

University of Southampton Research Repository
ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON
FACULTY OF PHYSICAL AND APPLIED SCIENCES
PHYSICS AND ASTRONOMY

The Low Frequency Array and the Transient and Variable Radio Sky

by

Martin Ellis Bell

**Thesis for the degree of Doctor of Philosophy
October 11, 2011**

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES

PHYSICS AND ASTRONOMY

Doctor of Philosophy

THE LOW FREQUENCY ARRAY AND THE TRANSIENT AND VARIABLE
RADIO SKY

by Martin Ellis Bell

This thesis addresses the topic of exploring and characterising the transient and variable radio sky, using both existing radio telescopes, and the next generation of radio facilities such as the Low Frequency Array (LOFAR). Studies of well known variable radio sources are presented in conjunction with blind searches of parameter space for unknown sources. Firstly, a three year campaign to monitor the low luminosity Active Galactic Nucleus NGC 7213 in the radio and X-ray bands is presented. Cross-correlation functions are used to calculate a global time lag between inflow (X-ray) and outflow (radio) events. Through this work the previously established scaling relationship between core radio and X-ray luminosities and black hole mass, known as the ‘fundamental plane of black hole activity’ is also explored with respect to NGC 7213.

Secondly, the technical and algorithmic procedures to search for transient and variable radio sources within radio images is presented. These algorithms are intended for deployment on the LOFAR telescope, however, they are heavily tested in a blind survey using data obtained from the VLA archive. Through this work an upper limit on the rate of transient events on the sky at GHz frequencies is placed and compared with those found from other dedicated transient surveys.

Finally, the design, operation and data reduction procedure for the Low Frequency Array, which will revolutionise our understanding of low frequency time domain astrophysics is explored. LOFAR commissioning observations are reduced and searched for transient and variable radio sources. The current quality of the calibration limits accurate variability studies, however, two unique LOFAR transient candidates that are not present in known radio source catalogues are explored (including multi-wavelength followup observations).

In the conclusion to this thesis the parameter space that future radio telescopes may probe - including the potential rates of such events - is presented. At the nano-Jansky level up to 10^7 transients $\text{deg}^{-2} \text{ yr}^{-1}$ are predicted, which will form an unprecedented torrent of data, followup and unique physics to classify.

CONTENTS

1	Introduction	1
1.1	Principles of radio astronomy	1
1.1.1	Aperture synthesis	5
1.2	The Low Frequency Array (LOFAR)	9
1.3	The LOFAR Transients Key Science Project	17
1.3.1	The LOFAR Radio Sky Monitor	17
1.3.2	Commensal observing	19
1.3.3	High time resolution mode	19
1.3.4	The Transient Buffer Boards (TBBs)	20
1.3.5	LOFAR imaging and transient detection algorithms	20
2	Transient and Variable Radio Sources	21
2.1	Radio emission mechanisms	21
2.1.1	Synchrotron	21
2.1.2	Bremsstrahlung	22
2.1.3	Black body	23
2.1.4	Spectral lines	24
2.1.5	Coherent emission mechanisms	25
2.2	The transient radio universe	25
2.2.1	Flare stars, active binaries, brown dwarfs and ultracool dwarfs	26
2.2.2	Neutron Stars	27
2.2.3	Planets and exoplanets	29
2.2.4	Intraday variability (IDV) and scintillation	30
2.2.5	γ -Ray Bursts (GRBs) and Radio Supernovae (RSNe)	31
2.2.6	Active Galactic Nuclei (AGN) and X-ray binaries	33
2.2.7	Gravitational wave events	35
2.2.8	Future prospects	36
2.3	Review of transient surveys to date.	36
2.4	This thesis	38
3	X-ray and radio variability in the low-luminosity active galactic nucleus NGC 7213	41

3.1	Introduction	42
3.2	NGC 7213 Background	44
3.3	Observations	45
3.3.1	ATCA data analysis	45
3.3.2	RXTE data analysis	46
3.4	Results	46
3.4.1	X-ray/radio light curves and cross-correlation	46
3.4.2	Radio/radio cross-correlation	52
3.4.3	X-ray/radio cross-correlation when flaring	52
3.4.4	Long term variability	53
3.4.5	Linear polarisation	55
3.5	Discussion	56
3.5.1	X-ray/radio jet connection	56
3.5.2	The fundamental plane of black hole activity	57
3.6	Conclusion	64
4	Testing and refining the LOFAR transients detection pipeline	65
4.1	Introduction	66
4.2	The prototype LOFAR transients detection pipeline	66
4.2.1	Input: LOFAR images and metadata	66
4.2.2	Source extraction	68
4.2.3	Quality control	69
4.2.4	MonetDB database and transient searches	71
4.2.5	Source classification and alerts	74
4.2.6	Monitoring	75
4.3	Visualisation	79
4.3.1	Interactive tools	79
4.3.2	Webpages	80
4.4	Considerations of single epoch transients	83
4.5	Conclusion	86
5	Archival radio transients	89
5.1	Testing the LOFAR transient detection pipeline on the Bower field	89
5.1.1	Automated VLA reduction script	90
5.1.2	Results and discussion	92
5.2	The X-ray binary outburst of <i>Swift</i> J1753.3-0127	98
5.3	VLA transient survey using the calibrator fields	99

5.3.1	VLA calibrator fields	101
5.3.2	VLA data analysis	103
5.4	Transient Search	105
5.5	Results	108
5.6	Snapshot rate upper limit	109
5.7	Predictions for LOFAR	115
5.8	Conclusion	117
6	Blind searches for transients in LOFAR commissioning data	119
6.1	Introduction	119
6.2	Test field B0329+54	119
6.2.1	The LOFAR Standard Imaging Pipeline (SIP)	123
6.3	Imaging results and transient search	128
6.4	Transient Candidate ILT J0320.3+5512	137
6.4.1	Multi-wavelength followup and catalogue search	139
6.4.2	ILT J0320.3+5512 Discussion	143
6.5	Transient Candidate ILT J0331.5+5712	146
6.5.1	Multi-wavelength followup and catalogue search	146
6.5.2	ILT J0331.5+5712 Discussion	150
6.6	Conclusion and Future Work	153
7	Conclusion	155
7.1	Summary of thesis and future work	155
7.2	Exploring SKA pathfinder parameter space	157
7.3	A final remark	164
APPENDICES		
A	Example Parsets	167
Bibliography		173

LIST OF TABLES

1.1	LOFAR telescope parameters	14
1.2	Comparison of next generation radio telescopes survey speeds to a sensitivity of 1 μ Jy and 1 mJy	16
3.1	Summary of the lag times which have been calculated using the discrete correlation function	48
3.2	Summary of archival observations found in the literature	54
5.1	Summary of the Gaussian fits produced from the LTrAP compared with those reported in Bower et al. (2007).	94
5.2	Summary of VLA calibrator observations used for an archival transient search.	102
5.3	Table of snapshot rates reported in the literature	113
6.1	LOFAR observations of the B0329+54 field obtained for LOFAR commissioning	122
7.1	Predicted rates of radio transients for SKA pathfinder instruments . .	162

LIST OF FIGURES

1.1	Example aperture synthesis arrays	3
1.2	Field of view against 1 day sensitivity for current and future facilities	4
1.3	Example UV plots for the VLA and WSRT arrays	6
1.4	Map showing the locations of the LOFAR international and core stations	11
1.5	Example LBA and HBA antennas	12
1.6	An example of the possible RSM configurations projected onto a coordinate grid.	18
2.1	A diagram depicting the evolution of the synchrotron shells within the jet of a black hole	23
2.2	The radio evolution of the radio outburst of the X-ray binary XTE J0421+56, or CI Cam	24
2.3	The intensity of the Lorimer burst as a function of radio frequency versus time	29
2.4	Examples of transient radio sources	32
2.5	Examples of transient radio sources	34
3.1	X-ray and radio lightcurves of NGC 7213	47
3.2	Discrete cross-correlation function (DCF) X-ray to 4.8 GHz and X-ray to 8.4 GHz	50
3.3	The z-transformed discrete cross-correlation (ZDCF) function for the radio lightcurves	51
3.4	Achival lightcurves of NGC 7213	53
3.5	The fundamental plane of black hole activity	59
3.6	The fundamental plane using the KFC06 LLAGN sample only and updated FGR10 BHXR sample	60
3.7	Flare 1 radio data points at 4.8 GHz uncorrected for lag	63
4.1	Flow diagram of the prototype LOFAR transients detection pipeline (LTrAP)	67
4.2	Flux per subband of the calibrator B0323.5+5510.	70

4.3	Example lightcurve of significantly variable source	73
4.4	Example (simulated) lightcurves of transient behaviour that LTraP currently struggles to correctly identify.	77
4.5	Example lightcurves of a wrongly classified single epoch transient	78
4.6	Screenshot of prototype LOFAR transients visualisation tool	80
4.7	Example webpage for a transient source generated from the visualisation library.	82
4.8	Simulated detection of a transient where the transient duration is less than the observation integration time	85
4.9	Graph showing the burst flux S_B as a function of burst time T needed to produce a variety of final image fluxes S_{Im}	86
5.1	Flow diagram showing the automated VLA data reduction pipeline (version 1)	92
5.2	Example images of Bower et al. (2007) field	93
5.3	Image and lightcurve of the transient detection RT 1984-05-02	95
5.4	Image and lightcurve of the transient detection RT 1986-01-15	96
5.5	Image of the field containing RT 1999-05-04 made using the automated VLA pipeline	97
5.6	Flow diagram showing the automated VLA data reduction pipeline (version 2)	99
5.7	Radio and X-ray lightcurves of the transient outburst of <i>Swift</i> J1753.5-0127	100
5.8	Histogram showing the time difference between observations for an archival VLA transient survey	104
5.9	Flow diagram showing the automated VLA data reduction pipeline (version 3)	106
5.10	An example image of 3C48 at 4.8 GHz produced by the automated VLA reduction pipeline	107
5.11	Histogram showing the measured noise, calculated in the same region, for all VLA calibrator images	110
5.12	Snapshot rate against flux density of transient detections reported in the literature	112
6.1	A map of the LOFAR station locations used in the commissioning observation L21641	121
6.2	A schematic of the steps involved in the LOFAR standard imaging pipeline	124

