source in the field. Instead we reduce the DDE-induced artefacts contributed by the strong sources. This results in the improved detection of weaker sources but without improving the dynamic range on the strongest sources. Furthermore, the

procedures are conceptually simple, computationally efficient and utilise existing tools.

The visibilities measured by every baseline are given by the measurement equation (Equation 4.1)

$$V_{pq} = G_p \left(\int E_p(\vec{l}) B_{pq}(\vec{l}) E_q^\dagger(\vec{l}) e^{2\pi i(u_{pq}l + v_{pq}m)} dldm \right) G_q^\dagger \quad (4.1)$$

where p and q are the antenna indices, G is the complex antenna gain, E_{pq} contains the direction dependent errors, and B_{pq} is the sky brightness distribution. The matrix E_{pq} can be modified to accommodate an arbitrary number of DDEs, including the w-term, polarisation leakage etc. These DDEs cause the apparent sky brightness distribution to vary as a function of time and direction. Standard imaging and self-calibration algorithms assume a time-invariant sky and hence use the time (and usually frequency) averaged flux of the source. This results in two separate, but related, errors - (a) closure errors during self-calibration and (b) systematic artefacts in the image plane.

Essentially, a point source of flux-density S should result in a fringe of amplitude S (but different phases) in all baselines. An array with symmetric antenna primary beams of differing sizes will require a direction-dependent self-calibration, without which residual visibility fringes (of constant amplitude) will be found across the array. On the other hand, in the case of an azimuthally asymmetric antenna primary beam and/or a time-dependent antenna pointing error, the fringe amplitude will differ from S as a function of time, both, across the array ($S' = S + \delta S_A(t)$), and within a baseline ($S' = S + \delta S_B(t)$). While subtracting a source of average flux density S , an imaging algorithm will leave behind residual sources of flux density $\delta S_A(t)$ in the image plane, and residual fringes of amplitude $\delta S_B(t)$ in the individual baselines. These residual point sources will result in dirty beam artefacts across the image. The residual baseline fringes will result in closure errors and incorrect calibration solutions.

In this chapter we describe two algorithms to reduce the residues of strong sources in the two different regimes, *viz.* by “snapshot” imaging in the image plane, and baseline defringing in the time-channel visibility plane.

4.2 Snapshot imaging

From Chapter 2, Section 2.2 the two extant methods to correct for an azimuthally asymmetric antenna primary beam (or more generally, time varying antenna primary beam effects) are the A-projection algorithm, or the source peeling method

with amplitude calibration. We propose here a new approach that measures and compensates for the variation of the source flux, without correcting the source structure.

Subtracting out the time-averaged flux of a source results in departures from Gaussianity in the residual visibilities because of the aforementioned time variability. Instead, the visibilities in the array will have time-variable residual point sources embedded in them. We estimated this effect by making a series of “snapshot” images using visibilities in 5-10 minute blocks within the several hours of a full synthesis observation. We measured the flux densities of the brightest 50 sources in these snapshots and investigated their variation with time, radial distance from the pointing centre, and the antenna parallactic and elevation angles. The visibilities of all the snapshot images were made from a self-calibrated data set which used the same sky model.

Two of the sources (J1453+3308 and J1158+2621) in Figures 4.1 and 4.2 were observed using the GMRT software backend (Roy *et al.* 2010) along with a pointing model that corrected for pointing offsets (Lal 2013). The final two sources (VIRMOSC and TAUBOO) in Figures 4.3 and 4.4 were observed using the older hardware correlator prior to the implementation of the pointing correction model.

4.2.1 Flux density variation with time

Figures 4.1 - 4.4 show the variation of source flux densities with time. The flux density of each snapshot, shown as a scatter plot of grey dots, is represented by its normalised deviation from the median of the source for the entire observing session. The blue line is a moving average of the scatter points. These plots illustrate that subtracting the session averaged flux density of a source still leaves behind ‘uncleaned’ components with a sizeable fraction of the original flux density. Some of the variations could just be due to the problem of estimating the flux densities using the automated source finding utility (e.g. pyBDSF; Mohan & Rafferty 2015). The increased dirty beam sidelobes in a snapshot can create spurious structures around a source. Nevertheless, fluctuations of the order of a few percent for dozens of sources ranging in flux densities from ~ 250 mJy to several Jansky can leave behind residuals of tens of milliJansky. This can seriously impact a campaign which seeks to generate images with RMS noise at the level of the confusion limit.

Two of our targets were observed prior to the use of the new pointing model. However, there seems to be no substantial difference in the flux density variation before and after the implementation of the pointing correction model.

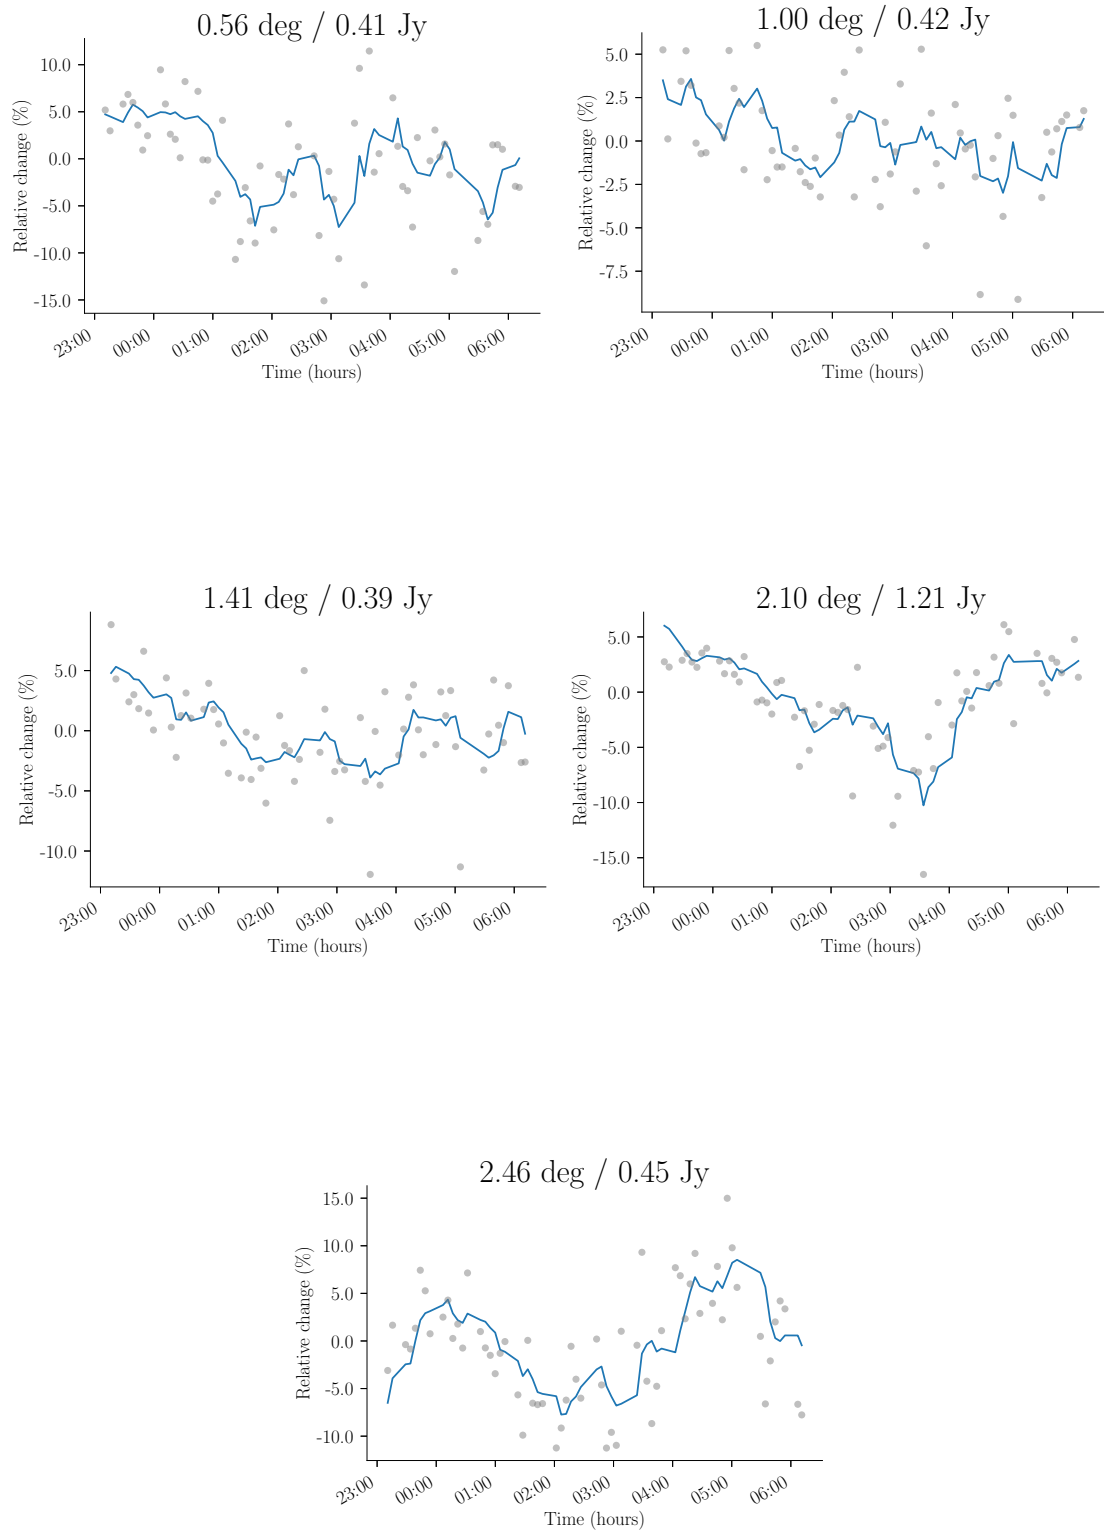


Figure 4.1: Percentage variation about the median flux density as a function of time for representative sources at different radial distances from the pointing centre in the J1453+3308 field. The distance of the source from the pointing centre and its flux density are listed above the plot.

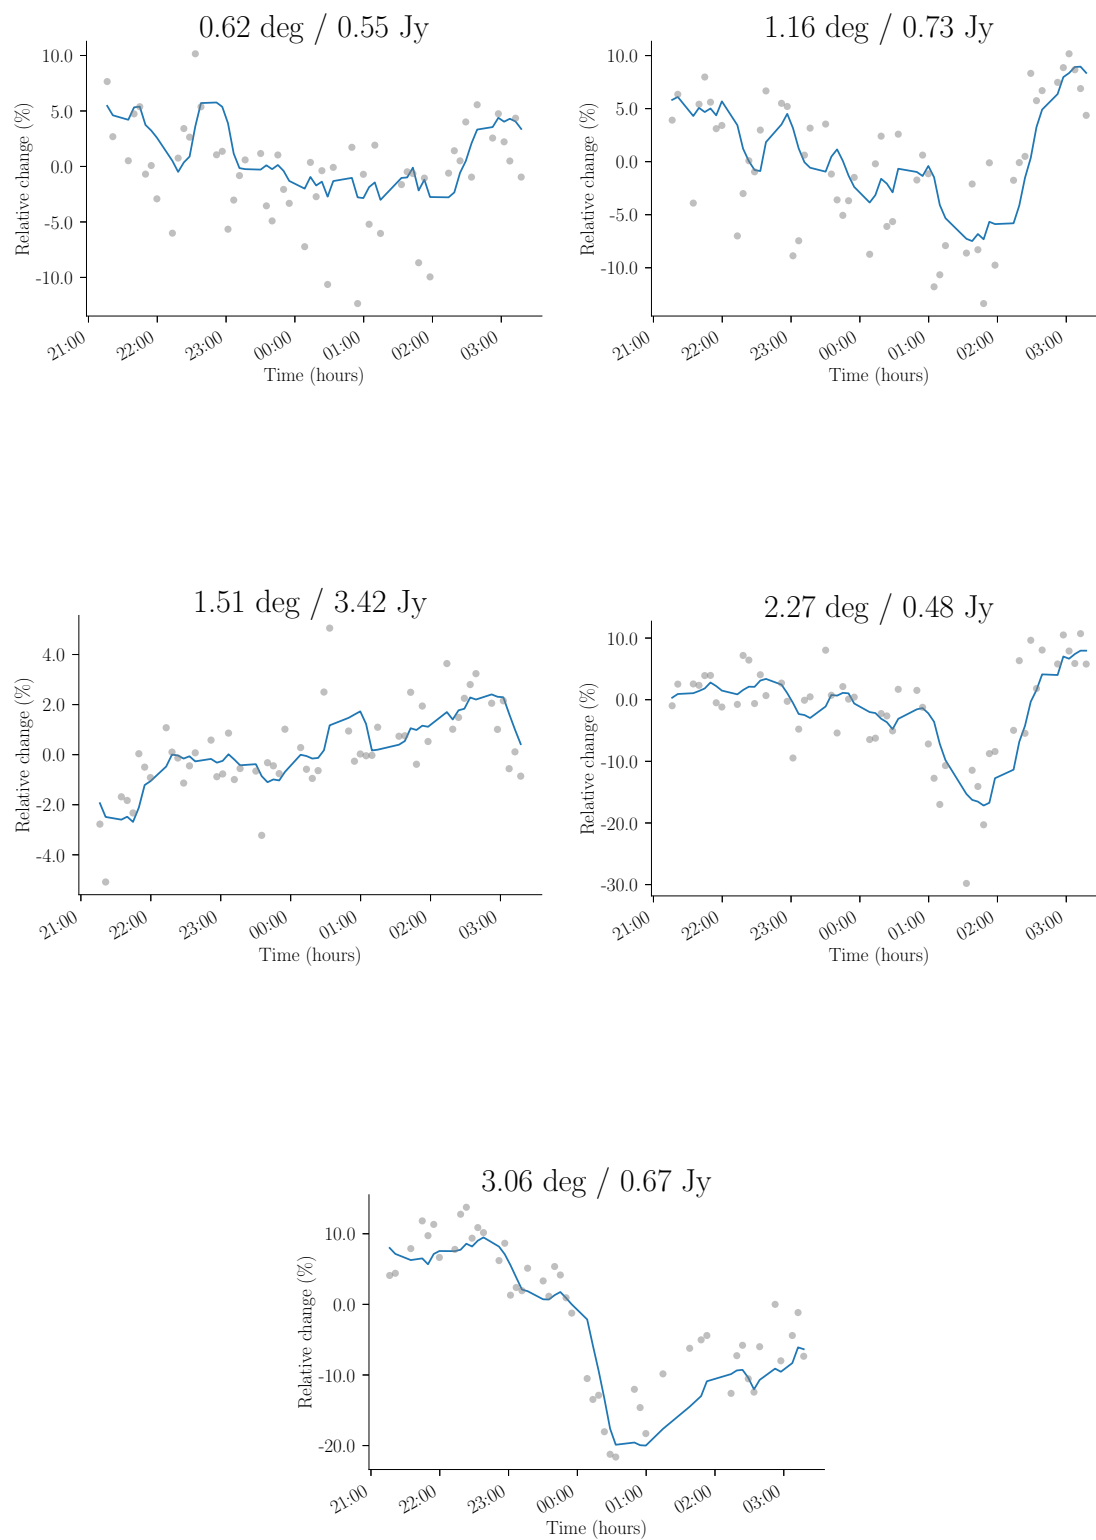


Figure 4.2: Percentage variation about the median flux density as a function of time for representative sources at different radial distances from the pointing centre in the J1158+2621 field. The distance of the source from the pointing centre and its flux density are listed above the plot.

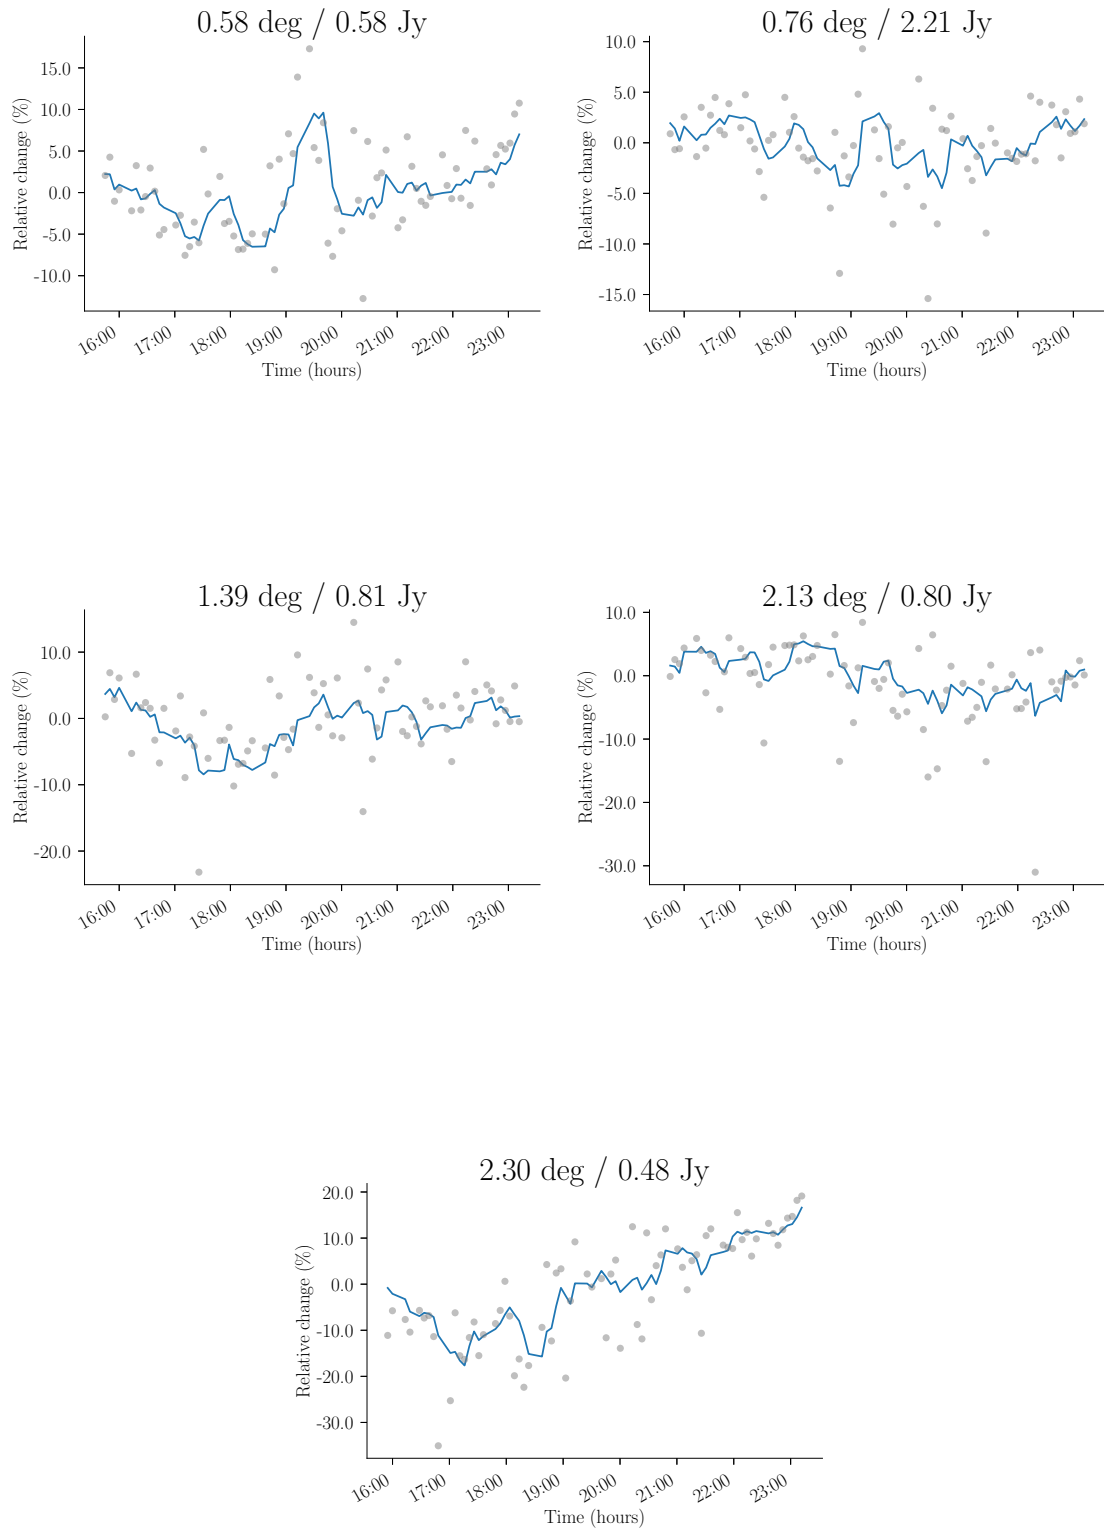


Figure 4.3: Percentage variation about the median flux density as a function of time for representative sources at different radial distances from the pointing centre in the VIRMOSC field. The distance of the source from the pointing centre and its flux density are listed above the plot.

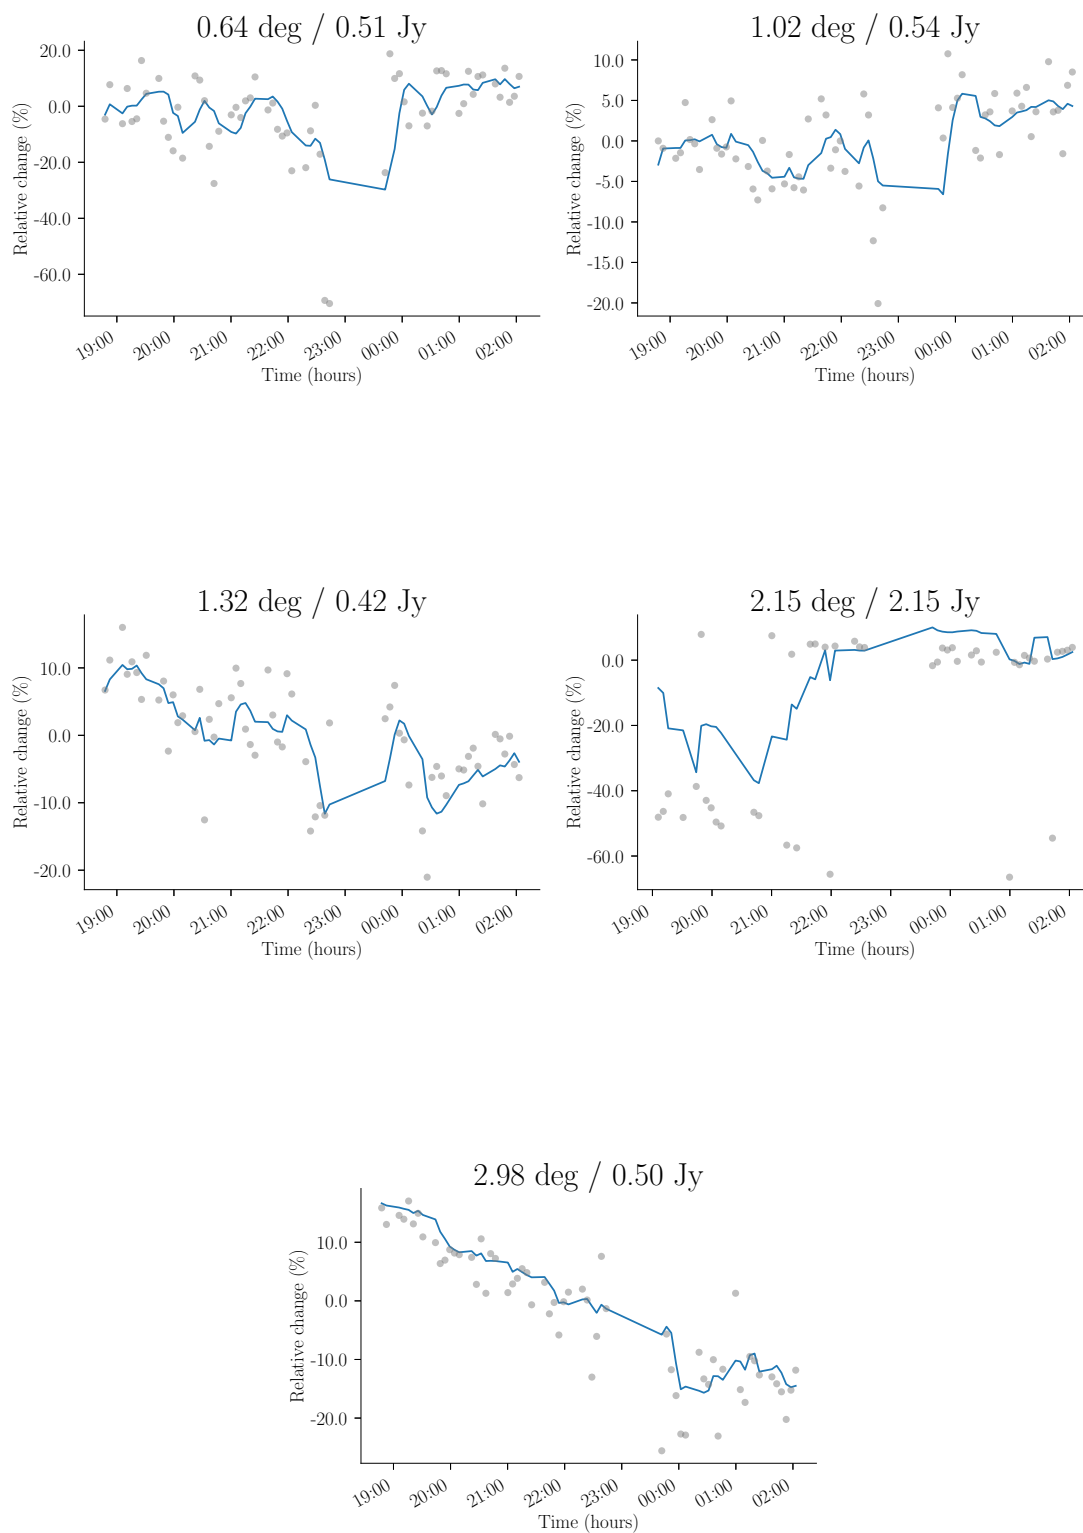


Figure 4.4: Percentage variation about the median flux density as a function of time for representative sources at different radial distances from the pointing centre in the TAUBOO field. The distance of the source from the pointing centre and its flux density are listed above the plot.

4.2.2 Flux density variation with radial distance

Suspecting that the primary beam pattern was the main source of DDE impacting our data we expected the source variation to increase with radial distance from the pointing centre. This is because the slope of the beam pattern increases with radial distance out to the half power point. Figure 4.5 shows that for three of the four fields tested this is indeed the case, albeit the effect is small. In the final field (TAUBOO, Figure 4.5d) the slope is negative. While primary beam effects should in principle cause a radially dependent pattern, ionospheric effects, which can have a complex pattern across the field of view, can mask it. Variable calibration errors, due to strong RFI or the ionosphere, can also wash out the effect due to the primary beam.

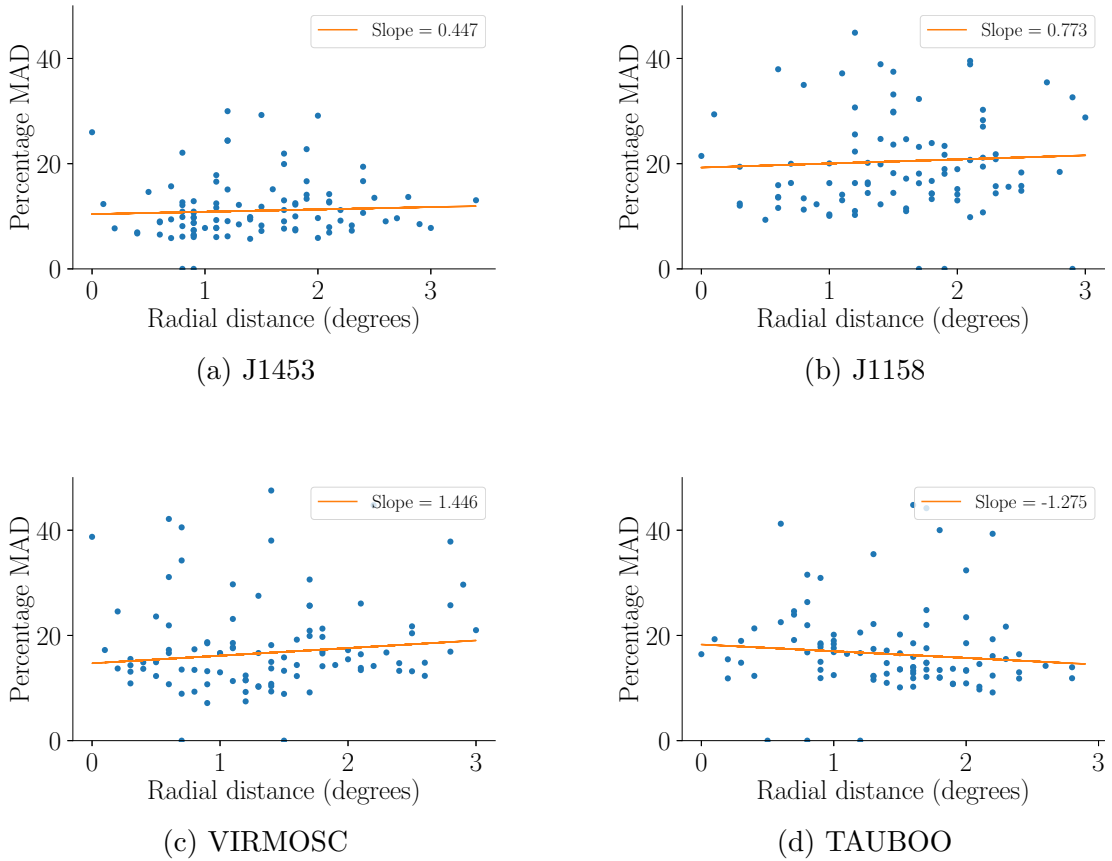


Figure 4.5: A plot of the normalised median absolute deviation (MAD) of flux densities of the sources within each field as a function of radial distance.

4.2.3 Flux density variation with parallactic angle

Primary beam azimuthal asymmetries may arise as a result of non-zero eccentricity in the antenna beam gain, or due to the presence of support structures in the path of the radiation (e.g. quadrupod holding up the feed).

Alt-azimuth mounted antennas rotate against the sky, as a function of time, while tracking a source. This results in the antenna beam pattern rotating around the pointing centre as a function of time. As the telescope tracks the field, an off-axis point source will trace a circular path around the pointing centre but the asymmetry of the beam will result in the receiver system detecting a variable amount of flux density from it.

The parallactic angle is the position angle of some ‘reference’ direction of the antenna on the sky. The change in the parallactic angle of the antenna will be the same as the amount of rotation of a source around the pointing centre. A schematic representation of the locus of a source around an elliptical beam in an alt-azimuth telescope and its apparent flux density is shown in Figure 4.6. The parallactic angle χ is obtained from (Perley *et al.* 1989)

$$\chi = \tan^{-1} \left(\frac{\cos(\lambda) \sin(h)}{\sin(\lambda) \cos(\delta) - \cos(\lambda) \sin(\delta) \cos(h)} \right) \quad (4.2)$$

where, δ is the source declination, λ is the latitude of the telescope, and h is the hour angle of the source.

Figures 4.8 - 4.11 show the relative variation of source flux as a function of parallactic angle. The 5 panels in each figure depict sources within successive annuli of 0.3° width from the pointing centre. All the sources in excess of 100 mJy have been shown in the plots. The Y-axis is the percentage flux density deviation for each source from its median value for the entire observation. The zero of the parallactic angle corresponds to the hour angle of the target location. The blue dots represent individual sources while the solid black line is their trend line obtained by a moving median filter.

Within each field the pattern seems to be broadly the same at all distances from the pointing centre. Even though the trend line is close to zero, the systematic shift of all the sources with parallactic angle suggests that the residuals (from improper subtraction of the source) can considerably increase the noise floor.

The expected variation of flux density with parallactic angle for an elliptical antenna beam is shown in Figure 4.7. While our observed patterns do indeed show systematic shifts as a function of parallactic angle we are unable to ascribe it to an elliptic beam pattern given the lack of similarity between the observed and the expected pattern.

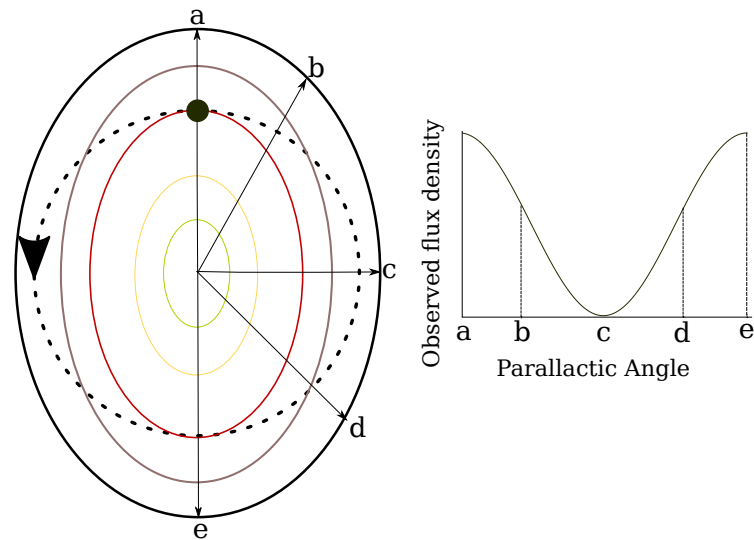


Figure 4.6: A schematic representation of the change in observed flux density of a source in an alt-azimuth telescope with an elliptical beam. The ellipses are contours of equal antenna beam gain; the dashed line is the locus of the source around the pointing centre; the arrow represents the direction of rotation of the antenna; and the directions a-e are the parallaxic angles. The plot shows the change in observed flux density at different parallaxic angles.




Figure 4.7: Representative percentage deviation of flux density from the median value for sources circling the pointing centre (at the same radius) in an alt-azimuth telescope with an elliptical antenna beam.

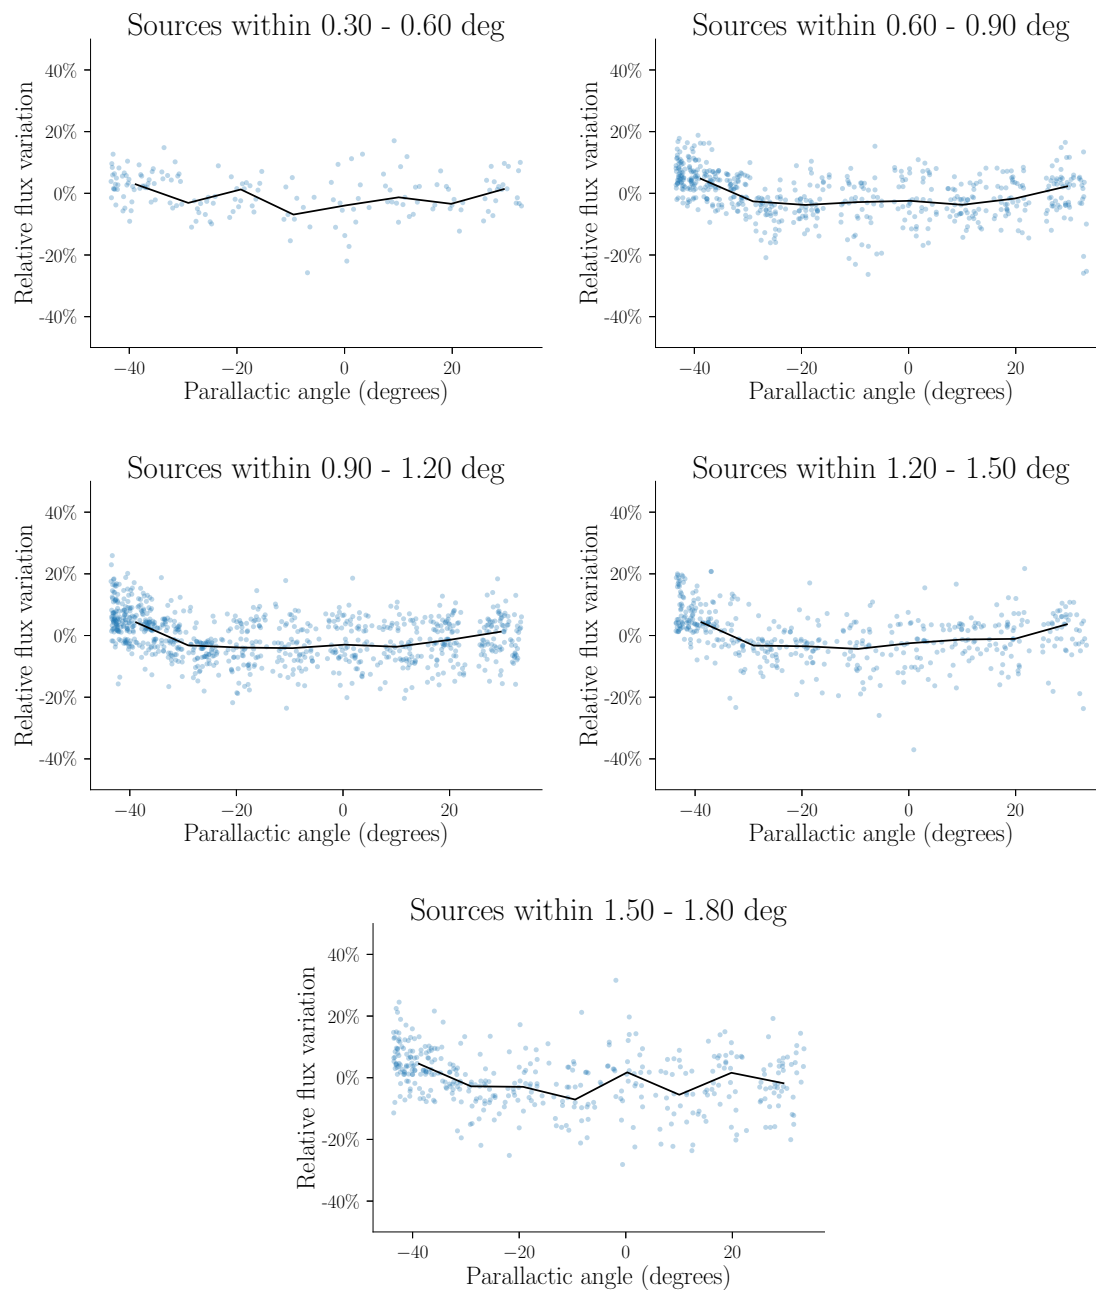


Figure 4.8: Percentage variation of the flux density about the median value as a function of parallactic angle for sources within concentric annuli in the J1453+3308 field.

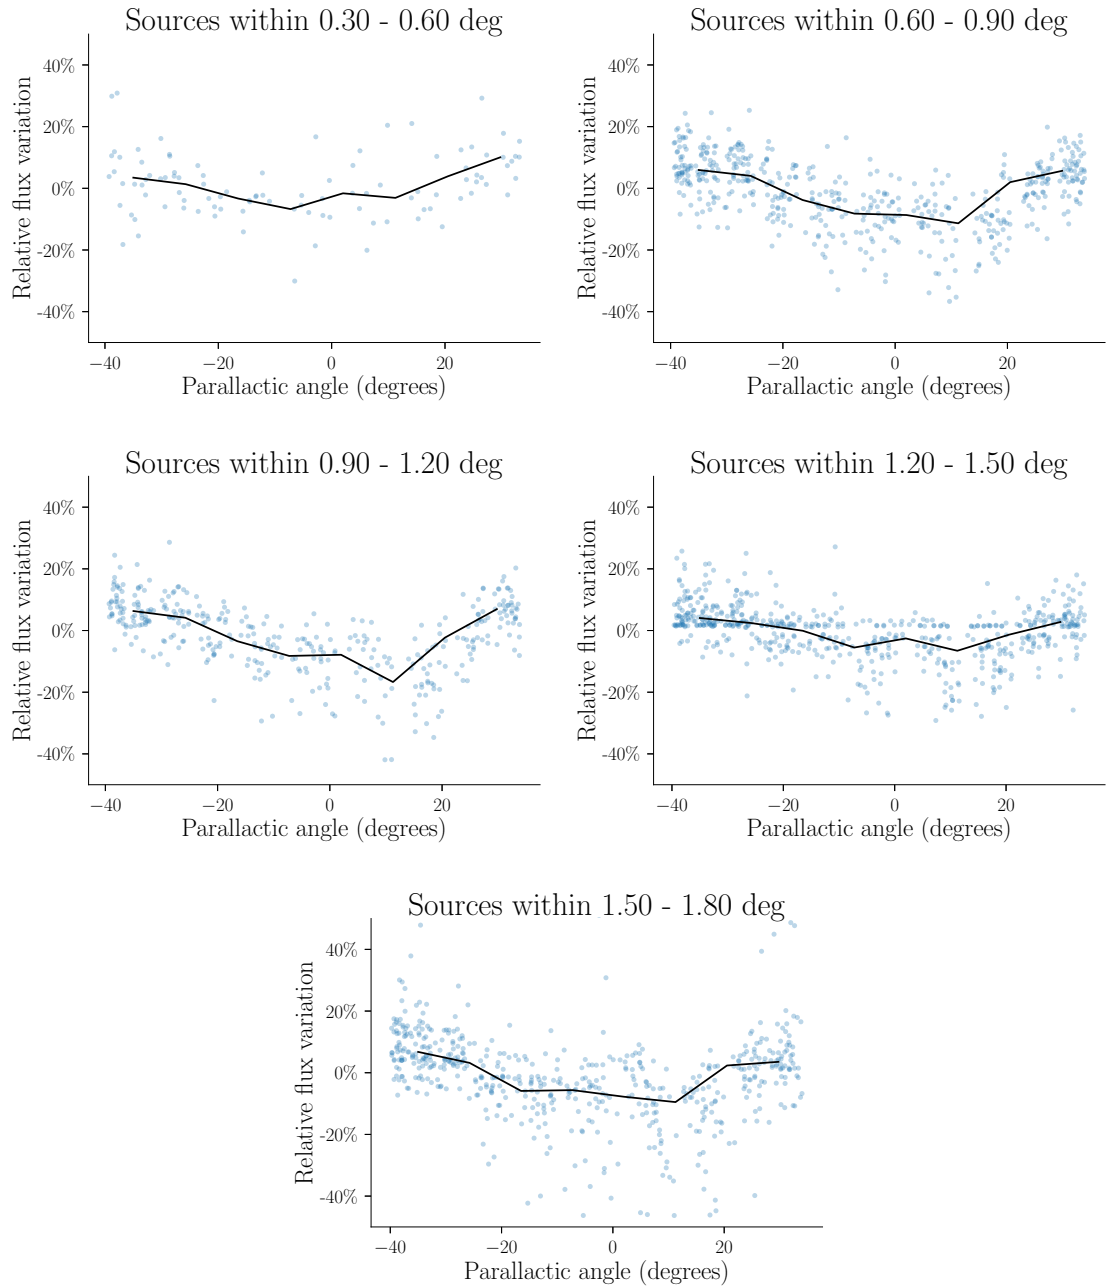


Figure 4.9: Percentage variation of the flux density about the median value as a function of parallactic angle for sources within concentric annuli in the J1158+2621 field.

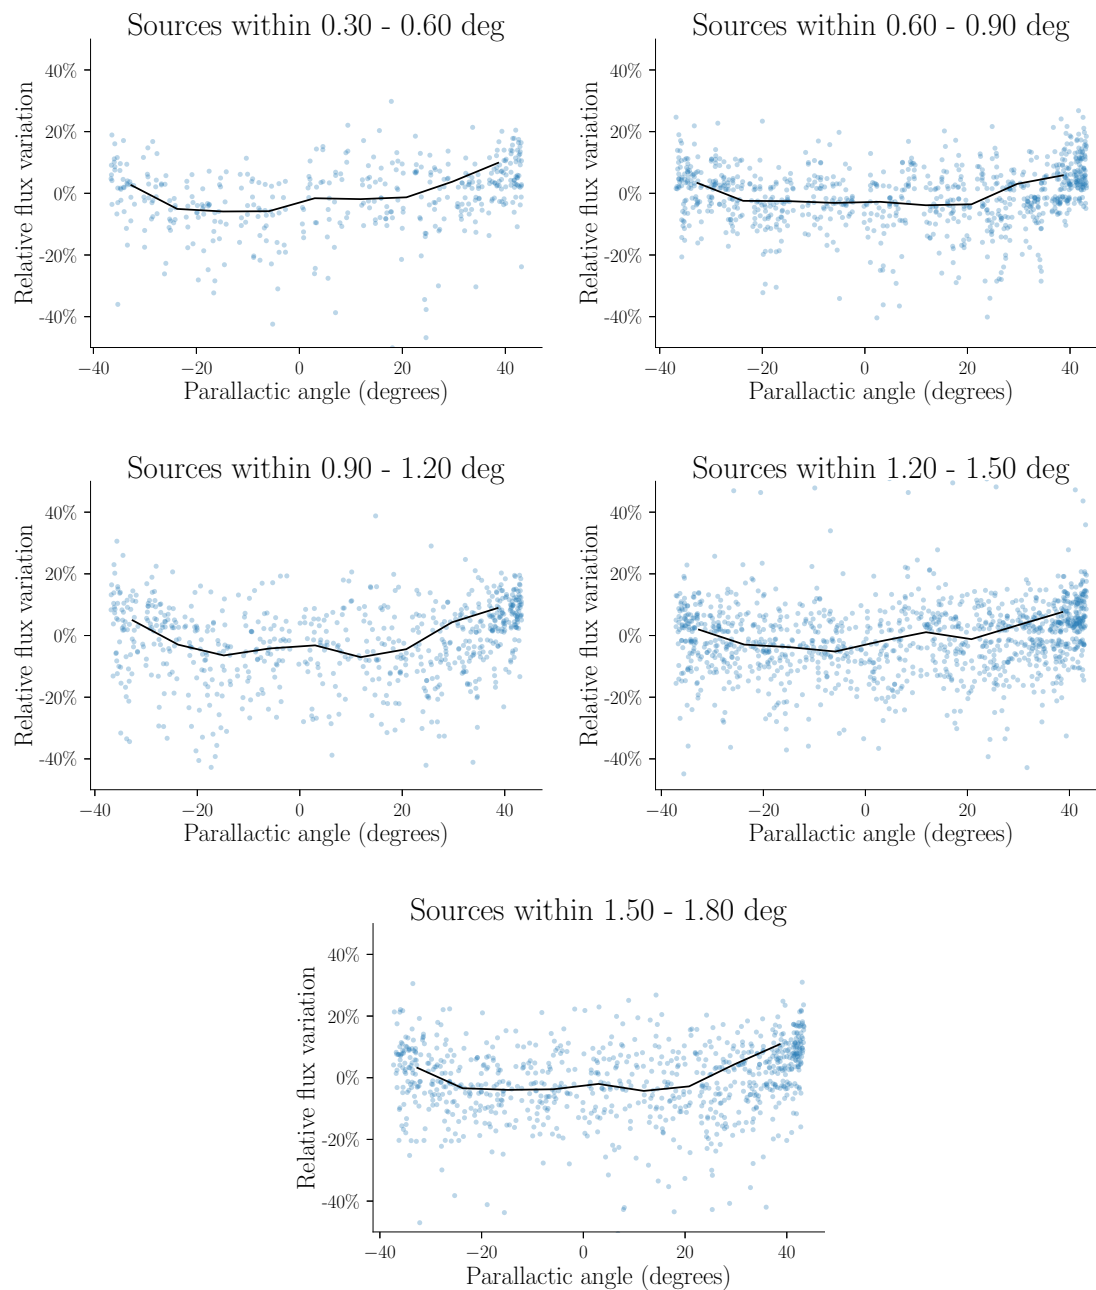


Figure 4.10: Percentage variation of the flux density about the median value as a function of parallactic angle for sources within concentric annuli in the VIRMOSEC field.