6.3	Plots showing the effectiveness of the LOFAR automated RFI flagging routines	127
6.4	LOFAR image of the field 0329+54 taken on December 3 rd 2010	130
6.5	Lightcurve of the three unsubtracted LOFAR calibrators	131
6.6	An image of the 2010-Dec-03 LOFAR observation	132
6.7	An image of the 2010-Oct-10 LOFAR observation	133
6.8	Comparison of the catalogued WENSS fluxes with the associated averaged LOFAR source fluxes	134
6.9	A ‘zoom in’ to the edge of the LOFAR 2010-Dec-03 observation showing the postions of the WENSS sources	135
6.10	The index of variability of all sources extracted in the B0329+54 images	136
6.11	Images showing the transient candidate ILT J0320.3+5512	138
6.12	Lightcurve of ILT J0320.3+5512	140
6.13	Plot showing the positional offset of the Gaussian fits (including errors), to the transient candidate ILT J0320.3+5512	141
6.14	Optical image of the ILT J0320.3+5512	142
6.15	Spectral energy distribution of the candidate source	144
6.16	Images of the transient candidate ILT J0331.5+5712	147
6.17	Lightcurve of the transient candidate ILT J0331.5+5712	148
6.18	Plot showing the positional offset of the Gaussian fits (including errors), to the transient candidate ILT J0331.5+5712	149
6.19	Optical image of the ILT J0331.5+5712	150
6.20	WSRT 1.4 GHz observation of the candidate transient	151
6.21	Images of the transient candidate ILT J0331.5+5712	152
7.1	Predicted parameter space that the SKA and SKA pathfinders may explore	159
7.2	Field of view against 1 day sensitivity for current and future facilities	160
7.3	An example Log N - Log S - Log t	163

DECLARATION

I, Martin Ellis Bell, declare that my thesis entitled *The Low Frequency Array and the Transient and Variable Radio Sky* and the work presented in the thesis are both my own, and have been generated by me as a result of my own original work, I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
 - ★ Investigating the disc-jet coupling in accreting compact objects using the black hole candidate Swift J1753.5-0127, 2010, MNRAS, 406, 1471.
 - ★ X-ray and radio variability in the low-luminosity active galactic nucleus NGC 7213, 2011, MNRAS, 411, 402.
 - ★ An automated archival VLA transients survey, 2011, MNRAS, 415, 2.
- parts of this work will soon be published as:
 - ★ Radio Transients: an antediluvian summary, Fender & Bell, 2011, Bulletin of the Astronomical Society of India, in prep.

Martin Ellis Bell, October 11, 2011

ACKNOWLEDGEMENTS

First and foremost I would like to thank Charli, my wonderful girlfriend. She has been there from the start of my University education 10 years ago - this thesis is dedicated to her. Secondly I would like to thank Rob Fender for giving me this opportunity, helping me through my studies, and drinking many beers - in many countries. Finally I would also like to thank my parents and family for their love and support and for staving off bankruptcy, ruin and hunger throughout my years as a student.

I would also like to thank my good friends (and bandmates) from Cornwall: Luke Toms, Mike Kitaruth, Phil Wilkes, Matt Truen, Niall Tomlinson and Ryan Jones. Thankyou for teaching me to not take myself too seriously (which wasn't too difficult!). We will make some noise one last time at my *viva* party. I have met some really amazing people over the last four years in Southampton. My previous office mates and compadres who include Matthew Schurch, Mark Peacock, Simone Scaringi, Dave Russell, Omar Jamil, Sadie Jones, Grace O'Connor, Tony Wilkinson, Seb Drave and Tana Joseph. Of which, many of them have *forced* me to party with them - a consequence of which being - that most Friday nights I had to sleep in the back of my car. I moved to Southampton shortly after this started to occur regularly. I would like to thank my current office mates Lee Townsend, Liz Bartlett (who shouts Jägerbomb at me any time she is drunk), Dan Calvelo (let us not forget the portrait from Bangkok) and Dan Plant (for a bit of early morning egging).

I would also like to thank some of my great collaborators I have worked over the years. First of all Michael Wise and Sera Markoff, for giving me a roof over my head in Amsterdam and introducing me to a really fat cat (sorry Gingit). Also many thanks to Michael for putting up with my relative deterioration in mental health at a summer school (at an undisclosed astronomical facility). A further mention needs to be given to Sera, who helped relieve my T-shirts of sleeves - for reasons neither her or I can remember - at a late night bowling session in Dwingeloo. Also, Rob for taking a photo of the scene and showing it to everybody the next day. Many thanks to the LOFAR team: John Swinbank, Casey Law, James Miller-Jones, Bart Scheers, Ben Stappers, Evert Rol, Ralph Wijers, Thijs Cohen and Jason Hessels - I have enjoyed many a fine bottle of scotch with you all. A special thanks also to Jess Broderick for the countless conversations we have had about LOFAR data reduction over the last 18 months.

I also like to thank the staff at the University of Southampton. Specifically, Tom Maccarone and Phil Uttley for their help and support over the years. I would also like to thank the 'girls' from the school office: Mandy, Wendy, Kim, Hannah and Jan for the countless questions I have needed to ask regarding (usually) badly behaving students.

Merastawhye - Kernow bys Vyken

Martin Ellis Bell, October 11, 2011

To 'C' for all the soup years

Astronomy? Impossible to understand and madness to investigate.

SOPHOCLES, C. 420 BCE

I don't know how people have time to read books, I am far too busy lying drunk face down in wet fields.

OLIVER REED (1938 – 1999)

When I had satisfied myself that no star of that kind had ever shone before, I was led into such perplexity, by the unbelievability of the thing, that I began to doubt the faith of my own eyes.

TYCHO BRAHE (1546 - 1601)

1

Introduction

1.1 Principles of radio astronomy

Since the serendipitous discovery of radio emission emanating from the Milky Way by Karl Jansky (1905-1950), radio astronomy has delivered paradigm shifting advances in astrophysics. These include the discovery of the Cosmic Microwave Background radiation (CMB; Penzias & Wilson 1965), the Pulsar (Hewish et al., 1968), and the first binary Pulsar - which provided the most stringent test of General Relativity to date (Taylor et al., 1979). All three of these discoveries were awarded the Nobel prize. One of the most difficult challenges, which has and always will make radio astronomy a difficult discipline, is manifested in the very nature of radio waves: wavelength. The typical radio wavelength is some million times greater than in the optical regime. Thus to achieve a radio image with the same resolution as the human eye (~ 1 arcsec) we would need a radio telescope with a diameter of ~ 1 kilometre, operating at a wavelength of 50cm (see Sir Martin Ryle's Nobel Lecture 1974).

There were obviously technological challenges to building parabolic radio reflectors of 1 km in diameter in the mid 20th Century and many still remain today. To reconcile the problem multiple (smaller) radio telescopes are placed some dis-

tance apart (in arrays) and a technique called aperture synthesis (Ryle & Hewish, 1960) is used to make a ‘synthesised’ single telescope. The collecting area of the array is less than a completely filled aperture (resulting in lower sensitivity for such arrays), but the greatest angular resolution is dictated by the maximum baseline length. By utilising greater numbers of parabolic reflectors or dipole elements the telescope sensitivity can be improved.

Radio interferometers that were constructed in the late 20th Century, such as the Westerbork Synthesis Radio Telescope (WSRT; Baars & Hooghoudt 1974), the Very Large Array (VLA; Hjellming & Bignell 1982), the Giant Metrewave Radio Telescope (GMRT; Swarup 1990) and the Australian Telescope Compact Array (ATCA; Frater et al. 1992), use a small number (6 - 27) of parabolic reflector antennas. The VLA, ATCA and WSRT telescope diameters are $\sim 25\text{m}$, the GMRT uses 42m dishes. The next generation of radio telescopes many use may thousands of smaller ($\sim 10\text{m}$ and less) parabolic reflectors, or dipole elements. See Figure 1.1 for photographs of the VLA (top panel), the ATCA (middle panel) and also an artists impression of the proposed Square Kilometer Array (SKA; bottom panel), which will potentially have a total collecting area of one square kilometre. In the parabolic case, decreasing the antenna diameter increases the field of view, while increasing the numbers of antennas and baseline length maximises the sensitivity and resolution. In the case of dipoles, the potential field of view is the whole sky, as they are omni-directional. They are also cheap to build, therefore many hundreds of thousands of antenna elements can be incorporated into an array. Figure 1.2 shows the field of view of current and future radio facilities, with respect to sensitivity. A defining characteristic of the future radio facilities is the instantaneous sensitivity, and wide-field of view they can achieve. I will discuss some of these facilities further in later sections.

One of the major difficulties with ‘many element’ aperture synthesis is the transportation of vast quantities of data over large distances and also the computational power needed to produce images from the data. Thus the revolution, for radio astronomy, comes in the form of high speed optical fibre networks (either dedicated or over the internet) which are used for data transport; and the use of (what is considered at the time of writing) supercomputers for correlation and data reduction. We are on the cusp of a major shift derived from these technological advances which may allow us to deliver Sir Martin Ryle’s one kilometre telescope (or larger). However, the next generation of radio telescopes will still rely on the principles of aperture synthesis.

In the following sections, I will briefly explore the mathematics pertaining to



Figure 1.1: Example aperture synthesis arrays. Top panel: The Very Large Array (VLA) which consists of $27 \times 25\text{m}$ dishes, with a maximum baseline length of $\sim 35\text{km}$, located in New Mexico, USA. Middle panel: The Australia Telescope Compact Array (ATCA) located in Narrabri, Australia, which contains $6 \times 25\text{m}$ dishes that can have a maximum baseline length of 6 km . Bottom panel: The proposed Square Kilometer Array (SKA) which may consist of many thousands of small parabolic reflectors and dipoles and will be located in either South Africa or Australia.