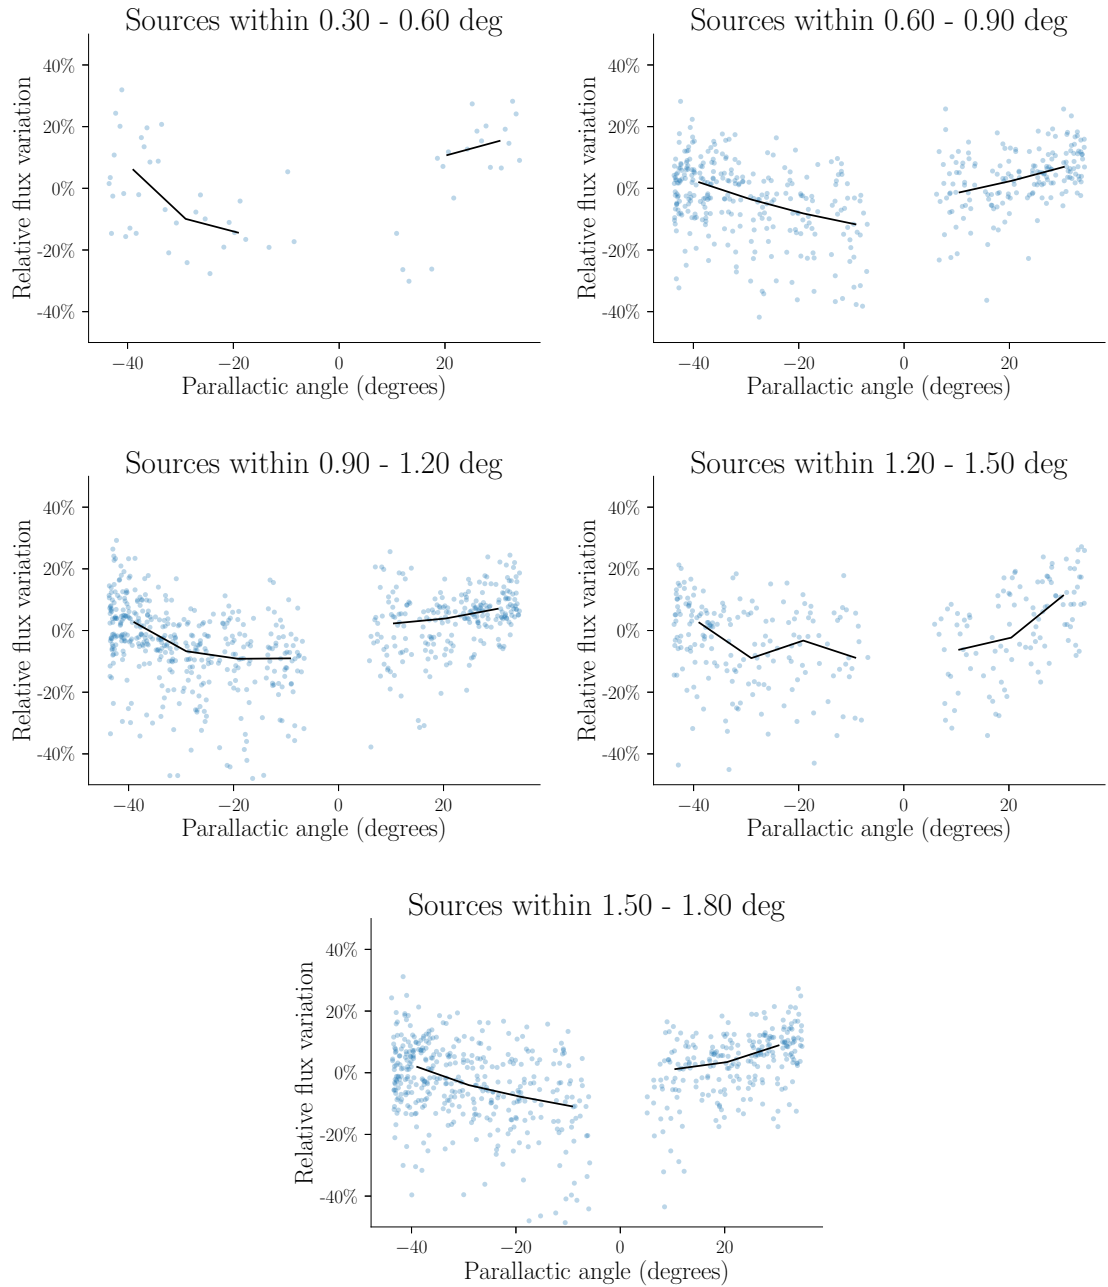


Figure 4.11: Percentage variation of the flux density about the median value as a function of parallactic angle for sources within concentric annuli in the TAUBOO field.

4.2.4 Flux density variation with elevation angle

The variation of antenna main lobe gain as a function of antenna elevation angle are usually called the “gain-elevation curves” (e.g., Greve *et al.* 1998). At lower elevations gravitational deformation can change the antenna structure and hence the gain. Figures 4.12 – 4.15 show the observed relative flux variation as a function of elevation angle. As for parallactic angle, the sources are separated into annuli of width 0.3° ; the blue points are the deviation of flux densities from the median values of individual sources in each snapshot; the solid black line is the trend line obtained from a moving median filter of width 10° .

4.2.5 Variable ionosphere

In addition to an asymmetric and variable primary beam gain, a time-variable non-isoplanatic ionosphere causes a different class of errors in the image plane. The ionosphere acts as a phase screen on the incoming radiation, and a differential change in the density of the ionosphere across the field of view can cause two different apparent errors - (a) an apparent shift in the position of the source and (b) a distortion of the source structure. The former error occurs when there is a linear gradient in the ionospheric phase screen across the array along a particular line of sight. However if there is a non-linear change in the ionospheric phase across the array, this can cause distortions in the apparent source structure. In such a situation, a time-average image will cause the source to appear ‘smeared out’.

Figures 4.16 – 4.19 show the histogram of the angular separation between the position measured in a snapshot and the mean position (measured from the time-averaged visibilities) for the same sources plotted in Figures 4.1 – 4.4. The median positional offset is typically smaller than 1 pixel (pixel size = $4.5''$ for all images), and much smaller than the beam size ($= 22''$). Even such small offsets can result in incorrect source subtraction and residuals sufficient to hinder the quality of the image we seek. These residuals can be taken care of in each baseline, or alternatively, the imaging process has to be modified to include an ionospheric correction (e.g., Intema *et al.* 2009).

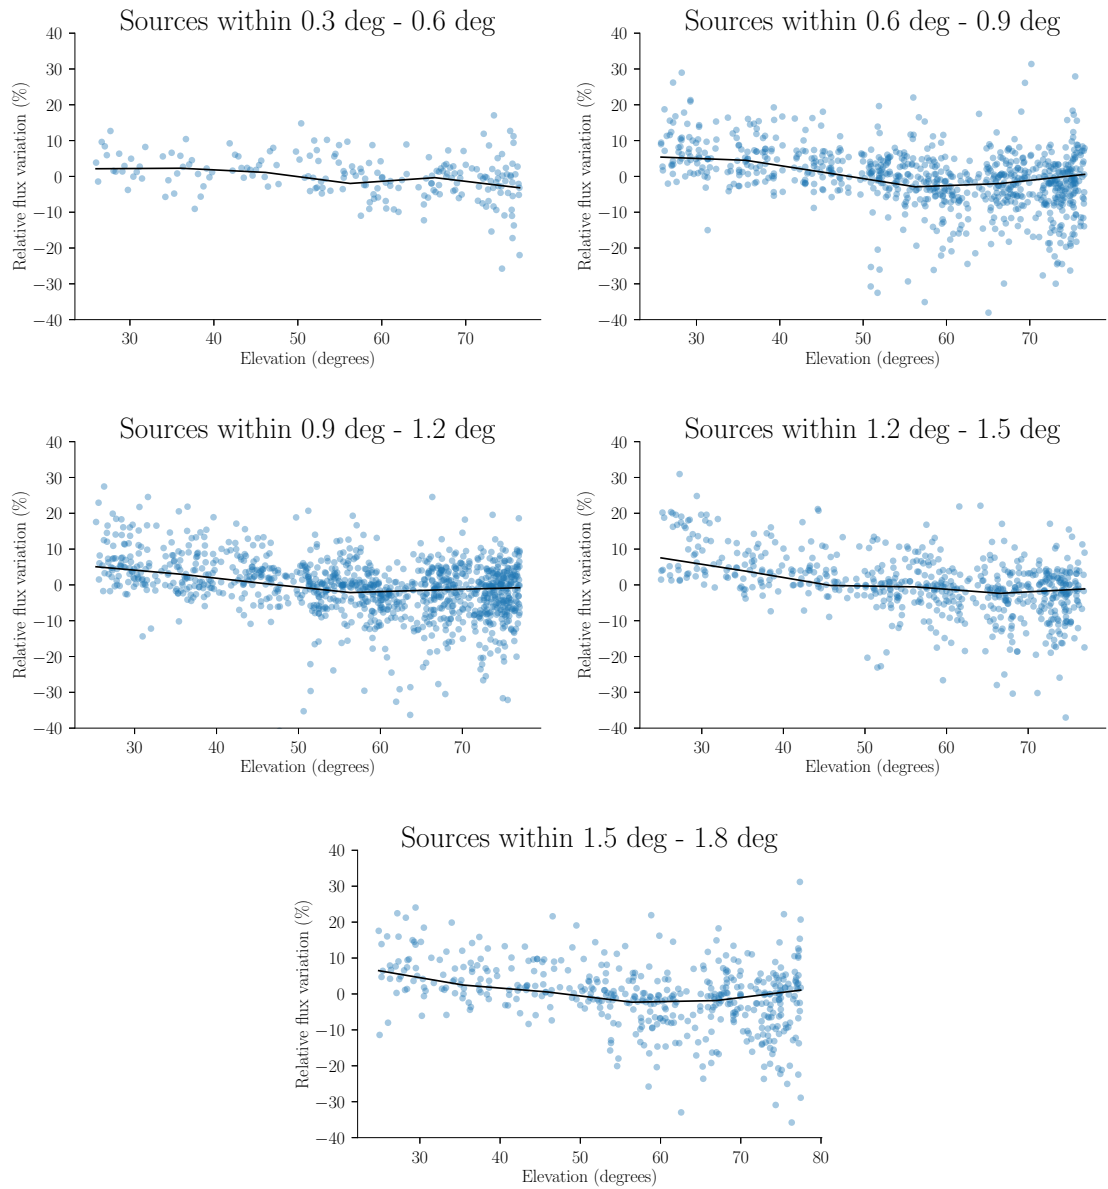


Figure 4.12: Percentage variation of the flux density about the median value as a function of elevation angle for sources within concentric annuli in the J1453+3308 field.

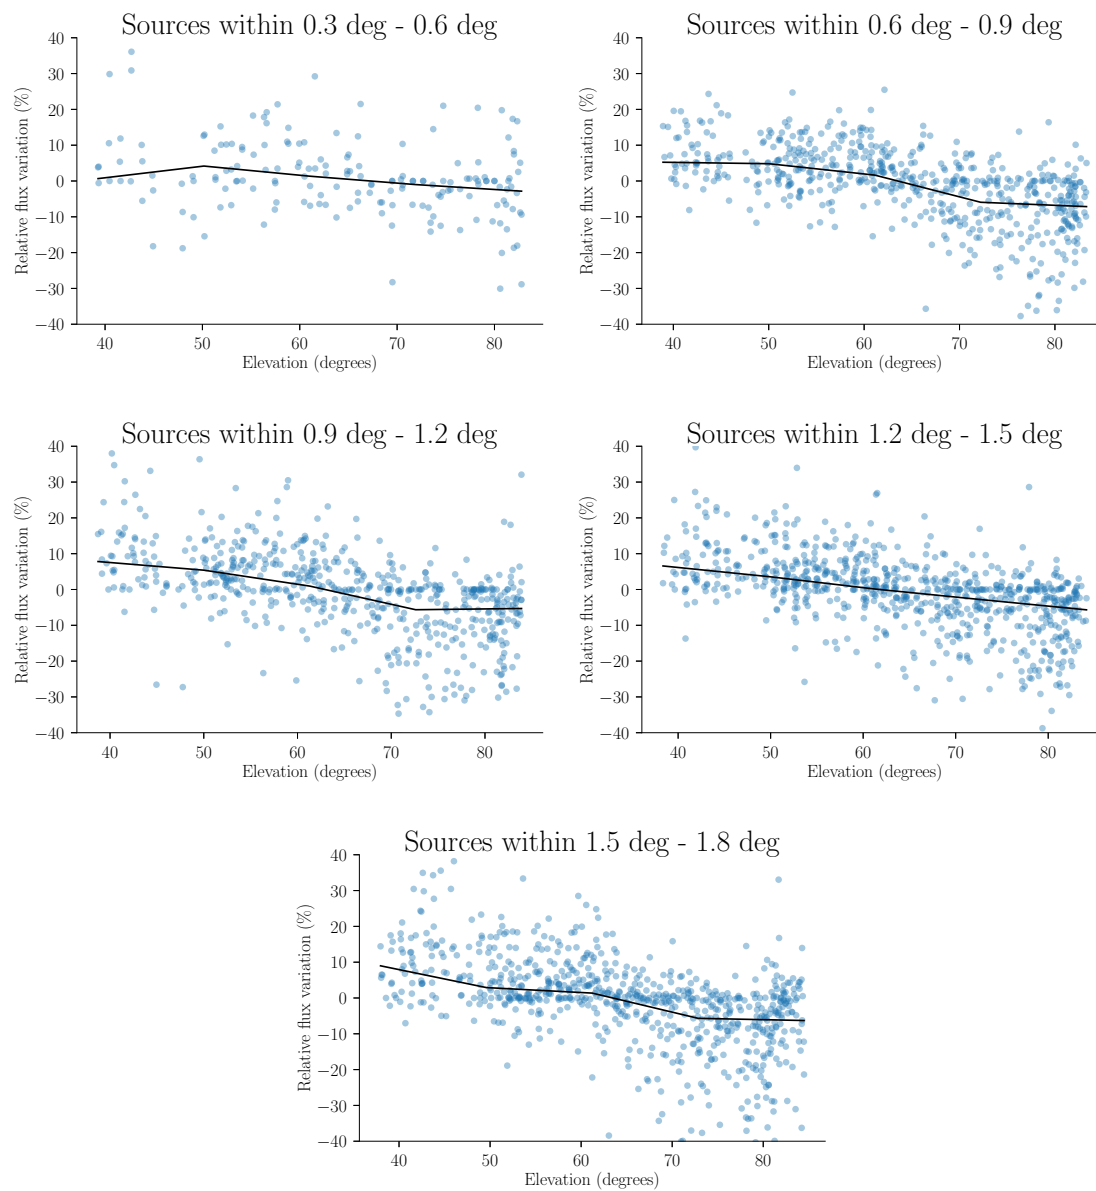


Figure 4.13: Percentage variation of the flux density about the median value as a function of elevation angle for sources within concentric annuli in the J1158+2621 field.

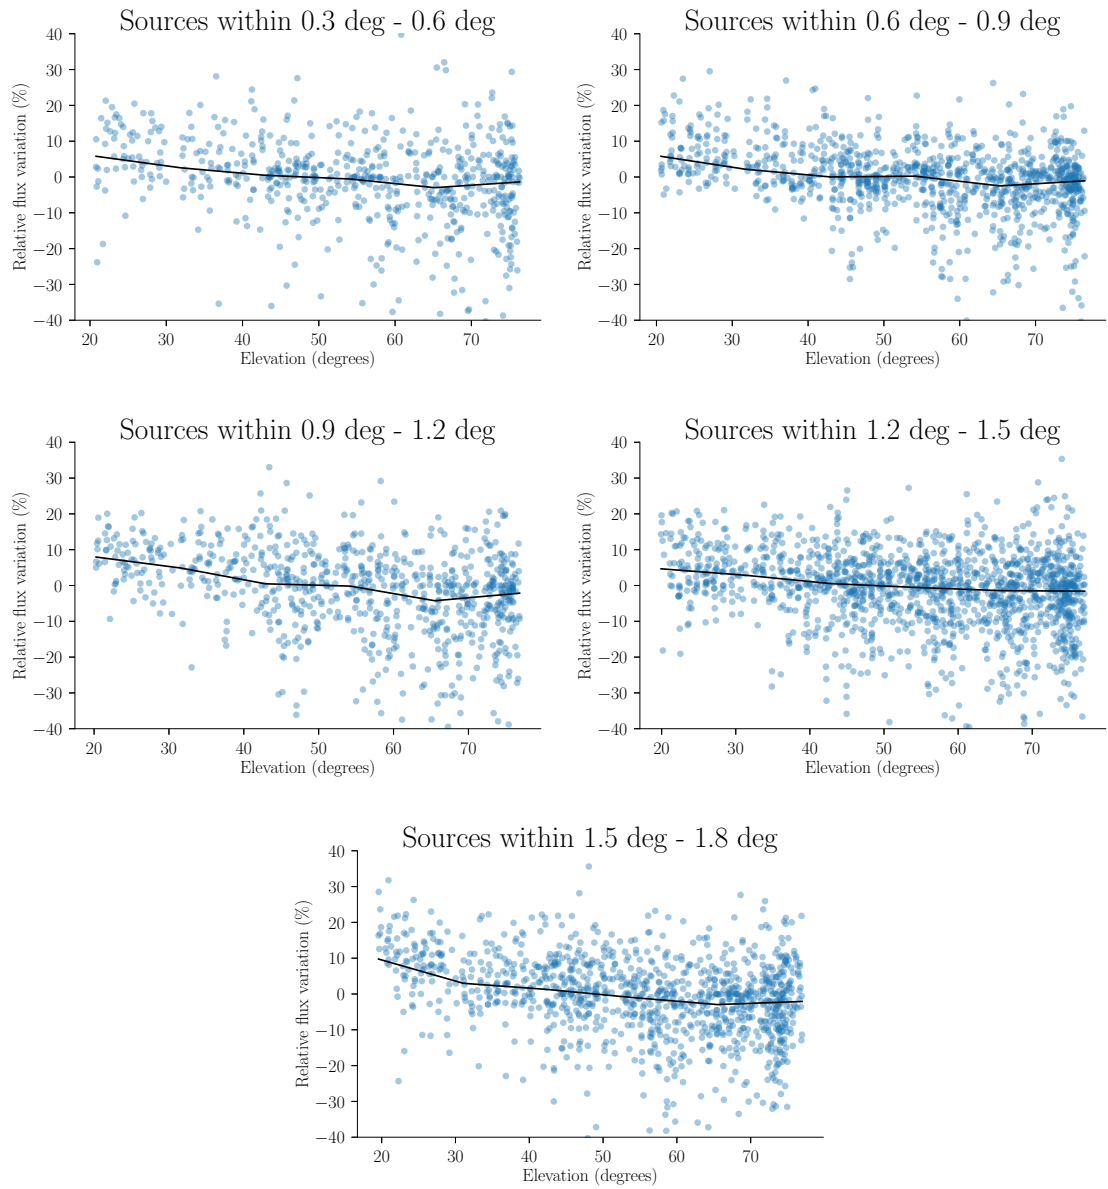


Figure 4.14: Percentage variation of the flux density about the median value as a function of elevation angle for sources within concentric annuli in the VIRMOSC field.

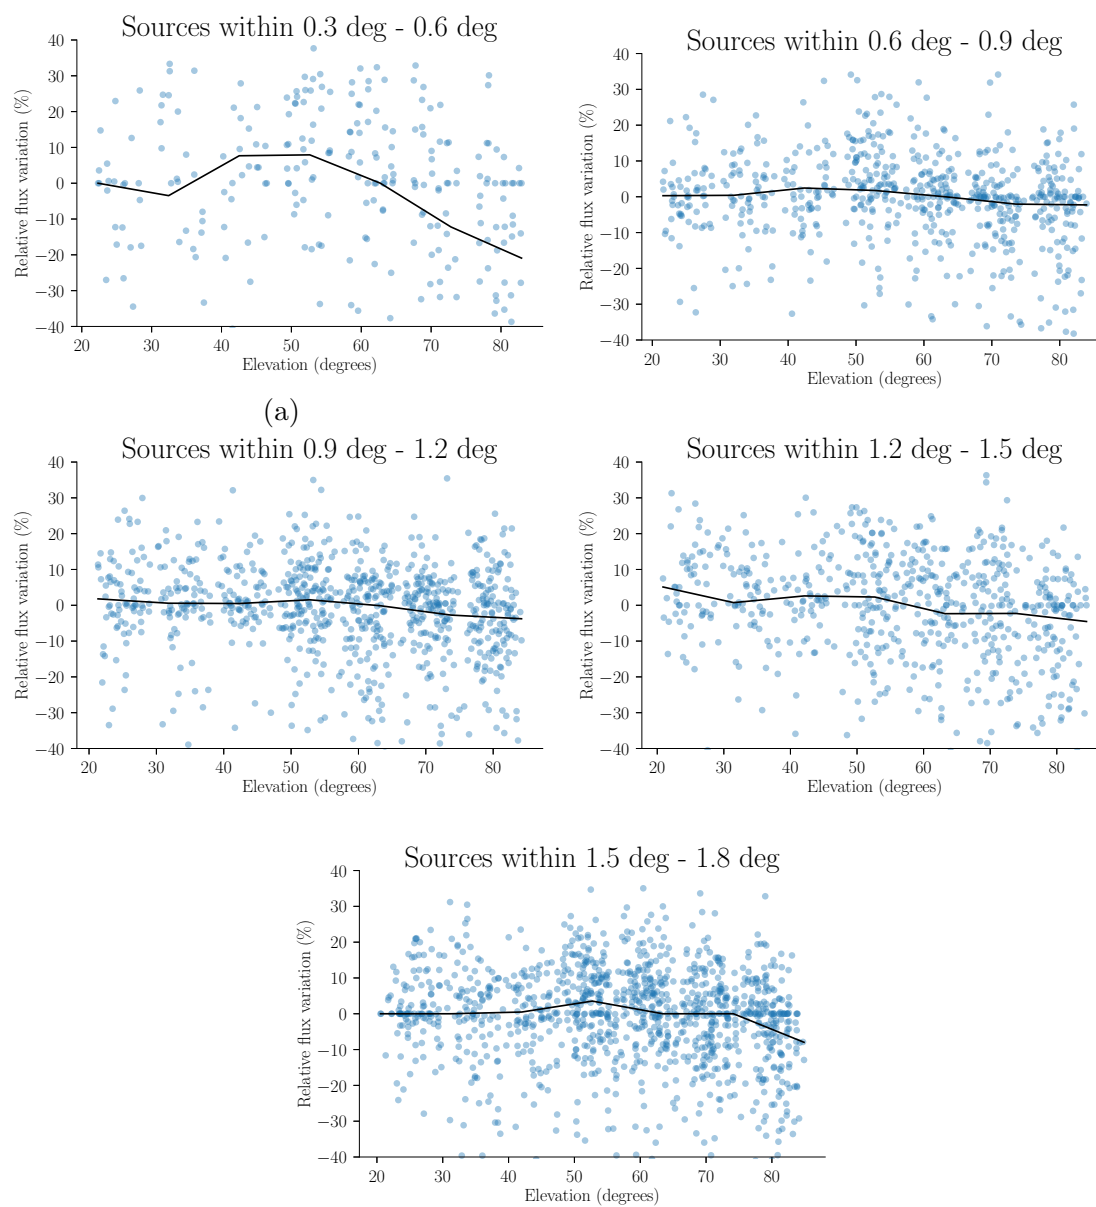


Figure 4.15: Percentage variation of the flux density about the median value as a function of elevation angle for sources within concentric annuli in the TAUBOO field.