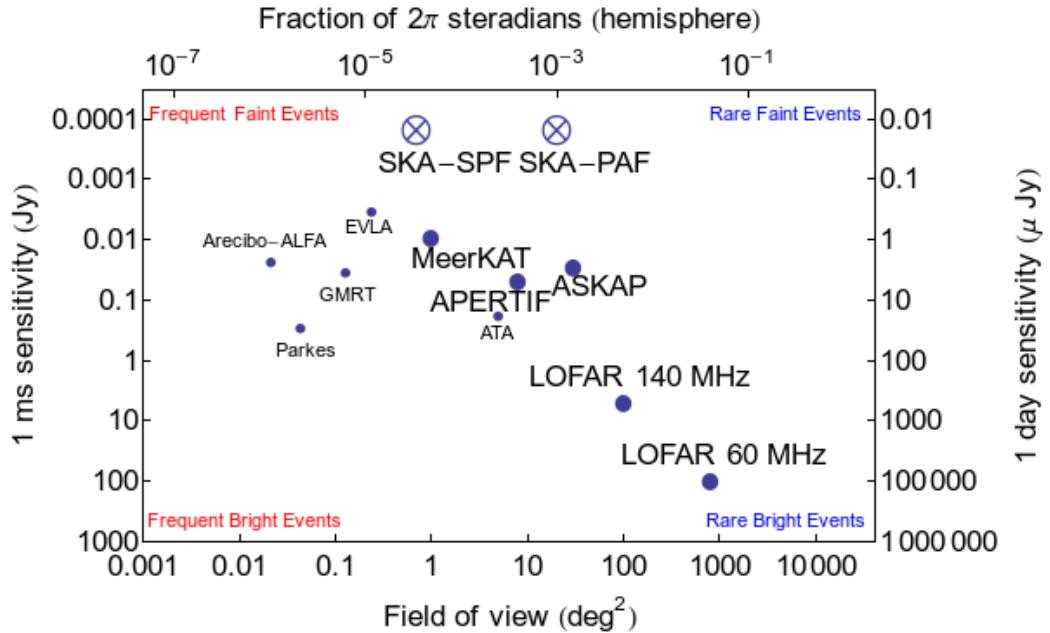


Figure 1.2: Field of View (FoV; deg^2) against 1 day sensitivity (μJy right axis) for current facilities (small circles) and new facilities (large circles). The Australian Square Kilometre Array Pathfinder (ASKAP) and the Karoo Array Telescope (MeerKAT) are both GHz facilities which are in the early stages of construction in Australia and South Africa respectively. ASKAP will utilise phased array feed technology to achieve a FoV of $\sim 30 \text{ deg}^2$, whilst MeerKAT will use a single pixel feed to achieve a FoV $\sim 0.5 \text{ deg}^2$: both facilities will be fully operational ~ 2014 . APERTure Tiles in Focus (APERTIF) is an upgrade to the WSRT telescope which will increase the FoV to $\sim 8 \text{ deg}^2$ (using phased array feeds), and will be completed by 2013. SKA-SPF denotes the FoV and sensitivity of a single pixel SKA feed. SKA-PAF denotes a phased array feed, which has a wider FoV. Figure courtesy of R. Fender.

aperture synthesis, as this is fundamental to the science presented in this thesis. I will also explore how the next generation of radio telescopes, such as The LOw Frequency ARray (LOFAR; de Vos et al. 2009) will need to advance on the strategies developed for the older generation of radio instruments to achieve wide-field, high dynamic range imaging. It is not only a technological shift which will define this new generation of radio telescope: it is also the transformational science they will provide. In the context of this thesis, I will concentrate on time domain or ‘transient’ radio science. This introduction reviews the current state of the art of transient radio surveys and also explore examples of known and predicted transient radio sources. It also presents the design, construction and operation of the LOFAR telescope.

1.1.1 Aperture synthesis

The principle of aperture synthesis relies heavily on the inverse Fourier transform. A large number of individual antennas (or telescopes) are separated by a variety of distances. The completeness of the synthesised aperture is dependent on the declination of the source being observed, and the antenna locations as a function of time: these components define the uv plane. With a good distribution of antennas, short timescale observations can produce good images. Alternatively, a method that is more common for East - West type arrays is using the Earth's rotation to synthesise a complete aperture. Both methods effectively measure Fourier components of the astronomical target source(s). The WSRT is an East - West array i.e. the antennas are positioned linearly in an East - West direction (see Figure 1.3). Therefore to achieve good spatial sensitivity, the Earth's rotation must be used to sample a large number of Fourier components. The VLA is a 'snapshot' array whereby the antennas are positioned in a Y-shape (see Figure 1.3) which achieves much better instantaneous spatial sensitivity. The Earth's rotation can still be used to fill in the VLA's synthesised aperture further, however, good images can be produced in short timescales. Figure 1.3 shows the array layout of both the VLA and WSRT; also shown is the instantaneous overhead snapshot uv coverage and the full synthesis (12 hour) uv coverage for a source at $\delta = 60^\circ$. It can be seen that WSRT relies heavily on the Earth's rotation to adequately fill in the uv plane, whilst the VLA can achieve relatively good coverage in a short time.

The astronomical radio waves arriving at each different antenna are corrected for geometric delay (due to different path length differences). Then, for each antenna pair within the telescope, at each uv position and integration time unit, the amplitude (A) and phase (θ) of the correlated signal is measured. We can then use both A and θ to calculate the complex visibility $V(u, v) = Ae^{i\theta}$: which is equal to the 2D Fourier transform of the source sky function $I(x, y)$ evaluated at each uv position. The full expression is shown in equation 1.1 below.

$$V(u, v) = \iint_{-\infty}^{+\infty} I(x, y) e^{2\pi i(ux+vy)} dx dy \quad (1.1)$$

To characterise the target source well, we would want to sample as many visibilities as possible at different spacial frequencies. We can then find the sampled sky function $I'(x, y)$ and produce an image by using the inverse Fourier transform:

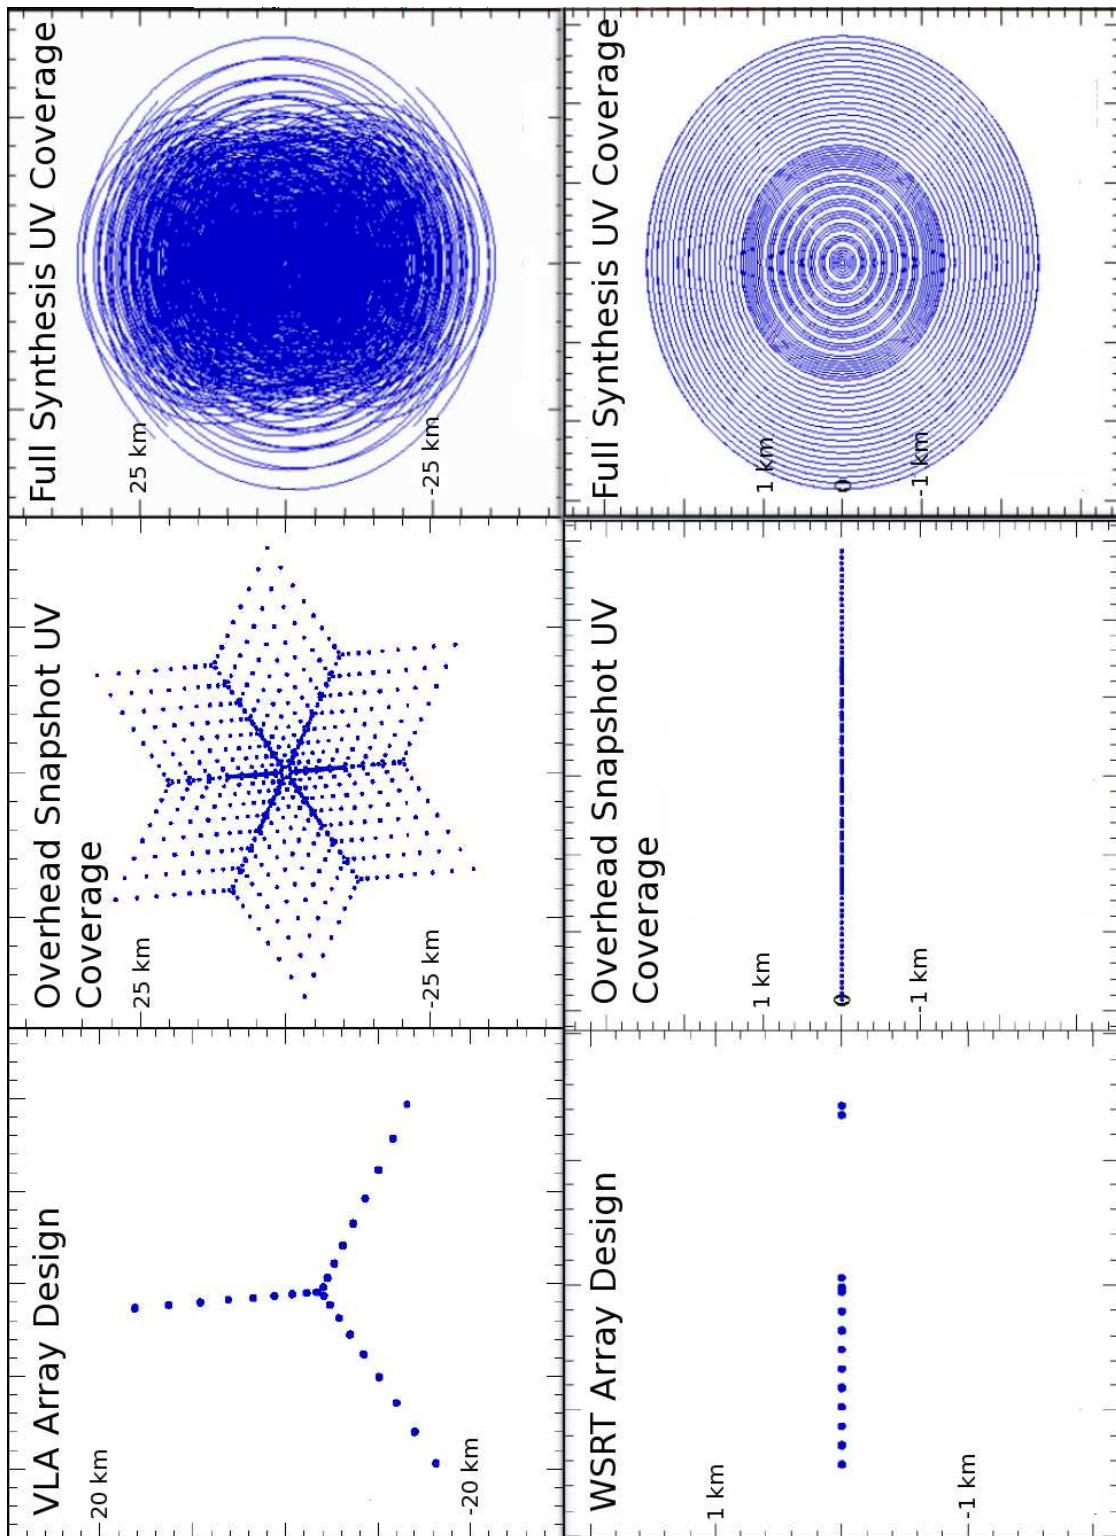


Figure 1.3: Top panel: The array layout of the VLA (left) with the corresponding overhead snapshot uv coverage (middle) and the full synthesis UV coverage for a source at $\delta = 60^\circ$. Bottom panel: The array layout of the WSRT (left) with the corresponding overhead snapshot uv coverage (middle) and the full synthesis UV coverage for a source at $\delta = 60^\circ$. Figure modified from http://192.96.5.2/synthesis_school/Miod_Array_Design.pdf

$$I'(x, y) = \iint_{-\infty}^{+\infty} V(u, v) e^{2\pi i(ux+vy)} du dv \quad (1.2)$$

Note, if the uv plane is completely sampled then $I'(x, y)$ is the true sky function. If the uv plane is sparsely filled then we yield an approximation of the sky function. This approximation can also be characterised by using the convolution theorem:

$$I'(x, y) = B(x, y) * I(x, y) \quad (1.3)$$

$B(x, y)$ is the Fourier transform of the uv function, which is typically referred to as the dirty beam. Equation 1.3 shows that what we are typically measuring with an interferometer, is the sky function $I(x, y)$, convolved with the response of our instrument $B(x, y)$.

The difficulty from this stage onwards is that we do not know the values of $V(u, v)$ for all possible uv components. There are a number of approaches to solving this problem. In one implementation, we can set all visibilities equal to zero where they have not been sampled, then perform the inverse Fourier transform to get a distorted image. We then use an approach called deconvolution which assumes that we have a number of point sources within our sky function. By assuming that these sources are point like, we can then deconvolve the effect of the dirty beam and hence ‘clean’ an image to reduce distortions (e.g. see Högbom 1974, Clark 1980).

To further complicate the situation, we must undertake a set of tasks collectively referred to as calibration. The term calibration encompasses a large number of corrections that must be applied to the visibilities to produce an adequate image. These corrections must account for: perturbations in instrumental stability (system temperature), atmospheric transmission (e.g. phase calibration), positional/pointing accuracy, polarisation (including cross-polarisation leakage), absolute flux scale calibration, bandpass calibration, bandwidth smearing and primary beam correction.

Typically for older interferometers, such as the VLA or ATCA, the calibration applied only makes approximations for *some* of the corrections described above. For example, a typical 1st generation calibration strategy is to measure the flux and phase of a nearby bright source, before and after the target observation. The derived calibration is then applied to the target source. This method relies on instrumental stability between the target and calibrator observations. A 2nd generation calibration strategy, sometimes referred to as SELFCAL, is a closed-loop system which continuously updates the calibration as a function of time and is therefore less dependent on instrumental stability (see Noordam & Smirnov 2010 for a discussion

of calibration strategies). Some of the corrections described above (such as amplitude and phase calibration) are degenerate moving away from the phase centre across the field of view. As discussed earlier the next generation of radio telescopes will potentially image very large fields of view, therefore a more complete description of equation 1.1 must be implemented, which will incorporate 3rd generation calibration techniques.

The Radio Interferometer Measurement Equation (RIME) was first derived by Hamaker et al. (1996) and is further discussed in Smirnov (2011a,b,c). The RIME takes *all* the terms needed for calibration and represents them as Jones matrices. These Jones matrices fully map out the transmission of the electromagnetic wave (in polarisation, amplitude and phase) through the interferometer system. They form multiple simultaneous equations which must be solved to completely calibrate all the perturbations in the system. The calibration matrices are split into two types, direction dependent and direction independent. Direction dependent equations have typically been ignored for interferometers such as the VLA and ATCA, but must be used for calibrating wide-field instruments. The matrices are combined with equation 1.1 to form a complete measurement equation:

$$V(u, v) = G_p \left(\iint_{-\infty}^{+\infty} X(p, q) e^{2\pi i(ux+vy)} dx dy \right) G_q \quad (1.4)$$

where

$$X_{p,q} = E_p I(x, y) E_q \quad (1.5)$$

$I(x, y)$ is a 2×2 brightness matrix which characterises the polarised sky brightness as a function of direction (x, y) . E_p and E_q represent a number of direction dependent Jones matrices which contain the corrections needed for antennas p and q . G_p and G_q represent the direction independent Jones matrices. If the telescope only has a small field of view, E_p and E_q can be neglected and the measurement equation returns to equation 1.3 - with the inclusion of the direction independent corrections G_p and G_q . The full implementation of RIME is computationally expensive and has probably slowed progress in radio astronomy. However, RIME is a complete description of the telescope system and must be considered when working with wide-field instruments.

The purpose of showing this theory is to place the technical work of this thesis in context. In this thesis I will use typical 1st and 2nd generation calibration and imaging techniques in Chapters 3 and 5. In Chapter 6 I will present the process of

producing wide-field images with LOFAR, for which direction dependent effects are apparent. A variety of new wide-field facilities will soon be available to sample the transient sky, and will make use of the techniques described above. In the radio band the Allen Telescope Array (ATA; Welch et al. 2009), the Murchison Wide Field Array (MWA; Lonsdale et al. 2009), the Long Wavelength Array (LWA; Ellingson et al. 2009) and LOFAR will soon begin or have already commenced operations. Other wide-field radio pathfinders such as the Karoo Array Telescope (MeerKAT; Booth et al. 2009) and the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al. 2008) are also being developed on the road to the Square Kilometer Array (SKA; Rawlings & Schilizzi 2011). Transient studies are a key science goal for all of these facilities.

1.2 The Low Frequency Array (LOFAR)

The LOFAR telescope is an international project (led by The Netherlands) which utilises many low-cost dipole antenna elements (organised into stations) that are constructed at different locations across Europe. There are eight international LOFAR stations located in the following countries: France ($\times 1$), Germany ($\times 5$), Sweden ($\times 1$) and the United Kingdom ($\times 1$), see Figure 1.4 for a photograph of the UK Chilbolton station (left). There is a central ‘core’ located in The Netherlands which consists of 22 individual stations, see Figure 1.4 for a photograph (right). There are also 18 Dutch remote stations with baseline lengths < 100 km (from the central core). The locations of the international and the Dutch core LOFAR stations are shown in Figure 1.4. As LOFAR operates in the relatively unexplored low frequency regime, very long baseline lengths are needed to achieve good resolution.

The LOFAR telescope design is a relatively unique concept as it does not use large steel dishes to focus radio waves, it uses many thousands of non-movable omni-directional dipole antennas, combined with the process of aperture synthesis and imaging. The concept of using dipoles in a telescope array is not so new. There has been a rich historical heritage in using dipoles and non-parabolic telescope design in the past. For example Karl Jansky’s radio telescope used simple dipoles on a steerable platform (circa. 1930). Other instruments such as the Clark Lake Teepee-Tee Telescope and the Ukrainian T-shaped Radio Telescope (UTR-2) have, or still use, non-parabolic telescope design. The difference for LOFAR (and other instruments discussed below) is the inclusion of long baseline aperture synthesis and imaging. Many of the older generation instruments were not capable of pro-

ducing images or incorporating signals from antennas spaced on kilometre baseline lengths. Telescopes such as the MWA and LWA have also adopted a non-parabolic design, similar to LOFAR.

Each LOFAR international station typically consists of 96 low band dipole antennas, and 96 high band tiles. The Dutch stations can vary in size from a half to a quarter that of the international stations. Signal processing techniques are used to add phase delays in the signal path (between different dipole elements) to ‘steer’ the telescope beam to a region of the sky. This process is typically referred to as beamforming and the result is referred to as the station beam. Once a station has been phased together, the station beams from two different stations are correlated to provide the *uv* components described in section 1.1

There are two distinct types of antennas: the Low Band Antenna (LBA) shown on the top panel of Figure 1.5, which operate in the frequency range 10 to 90 MHz (but are mostly effective in the range 30 - 80 MHz); and the High Band Antenna (HBA) shown on the bottom panel of Figure 1.5, which operate from 110 to 250 MHz. The gap in observing frequency is due to the FM radio band, which will saturate all observations at those frequencies. The station design and subsequent beam forming of the LBA and HBA are slightly different. The HBA are arranged into tiles which consist of 16 folded dipoles arranged into 4×4 blocks. These 4×4 dipole blocks are then phased together using an analogue beam former at the tile level - which is referred to as the tile beam. This restricts the HBA field of view before the data have been digitised. The station beams are then formed digitally from the data stream received from the tiles. Note, there are 96 dipoles in an international LBA station and 96 *tiles* in the HBA. As each HBA tile consists of 16 dipoles, a total of 1536 dipoles are in fact present in the HBA station. The wavelength is smaller at the HBA frequency range than at the LBA range, therefore a larger number of dipoles can be packed together in a smaller space in the HBA. This increase in the number of dipoles comes with the consequence of increased data rates. The HBA dipoles are therefore phased together at the station level (using the analogue beam former) to reduce the data volume, which comes with the tradeoff of reducing the field of view.