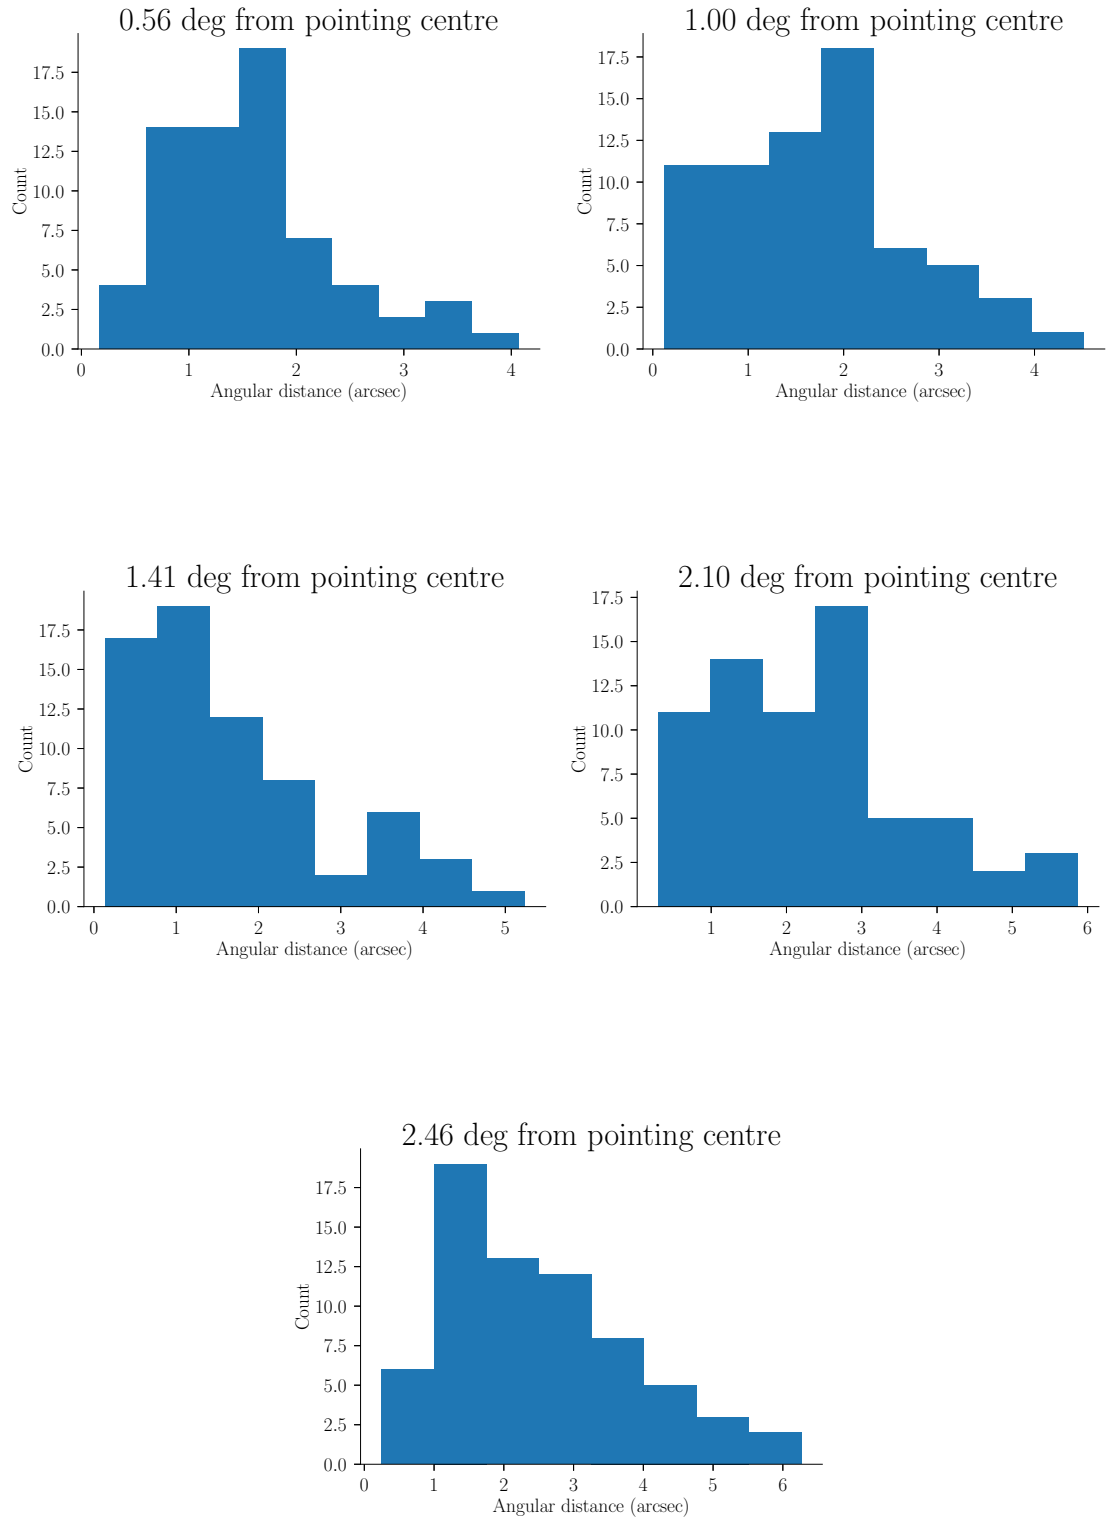


Figure 4.16: Histogram of the source position jitter for representative sources between 0.5° and 2.5° away from the pointing centre in the J1453+3308 field.

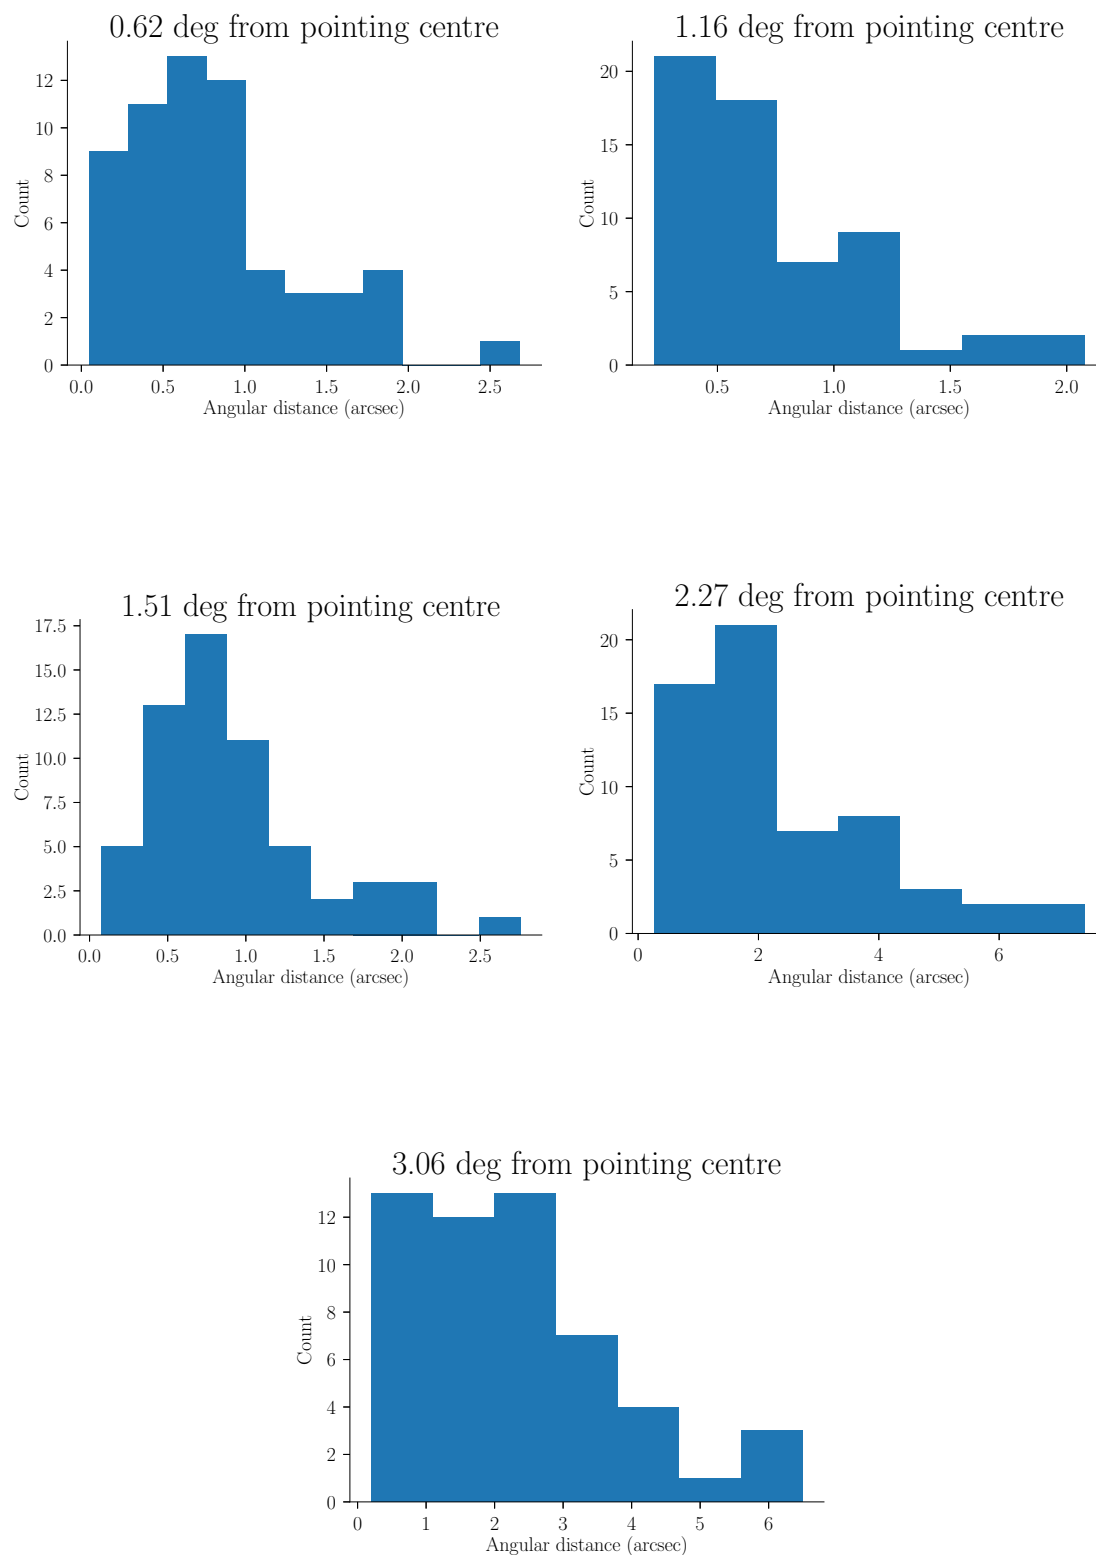


Figure 4.17: Histogram of the source position jitter for representative sources between 0.5° and 2.5° away from the pointing centre in the J1158+2621 field.

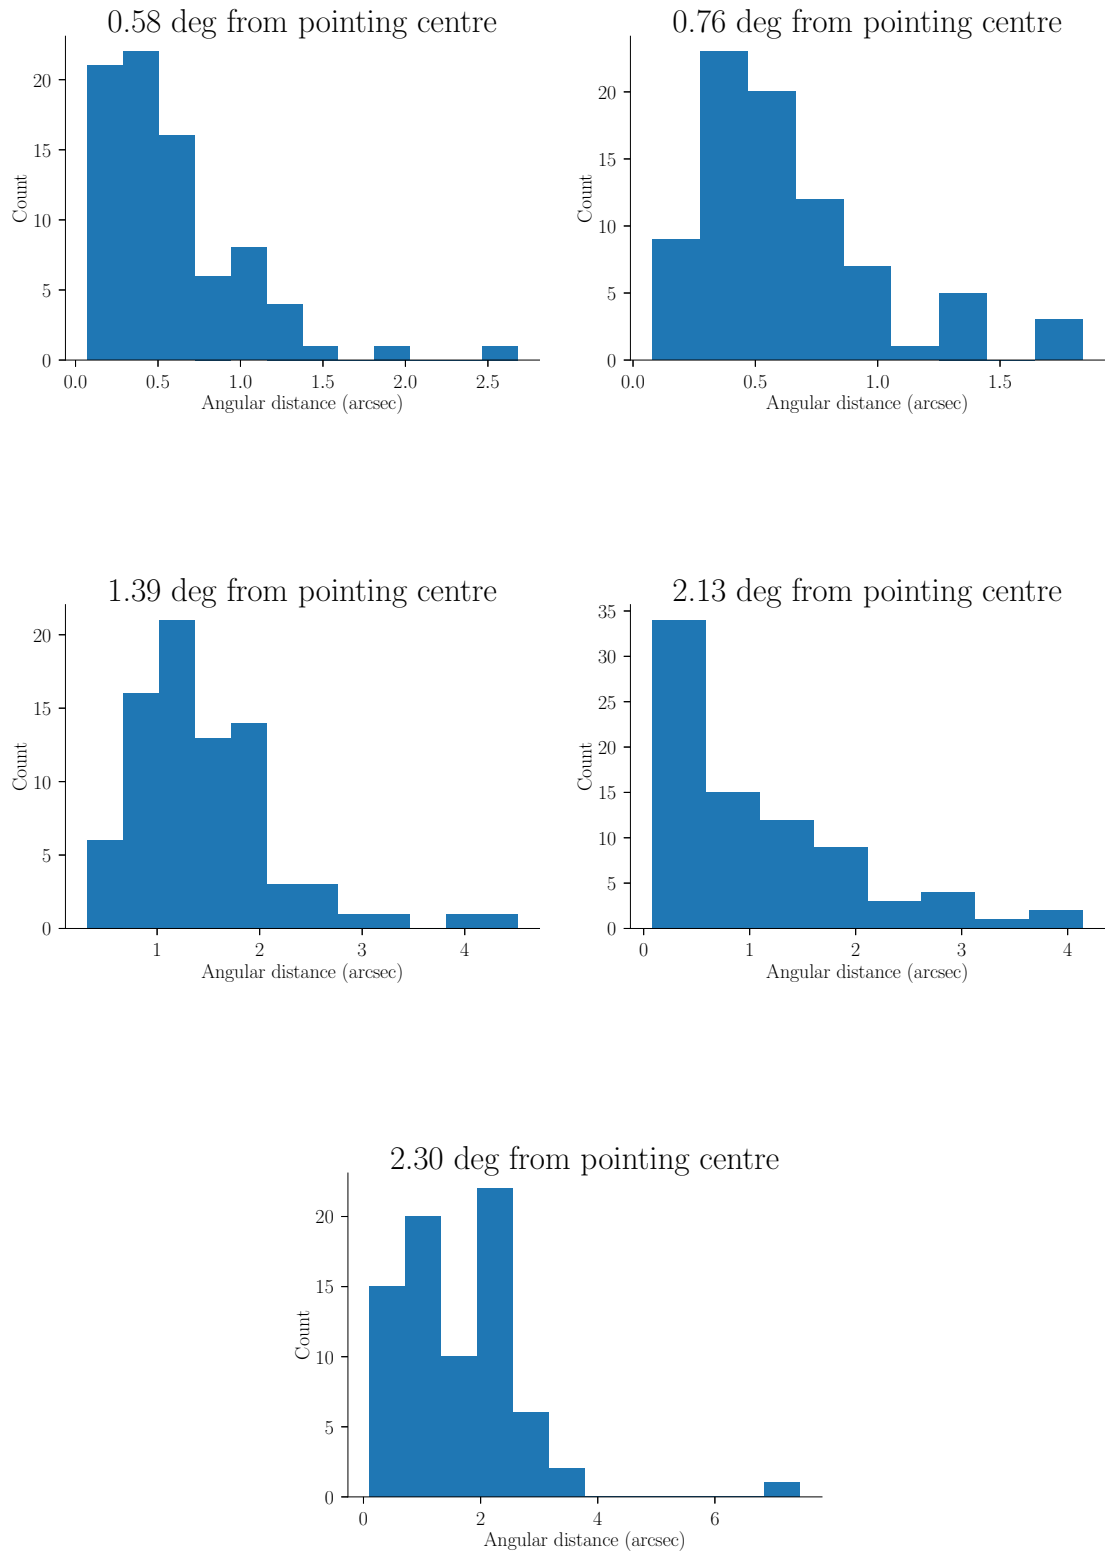


Figure 4.18: Histogram of the source position jitter for representative sources between 0.5° and 2.5° away from the pointing centre in the VIRMOSC field.

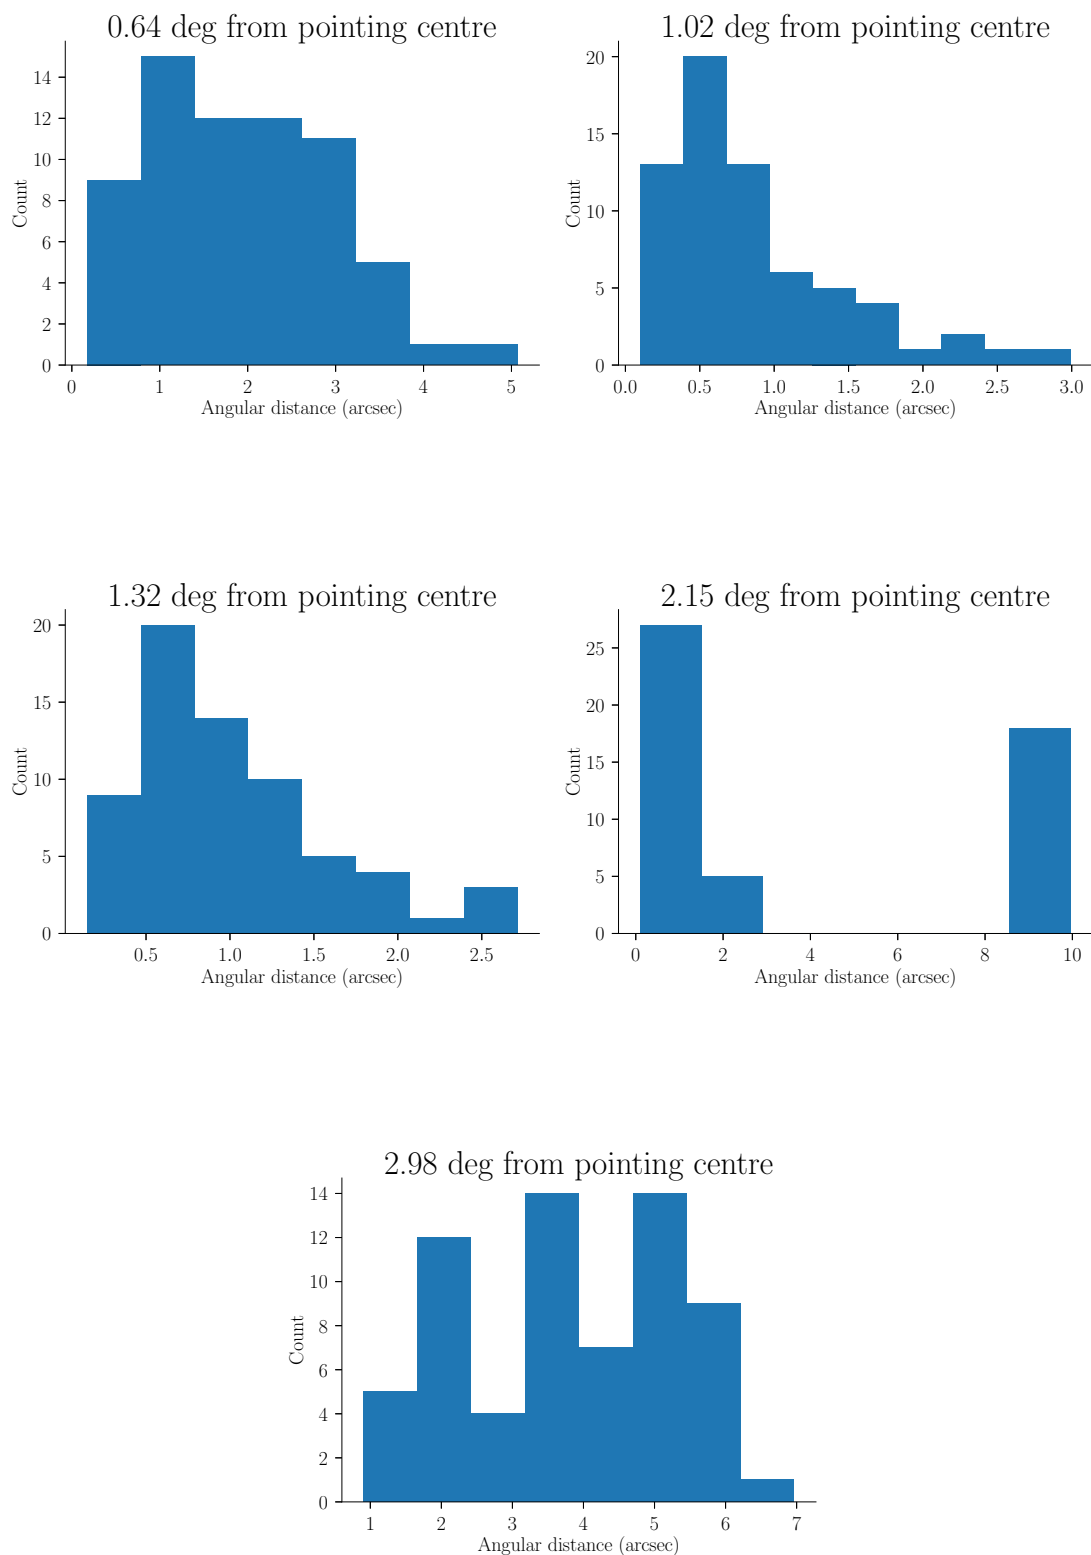


Figure 4.19: Histogram of the source position jitter for representative sources between 0.5° and 2.5° away from the pointing centre in the TAUBOO field.

4.2.6 Subtracting residual source flux variation

The previous subsections have shown that the observed source flux densities vary as a function of time, parallactic angle, elevation angle and ionosphere. Though we are unable to model the variation in terms of specific causative factors or processes, we started this exercise with the intention of mitigating the consequences without regard to the causes. Even if we were in a position to identify and model the causes with some degree of accuracy we may still not be able to eliminate them in the absence of other information — like accurate antenna beam patterns, a good ionospheric model, etc. Therefore, we provide a prescription for mitigating the consequences in order to improve the detectability of fainter sources.

Snapshot images are made using the residual (source subtracted) visibilities, to estimate the residual flux at the locations of the brightest 50 sources. The procedure is as follows:

1. Make an image using the entire (amplitude) calibrated data, and subtract this model from the visibilities to create the residual visibilities.
2. Identify a set of 25-100 of the brightest sources in the field
3. Make 5-10 min snapshot images from the residual visibilities and CLEAN strictly within small boxes at the locations of the bright source set.
4. Subtract the snapshot model from the original visibilities but only within the corresponding time range *i.e.*, the same time range used to create the snapshot image. The standard CASA task 'uvsub' subtracts every component from the entire visibility data. Therefore, we developed a CASA task which subtracts the components only from visibilities within the appropriate time range.
5. Add back the original, time-averaged model to the 'corrected' residual visibilities to obtain the error-mitigated visibilities.

This method is essentially sampling the measurement equation

$$V^R = G_p \left(\int \phi_p A_p \bar{B}_{pq}^R A_q^\dagger \phi_q^\dagger e^{-2\pi i(u_{pq}l + v_{pq}m)} dl dm \right) G_q^\dagger \quad (4.3)$$

every 5 minutes; where V^R are the residual visibilities, A is the antenna primary beam Jones matrix, ϕ is the ionospheric phase Jones matrix, and \bar{B}^R is the residual observed brightness distribution, which includes any other DDEs that may be present. The lm dependence is implied.

4.3 Baseline based defringing

The Snapshot source subtraction method described in the previous section subtracts from a residual source which is present in all the residual visibilities. Implicit in the snapshot imaging procedure is the assumption that within the snapshot interval every baseline sees an identical sky. As we have discussed earlier, this is not valid in many situations (a non-isoplanatic ionosphere, time dependent antenna pointing errors, etc.). This will result in the same source being seen at different levels of flux density by different baselines, even after the application of the Snapshot procedure. These errors will appear as fringes in the residual baseline visibilities. Some examples of these residual source fringes may be seen in Figure 4.20. These baseline-based errors will contribute toward the closure errors during self-calibration, and result in systematic artefacts in the image plane.

We propose here a method to mitigate these baseline-based errors, and in doing so simultaneously mitigate the errors due to a non-isoplanatic ionosphere as well as with an asymmetric antenna primary beam.

- Calculate the visibility fringe pattern corresponding to residual point sources at the location of the brightest sources and for a particular baseline-time-frequency.
- Fit for these model fringe patterns in the residual visibilities of each baseline-time-frequency while keeping the flux density and the phase offsets of the sources as free parameters. The phase offsets correspond to assuming that the residual sources are very close to but not necessarily exactly at the position of the cleaned bright sources.
- The fit is attempted within a time window of 5 minutes for sufficient signal-to-noise.
- The fit is discarded if it does not have a sufficiently large amplitude — more than 50 times the image noise — to avoid subtraction of any faint source structure. This is also ensured by limiting the model to the close vicinity of only a few of the strongest sources in the field.
- Subtract the fit from the residual visibility.

A source fringe will typically be a sinusoid with amplitude A , angular frequency ω and initial phase ϕ . It can be decomposed into the sum of a sine and cosine, with amplitudes a and b .

$$A \sin(\omega t + \phi) = a \sin(\omega t) + b \cos(\omega t) \quad (4.4)$$

The fit is then performed for the parameters a and b , which are related to the amplitude A and initial phase ϕ as

$$A = \sqrt{a^2 + b^2} \quad (4.5)$$

$$\phi = \arctan\left(\frac{b}{a}\right) \quad (4.6)$$


We used a linear least squares linear solver to fit for the values of a and b , and hence calculate the original amplitude and initial phase.

All the off-axis sources will exhibit an apparent flux variation, and will have associated systematic fringes in the residual baselines. However, the fringes of the fainter sources are usually well below the noise in the baseline visibilities. In practice, most of the residual fringes which impact the image quality are typically due to the brightest handful of sources. When we first addressed this problem we subtracted only the strongest source, which did not result in much improvement. Subsequently we tried to simultaneously fit model fringes of up to 20 sources, which also did not work because of the complex pattern of the combination of fringes, and the increase in the number of degrees of freedom. In our experience the fringes associated with the brightest 3-5 sources results in an improvement of about 20-30% in RMS noise in the vicinity of the strongest sources, and an improvement of a few percent elsewhere.


The standard GMRT integration time of 16.8s yields 18 time points in a 5 minute window. Each fit has 2 degrees of freedom, and therefore fitting for 5 sources results in 10 degrees of freedom over 18 data points, which is not ideal for noisy data. One will either have to use shorter integration times or reduce the number of sources in the model. One could consider fitting multiple spectral channels simultaneously with the same model. This did not work well in practice; perhaps it requires appropriate modelling of the spectral index due to the source and the frequency-dependent primary beam size.

Figures 4.20–4.22 show the residual and the fitted model fringes in the baseline time-channel plane. One can clearly see the correspondence between the structures in the two, even by eye. The quality and the nature of the fits may also be evaluated from the single channel fits in Figure 4.23 (good fits) and in Figure 4.24 (bad fits). In most cases, a poor fit usually comes with a small fit amplitude that gets rejected by our amplitude selection criterion. However, in a handful of cases the model fit, though poor, was subtracted because the amplitude was large. While we have not evaluated the impact of these cases we believe that their deleterious impact is minimal. This is because our model fits are independent from


channel-to-channel and baseline-to-baseline. Therefore, the probability of such a situation coherently introducing an artificial source, or erasing a real source, is vanishingly small. At worst we expect it to marginally increase the image noise.



(a)




(b)




(c)


Figure 4.20: Plot of the residual fringe in the time-channel plane of a baseline, and the fitted model. The colour indicates the amplitude of the real component of the visibilities.



(a)




(b)

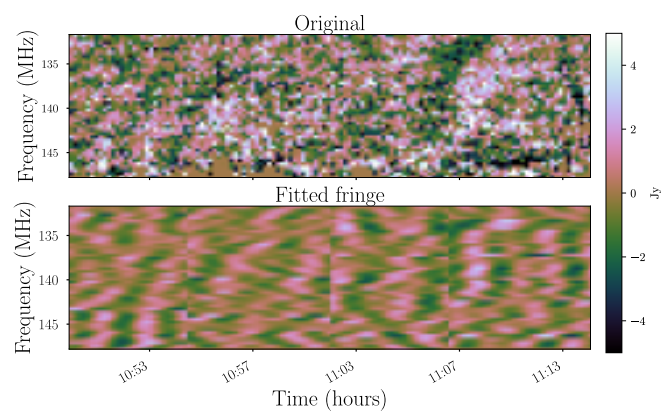


(c)


Figure 4.21: Residual fringe in the time-channel plane of a baseline, and the fitted model. The colour indicates the amplitude of the real component of the visibilities.



(a)



(b)



(c)

Figure 4.22: Residual fringe in the time-channel plane of a baseline, and the fitted model. The colour indicates the amplitude of the real component of the visibilities.

4.4 Discussion

The efficacy of the two algorithms described in this chapter can be seen in the reduction of artefacts around sources shown in Figure 4.25. However, in some cases the improvements are marginal (e.g. bottom panels in Figure 4.26) as the artefacts around the source are “frozen-in” to the self-calibration model. This is an unwanted side-effect of the CLEAN algorithm (e.g., Högbom 1974). Briefly, CLEAN is an iterative deconvolution algorithm that identifies the largest peak in the residual image and subtracts a fraction of the peak along with the associated dirty beam pattern. During this process, if the systematics around a strong source are large enough they will register as legitimate sources during deconvolution and will then be added to the clean image. This freezes the systematic errors into the visibilities as they will be considered as legitimate structure during the self-calibration process. One solution to this issue is to define a box within which CLEAN operates, that includes the legitimate source but excludes the artefacts around it. While this is feasible for a handful of sources, at lower frequencies the number of bright sources with noticeable artefacts around them grow too large to be handled manually. Identification and boxing of bright sources will have to be handled in an automated manner using source detection algorithms like pyBDSF (Mohan & Rafferty 2015). The box so defined is in general significantly larger than the source itself, and is usually large enough to include at least some of the surrounding artefacts.

We attempted to solve the direction-dependent errors by tackling them in the time domain. This works because the net effect of the DDEs is to induce an apparent time-dependent variation of source flux densities. Since CLEAN and its variants assume a time-invariant sky, this causes errors during imaging and self-calibration. We reduced the requirement of time invariance from the entire span of the observations to just 5-10 minute intervals. The other methods commonly used to correct for DDEs attempt to modify the imaging algorithm itself, by making the algorithm “smarter” and able to track the apparent variability in either the visibility amplitude, phase, or both for a particular source in the sky by providing additional information in the form of the beam shape, or the ionospheric phase variability across the field etc. These other algorithms are conceptually quite complex, and computationally intensive.

It must be emphasised that our procedures do not improve the dynamic range of the bright sources at all. Nor do they undo the damage to the structure of faint sources caused by DDEs. What they do achieve, without much increase in computation or algorithmic complexity, is the reduction of systematic patterns in the image caused by improper subtraction of the bright sources. This should

further improve the calibration solutions and hence image quality. We suggest that this is a very useful intermediate step while telescopes await improved and robust versions of algorithms which will target the root causes of DDEs.

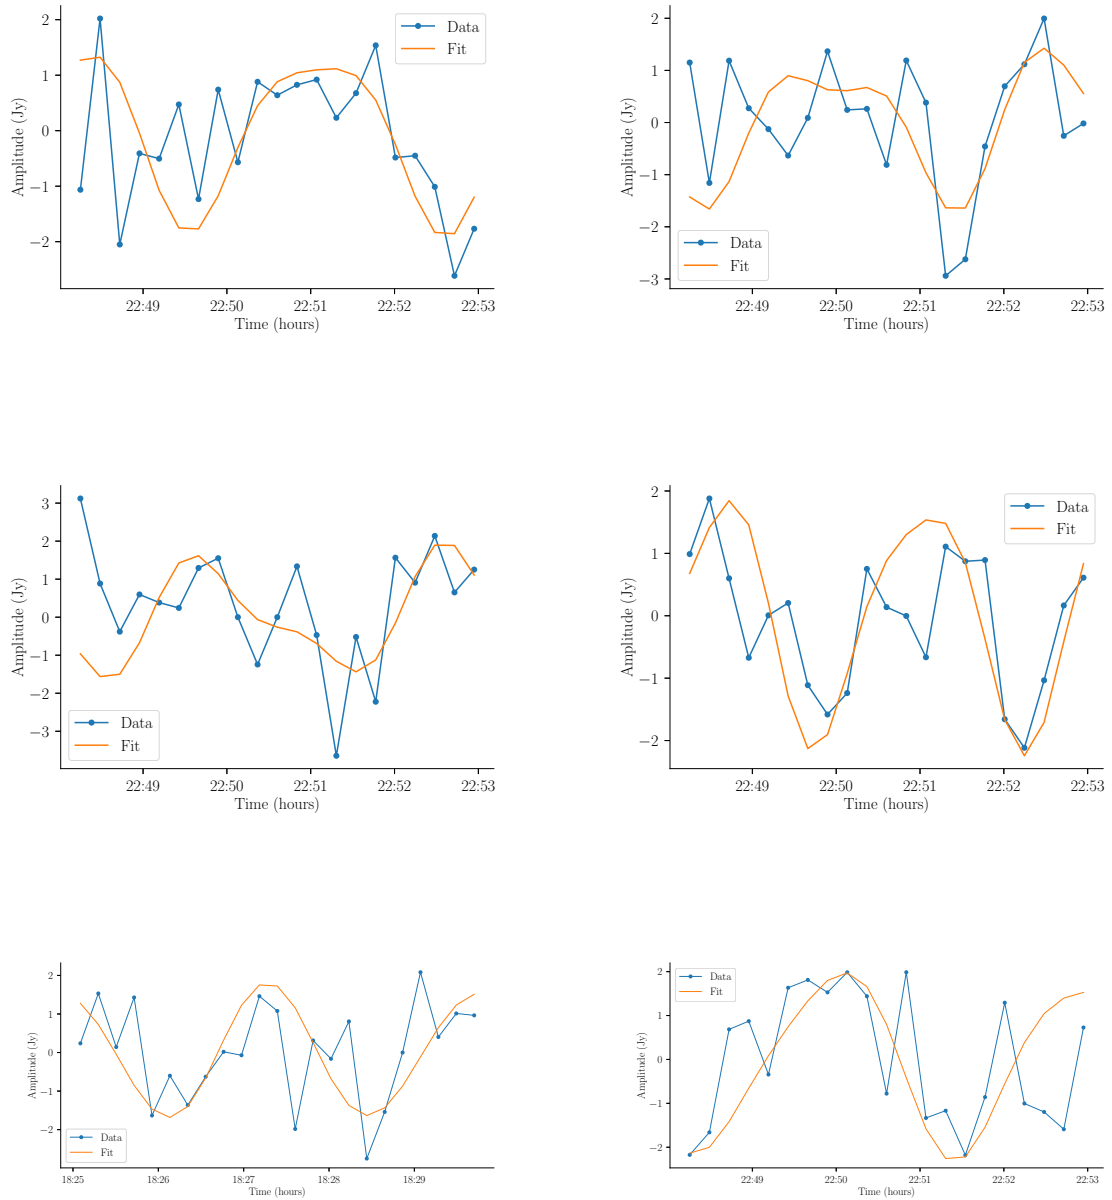


Figure 4.23: Good fits of the model fringe (orange line) to the observed data (blue line) from a single baseline-channel.

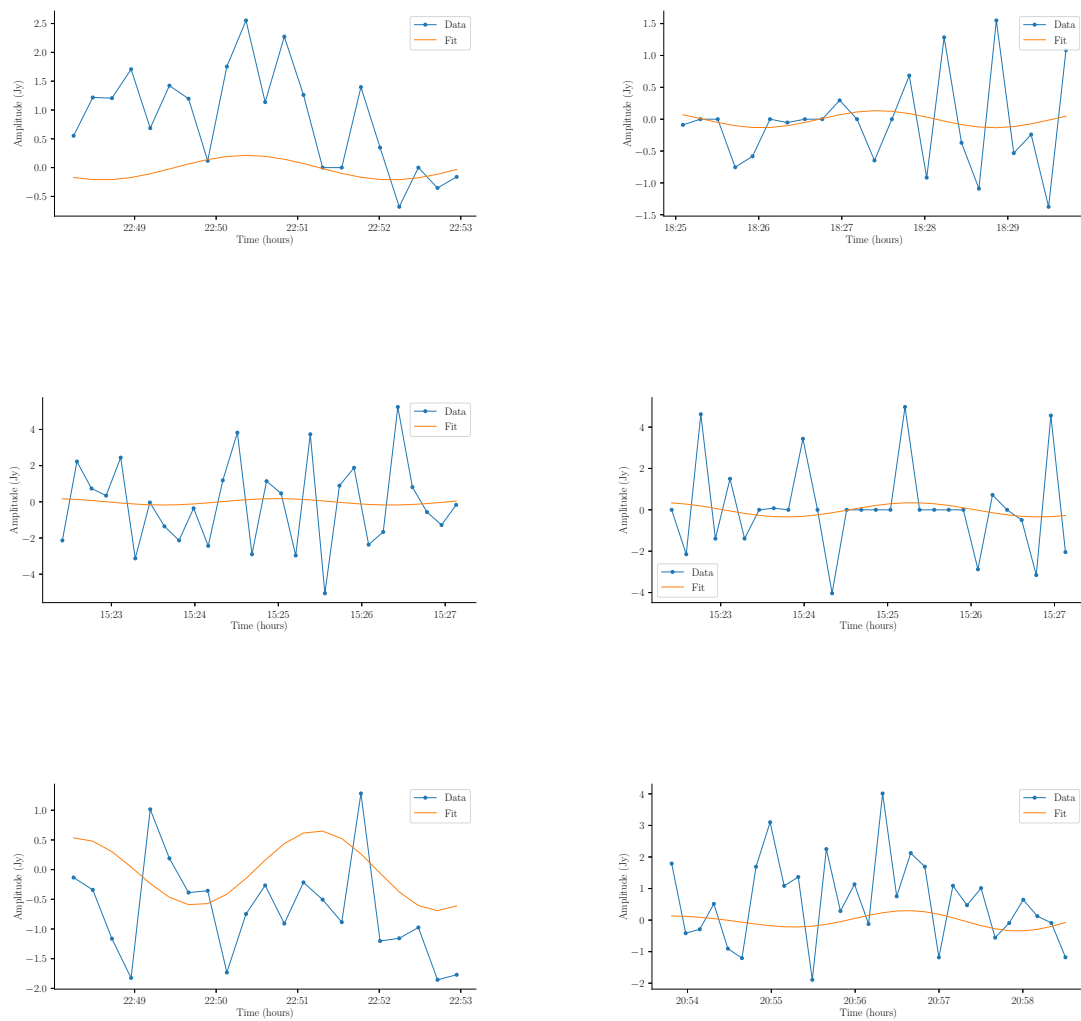


Figure 4.24: Poor fits of the model fringe (orange line) to the observed data (blue line) from a single baseline-channel.

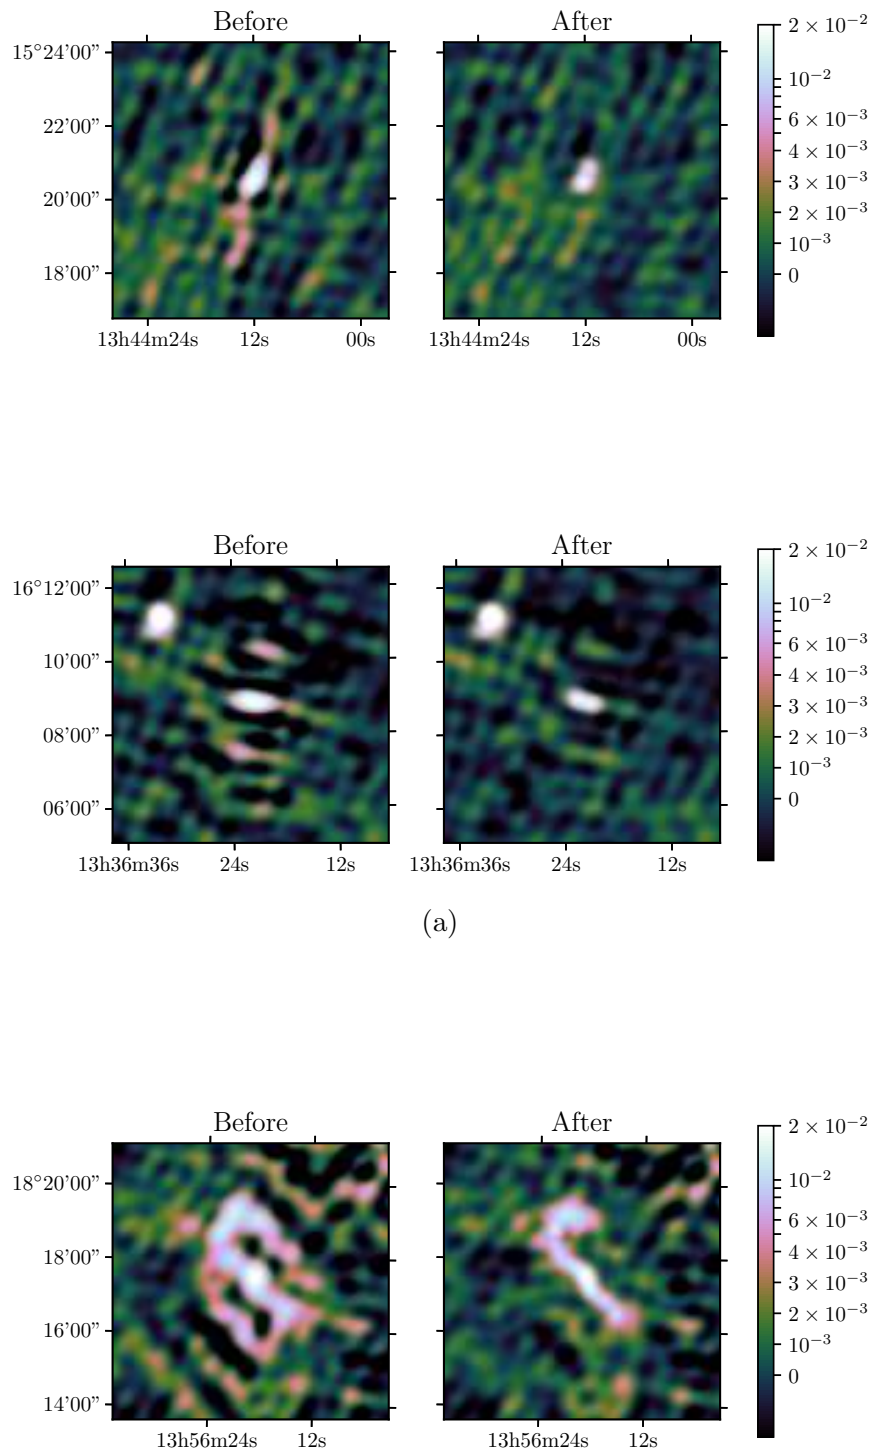


Figure 4.25: Examples of improvement to source structure after the DDE mitigation algorithms were applied.

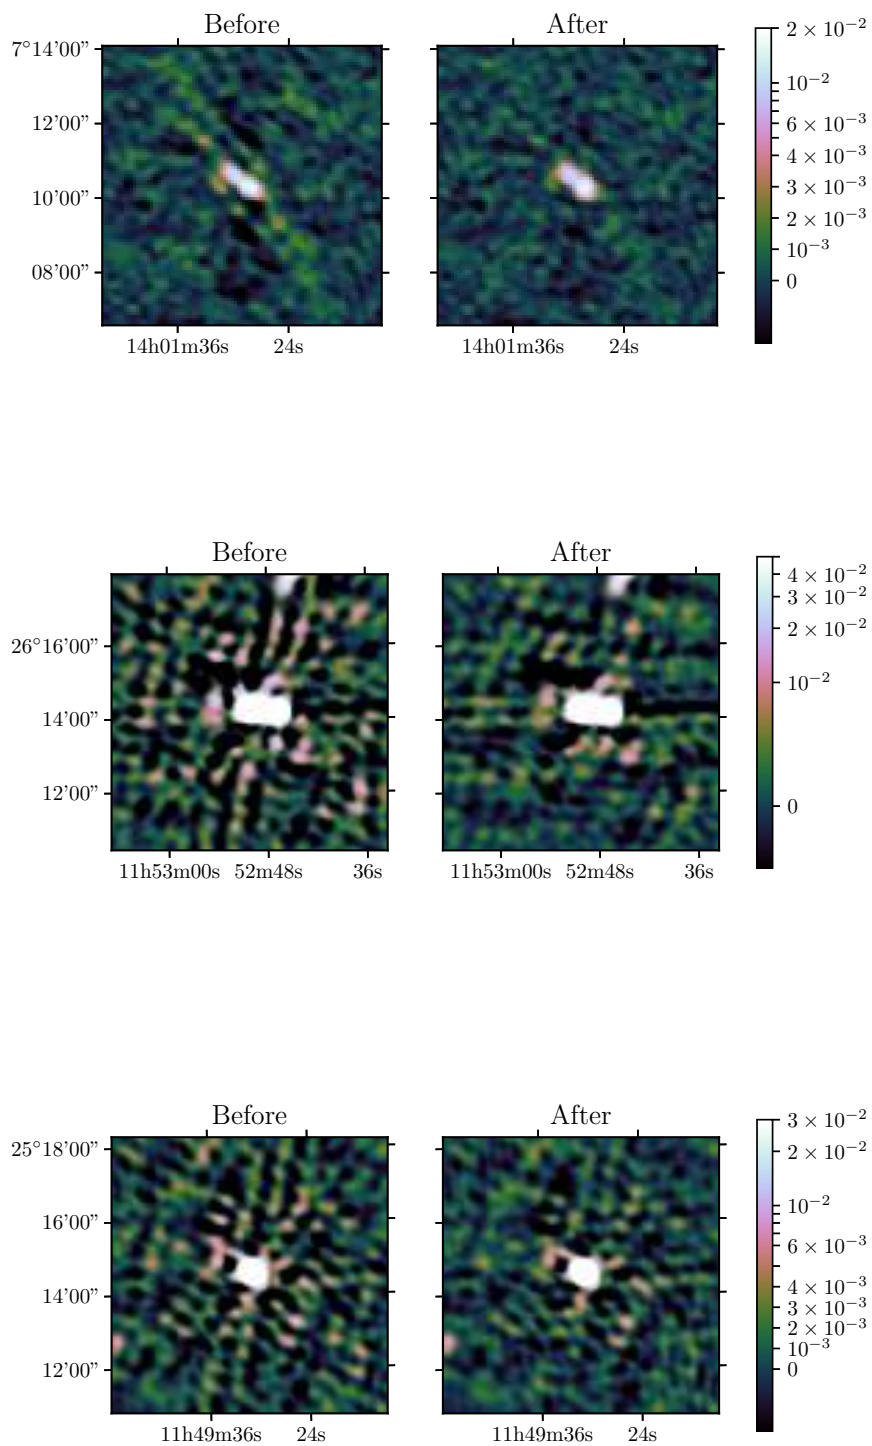


Figure 4.26: Examples of improvement to source structure after the DDE mitigation algorithms were applied.

Chapter 5

Evolution of spectral index in restarting AGN

5.1 Introduction

Active galactic nuclei (AGNs) are regions at the core of galaxies that have a much higher than average luminosity across large parts of the electromagnetic spectrum (Osterbrock 1991). AGNs are classified as ‘radio-loud’ or ‘radio-quiet’ depending on their luminosity at radio wavelengths. Often, the energy generated in the central regions of the AGN is transported far outside the host galaxy and radiated in the form of radio waves from extended lobes of synchrotron plasma. In this chapter we present a study of radio-loud AGNs which show multiple epochs of activity.

Around 20% of all AGNs are radio loud, meaning that they have a radio (5 GHz) to optical (B-band) flux ratio greater than 10 (Kellermann *et al.* 1989). Radio loud AGNs are typically described by the Fanaroff-Riley classification (Fanaroff & Riley 1974), in which they are either ‘Type I’ (FR I) or ‘Type II’ (FR II). This classification differentiates the radio source on the basis of their luminosity — FR I radio sources occur with luminosity $L < 2 \times 10^{25} \text{ W Hz}^{-1} \text{ sr}^{-1}$ (at 178 MHz) while FR II sources are more luminous than this. FR I sources typically reduce in brightness as a function of distance from the core, whereas FR II sources have an ‘edge brightened’ morphology in which the brightness toward the edge of the radio lobes is comparable to or larger than that of the core.

The differences in luminosity and morphology of these two sources seem to be a consequence of the amount of energy and the efficiency of its transport from the core to the lobes. FR I galaxies, which have a lower luminosity, have jets that decelerate while traversing the galactic medium. The jets are initially relativistic (Biretta *et al.* 1995), and the entrainment of the galactic medium across

the turbulent interface of the jet leads to the deceleration of the jet to subsonic speeds (Laing & Bridle 2002). On the other hand the jets of FR II type sources remain supersonic until they terminate at the lobes (Wardle & Aaron 1997). The termination location of these jets are often referred to as ‘hotspots’, since they are associated with a much higher luminosity than the rest of the lobe. Hotspots are the site of active energisation in a radio lobe and hence have a flatter spectral index (Blundell *et al.* 1999). The hotspots are usually present at the edge of the lobe, and represent the site at which the shock that defines the outer edge of the lobe is forming (Miley 1980). Thus the FR II type sources have two well defined radio lobes that are (usually) terminated by the hotspot, and the lobes are connected to the core via the radio jet. Such sources are also known as ‘classical doubles’. To summarise, FR I sources have (a) lower luminosity, (b) subsonic jets and (c) no well defined lobe structure whereas FR II sources have (a) a larger overall luminosity, (b) supersonic jets to several hundred kiloparsecs (or even a few megaparsecs in some instances) and (c) well defined lobes with hotspots.

There is evidence that radio-loud AGNs go through phases of quiescence, and exhibit multiple epochs of activity. Some studies indicate that long-lived AGNs are the exception, rather than the rule (Franceschini *et al.* 1998). Therefore if most radio-loud AGNs exhibit some form of restarting behaviour, where they periodically cycle through active and quiescent phases, one must observe a sample of AGNs that are currently in their restarting phase. For example, there is some evidence that Hercules A has gone through several epochs of activity (Gizani *et al.* 2005; Gizani & Leahy 2003).