The LBAs are arranged in a pseudorandom manner (within a station) and the signals are phased together digitally to form the station beam. Note, as there is no analogue beam former for the LBAs, they essentially view the whole sky all the time. Bright sources, such as Cassiopeia A (which are extremely bright at low frequencies \sim kJy), are therefore always present in observations. Signal processing techniques must be used to subtract the effects of these bright sources to produce

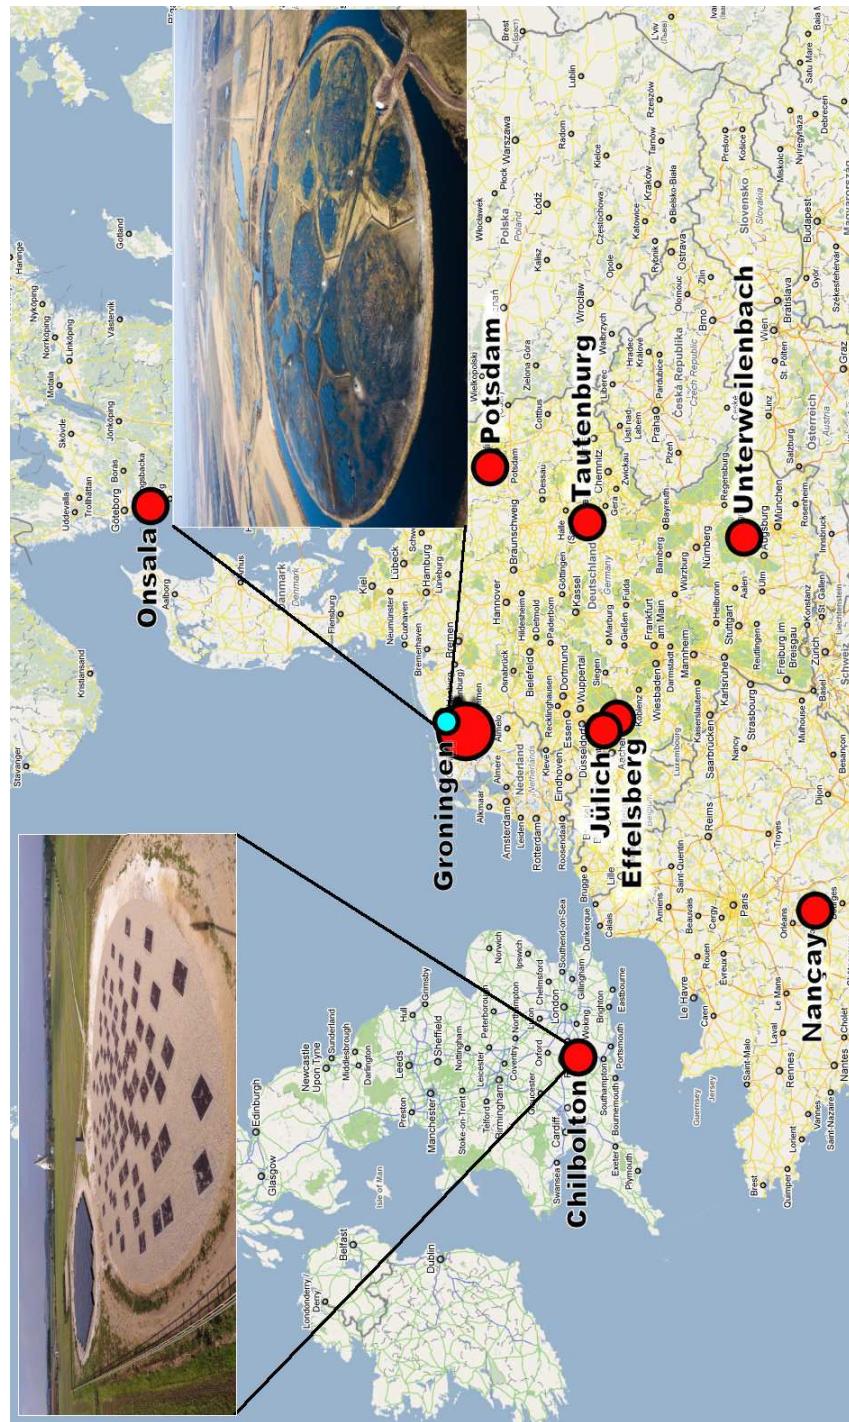


Figure 1.4: Map showing the locations of the LOFAR international and core stations. A photograph is shown (left) of the recently constructed Chilbolton UK station (credit: LOFAR UK). The large circle represents the central core station, which comprises of a number of smaller individual stations. It also encompasses the Dutch remote stations which offer intermediate baseline lengths with respect to the core and international stations. A photograph is shown (right) of the central core station or ‘Superterp’ (credit: Astron).



Figure 1.5: Top panel: The first Low Band Antenna (LBA) to be constructed at the LOFAR UK Chilbolton station. The antenna consists of a low noise amplifier mounted on a piece of tubing, which is secured via the XX and YY polarised dipole wires to a reflective ground plane (credit: LOFAR UK). Bottom panel: Two High Band Antennas (HBA) which are typically deployed in 4×4 tiles within a protective casing (credit: Astron).

adequate images.

The number of station beams formed by LOFAR is simply dictated by hardware and software constraints. For example one single station beam can be made with a bandwidth of 48 MHz, or up to 244 beams can be formed with a bandwidth of $48/244 = 0.2$ MHz per-beam. All beam formed data collected via the stations is transported to a central processing hub in Groningen (shown on Figure 1.4 in blue) which utilises an IBM BlueGene/P supercomputer to produce the correlated visibilities. The data volumes are large; a 12 hour dataset with 25 stations using 1 second visibility averaging produces 15 TB of data. (McKean et al., 2011). Further details of the design and signal processing of LOFAR see de Vos et al. (2009). A full description will also soon be available by Haarlem et al. (in prep).

The RMS thermal noise for the LOFAR array can be calculated using (see de

Vos et al. 2009):

$$\sigma = \frac{W}{\sqrt{2(2\delta v \delta t) \left\{ \frac{N_c(N_c-1)/2}{S_{core}^2} + \frac{N_c N_r}{S_{core} S_{remote}} + \frac{N_r(N_r-1)/2}{S_{remote}^2} \right\}}} \quad (1.6)$$

Where σ denotes the noise measured in the final image (Jy). δv is the operating bandwidth (MHz). δt is the integration time (s). N_c is the number of core stations and N_r is the number of remote stations. S_{core} and S_{remote} denote the system equivalent flux densities (SEFDs) for the core and remote station respectively. W is a factor that accounts for the weighting scheme used in imaging (i.e. Natural, Robust etc.). At 150 MHz these values are $S_{core} = 2.8$ kJy, $S_{remote} = 1.4$ kJy and $W = 1.3$. It should be noted that at the time of writing LOFAR was still being constructed, therefore these values are subject to change. The reader should refer to de Vos (2009) for the predicted values of the SEFDs to be used in calculating the RMS noise. If only one type of station is used in the telescope setup then equation 1.7 (below) can be implemented.

$$\sigma = \frac{S_{sys}}{\sqrt{N(N-1)2\delta v \delta t}} \quad (1.7)$$

Note, the calculation above only accounts for the thermal or theoretical noise limit, other factors such as dynamic range effects, bright sources in the side lobes, ambient temperature and RFI etc can play a role in increasing the noise. The resolution of LOFAR can be calculated using $\theta \sim \lambda/D$. Typical values of the sensitivity and resolution as a function of frequency for the core and full LOFAR array are given in Table 1.1.

Due to LOFAR's large field of view it is extremely well placed to survey large areas of sky. This also makes LOFAR an excellent instrument to explore time domain radio astrophysics. To place the parameters summarised in Table 1.1 in context, I can compare these parameters with those expected for the Expanded Very Large Array (EVLA), which is an older generation interferometer with upgraded receivers. For the EVLA in A-configuration at 5 GHz, a field of view (FoV) of 0.07 deg^2 is realised. The sensitivity would be $\sigma_{EVLA} = 2.0 \mu\text{Jy}$, in a 12 hour synthesis (assuming a 1 GHz bandwidth). When compared with LOFAR a FoV of 28.8 deg^2 is achievable at 120 MHz with a sensitivity of $\sigma_{LOF} = 110 \mu\text{Jy}$, using the core station only in a 12 hour synthesis with 48 MHz bandwidth.

Although the LOFAR sensitivity does not seem as competitive as the EVLA, it is important to consider the spectral indexes of radio sources. For example, a typical synchrotron source with $\alpha = -0.7$ (where $S_\nu \propto \nu^{+\alpha}$) at 5 GHz would increase in flux

Table 1.1: Typical sensitivity (σ), resolution (θ) and field of view (Ω) as a function of frequency, for the full array (including all international stations) and the Dutch core station only. I assume the core consists of 22 stations, and the full array consists of 22 core and 26 remote stations. The sensitivities have been calculated using a bandwidth of 48 MHz. The field of view for each configuration has been taken from the LOFAR technical website¹. Note, the system equivalent flux densities (SEFDs) used to calculate the parameters below are based on predicted values. The sensitivities may therefore change (slightly) once the array is completed and the SEFD is properly parametrised.

Freq. (MHz)	$\sigma_{1\text{sec}}$ (mJy) Core	$\sigma_{12\text{hours}}$ (mJy) Core	$\sigma_{1\text{secs}}$ (mJy) Full	$\sigma_{12\text{hours}}$ (mJy) Full	θ ($''$) Core	θ ($''$) Full	Ω (deg 2) Core
15	2980	14.3	1350	6.5	1650	48	1676
30	549.3	2.64	248.6	1.20	825	25	419
75	314.8	1.51	142.4	0.69	330	10	67
120	22.2	0.10	6.5	0.0314	206	6	28.8
200	22.8	0.11	6.6	0.0316	125	3.5	10.1

density by approximately an order of magnitude at 150 MHz. Therefore to achieve a source detection of 10σ with LOFAR with a source spectrum of $\alpha = -0.7$, the equivalent detection with the EVLA would yield a 55σ detection, which is a factor of five improvement. However, in terms of FoV approximately 400 times the EVLA FoV can be imaged with LOFAR, in this example.