We are interested in the class of restarting AGN that show two clear epochs of activity known as ‘double double radio galaxies’ (DDRGs) since they show two clear pairs of double-lobes that are morphologically similar to FR II sources. The outer pair of lobes are from the initial epoch of jet activity and are no longer being actively energised by the radio jets, whereas the inner lobes are formed in the newer epoch of activity (Schoenmakers *et al.* 2000). During the initial phase of activity, the radio jet impacts upon the inter-galactic medium (IGM) and the momentum transport is balanced by the ram pressure of the IGM. This leads to the formation of a shock, since the jet is supersonic, and once the jet has passed through this shock it inflates a cocoon of low density, high pressure gas that protects the jet flow from turbulence (Blandford & Rees 1974). In the case of a DDRG the outer lobes are no longer being energised by the jets and therefore the hotspot and the shock will dissipate once the jet has turned off. If the jet is turned off at time t_0 at the core, the lobe will continue to be energised until time t_1 which is the jet travel time between the core and the hotspot. After time t_1 the lobe will no longer evolve as an FR II source, but rather as a ‘dead’ radio source (Brocksopp *et al.*

2010). Once the jets are restarted, it is observed that in almost all cases the inner jets also show an edge-brightened morphology. Before discussing the evolution of these dead lobes, which is the focus of this chapter, we will first discuss the possible scenarios that can cause the restarting behaviour in the AGN.

5.1.1 Quenching of the radio jet

In most of the observed DDRGs there exists a relatively high degree of symmetry of the inner (newer) lobes, regardless of the degree of symmetry of the outer lobes. The fact that both pairs of lobes are on either side of the core, as well as the degree of symmetry in the inner core points toward processes at the nucleus of the radio source that trigger the formation of DDRGs (Schoenmakers *et al.* 2000). We summarise here the three scenarios outlined in Schoenmakers *et al.* (2000) to explain the quenching of the radio jet -

1. A change in the jet axis - If the jet axis abruptly changes, then the outflow will again impact on the IGM and start forming lobes with an edge-brightened morphology. This is not particularly likely in the case of DDRGs since most DDRGs have a very good alignment ($< 30^\circ$) between the axes of the outer and inner lobes. However this is a possible scenario to explain X-shaped radio sources (Capetti *et al.* 2002; Lal & Rao 2006).
2. Backflow instability - When the head of the jet starts to slow down from supersonic to subsonic speeds, instabilities are created in the backflow of the lobe. If the instabilities are strong enough they can eventually pinch off the jet channel, cutting off the lobe from the jet flow. Subsequently this will give rise to a new shock front which may lead to a lobe- and hotspot-like structure. While this adequately explains the alignment between the inner and outer lobes, it does not explain the symmetry of the inner lobes. In this model for the inner lobes to be symmetric the pinching off of the outer lobes must happen at similar times on either side of the core. This is very unlikely, since the medium through which each jet is propagating is different and therefore the properties of the instabilities will also be different.
3. Jet interruption - The jet power of FR II sources is not constant, and varies as a function of time. If the jet is completely quenched, the jet channel will collapse since there is no more active material to balance the pressure against the external medium. In such a scenario if the jets are to restart, they will again impact against some external medium and consequently lead to the formation of shock fronts and hotspots (*i.e.*, FR II characteristics).

Simulations indicate that the lobe is overpressured compared to the surrounding medium (Reynolds *et al.* 2002), but the jet material is less dense than the surrounding medium. Therefore once the jet is switched off (accompanied by a possible collapse of the jet channel itself), the lobe will expand until pressure is equalised and subsequently will enter a ‘bouyant’ phase where the lobe material will continue to move away from the core, while carrying with it entrained material from the surrounding medium. These simulations indicate that the surrounding IGM (which is relatively denser than the dead lobe) will re-occupy the regions between the core and the lobe, particularly since the lobe will tend to move away from the core. This scenario seems to be the only likely scenario that can explain the edge-brightened morphology in the restarting jets. If the jets were instead propagating into the rarer lobe plasma, Schoenmakers *et al.* (2000) suggest that the appearance of an edge-brightened morphology is instead due to the re-energisation of old lobe plasma.

5.1.2 Restarting the radio jet

The mechanism of the restarting jet is still poorly understood. The radio jet itself is thought to arise from the accretion disk surrounding the central black hole in a galaxy (Blandford & Payne 1982; Chakrabarti & Bhaskaran 1992). In these models, the jet collimation is provided by the angular momentum ‘outflow’ from the accretion disk, and is accompanied by a corresponding loss of angular momentum of the disk itself. From the previous section, various possibilities are put forward to explain the quenching of the radio jets. Assuming that the radio jet is quenched by some physical process, Kaiser *et al.* (2000) propose that one particular scenario for the restarting of the radio jet is the infall of a large amount of matter onto to the central black hole. Under the assumption that the angular momentum of the radio jet determines the jet axis, such an infall can significantly alter the angular momentum and therefore there is no reason to expect the axis of the new jet to be aligned with the old. A situation where the new jets are highly misaligned from the old jets could result in the formation of the ‘X-shaped’ radio sources (e.g., Dennett-Thorpe *et al.* 2002; Lal & Rao 2006). These sources are known to have jet powers that are close to the FRI-FR II cutoff (*i.e.*, $5 \times 10^{25} \text{ W Hz}^{-1} \text{sr}^{-1}$), which is similar to the powers of most known DDRGs (Saikia & Jamrozy 2009). While this is indicative, it is in no way conclusive that the two processes are related.

5.1.3 Time scales of quiescence

There are two methods used to estimate the age of the inner and outer lobes of the DDRGs and subsequently estimate the timescale of quiescence. The ages can be estimated in the following ways:

1. **Dynamical age** - This is the age of the source, assuming the expansion speed of the radio hotspots and given the linear size of the source. The linear size is straightforward to estimate from the measured angular size of the source, but the expansion speed of the hotspots is an assumption. Typical speeds of expansion are usually of the order of a few percent of the speed of light (e.g., Scheuer 1995).
2. **Spectral age** - The spectral age measures the time since the last injection of energy into a particular location. In a classical double FR II source, the hotspot location in the lobe is assumed to be the site at which energy is dumped into the surroundings. The spectral age is determined by fitting the $\log \nu - \log S$ curve, and determining the ‘break frequency’ which determines the age of the source. There have been critiques of this method, questioning the accuracy of the resulting ages (Blundell & Rawlings 2001), but it still remains one of the most popular ways to estimate the age of the source.

From the estimates of the age of the outer and inner lobes, constraints can be put on the period of quiescence. However the spectral and dynamical ages do not always agree, and can differ by up to a factor of two or more (Machalski *et al.* 2009). This could be due to re-acceleration of the radiating electrons which results in a smaller measured value of the spectral age as compared to the dynamical age.

5.2 Evolution of the spectral index structure of dying radio lobes

We investigated the evolution of the spectral index structure in the outer lobes of DDRGs. The outer lobes are over-pressured, adiabatically expanding plasma that are no longer being injected with energy from the radio jet. The radiative lifetime of such plasma is thought to be of the order of $10^6 - 10^7$ yr (Blundell & Rawlings 2000) which means that in the absence of other physical effects the radio lobes will fade in a few tens of Myr. However this is contradicted by the measured spectral ages of sources which can be of the order of a few 100 Myr. Blundell & Rawlings (2000) suggests that this issue can be resolved by considering either (a) A lobe magnetic field that is substantially lower than equipartition, or

(b) Turbulent mixing of the lobe plasma, such that electrons of different ages and from different regions of the lobe mix.

The evolution of the spectral index structure of a lobe undergoing adiabatic expansion (and possibly turbulent mixing) is not well understood. One expects that after the radio jets are switched off and the last energised electrons cross the terminating shock of the jet, that the hotspot will quickly dissipate. This is in general expected, since the highest energy electrons radiate away the most energy. The spectral index structure for a classical FR II type galaxy varies smoothly, with the flattest spectra found at and around the hotspot and the spectrum steadily steepening as one moves away from it. *The question then is how long would the spectral index structure still indicate the location of the hotspot?* If the spectral index still has a definite structure with respect to the location of the old hotspot (*i.e.*, the ‘warmspot’) this gives us some indication as to the level of mixing that happens within the lobe. A combination of the preferential radiation from high energy particles, the adiabatic expansion and the turbulent mixing will eventually result in a homogenisation of the spectral index structure of the entire lobe. Therefore studying the evolution of the spectral index structure as a function of time can provide insight into the timescales of these phenomena.

We were hoping to apply the wisdom from this study to other systems like radio halos which have no discernible hotspots nor any other obvious sources of energisation, but nevertheless have to have been subjected to some process of particle reacceleration. We expected to generate high quality spectral index images using the improved imaging techniques described in the previous chapters, but our images showed no systematic trend in spectral index away from the putative hotspots.

5.2.1 Sample of double-double radio galaxies

There are about 100 DDRG candidates known at present (Nandi & Saikia 2012), of which some 50 have been confirmed to exhibit restarting activity; the remaining require additional higher resolution images for confirmation. A study of the evolution of the spectral index structure of the dead outer lobes require that the lobe must be many times larger (in angular size) than the synthesised beam size. Therefore we selected DDRGs whose radio lobes were known to be at least twice as wide as the GMRT 150 MHz synthesised beam. The lengths of the lobes were typically much larger than this. The sample was also selected such that the galaxies are at different stages of evolution after the initial jet was switched off. The selected sample of 9 DDRGs were observed at 3 frequencies using the GMRT: 150, 325, and 610 MHz bands with either 16 MHz or 32 MHz bandwidth. The 610 MHz

and 325 MHz data were obtained from the GMRT online archive (except for one DDRG, 3C293, for which no 325 MHz data existed). The 150 MHz observations (and of 3C293 at 325 MHz) were obtained through the standard call for telescope proposal process.

Some of the galaxies did have existing observations at 150 MHz but the data had been observed using the ‘standard’ 16s integration time of the GMRT. Effective excision of the strong RFI at 150 MHz (e.g Athreya 2009), essential for making high quality images, requires much shorter integrations. Therefore, we observed all the targets at 150 MHz using 2s integration (which is the lowest ‘standard’ integration at the GMRT). Telescope time was granted to us in Cycle 27 (proposal code: 27_063) and Cycle 29 (29_062). Table 5.1 summarises some of the source properties, such as the redshift and source size (linear and angular).

Of the 9 DDRGs observed for our sample, 3C452 turned out to be a classical double (FR II) type source, rather than a DDRG. The images of Sirothia & Wiita (2013) had shown a pair of very steep spectrum relic lobes straddling the active lobes. After this source had been observed by us the work of other investigators (Harwood *et al.* 2016; van Haarlem *et al.* 2013) showed no evidence of a relic radio lobe. Our image at 150 MHz of 3C452 also show no signs of any relic emission outside the FR II lobes. Our image of 3C452 is shown in Figure 5.1.




Figure 5.1: Image of 3C452, observed with the GMRT at 150 MHz as part of the sample selected for this thesis. There is no evidence of any diffuse, relic emission beyond the lobes of the galaxy.

Table 5.1: The list of DDRGs observed, with their associated redshifts, measured angular sizes and derived linear sizes. The linear sizes were derived assuming a cosmology of $H_0 = 71.9 \text{ km s}^{-1}\text{Mpc}^{-1}$, $\Omega_M = 0.26$ and $\Omega_\Lambda = 0.74$. The redshifts were obtained from Saikia & Jamrozy (2009) except where noted. The angular sizes reported here were measured from the 150 MHz maps of each source. The angular sizes were measured by drawing a straight line between the edges of the outer lobes.

(1) Redshift obtained from Sirothia & Wiita (2013)

(2) Redshift obtained from Bagchi *et al.* (2014)

Source Name	Redshift	Angular size	Linear size	Position	
		(arcsec)	(Mpc)	Right Ascension	Declination
J0041+3224	0.45	216	1.218	00h41m46.1627s	+32d24m52.1724s
J0116-4722	0.1461	680	1.701	01h16m25.005201s	-47d22m39.96212s
J1158+2621	0.1121	290	0.579	11h58m20.1019s	+26d21m11.98s
3C293	0.045	230	0.198	13h52m17.8s	+31d26m46s
J1453+3308	0.076	383	1.463	14h53m02.8595s	+33d08m41.2449s
J1548-3216	0.1082	517	0.998	15h48m58.0302s	-32d16m57.3793s
J1835+6204	0.5194	275	1.688	18h35m10.8682s	+62d04m09.0107s
3C452 ⁽¹⁾	0.0811	270	0.416	22h43m32.808s	39d25m27.56s
J2345-0449 ⁽²⁾	0.0755	1141	1.595	23h45m32.0062s	-04d49m58.9597s

We present below a brief description of each of the other 8 sources, with notes on their morphology and other salient features.

J0041+3224 – This is a DDRG at a redshift of 0.45, and a linear size of ~ 1.2 Mpc oriented in the NW-SE direction. Both the inner and outer lobes are remarkably collimated and do not show much backflow at any of the observed frequencies. In addition to this neither the core nor the jets are visible. The lobes of the DDRG show a substantial degree of asymmetry - the distance between the outer and inner southern lobes is $\sim 1'$ while the distance between the outer and inner northern lobes is $\sim 1.5'$. The 150 MHz image is shown in Figure 5.2.




Figure 5.2: Stokes I image of J0041+3224 at 150 MHz. The resolution of the image is $20'' \times 20''$. Contour levels are $[-5, -3, 3, 5 \dots] \times \sigma$ where σ is the local RMS.

J0116-4722 – This is a giant radio source with a very large linear size (~ 1.7 Mpc), at a redshift of 0.146. The southern inner lobe is very bright ($S \simeq 1$ Jy at 150 MHz) and the northern inner lobe is not visible at 150 MHz. The northern inner lobes gets progressively brighter as the frequency of observation increases.

Particularly, it could mean that the turnover frequency for the synchrotron spectrum lies between 325 MHz and 150 MHz. The outer lobes are large and diffuse with no obvious hotspot in either lobe, although the possible location of the old hotspot could be inferred from the observed flux density distribution. This image was created from the re-analysis of the TGSS archival observations (Intema *et al.* 2017) rather than the full synthesis observation granted during Cycle 29 of the GMRT. This is because the data acquired from our observation suffered from severe ionospheric scintillation, and calibration was almost impossible. The calibrated image quality was significantly worse than the TGSS snapshot image, and therefore we opted to continue analysis on the TGSS image. Figure 5.3 shows the 150 MHz map of the source.




Figure 5.3: Stokes I image of J0116-4722 at 150 MHz. The resolution of the image is $20'' \times 20''$. This image is made from the TGSS archival observations.

J1158+2621 – This source is at a redshift of 0.1121, and is among the smaller sources in the sample with a linear size of ~ 580 kpc. It is a bright source, with a total flux density of ~ 6.8 Jy at 150 MHz. Like most of the sources in this sample,

it is oriented in the NW-SE direction. The inner double is visible at all observed frequencies, and are well aligned with the outer double. Figure 5.4 shows the 150 MHz Stokes I map of the source.

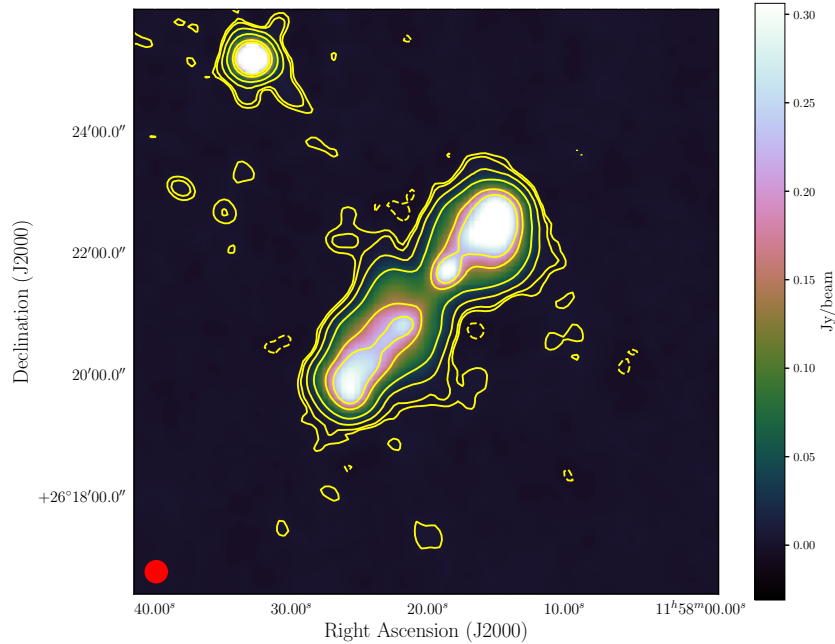


Figure 5.4: Stokes I image of J1158+2621 at 150 MHz. The resolution of the image is $20'' \times 20''$. Contour levels are $[-5, -3, 3, 5 \dots] \times \sigma$ where σ is the local RMS.

3C293 – This is a very bright source ($S \simeq 15$ Jy at 150 MHz) where the inner lobes and the core are not separable at any of the observed frequencies (*i.e.*, 150, 325 and 610 MHz). We refer here to both the inner lobes and the core as simply the core. The outer lobes are significantly fainter ($S \simeq 4$ Jy) at 150 MHz, and have a steeper spectral index than the core, and the lobes quickly fade at higher frequencies. Even at 325 MHz the artefacts close to the core are strong enough to completely wash out any structure from the relic emission. The lobes are almost completely invisible at 610 MHz. Therefore we have excluded this source from any further analysis since (a) The spectral index maps will be highly contaminated by systematic errors associated with the very bright core and (b) Reliable spectral

index maps cannot be obtained between 150 MHz and 325 MHz.




Figure 5.5: Stokes I image of 3C293 at 150 MHz. The outer lobes of this source are significantly fainter than the core. The inner lobes and core are both unresolved at all the observed frequencies.

J1453+3308 – A giant radio source at a redshift of 0.076, with a total linear size of ~ 1.46 Mpc. The radio core is not visible at any of the observed frequencies. The inner lobes are slightly misaligned from the outer lobes, although only to within a few degrees. There is no clear evidence of a hotspot in either of the outer lobes.

J1548-3216 – This is a giant radio source of linear size ~ 1 Mpc at a redshift of 0.108. The inner doubles are resolved at all observed frequencies, although the core is not detected at any. Although the inner and outer lobes seem to be aligned, the terminal point of the inner and outer SW lobe are offset by a few degrees. The SW lobe shows an interesting morphology with a continuous flux density distribution that is bent between the inner and outer lobes. The possible explanations for this need to be further investigated. Figure 5.7 shows the source at 150 MHz.




Figure 5.6: Stokes I image of J1453+3308 at 150 MHz. The resolution of the image is $20'' \times 20''$

J1835+6204 – One of the largest sources in the sample, with a linear size of ~ 1.7 Mpc, at a redshift of 0.5194. The inner and outer doubles of this source are highly collimated, with no measureable offset in the alignment. The radio core is not detected at any of the observed frequencies. At a 150 MHz, the southern inner and outer lobes are interestingly of comparable flux density, although the inner and outer northern lobes differ by almost a factor of 2. This trend continues to higher frequencies, although at 610 MHz the southern inner lobe is brighter than the corresponding outer lobe.

J2345-0449 – The only radio source in our sample where the host galaxy is a spiral (Bagchi *et al.* 2014). For reasons that are not yet understood, the quality of the final image was very poor. While the fluxes of the point sources were accurate, most of the diffuse emission appeared to have been lost. At 150 MHz the outer lobes have a combined flux of ~ 0.45 Jy. However at higher frequencies, the strong point sources in the vicinity of the DDRG showed very large systematic artefacts (most likely due to an active ionosphere) and while calibrating those artefacts (using the SPAM package, described in the next section) most of the diffuse emission was lost. Therefore this was also excluded from any further analysis.




Figure 5.7: Stokes I image of J1548-3216 at 150 MHz. The resolution of the image is $40'' \times 24''$, since a large fraction of the UV coverage was lost during imaging and self-calibration.

5.2.2 Image Processing

The data at 150 MHz were acquired during GMRT cycles 27 and 29 (observing code: 27_063, 29_062), with one target (3C293) also observed at 325 MHz. These data were recorded with an integration time of 2s, and a bandwidth of 16 MHz using the software correlator of the GMRT (Roy *et al.* 2010). The archival data were recorded with an integration time of 16s, and a bandwidth of either 16 MHz or 32 MHz. Both the new observations as well as the archival data were stored in the LTA format used at the GMRT observatory. All the target sources were put through a similar data analysis process, with the initial bandpass and phase calibration done through AIPS (Greisen 2003), and the imaging and self-calibration done through CASA (McMullin *et al.* 2007). There were however small differences in the processing procedure for different sources, such as the solution intervals for self-calibration, the number of iterations used for imaging etc.

The general analysis procedure used is described below:

1. After converting the LTA file to a FITS file, the RfiX algorithm (Athreya 2009) was run on the un-calibrated flux calibration, phase calibrator and target respectively.




Figure 5.8: Stokes I image of J1835+6204 at 150 MHz. The resolution of the image is $22'' \times 22''$.

2. The bandpass calibrator (typically the same as the flux calibrator) was used to obtain the bandpass solutions. The solutions were manually inspected for deviant points, which were flagged, and the “good” solutions were smoothed with a sliding median window. This procedure results in a greater sensitivity to RFI affected channels later on in the analysis procedure.
3. The bandpass solutions were applied to the phase calibrator and the target, and the flux scale was bootstrapped from the flux calibrator to the phase calibrator, and subsequently from the phase calibrator to the target.
4. The “zeroth” order calibrated target was then used to perform several rounds of imaging, self-calibration and residual flagging. Typically, 6-8 rounds of imaging and (phase only) self-calibration were performed after which the residual visibilities were obtained. The residual visibilities were flagged with the IPFLAG algorithm (see Chapter 3) and the flags transferred back to the original (zeroth order) calibrated file.
5. The entire procedure of imaging, self-calibration and residual flagging was repeated 2-3 times until there was no further improvement in image sensitivity. The final image was amplitude calibrated, and imaged using multi-term,




Figure 5.9: Stokes I image of J2345-0449 at 150 MHz. The resolution of the image is $22'' \times 22''$.

multi-frequency synthesis (Rau & Cornwell 2011).

6. If the image still contained significant artefacts around the strong sources, the final calibrated file was then run through the SPAM pipeline (Intema *et al.* 2009). The SPAM algorithm solves for the ionospheric phase in multiple directions and fits for the ionospheric phase screen. We achieved far better results by running SPAM on our final calibrated product rather than the initial, uncalibrated LTA file. This is probably down to more effective residual flagging using IPFLAG, as well as more careful calibration (since the calibration was done semi-automatically, as opposed to a fully automatic calibration in SPAM).

The images of each DDRG used identical parameters at each frequency, *i.e.*, the same cell size and restoring beam, as well as the image weighting scheme. All the images used the Briggs weighting (Briggs *et al.* 1999) with the robust parameter set to 0 in order to achieve a reasonable balance between sensitivity to diffuse emission and resolution. The SPAM step of the analysis was not applied on every DDRG, as in some cases running the data through SPAM resulted in the loss of a significant amount of extended emission (such as in the case of J2345-0449). The direction dependent error mitigation described in Chapter 4 was not used for all the DDRGs since the algorithm was still under development while the DDRGs were being analysed. We found that in most cases we achieved an acceptable

Table 5.2: The list of DDRGs used for the analysis and their spectral ages (in Myr).

Source Name	Spectral Age (Myr)	Reference
J0041+3224	26	Konar <i>et al.</i> 2012
J0116-4722	236	Konar <i>et al.</i> 2013
J1158+2621	113.5	Konar <i>et al.</i> 2013
J1453+3308	52.5	Konar <i>et al.</i> 2006
J1548-3216	26	Machalski <i>et al.</i> 2010
J1835+6204	22	Konar <i>et al.</i> 2012

sensitivity without any direction-dependent calibration and in cases where we did not, SPAM removed a lot of the systematic artefacts.

The flux scales that were obtained by bootstrapping the bandpass calibrator to a standard scale were typically found to have a $\sim 10 - 20\%$ error. In order to more accurately determine the spectral indices, we tied the overall flux scale to the derived flux density between the TGSS-ADR (Intema *et al.* 2017) and NVSS (Condon *et al.* 1998) sky surveys. We compared the fluxes of the common sources at each observed frequency (in the 150, 325 and 610 MHz bands) to the predicted flux density from these two surveys. This ensures that any flux scaling discrepancies between the frequencies will not affect the spectral index structure since all the frequencies are tied to the same reference scale.

Each of the target sources were observed at 3 frequencies (150, 325 and 610 MHz) and processed through the analysis procedure described above. In most cases the 610 MHz was observed in the ‘dual-frequency’ mode of the GMRT wherein the RR polarisation channel records at 610 MHz and the LL polarisation channel records at 240 MHz. In these cases the 240 MHz data were discarded since the sensitivity is typically poor due to the presence of only a single polarisation. The spectral index maps were created using a matched resolution and UV coverage between all three frequencies. The lower UV cutoff was determined by coverage at 610 MHz and the upper UV cutoff by the coverage at 150 MHz. We used a cell size of $4.5''$ and a restoring beam of $22''$. The spectral index maps were created from these matched resolution images, after thresholding all the pixels below $\mu + 1.5 \times \sigma$ where μ and σ are the robust mean and standard deviation over all the pixels in the central region encompassing the DDRG structure. A larger sigma cutoff resulted in fewer pixels surviving the masking process and hence this affected the structure of the spectral index maps. The maps were made between 150 MHz and 325 MHz, and between 325 MHz and 610 MHz independently.

5.3 Results and Discussion

5.3.1 Hotspot Compactness

Section 5.1.3 briefly describes the two methods of deriving the age of a radio source. The methods work for both active and dead galaxies although the dynamic ageing method will result in a significant underestimate of the source age if the deceleration of the radio jet is not taken into account (e.g. Falle 1991; Laing & Bridle 2002). There are several caveats to be used while using a spectral ageing model (Blundell & Rawlings 2000, 2001); the two biggest sources of systematic errors in the spectral age estimation are (a) the lobe age is significantly larger than the radiative lifetime of particles, and (b) the magnetic field in the lobes are not well characterised. The errors that arise when attempting to calculate the spectral ages of older radio lobes are because re-energisation processes are typically not taken into account. While the dynamical age will result in a reasonably accurate estimate, spectral ages derived in such situations can be a factor of 2 (or more) different from the dynamical ages. Accurately deriving the spectral age involves assuming the injection spectrum, the magnetic field, the evolution of the magnetic field (due to adiabatic expansion of the lobe), particle re-energisation through turbulent processes as well as other losses such as inverse compton scattering of the microwave background. There are methods that attempt to take care of some of these issues (e.g., Kaiser *et al.* 1997) however a common assumption is that of a constant magnetic field. This will result in a systematic under-estimate of the spectral age of a source.

We investigated the utility of hotspot size as a proxy for the spectral age of the lobe since the former is much easier to estimate than the latter. In an active FR II radio source, the hotspot is the location at which energy is transferred between the radio jet and the lobe, and is often the brightest region in the lobe. Since the pressure in the hotspot is related to the energy density, and hence its brightness (Blundell *et al.* 1999), the hotspot should be typically over-pressured relative to the surrounding medium. Therefore, when the radio jet is switched off, and the lobe evolves adiabatically, the hotspot should expand into the surrounding medium, increasing in size and decreasing in luminosity until it homogenises with the lobe. Therefore, in principle one should be able to use the hotspot size as a proxy for the age of the source. However, given that the hotspot sizes in active FR II sources scale linearly with the size of the source (Hardcastle *et al.* 1998), one would need to normalise the measured size of the hotspot by the size of the lobe.

We define a quantity “hotspot FWHM” which is the average width of the region around the peak where the average brightness is half that of the peak. We

also plotted a “normalised hotspot FWHM” which takes into account the increase in hotspot size with the overall linear size of the source. Since the outer lobes of DDRGs often do not show the presence of a classical hotspot-like structure we assumed that the brightest pixel was the location of the erstwhile hotspot in order to derive the hotspot FWHM. Figures 5.10b and 5.10 show the evolution of hotspot size with spectral age of the source (from literature). The scatter shows the individual data points, and the line is the best fit to the scatter. The trends at 150 MHz and 325 MHz are similar, although the scatter in the latter is smaller. The estimates of hotspot size and compactness at 325 MHz should be more accurate since its resolution is higher. While we indeed measure some (expected) positive correlation we cannot claim that the hotspot FWHM faithfully reflects the spectral age to any great accuracy.

5.3.2 Spectral index structure

There are very few confirmed examples of dead radio sources in the literature, and therefore there have not been many studies on the longevity of dead sources. The radiative lifetime of a population of synchrotron electrons that are no longer being injected with energy is thought to be of the order of $\sim 10^7$ years (Blundell & Rawlings 2000). On the other hand, the typical measured spectral ages of sources are of the order of $10^7 - 10^8$ years which is an order of magnitude larger than the radiative lifetimes. This discrepancy has been explained by invoking population mixing and/or reacceleration of electrons (Blundell & Rawlings 2001). The evolution of spectral index structure in a radio lobe should give us an estimate of the timescale of the homogenisation of a radio lobe.

Beyond the question of the evolution of the spectral index structure, which is interesting in itself, we were also interested in the possibility of inferring the location of re-energisation of synchrotron electrons from the spectral index maps alone. Essentially, we were hoping to use DDRGs to estimate the timescale over which the brightest region of a radio lobe (*i.e.*, the hotspot) remains also its region of flattest spectral index. The wisdom from such an exercise can be utilised in the case of galaxy cluster radio halos, which are diffuse emission which have no obvious source (or locality) of energisation (e.g., Buote 2001). Characterising the spectral index as a tracer for energisation is markedly more straightforward in the case of DDRGs since the direction and source of energisation are known (*i.e.*, the radio jets), as is the location of energisation, *i.e.*, the hotspots, even if they have become more diffuse after the cessation of the radio jet.

In order to perform such a study, one requires high resolution spectral index maps at low frequencies (Figure 5.11). This is because the radiative lifetimes of

synchrotron electrons increase rapidly with decrease of emission frequency. There have been very few studies that have used low frequency ($\nu \leq 325$ MHz) spectral index images and all such have only worked with the integrated spectrum (e.g., Joshi *et al.* 2011; Konar *et al.* 2006).

Identifying the “dead” hotspots (or warmspots) turned out to be more difficult than anticipated. We used the peak in the locally smoothed lobe as the putative warmspot, and used it as the centre of the radial profile of spectral index. These plots are shown in Figures 5.12 - 5.17. The plots show either very little variation of the spectral index with radial distance or have the flattest spectral index at some distance from the putative warmspot. Both these patterns suggest that the old hotspot has essentially homogenised within the lobe even for the youngest DDRGs which are about 25 Myr old.

5.4 Conclusion

We have observed a sample of 9 double-double radio galaxies (of which only 6 were useful) to study whether the spectral index structure can predict the previous location of energisation in a dead radio source (*i.e.*, the warmspot). We also investigated if warmspot sizes are a good proxy for spectral ages of a dead radio lobe. We found that while there is some positive correlation between warmspot sizes and spectral age it has a lot of scatter which makes it less useful for the purpose. We also found that dead hotspots tend to homogenise within the lobe inside 25 Myr, and therefore dead sources older than this will not be expected to show any spectral index signatures of the location of energisation of synchrotron electrons.

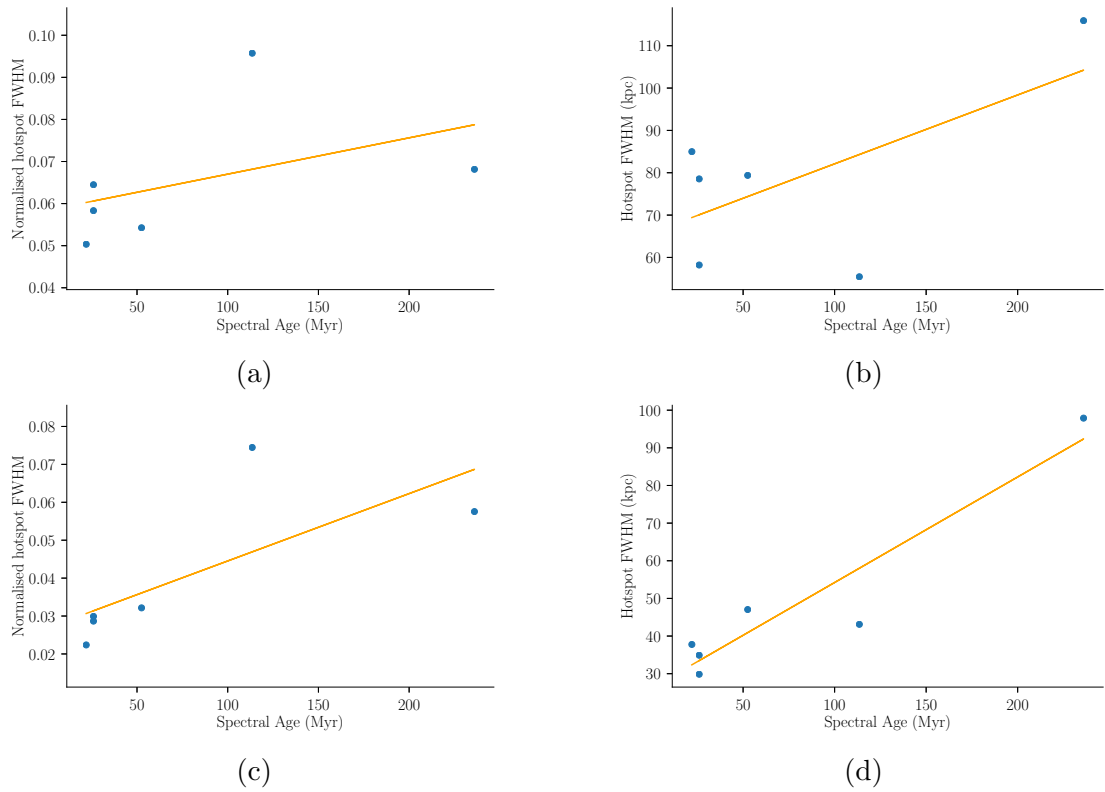




Figure 5.10: The hotspot FWHM (normalised and linear) measured at 150 MHz (top) and 325 MHz (bottom) for different DDRGs vs. their spectral ages (as reported in the literature). Left: The normalised hotspot FWHM (the hotspot size normalised by the source size) as a function of spectral age. Right: The hotspot FWHM as a function of spectral age.




(a) J0041+3224




(b) J0116-4722

Figure 5.11: Spectral index map of J0041+3224 and J0116-4722. In this and all the following spectral index maps, a 3σ threshold has been applied prior to the calculation of the map in order to retain only reliable features.




(c) J1158+2621




(d) J1453+3308

Figure 5.11: Spectral index map of J1158+2621 and J1453+3308



(e) J1548-3216



(f) J1835+6204

Figure 5.11: Spectral index map of J1548-3216 and J1835+6204




Figure 5.12: The radial spectral index profile of J0041+3224 between 150–325 MHz and 325–610 MHz. Each plot contains the radial spectral index profiles from both lobes of the DDRG.

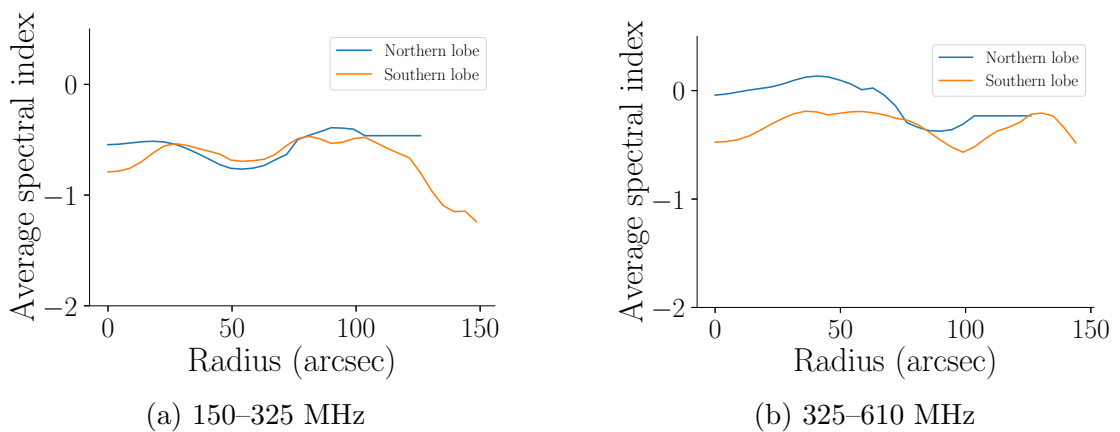


Figure 5.13: The radial spectral index profile of J0116-4722 between 150–325 MHz and 325–610 MHz. Each plot contains the radial spectral index profiles from both lobes of the DDRG.

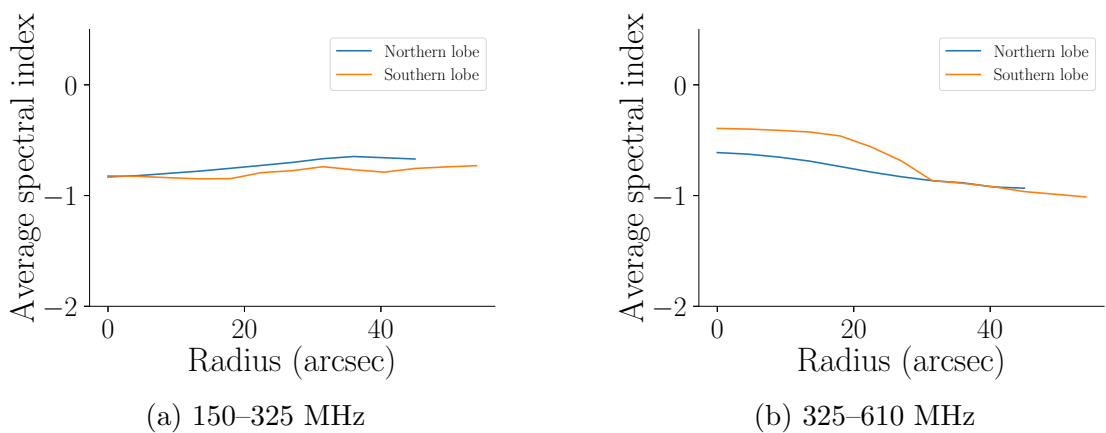


Figure 5.14: The radial spectral index profile of J1158+2621 between 150–325 MHz and 325–610 MHz. Each plot contains the radial spectral index profiles from both lobes of the DDRG.

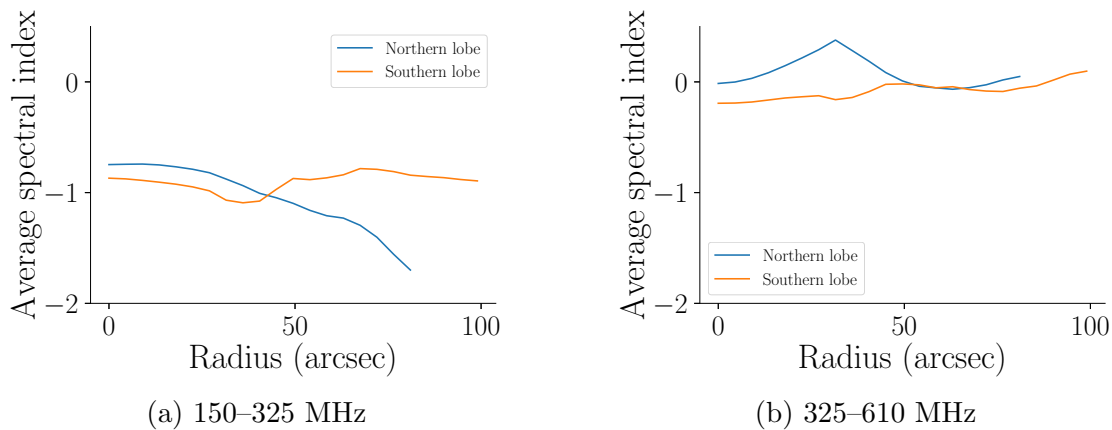


Figure 5.15: The radial spectral index profile of J1453+3308 between 150–325 MHz and 325–610 MHz. Each plot contains the radial spectral index profiles from both lobes of the DDRG.

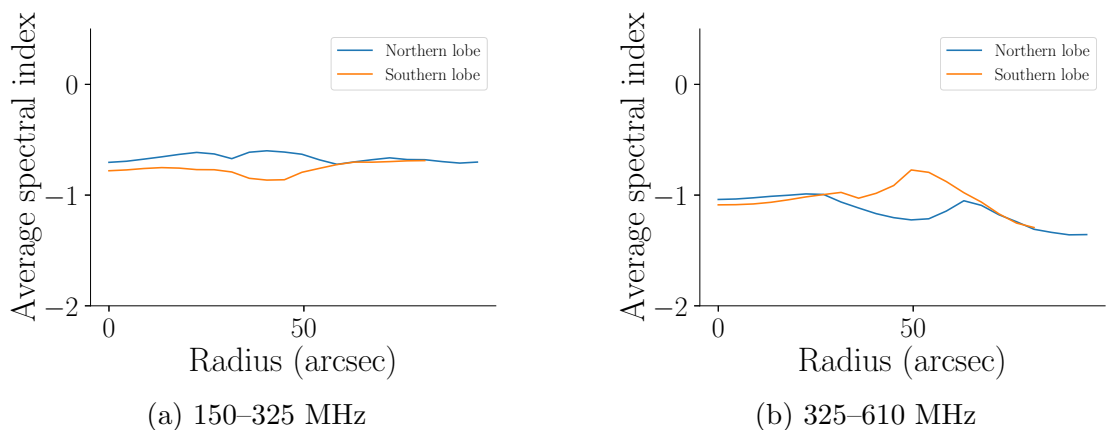


Figure 5.16: The radial spectral index profile of J1548-3216 between 150–325 MHz and 325–610 MHz. Each plot contains the radial spectral index profiles from both lobes of the DDRG.

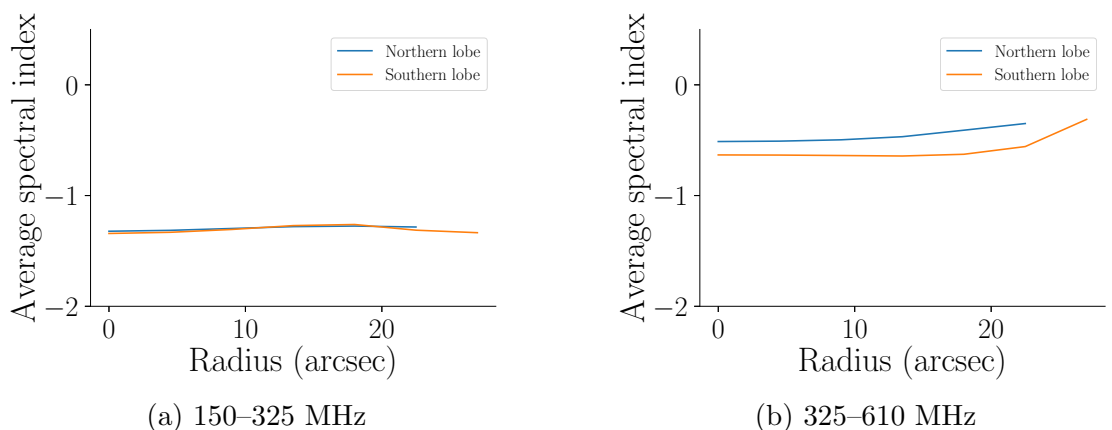


Figure 5.17: The radial spectral index profile of J1835+6204 between 150–325 MHz and 325–610 MHz. Each plot contains the radial spectral index profiles from both lobes of the DDRG.

Chapter 6

Summary

We have described algorithms to obtain high sensitivity, low frequency images with the GMRT. The two main sources of systematic errors in interferometric imaging data are radio frequency interference and direction dependent calibration errors (DDEs) respectively. Although DDEs only contribute errors to the level of a few percent of source flux, it is crucial to eliminate or calibrate for these errors in order to achieve thermal noise limited images. At the GMRT 150 MHz (using the software narrowband correlator) we have managed to achieve an rms noise of ~ 0.40 mJy in several images which is only about a factor of two higher than the theoretical confusion limit at 150 MHz. We have otherwise routinely reached a noise level of below 0.7 mJy with the application of the techniques described in this thesis. While the algorithms were developed for and tested on GMRT data there is no “telescope specificity” in them and they can be used for any interferometer.

The first set of algorithms — IPFLAG — were developed to identify and flag intermittent RFI in the binned UV plane (GRIDflag) and in the baseline time-channel plane (TCflag). GRIDflag exploits the within-cell redundancy of visibilities in the binned UV plane to statistically identify and flag RFI corrupted visibilities. This has the advantage that flagging independently within each UV-bin will tend to preserve UV coverage, as at least some visibilities remain within each UV-bin at the end of the process. TCflag subtracts the residual fringes in the time-channel (TC) plane of a single baseline in order to obtain more sensitive RFI thresholds. The baselines are first “defringed” by clipping in the two-dimensional Fourier space, the group-delay – delay-rate plane, before identifying the thresholds. The flags are then applied on the original data. The application of these procedures to several different fields, with both compact and diffuse emission yielded a consistent improvement of 20-50%. It also resulted in a substantial reduction in the spurious patterns in the image, and an increase in the number of detected sources. The GRIDflag algorithm is expected to work well for deep imaging programmes that involve repeated multi-epoch observations of the same fields. The

repeated sampling within a UV-cell is expected to facilitate the process of identifying and separating RFI-corrupted visibilities. However, wide bandwidth data sets will require some more care as the source and primary beam-induced spectral index and noise characteristics can change substantially across the band. Understanding these issues and adapting these algorithms for wide bandwidth imaging should be the next step.

Direction-dependent errors from various sources — primary beam asymmetries, antenna pointing errors and non-isoplanatic ionosphere — have adversely impacted image quality particularly at low radio frequencies. Various workers have been attempting the mitigation of these effects with different degrees of success. We have developed an alternative approach of addressing the consequences of these errors rather than the root cause. Essentially, the DDEs result in an apparent time variability in the sky intensity as viewed by the interferometer. The standard imaging and calibration algorithms assume that the sky is invariant over the entire duration of the observation. We realised that this inappropriate assumption results in inaccurate subtraction of the sources. As a consequence the strongest sources leave behind “uncleaned” residual components which spread dirty-beam structures across the image. We addressed this very simply by cleaning the 50-100 strongest sources in independent data sets of 5-10 minute intervals, i.e. we no longer required the source flux to remain constant across the entire observation. While this does make the images useless for studying the strongest sources they do reduce the systematic patterns associated with the strongest sources, thus reducing the noise floor and improving detectability of faint sources.

These DDEs also result in the different baselines detecting a source with different flux densities. As a consequence, the assumption of constancy of flux density results in residual source patterns as a function of time and baselines. We again solved this problem by searching for such patterns (by using the model for the brightest 3-5 sources) and fitting them out in the residual baseline visibilities.

The challenge of generating high quality images with the GMRT at 150 MHz has now many tools including RfiX (RFI excision), SPAM (ionosphere non isoplanaticity), the algorithms presented in this thesis, and the real time RFI excision (at Nyquist sampling) being currently developed by the GMRT team. The one major lacuna in this tool kit is the lack of weighting of GMRT data. We have plans to use residual statistics to weight GMRT data during calibration and imaging. This is because the “noise” in GMRT data perhaps has a greater contribution from RFI and DDEs than from system temperature. We hope that with all these the GMRT 150 Mhz observations will routinely produce images with confusion limited noise.

Finally, we applied these techniques to image a class of restarting radio galax-

ies, the double-double radio galaxies (DDRGs) and mapped the spectral index structures of the dead outer lobes to attempt to identify the location of energisation (i.e. dead hotspot) from the spectral index maps alone. We found no clear trend in the DDRGs. It seems that the spectral index structure seen in active radio sources (flattest at the hotspot and steepening away from it) becomes homogenised within 25 Myr, which corresponds to the ages of the youngest DDRGs in our sample.

Appendix A

Spectral leakage

Typically, the Fourier transform is taken using the Fast Fourier Transform (FFT) algorithm (Cooley & Tukey 1965) which requires data in a uniformly sampled grid (N dimensional) and will return the resulting Fourier domain information on a uniform grid. Typically, for one dimensional time series data this gridding introduces what is known as “spectral leakage” into the Fourier domain (Harris 1978).

There are many ways to understand the phenomenon of spectral leakage. The FFT algorithm operates on finite-length data, which can be thought of as a rectangular windowing function applied to the input data.


$$D_o = D_t \times R_t \tag{A.1}$$

where D_o is the observed data, D_t is the “true” data, multiplied by R_t the rectangular sampling function which is given by


$$R_t = \begin{cases} 1 & -T/2 < t < T/2 \\ 0 & \text{otherwise} \end{cases}$$

Under a Fourier transformation, this multiplication turns into a convolution (due to the convolution theorem Bracewell 1986). The Fourier transform of the rectangular window is the sinc function, which extends from $-\infty$ to ∞

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \tag{A.2}$$



(a)



(b)

Figure A.1: (a) Fourier transform of a signal with an integer number of periods within the window (10 periods). The values in the bins adjacent to the peaks are of the order of numerical precision. (b) Fourier transform of a signal with a non-integer number of periods with the window (10.1 periods). The values in the bins adjacent to the peak are several orders of magnitude larger than the numerical precision.