To further explore the competitiveness of LOFAR as a survey instrument in Table 1.2 I compare the survey speed of LOFAR with current and future instruments. The survey speed has been calculated using:

$$SS = \frac{1}{\tau} \times FoV \left[\text{deg}^2 \text{hr}^{-1} \right] \quad (1.8)$$

Where τ is the time taken to reach a given sensitivity in hours, and FoV is the field of view of the instrument. Each instrument will obviously reach a given sensitivity in a different amount of time, I therefore normalise each instrument by the amount of time taken to reach $1 \mu\text{Jy}$ and 1mJy (see Table 1.2). I then calculate how many deg 2 (using the FoV) each instrument could survey in one hour, to define the survey speed in units of deg $^2 \text{hr}^{-1}$. The LOFAR survey speed at 120 MHz using the full array is a factor of 100 faster than the VLA. However, some of the next generation of radio telescopes, such as the SKA, are extremely competitive. In this calculation I have used raw sensitivity only; as I demonstrated in the discussion earlier, a source with a typical spectral index of $\alpha = -0.7$ at 5 GHz would increase in flux by a factor of ~ 10 at 150 MHz. This corresponds to a factor of 100 increase in survey speed to detect a population of sources with this spectral index, because $\tau \propto 1/(10\sigma)^2$.

The spectral index should be taken into account when comparing survey speeds at different wavelengths, as should the Log N - Log S of the known radio source population. I will return to a discussion of the known radio source and transient Log N - Log S later in this thesis. I will also discuss the concept of the transient source Log N - Log S - Log t, which is the time domain consideration of the standard Log N - Log S.

Table 1.2: Comparison of next generation radio telescope survey speeds. $\sigma_{12\text{hours}}$ gives the sensitivity of each instrument in a 12 hour integration. FoV gives the field of view of each instrument. τ gives the times needed to reach a sensitivity of 1 μJy and 1 mJy respectively. The survey speed gives the area of sky that could be surveyed per hour to a sensitivity of 1 μJy or 1 mJy.

Telescope	$\sigma_{12\text{hours}}$ (μJy)	FoV (deg^2)	$\tau = 1\mu\text{Jy}$ (hours)	Survey speed to 1 μJy ($\text{deg}^2 \text{ hr}^{-1}$)	$\tau = 1\text{mJy}$ (seconds)	Survey speed to 1 mJy ($\text{deg}^2 \text{ hr}^{-1}$)
APERTIF (1.4 GHz)	10	8	1200	0.006	4.3	6666
SKA (Mid 1.4 GHz)	0.015	1	0.0027	370	9.7×10^{-6}	3.7×10^8
SKA (Low $< 300\text{MHz}$)	0.4	200	1.9	104	6.9×10^{-3}	1×10^8
MeerKAT (1.4 GHz)	1.8	1	38	0.025	0.14	2.5×10^4
ASKAP (1.4 GHz)	10.2	30	1248	0.024	4.5	2.4×10^4
LOFAR (120 MHz Full)	31.4	28.8	1×10^4	0.0024	42.6	2.4×10^3
LOFAR (30 MHz Full)	1200	419	1.7×10^7	2.4×10^{-5}	6.2×10^4	24.2
EVLA (5 GHz)	2.0	0.07	47	1.4×10^{-3}	0.17	1500
EVLA (74 MHz)	5000	322	3.3×10^8	2.0×10^{-10}	1.2×10^6	2.0×10^{-4}
ATCA (5 GHz)	5.3	0.07	300	2.3×10^{-4}	1.0	2.3×10^2
GMRT (153 MHz)	15.5	10	2700	3.7×10^{-3}	9.7	3.7×10^3
ATA ^a (1.4 GHz)	50	2.45	3.0×10^4	8.1×10^{-5}	100	81

^a ATA - Allen Telescope Array

1.3 The LOFAR Transients Key Science Project

There are six LOFAR key science projects: The Epoch of Reionisation, Extragalactic Surveys, Cosmic Rays, Solar Physics, Cosmic Magnetism, and Transients (see <http://www.lofar.org> for further details). Each of these projects will benefit from dedicated observing time with LOFAR to achieve specific science goals. The Transients Key Science Project (TKSP; Fender et al. 2006) will aim to detect and monitor a broad range of transient astrophysical phenomena using LOFAR. The TKSP can be broadly split into two specialist areas: high time resolution and imaging. The high time resolution team (sometimes referred to as the Pulsar team) will probe transient phenomena with characteristic timescales $1\mu\text{sec} \leq \delta t \leq 1\text{sec}$. The imaging team will probe transient phenomena with timescales $1\text{sec} \leq \delta t$ to many years. A variety of observational and transient detection techniques will be employed by both teams, and I will explore these in detail in the following sections. In this thesis I am primarily concerned with image plane searches for transients, and the associated algorithmic detection techniques. The division between high time resolution and image plane searches is merely a division of observation and detection method. As I will discuss in greater detail later, certain classes of transients could be prominent in both high time resolution and image plane searches. For example, a Pulsar would typically fall into the category of high time resolution, however, pulsars are known to scintillate and should also be detected as highly variable sources in image plane searches.

1.3.1 The LOFAR Radio Sky Monitor

The Radio Sky Monitor (RSM) holds enormous scientific potential and will usher in a new era of all-sky radio astrophysics. The RSM will use a number of different overlapping beam configurations to tile up a large fraction of the sky to search for transients. This wide-field mode will open up the type of serendipitous discovery space that has been available to X-ray instruments, such as the All Sky Monitor (ASM) on the Rossi X-ray Timing Explorer (RXTE), for some time. Utilising LOFAR's wide-field radio advantage, coupled with a potential daily monitoring cadence, the low-frequency radio parameter space can be explored and characterised.

As discussed earlier, the available 48 MHz LOFAR bandwidth can be divided up to form a maximum of 244 beams. Figure 1.6 shows two possible configurations, at both 120 MHz (left panel) and 30 MHz (right panel) for surveying large areas of the sky. On the left panel, I show using black circles a Full Width Half Maximum

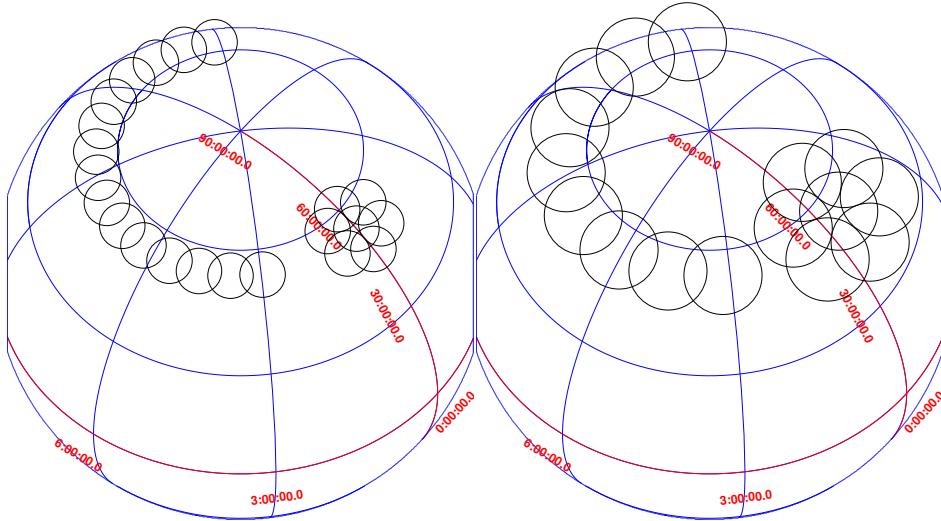


Figure 1.6: An example of the possible RSM configurations projected onto a coordinate grid. Left panel: Black circles represent the FWHM of the HBA at 120 MHz ($\theta = 5.3^\circ$). The hexagonal pattern shows seven beams which overlap by a factor of $\text{FWHM}^{0.5}$, which yields a total FoV ~ 0.12 steradians. The LOFAR zenith track is also shown which can be tiled up using ~ 30 overlapping beams. Right panel: The same configuration but for the LBA at 30 MHz, each circle represents a FWHM of $\theta = 9.1^\circ$. When tiled together using seven beams, a FoV of ~ 0.35 steradians is realised. Approximately 20 LBA beams can tile up the entire zenith.

(FWHM) of 5.3 degrees, which corresponds to the radius of the LOFAR HBA FoV at 120 MHz. By overlapping these fields of view by a factor of $\text{FWHM}^{0.5}$ (see Condon et al. 1998), using seven overlapping beams, we can survey an area of sky ~ 0.12 steradians (see hexagonal shape shown in Figure 1.6). Alternatively, in this figure I also show a zenith monitoring track. The LOFAR zenith is at $\sim 54^\circ$ North, by placing consecutive overlapping beams along this track, approximately 30 HBA beams can tile up the entire zenith. Observing at the zenith offers extra beam stability and optimises sensitivity (Fender et al., 2006). A combination of overlapping beams can be used to track the zenith fields for a given amount of time, using Earth's rotation. The right panel of Figure 1.6 shows the same concept for the LBA at 30 MHz. At 30 MHz the LBA has a FWHM of 9.1 degrees; when this is tiled together with seven beams an area of 0.35 steradians is realised.

To best probe the parameter space available, images will be created on logarithmically spaced ranges of timescales. The images will then be searched for temporal changes in source brightness, and sources which are not found in current radio or multi-wavelength catalogues. Also, in a hybrid type observing mode, LOFAR could dedicate a large number of its beams to the RSM (including the zenith fields),

while a smaller number of beams perform daily monitoring of known radio transient sources, such as black hole X-ray binaries or active galactic nuclei.

1.3.2 Commensal observing

In parallel to dedicated TKSP observations, commensal observations can offer a large amount of data for transient searching. For example, the planned LOFAR Million Source Sky Survey (MS³) will survey the low frequency sky. The survey will not only offer a deep census of the LOFAR sky, but all data obtained through this survey can be searched for transient and variable radio sources, on a variety of timescales. The MS³ survey will potentially be primarily concerned with producing one image of a given field over the total integration time (~ 15 mins). The TKSP can divide this data up into (on the smallest scale) one second chunks to search for transients. In the future deeper surveys will be performed and the TKSP anticipate to search this data for transients. Furthermore, high time resolution data can be obtained simultaneous with imaging data, therefore all pulsars surveys (which are part of the TKSP mandate) are excellent candidates for commensal image plane transient searches. If a transient is detected via commensal observations, internal triggers can be used to re-point LOFAR, or via external triggers, to alert other facilities to obtain multi-wavelength followup.

1.3.3 High time resolution mode

High time resolution mode (sometimes referred to as Pulsar mode) will give LOFAR the capabilities of probing time-scales $1\mu\text{sec} \leq \delta t \leq 1\text{sec}$. The telescope setup, data reduction and data products are different from those produced in imaging mode. The data products will be time-series data rather than images. Specialist high time resolution algorithms will search for fast periodic and intermittent periodic signals, as well as single bursts. Also, known pulsars will be monitored to explore their low frequency profiles and properties.

Two different signal processing techniques, known as coherent and incoherent station addition, are used to combine the signals from different stations to form times-series data. The coherent station addition adds all the signals from the stations together, explicitly correcting for the phase difference between the station signals, from a particular direction. The incoherent addition combines the signals from the stations without correcting for the phase difference. There are trade offs to using each of these signal processing techniques. Coherent addition yields a better signal

to noise, but has a small field of view; this mode will mostly be used to monitor known pulsars due to its precision. Incoherent addition gains a larger field of view, but loses sensitivity, therefore this mode will be largely used for surveys. For further information on high time resolution mode and early results from observations of known pulsars see Stappers et al. (2011).

1.3.4 The Transient Buffer Boards (TBBs)

One of the unique properties of LOFAR is the Transient Buffer Boards (TBBs). The TBBs essentially consist of Random Access Memory (RAM) located at each station, that has the ability to record raw voltages for up to 1.3 seconds, at full bandwidth (Stappers et al., 2011). This gives LOFAR the ability to reprocess visibilities in greater detail after an event has occurred. The amount of time the TBBs can record is simply scaled by the amount of RAM available (and is upgradeable). Larger amounts of data can also be recorded if the bandwidth is reduced, which is defined by:

$$T = 1.3 \text{ seconds} \times \frac{100 \text{ MHz}}{B_r} \quad (1.9)$$

B_r is the reduced bandwidth (Stappers et al., 2011). Using a 4 MHz bandwidth, which is potentially the amount which will be assigned to each RSM beam, 32.2 seconds of raw data may be recorded. This functionality is extremely useful for potentially receiving transients triggers from other instruments, or LOFAR itself, and beamforming at the region of interest to produce an image(s). See Singh et al. (2008) for further information on the TBBs. Note, this mode will primarily be used (at least initially) by the Cosmic Ray Key Science Project (see Falcke & Key Science Project (2007) for more details).

1.3.5 LOFAR imaging and transient detection algorithms

To produce images and detect transients on \sim one second timescales, all data reduction must be handled by dedicated pipelines. Automated data reduction pipelines, for both imaging and transient detection, are currently in a state of heavy testing and commissioning. Chapter 4 describes in detail the LOFAR transient detection pipeline (LTrAP) and the testing that has been conducted on the datasets described in this thesis. Chapter 6 describes the LOFAR standard imaging pipeline (SIP), and the application of the LTrAP to LOFAR commissioning data.

All energy flows according to the whims of the great magnet.

HUNTER S. THOMPSON (1937 - 2005)

2

Transient and Variable Radio Sources

In this chapter I will review some of the prominent radio emission mechanisms, as they are fundamental to the science presented in this thesis. I will also review the types of known transient and variable radio sources. In the conclusion to this chapter I will review the current state of the art of transient radio surveys.

2.1 Radio emission mechanisms

2.1.1 Synchrotron

Synchrotron radiation is produced by the acceleration of electrons (or positrons) in a helical path due to magnetic field lines. The velocity of the spiralling electrons is relativistic. Synchrotron is closely related to cyclotron radiation, where the velocity of the spiralling electrons is non-relativistic, and also gyro-synchrotron where the velocity is mildly relativistic.

The spectrum of a synchrotron source is characterised by the underlying electron population. For the non-thermal case, the electron population can be described by a power-law distribution of electron energies i.e. $N(E) = N_0 E^{-\delta}$, then $S_\nu \propto \nu^{-\frac{\delta-1}{2}}$ (where S_ν is the flux density in Janskys). This typically results in optically thin spectrums of $\alpha \sim -0.7$ (where $S_\nu \propto \nu^{+\alpha}$). Synchrotron emission can also

be absorbed, which results in a turnover of the spectrum with $\alpha \sim +2.5$ at lower frequencies, for the non-thermal case. If the distribution of electron energies can be described as a thermal distribution, then the resulting spectrum at low frequencies yields $\alpha \sim +2$ (further discussion will follow on thermal radiation). The peak flux of the radio emission is found at the frequency at which the emission switches from optically thick to thin.

Furthermore, the electron populations and the environments in which synchrotron radiation is produced can evolve. For example, in the radio emission produced by an expanding shell of plasma in the jet of a black hole (either an X-ray binary or active galactic nuclei), the electron population evolves from being dense to less-dense, as a function of time. Other factors such as changes in the magnitude of the magnetic field also affect the evolution of emission as a function of time (see van der Laan 1966 for further details). Figure 2.1 shows a simplified cartoon of this physical process. At a given frequency the shell of material will initially be optically thick (which will yield a spectrum of $\alpha \sim +2.5$ in the non-thermal case). As the shell expands the flux rises as the shell becomes optically thin ($\alpha \sim -0.7$). This process is repeated at each different observing frequency. In active galactic nuclei and X-ray binaries the broad-band radio emission can commonly be described as flat spectrum ($\alpha \sim 0$). If we sum up all the synchrotron peaks (over a range of frequencies) we arrive at a flat spectrum (see Figure 2.1). In this simple model, I have ignored the consequence of ‘shocks’. I will return to this topic including a further discussion on jets in active galactic nuclei and X-ray binaries in Chapter 3.

To further demonstrate the model, Figure 2.2 shows lightcurves at a number of different frequencies, of an outburst from the X-ray binary XTE J0421+56, or CI Cam (see Belloni et al. 1999 and Fender et al. 2006). At the early stages the outburst of the source has already peaked at higher frequencies. The lower frequency emission, specifically 0.33 GHz, starts to rise later due to the delay time for the expanding synchrotron shell to become optically thin at lower frequencies. At all frequencies, the decay of the emission (after the peak) is due to adiabatic expansion losses as the ejecta expands.

2.1.2 Bremsstrahlung

Bremsstrahlung emission (sometimes referred to as free-free emission) results from the acceleration of electrons within a plasma. Electrons are deflected by the ionised charge of a heavier ion and the resulting change in energy releases radiation. The resulting spectrum at high (radio) frequencies is relatively flat, then moving to low

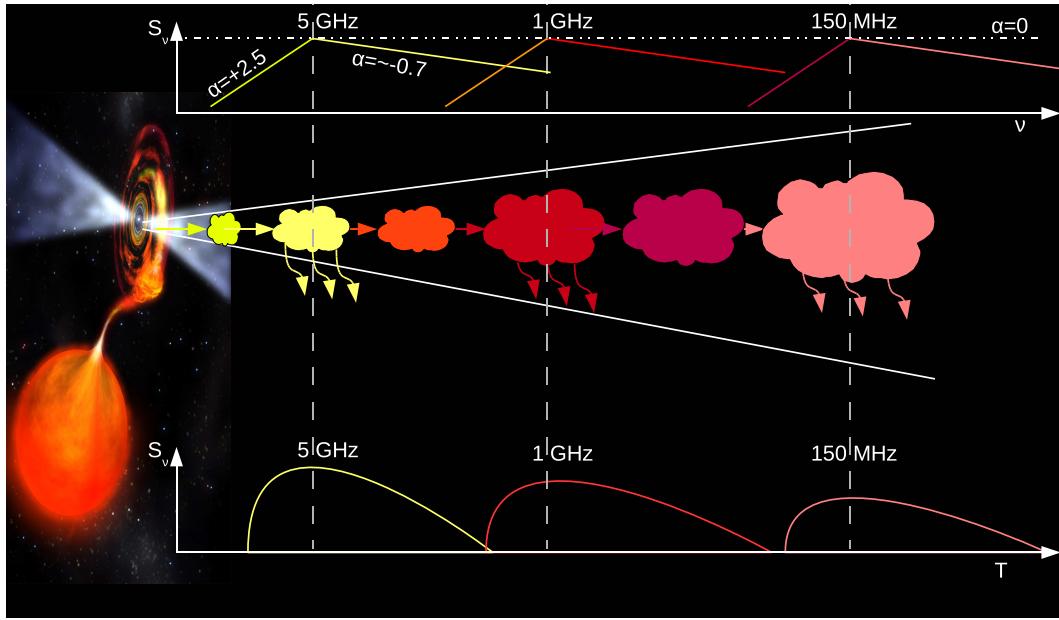


Figure 2.1: A diagram depicting the evolution of the synchrotron shells within the jet of a black hole. At a given frequency, as the shell expands, the radio emission will rise whilst the spectral index evolves from optically thick to thin. The summation of the spectra over a broad frequency range, typically results in a flat spectrum. This model does not consider ‘shocks’ within the jet which will be discussed later in this thesis.

frequencies there is a turnover after which the spectrum is described by $\alpha = +2$. Bremsstrahlung is very similar to synchrotron radiation, the difference being that synchrotron requires an ambient magnetic field: Bremsstrahlung utilises a charged plasma. Bremsstrahlung is therefore common in hot ionised environments, such as the hot intra-cluster gas of galaxy clusters.

2.1.3 Black body

Black body radiation arises from the oscillations of charged particles within an object. The flux density of a black body source increases with λ^{-2} (in the Raleigh-Jeans approximation), which yields a spectrum of $\alpha = +2$. Most astronomical sources emit via the black body process and it is also observed whenever a thermal population of electrons (or protons) is optically thick (e.g. thermal synchrotron). The quiet sun, the moon and the cosmic microwave background radiation are examples of black body emitters.