If the input signal to the FFT has an exactly integral number of periods, then the sinc function lines up such that the nulls of the sinc function coincide with the centres of each adjacent bin. If the input signal does not have an exactly intergral number of periods, then the nulls of the sinc function no longer line up with the centres of the bin, causing a non-zero contribution to several of the adjacent Fourier bins. This phenomenon is called “spectral leakage”.

[A.1](#) shows the effect of Fourier transforming a signal which does not have an integral number of periods within the sampling window. [A.1a](#) shows the spectrum of an input signal with exactly 10 periods. The peaks are symmetric about the centre as expected, and correspond to the tenth bin, and the rest of the bins have values on the order of numerical noise. However, [A.1b](#) shows the FFT spectrum of a signal without an integral number of periods within the sampling window - the input signal in this case has 10.1 periods within the window. There is clearly a component in several of the adjacent Fourier bins at a level several magnitudes larger than numerical noise.

Bibliography

1. Athreya, R. A new approach to mitigation of radio frequency interference in interferometric data. *The Astrophysical Journal* **696**, 885 (2009).
2. Bagchi, J., Vivek, M., Vikram, V., Hota, A., Biju, K., Sirothia, S., Srianand, R., Jacob, J., *et al.* Megaparsec relativistic jets launched from an accreting supermassive black hole in an extreme spiral galaxy. *The Astrophysical Journal* **788**, 174 (2014).
3. Barnbaum, C. & Bradley, R. F. A New Approach to Interference Excision in Radio Astronomy: Real-Time Adaptive Cancellation. *The Astronomical Journal* **116**, 2598. ISSN: 1538-3881 (1998).
4. Basseville, M. & Nikiforov, I. V. *Detection of abrupt changes: theory and application* (Prentice Hall Englewood Cliffs, 1993).
5. Bhat, N., Cordes, J., Chatterjee, S. & Lazio, T. Radio frequency interference identification and mitigation using simultaneous dual-station observations. *Radio science* **40** (2005).
6. Bhatnagar, S., Cornwell, T., Golap, K. & Uson, J. M. Correcting direction-dependent gains in the deconvolution of radio interferometric images. *Astronomy & Astrophysics* **487**, 419–429 (2008).
7. Biretta, J., Zhou, F. & Owen, F. Detection of proper motions in the M87 jet. *The Astrophysical Journal* **447**, 582 (1995).
8. Blandford, R. & Payne, D. Hydromagnetic flows from accretion discs and the production of radio jets. *Monthly Notices of the Royal Astronomical Society* **199**, 883–903 (1982).
9. Blandford, R. & Rees, M. A ‘twin-exhaust’ model for double radio sources. *Monthly Notices of the Royal Astronomical Society* **169**, 395–415 (1974).
10. Blundell, K. M. & Rawlings, S. The spectra and energies of classical double radio lobes. *The Astronomical Journal* **119**, 1111 (2000).
11. Blundell, K. M. & Rawlings, S. *Spectral ageing: a new age perspective in Particles and Fields in Radio Galaxies Conference* **250** (2001), 363.

12. Blundell, K. M., Rawlings, S. & Willott, C. J. The nature and evolution of classical double radio sources from complete samples. *The Astronomical Journal* **117**, 677 (1999).
13. Bracewell, R. N. *The Fourier Transform and its applications* (1986).
14. Briggs, D. S., Schwab, F. R. & Sramek, R. A. *Imaging in Synthesis Imaging in Radio Astronomy II* **180** (1999), 127.
15. Briggs, F. H., Bell, J. F. & Kesteven, M. J. Removing Radio Interference from Contaminated Astronomical Spectra Using an Independent Reference Signal and Closure Relations. *The Astronomical Journal* **120**, 3351. ISSN: 1538-3881 (2000).
16. Brocksopp, C., Kaiser, C., Schoenmakers, A. & de Bruyn, A. Double–double radio galaxies: further insights into the formation of the radio structures. *Monthly Notices of the Royal Astronomical Society* **410**, 484–498 (2010).
17. Buch, K. D., Gupta, Y., Bhatporia, S., Nalawade, S., Naik, K. & Ajithkumar, B. *Real-time RFI excision for the GMRT wideband correlator in Radio Frequency Interference (RFI)* (2016), 11–15.
18. Buote, D. A. On the origin of radio halos in galaxy clusters. *The Astrophysical Journal Letters* **553**, L15 (2001).
19. Capetti, A., Zamfir, S., Rossi, P., Bodo, G., Zanni, C. & Massaglia, S. On the origin of X-shaped radio-sources: New insights from the properties of their host galaxies. *Astronomy & Astrophysics* **394**, 39–45 (2002).
20. Chakrabarti, S. K. & Bhaskaran, P. On the origin, acceleration and collimation of bipolar outflows and cosmic radio jets. *Monthly Notices of the Royal Astronomical Society* **255**, 255–260 (1992).
21. Clark, B. An efficient implementation of the algorithm 'CLEAN'. *Astronomy and Astrophysics* **89**, 377 (1980).
22. Cohen, A., Lane, W., Cotton, W., Kassim, N., Lazio, T., Perley, R., Condon, J. & Erickson, W. The VLA low-frequency sky survey. *The Astronomical Journal* **134**, 1245 (2007).
23. Condon, J. J., Cotton, W., Greisen, E., Yin, Q., Perley, R., Taylor, G. & Broderick, J. The NRAO VLA sky survey. *The Astronomical Journal* **115**, 1693 (1998).
24. Cooley, J. W. & Tukey, J. W. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation* **19**, 297–301. ISSN: 0025-5718, 1088-6842 (1965).

25. Cornwell, T. J. Multiscale CLEAN deconvolution of radio synthesis images. *IEEE Journal of Selected Topics in Signal Processing* **2**, 793–801 (2008).
26. Cornwell, T. & Perley, R. Radio-interferometric imaging of very large fields—The problem of non-coplanar arrays. *Astronomy and Astrophysics* **261**, 353–364 (1992).
27. Cornwell, T., Perley, R., Golap, K. & Bhatnagar, S. Rfi excision in synthesis imaging without a reference signal. *EVLA memo series* **86** (2004).
28. Cotton, W. D., Condon, J. J., Perley, R. A., Kassim, N., Lazio, J., Cohen, A., Lane, W. & Erickson, W. C. *Beyond the isoplanatic patch in the VLA Low-frequency Sky Survey in Ground-based Telescopes* **5489** (2004), 180–190.
29. DeBoer, D. R., Gough, R. G., Bunton, J. D., Cornwell, T. J., Beresford, R. J., Johnston, S., Feain, I. J., Schinckel, A. E., Jackson, C. A., Kesteven, M. J., *et al.* Australian SKA pathfinder: A high-dynamic range wide-field of view survey telescope. *Proceedings of the IEEE* **97**, 1507–1521 (2009).
30. Dennett-Thorpe, J., Scheuer, P., Laing, R., Bridle, A., Pooley, G. G. & Reich, W. Jet reorientation in active galactic nuclei: two winged radio galaxies. *Monthly Notices of the Royal Astronomical Society* **330**, 609–620 (2002).
31. Ellingson, S. W., Bunton, J. D. & Bell, J. F. Removal of the GLONASS C/A Signal from OH Spectral Line Observations Using a Parametric Modeling Technique. *The Astrophysical Journal Supplement Series* **135**, 87 (2001).
32. Falle, S. Self-similar jets. *Monthly Notices of the Royal Astronomical Society* **250**, 581–596 (1991).
33. Fanaroff, B. & Riley, J. The morphology of extragalactic radio sources of high and low luminosity. *Monthly Notices of the Royal Astronomical Society* **167**, 31P–36P (1974).
34. Franceschini, A., Vercellone, S. & Fabian, A. Supermassive Black Holes in Early-Type Galaxies: Relationship with Radio Emission and Constraints on the Black Hole Mass Function. *Monthly Notices of the Royal Astronomical Society* **297**, 817–824 (1998).
35. Fridman, P. Statistically stable estimates of variance in radio-astronomy observations as tools for radio-frequency interference mitigation. *The Astronomical Journal* **135**, 1810 (2008).
36. Fridman, P. & Baan, W. *RFI* mitigation methods in radio astronomy. *Astronomy & Astrophysics* **378**, 327–344 (2001).

37. Gizani, N. A., Cohen, A. & Kassim, N. First results of the 74-MHz Very Large Array–Pie Town link. Hercules A at low frequencies. *Monthly Notices of the Royal Astronomical Society* **358**, 1061–1068 (2005).
38. Gizani, N. A. & Leahy, J. A multiband study of Hercules A—II. Multifrequency Very Large Array imaging. *Monthly Notices of the Royal Astronomical Society* **342**, 399–421 (2003).
39. Golap, K., Cornwell, T., Perley, R. & Bhatnagar, S. *Post Correlation RFI Excision in From Clark Lake to the Long Wavelength Array: Bill Erickson’s Radio Science* **345** (2005), 366.
40. Greisen, E. W. AIPS, the VLA, and the VLBA. *Astrophysics and Space Science Library* **285**, 109–126 (2003).
41. Greve, A., Neri, R. & Sievers, A. The gain–elevation correction of the IRAM 30–m telescope. *Astronomy and Astrophysics supplement series* **132**, 413–416 (1998).
42. Gupta, Y., Ajithkumar, B., Kale, H., Nayak, S., Sabhapathy, S., Sureshkumar, S., Swami, R., Chengalur, J., Ghosh, S., Ishwara-Chandra, C., *et al.* The upgraded GMRT: opening new windows on the radio Universe. *Current Science* **113**, 707 (2017).
43. Hamaker, J., Bregman, J. & Sault, R. Understanding radio polarimetry. I. Mathematical foundations. *Astronomy and Astrophysics Supplement Series* **117**, 137–147 (1996).
44. Hancock, P. J., Murphy, T., Gaensler, B. M., Hopkins, A. & Curran, J. R. Compact continuum source finding for next generation radio surveys. *Monthly Notices of the Royal Astronomical Society* **422**, 1812–1824 (2012).
45. Hardcastle, M., Alexander, P., Pooley, G. G. & Riley, J. FR II radio galaxies with z_j 0.3–I. Properties of jets, cores and hotspots. *Monthly Notices of the Royal Astronomical Society* **296**, 445–462 (1998).
46. Harris, F. J. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE* **66**, 51–83. ISSN: 0018-9219 (Jan. 1978).
47. Harwood, J. J., Croston, J. H., Intema, H. T., Stewart, A. J., Ineson, J., Hardcastle, M. J., Godfrey, L., Best, P., Brienza, M., Heesen, V., *et al.* FR II radio galaxies at low frequencies—I. Morphology, magnetic field strength and energetics. *Monthly Notices of the Royal Astronomical Society* **458**, 4443–4455 (2016).

48. Hein, G. W., Godet, J., Issler, J.-L., Martin, J.-C., Lucas-Rodriguez, R. & Pratt, T. *Status of Galileo frequency and signal design in in CDROM Proc. ION GPS* (2002).
49. Hellbourg, G., Chippendale, A. P., Kesteven, M. J. & Jeffs, B. D. *Reference antenna-based subspace tracking for RFI mitigation in radio astronomy in Signal and Information Processing (GlobalSIP), 2014 IEEE Global Conference on* (IEEE, 2014), 1286–1290. <http://ieeexplore.ieee.org/abstract/document/7032330/> (2017).
50. Hofmann-Wellenhof, B., Lichtenegger, H. & Wasle, E. *GNSS—global navigation satellite systems: GPS, GLONASS, Galileo, and more* (Springer Science & Business Media, 2007).
51. Högbom, J. Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement Series* **15**, 417 (1974).
52. Intema, H., Jagannathan, P., Mooley, K. & Frail, D. The GMRT 150 MHz all-sky radio survey—First alternative data release TGSS ADR1. *Astronomy & Astrophysics* **598**, A78 (2017).
53. Intema, H., Van der Tol, S., Cotton, W., Cohen, A., Van Bemmelen, I. & Röttgering, H. Ionospheric calibration of low frequency radio interferometric observations using the peeling scheme—I. Method description and first results. *Astronomy & Astrophysics* **501**, 1185–1205 (2009).
54. Ishwara-Chandra, C., Sirothia, S., Wadadekar, Y., Pal, S. & Windhorst, R. Deep GMRT 150-MHz observations of the LBDS-Lynx region: ultrasteepest spectrum radio sources. *Monthly Notices of the Royal Astronomical Society* **405**, 436–446 (2010).
55. Jarvis, M. J., Taylor, A., Agudo, I., Allison, J. R., Deane, R., Frank, B., Gupta, N., Heywood, I., Maddox, N., McAlpine, K., *et al.* The MeerKAT International GHz Tiered Extragalactic Exploration (MIGHTEE) Survey. *arXiv preprint arXiv:1709.01901* (2017).
56. Jonas, J. L. MeerKAT—The South African array with composite dishes and wide-band single pixel feeds. *Proceedings of the IEEE* **97**, 1522–1530 (2009).
57. Joshi, S., Nandi, S., Saikia, D., Ishwara-Chandra, C. & Konar, C. A radio study of the double–double radio galaxy 3C293. *Monthly Notices of the Royal Astronomical Society* **414**, 1397–1404 (2011).

58. Kaiser, C. R., Dennett-Thorpe, J. & Alexander, P. Evolutionary tracks of FR II sources through the PD diagram. *Monthly Notices of the Royal Astronomical Society* **292**, 723–732 (1997).
59. Kaiser, C. R., Schoenmakers, A. P. & Röttgering, H. J. Radio galaxies with a ‘double-double’ morphology—II. The evolution of double-double radio galaxies and implications for the alignment effect in FR II sources. *Monthly Notices of the Royal Astronomical Society* **315**, 381–394 (2000).
60. Kellermann, K., Sramek, R., Schmidt, M., Shaffer, D. & Green, R. VLA observations of objects in the Palomar Bright Quasar Survey. *The Astronomical Journal* **98**, 1195–1207 (1989).
61. Kocz, J., Briggs, F. & Reynolds, J. Radio frequency interference removal through the application of spatial filtering techniques on the parkes multi-beam receiver. *The Astronomical Journal* **140**, 2086 (2010).
62. Konar, C., Hardcastle, M., Jamrozy, M., Croston, J. & Nandi, S. Rejuvenated radio galaxies J0041+ 3224 and J1835+ 6204: how long can the quiescent phase of nuclear activity last? *Monthly Notices of the Royal Astronomical Society* **424**, 1061–1076 (2012).
63. Konar, C., Hardcastle, M., Jamrozy, M. & Croston, J. Episodic radio galaxies J0116- 4722 and J1158+ 2621: can we constrain the quiescent phase of nuclear activity? *Monthly Notices of the Royal Astronomical Society* **430**, 2137–2153 (2013).
64. Konar, C., Saikia, D., Jamrozy, M. & Machalski, J. Spectral ageing analysis of the double–double radio galaxy J1453+ 3308. *Monthly Notices of the Royal Astronomical Society* **372**, 693–702 (2006).
65. Laing, R. A. & Bridle, A. Dynamical models for jet deceleration in the radio galaxy 3C 31. *Monthly Notices of the Royal Astronomical Society* **336**, 1161–1180 (2002).
66. Lal, D. V. GMRT Observer’s Manual (2013).
67. Lal, D. V. & Rao, A. P. Giant Metrewave Radio Telescope observations of X-shaped radio sources. *Monthly Notices of the Royal Astronomical Society* **374**, 1085–1102 (2006).
68. Leshem, A., Veen, A.-J. v. d. & Boonstra, A.-J. Multichannel Interference Mitigation Techniques in Radio Astronomy. *The Astrophysical Journal Supplement Series* **131**, 355. ISSN: 0067-0049 (2000).

69. Lonsdale, C. *Configuration considerations for low frequency arrays in From Clark Lake to the Long Wavelength Array: Bill Erickson's Radio Science* **345** (2005), 399.
70. Machalski, J., Jamrozy, M. & Konar, C. Spectral ageing analysis and dynamical analysis of the double-double radio galaxy J1548–3216. *Astronomy & Astrophysics* **510**, A84 (2010).
71. Machalski, J., Jamrozy, M. & Saikia, D. A multifrequency study of giant radio sources–III. Dynamical age versus spectral age of the lobes of selected sources. *Monthly Notices of the Royal Astronomical Society* **395**, 812–822 (2009).
72. McMullin, J. P., Waters, B., Schiebel, D., Young, W. & Golap, K. *CASA Architecture and Applications* in. *Astronomical Data Analysis Software and Systems XVI*. **376** (Oct. 1, 2007), 127. <http://adsabs.harvard.edu/abs/2007ASPC..376..127M> (2017).
73. Middelberg, E. Automated editing of radio interferometer data with PIEFLAG. *Publications of the Astronomical Society of Australia* **23**, 64–68 (2006).
74. Miley, G. The structure of extended extragalactic radio sources. *Annual review of astronomy and astrophysics* **18**, 165–218 (1980).
75. Mitchell, D., Sault, R. & Kesteven, M. *Post correlation versus real-time adaptive RFI cancellation in XXVIIth General Assembly of the International Union of Radio Science, Maastricht, The Netherlands* (2002).
76. Mohan, N. & Rafferty, D. PyBDSF: Python Blob Detection and Source Finder. *Astrophysics Source Code Library* (2015).
77. Nandi, S. & Saikia, D. Double-double radio galaxies from the FIRST survey. *Bull. Astr. Soc. India* **40**, 121–137 (2012).
78. Nita, G. M., Gary, D. E., Liu, Z., Hurford, G. J. & White, S. M. Radio Frequency Interference Excision Using Spectral-Domain Statistics. *Publications of the Astronomical Society of the Pacific* **119**, 805. ISSN: 1538-3873 (2007).
79. Noordam, J. E. *LOFAR calibration challenges in Ground-Based Telescopes* **5489** (2004), 817–826.
80. Noordam, J. E. & Smirnov, O. M. The MeqTrees software system and its use for third-generation calibration of radio interferometers. *Astronomy & Astrophysics* **524**, A61 (2010).

81. Offringa, A. R., de Bruyn, A. G., Biehl, M., Zaroubi, S., Bernardi, G. & Pandey, V. N. Post-correlation radio frequency interference classification methods. *Monthly Notices of the Royal Astronomical Society* **405**, 155–167. ISSN: 0035-8711 (June 2010).
82. Offringa, A. R., Gronde, J. J. v. d. & Roerdink, J. B. T. M. A morphological algorithm for improving radio-frequency interference detection. *Astronomy & Astrophysics* **539**, A95. ISSN: 0004-6361, 1432-0746 (2012).
83. Okabe, A. *Spatial tessellations* (Wiley Online Library, 1992).
84. Osterbrock, D. E. Active galactic nuclei. *Reports on Progress in Physics* **54**, 579 (1991).
85. Pearson, T. & Readhead, A. Image formation by self-calibration in radio astronomy. *Annual review of astronomy and astrophysics* **22**, 97–130 (1984).
86. Pen, U.-L., Chang, T.-C., Hirata, C. M., Peterson, J. B., Roy, J., Gupta, Y., Odegova, J. & Sigurdson, K. The GMRT EoR experiment: limits on polarized sky brightness at 150 MHz. *Monthly Notices of the Royal Astronomical Society* **399**, 181–194. ISSN: 0035-8711 (Oct. 11, 2009).
87. Pen, U.-L., Chang, T.-C., Hirata, C. M., Peterson, J. B., Roy, J., Gupta, Y., Odegova, J. & Sigurdson, K. The GMRT EoR experiment: limits on polarized sky brightness at 150 MHz. *Monthly Notices of the Royal Astronomical Society* **399**, 181–194 (2009).
88. Perley, R. A. *Wide field imaging II: imaging with non-coplanar arrays in Synthesis Imaging in Radio Astronomy* **6** (1989), 259.
89. Perley, R. A., Schwab, F. R. & Bridle, A. H. Synthesis imaging in radio astronomy (1989).
90. Rau, U. & Cornwell, T. J. A multi-scale multi-frequency deconvolution algorithm for synthesis imaging in radio interferometry. *Astronomy & Astrophysics* **532**, A71 (2011).
91. Reynolds, C. S., Heinz, S. & Begelman, M. C. The hydrodynamics of dead radio galaxies. *Monthly Notices of the Royal Astronomical Society* **332**, 271–282 (2002).
92. Röttgering, H. LOFAR, a new low frequency radio telescope. *New astronomy reviews* **47**, 405–409 (2003).
93. Roy, J., Gupta, Y., Pen, U.-L., Peterson, J. B., Kudale, S. & Kodilkar, J. A real-time software backend for the GMRT. *Experimental Astronomy* **28**, 25–60 (2010).

94. Saikia, D. & Jamrozy, M. Recurrent activity in Active Galactic Nuclei. *Bull. Astr. Soc. India* **37**, 63–89 (2009).
95. Scaife, A. M. & Heald, G. H. A broad-band flux scale for low-frequency radio telescopes. *Monthly Notices of the Royal Astronomical Society: Letters* **423** (2012).
96. Scheuer, P. Lobe asymmetry and the expansion speeds of radio sources. *Monthly Notices of the Royal Astronomical Society* **277**, 331–340 (1995).
97. Schoenmakers, A. P., De Bruyn, A., Röttgering, H., Van Der Laan, H. & Kaiser, C. Radio galaxies with a ‘double-double morphology’-I. Analysis of the radio properties and evidence for interrupted activity in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society* **315**, 371–380 (2000).
98. Shurygin, A. Applied Statistics. *Robustness. Estimates. Prediction. Financy i Statistika, Moscow* (2000).
99. Sirothia, S. & Wiita, P. J. Discovery of giant relic radio lobes straddling the classical double radio galaxy 3C452. *The Astrophysical Journal Letters* **765**, L11 (2013).
100. Sokolowski, M., Wayth, R. B. & Lewis, M. *The statistics of low frequency radio interference at the Murchison Radio-astronomy Observatory in 2015 IEEE Global Electromagnetic Compatibility Conference (GEMCCON) 2015 IEEE Global Electromagnetic Compatibility Conference (GEMCCON)* (Nov. 2015), 1–6. doi:[10.1109/GEMCCON.2015.7386856](https://doi.org/10.1109/GEMCCON.2015.7386856).
101. Swarup, G., Ananthakrishnan, S., Kapahi, V., Rao, A., Subrahmanya, C. & Kulkarni, V. The giant metre-wave radio telescope. *Current science* **60**, 95–105 (1991).
102. Thompson, A. R., Moran, J. M. & Swenson Jr, G. W. *Interferometry and synthesis in radio astronomy* (John Wiley & Sons, 2008).
103. Tingay, S. J., Goeke, R., Bowman, J. D., Emrich, D., Ord, S., Mitchell, D. A., Morales, M. F., Boller, T., Crosse, B., Wayth, R., *et al.* The Murchison widefield array: The square kilometre array precursor at low radio frequencies. *Publications of the Astronomical Society of Australia* **30** (2013).
104. Van Der Tol, S., Jeffs, B. D. & van der Veen, A.-J. Self-calibration for the LOFAR radio astronomical array. *IEEE Transactions on Signal Processing* **55**, 4497–4510 (2007).

105. Van Der Veen, A.-J. & Boonstra, A.-J. *Spatial filtering of RF interference in radio astronomy using a reference antenna* in *Acoustics, Speech, and Signal Processing, 2004. Proceedings.(ICASSP'04). IEEE International Conference on* **2** (2004), ii–189.
106. Van Weeren, R., Williams, W., Hardcastle, M., Shimwell, T., Rafferty, D., Sabater, J., Heald, G., Sridhar, S., Dijkema, T., Brunetti, G., *et al.* LOFAR facet calibration. *The Astrophysical Journal Supplement Series* **223**, 2 (2016).
107. Van Haarlem, M. á., Wise, M., Gunst, A., Heald, G., McKean, J., Hessels, J., De Bruyn, A., Nijboer, R., Swinbank, J., Fallows, R., *et al.* LOFAR: The low-frequency array. *Astronomy & Astrophysics* **556**, A2 (2013).
108. Wardle, J. & Aaron, S. How fast are the large-scale jets in quasars? Constraints on both Doppler beaming and intrinsic asymmetries. *Monthly Notices of the Royal Astronomical Society* **286**, 425–435 (1997).
109. Weber, R., Faye, C., Biraud, F. & Dansou, J. Spectral detector for interference time blanking using quantized correlator. *Astronomy and Astrophysics Supplement Series* **126**, 161–167 (1997).
110. Winkel, B., Kerp, J. & Stanko, S. RFI detection by automated feature extraction and statistical analysis. *Astronomische Nachrichten* **328**, 68–79 (2007).