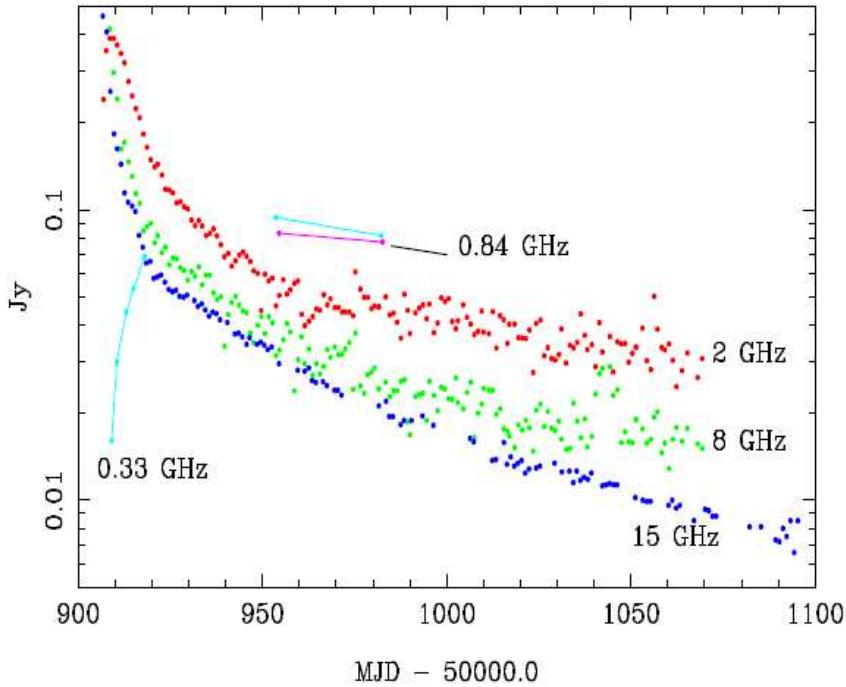


Figure 2.2: The radio evolution of the radio outburst of the X-ray binary XTE J0421+56, or CI Cam. Observations of the source were conducted with the Ryle Telescope (15 GHz), the Green Bank Interferometer (2 and 8 GHz) and the WSRT (0.8 and 0.33 GHz). The lightcurves show the late arrival time of the lower frequency (0.33 GHz) synchrotron emission. Figure and description taken from Fender et al. (2006).

2.1.4 Spectral lines

Spectral line processes are common across the electromagnetic spectrum. One of the most prominent and abundant spectral line processes in the radio regime, is the 21cm line. The 21cm line arises from a hyperfine spin-flip transition in neutral hydrogen, resulting in the release of a photon at 21cm, or 1.4 GHz. This type of radiation is produced at a discrete wavelength, therefore it can also be used to measure velocities using Doppler shifts. The 21cm has been used to map galactic rotation with our galaxy and others. It is also a potential probe of cosmology, a highly red-shifted 21cm line, at the time of the epoch of reionisation, would emit at a frequency <200 MHz. The LOFAR Epoch of Reionisation Key Project hope to study the density, temperature and velocity profiles of the heavily shifted 21cm line to learn more about the early universe.

2.1.5 Coherent emission mechanisms

The electron cyclotron maser is one example of a coherent radio emission mechanism. Electrons are accelerated along magnetic field lines (typically in a planetary or solar environment), these electrons then produce photons via the cyclotron process. The cyclotron photons are then subsequently amplified by the maser (Microwave Amplification through Stimulated Emission of Radiation) process. This bulk flow of electrons provides a population inversion in the electron distribution, which provides a ‘pump’ for the maser. The input photons stimulate greater numbers of photons to be produced, and provides an amplification of the original cyclotron radiation. The photons emitted by the maser process have the same (or similar) properties to the input photons, and since all photons are travelling in the same direction, they appear to come from a very compact region. The brightness temperature of a source is inversely proportional to the source size, therefore electron cyclotron masers (and coherent sources in general) have been shown to have extremely high brightness temperatures. Also, this type of emission has been shown to have high degrees of circular polarisation. Other types of coherent process exist which follow a similar theme to the electron cyclotron maser i.e. bulk flow of electrons - for further information on coherent processes see Dulk (1985). In the next section I will discuss the known types of transient radio sources and explore which emission mechanisms dominate.

2.2 The transient radio universe

The scope and depth of transient radio science is vast: by utilising the time domain we can gain unique insight into such objects as neutron stars and white dwarfs in binary systems, relativistic accretion and consequent jet launch around black holes, distant gamma-ray burst afterglows, supernovae and active galactic nuclei (AGN), to name a few. The distances to these objects, as well as the timescale for transient behaviour, varies dramatically. For example, giant kilo-Jansky micro-second radio pulses have been observed from the relatively nearby Crab Pulsar (e.g. see Bhat et al. 2008). In contrast, month timescale (and longer) variations are commonly observed in the radio emission produced by powerful jets driven by accretion onto super-massive black holes in distant AGN (Aller et al. 1985; de Vries et al. 2004; Bell et al. 2010; Jones et al. 2010). Through studying the transient and variable nature of these exotic and energetic objects, we obtain an unprecedented laboratory to probe extreme physics, on a variety of distance scales and in a number of different

environments. I will comment further on some of the broader scientific motivation to studying transients later in this chapter.

This thesis is concerned with the search, detection and scientific interpretation of transient and variable radio sources at both MHz and GHz frequencies. I will therefore review some of the known and predicted types of transient radio sources at both frequencies. I will put an emphasis on the expected differences between GHz and MHz populations (if appropriate), and also the typical timescales for transient behaviour and expected flux densities.

2.2.1 Flare stars, active binaries, brown dwarfs and ultracool dwarfs

Flare stars are typically late type stars (often referred to as dM or dMe stars; Bastian 1990) that undergo transient outbursts which increase the flux density by a factor of ~ 500 (from mJy quiescent levels), at a range of radio frequencies (Strassmeier et al., 1993). The outbursts are often characterised at both MHz and GHz frequencies by: (i) a high degree (up to 100%) of circular polarisation (Jackson et al. 1989; Lang et al. 1983; Gudel et al. 1989); (ii) a high spectral index, typically $\alpha \sim -2.5$ (Lovell 1969; Spangler et al. 1974); (iii) high brightness temperatures, typically $T_b > 10^{12} K$ (Lovell 1969; Spangler & Moffett 1976); (v) bursts that can last from tens of seconds to tens of minutes (Bastian 1990), sometimes up to days (Kundu et al. 1988; Slee et al. 2003). Due to the high brightness temperatures, steep spectral index and short timescale of variability a coherent emission mechanism is favoured over gyro-synchrotron $<\sim 5$ GHz. Above 5 GHz gyro-synchrotron can not be ruled out because the brightness temperature does not exceed $T_B = 10^{12}$. A sub-division of the flare star class incorporates active binaries (RS CVn and BY Dra Binaries) which share similar characteristics to standard flare stars (see Mutel & Weisberg 1978; Fix et al. 1980; van den Oord & de Bruyn 1994 and references therein). Based on the flare star radio burst statistics of Abada-Simon & Aubier (1997) and subsequent interpolation by Fender¹ et al. (2007) to LOFAR considerations, flare stars will be abundant in routine monitoring with LOFAR. Due to LOFARs one second cadence and wide-field of view, a broadband census of nearby flare star activity can be achieved.

Brown dwarfs are the smaller main sequence ‘non-starter’ cousins of flare stars; they are distinguished from planets via their ability to burn deuterium. Less is

¹The LOFAR Transients Key Project (TKSP) Project Plan

known about brown dwarfs due to the low surface brightness, which makes detection at some frequencies difficult. A small number of brown dwarfs (late M and L type) have produced radio bursts ($\sim < 1$ mJy between 5 and 8 GHz) with moderate amounts of circular polarisation. The bursts are potentially consequences of magnetic activity on the surface of the star (Berger et al. 2001; Berger 2002). Synchrotron radiation has been used in some cases to explain the emission properties at GHz frequencies (Berger et al., 2001), but more recently Hallinan et al. (2008) have shown that coherent bursts are the dominant mechanism.

Ultracool dwarfs straddle the stellar/substellar boundary between Brown dwarfs and stars (spectral type $\geq M7$; Kirkpatrick et al. 1997). Ultracool dwarfs have a much lower surface temperature and magnetic field than their higher mass stellar counterparts. Despite the lower magnetic fields these stars have been shown to be relatively bright sources at radio frequencies (see Burgasser & Putman 2005 and Phan-Bao et al. 2007). Of greater interest is the fact that some of these sources can produce *periodic* coherent bursts - which has dubbed them ‘dwarf pulsars’. The ultracool dwarf TVLM 513-46546 has been shown to produce periodic bursts of $p = 1.96$ hours, which is consistent with the orbital period of the star (see Hallinan et al. 2007 and Berger et al. 2005). With fluxes of milli-Janskys at GHz frequencies (potentially much brighter at lower frequencies) and periods of \sim hours, LOFAR image plane searches, coupled with periodicity searches, could potentially detect greater numbers of ultracool dwarfs.

2.2.2 Neutron Stars

Pulsars emit pulsed emission across the electromagnetic spectrum, with pulse periods typically less than a second. The pulsed emission is a consequence of a rapidly rotating neutron star, which creates beamed radio emission (at the poles) that traverses our telescopes line of sight (see Figure 2.5 for a schematic of a Pulsar). Pulsars typically have a steep spectral index 0 to -4 (Maron et al., 2000), with flux densities ranging from 0.1 mJy to 5 Jy at 400 MHz (Stappers et al., 2011). As well as emitting pulsed emission, pulsars have also been shown to emit giant radio pulses. For example, the Crab pulsar (PSR B0532+21) has emitted a number of $\sim 10^5$ Jy pulses (at 430 MHz) lasting just $100\mu\text{s}$ (Hankins et al. 2003, Cordes, Lazio, & McLaughlin 2004).

Currently there are ~ 2000 known radio pulsars (Lorimer et al. 2006; Stappers et al. 2011). Based on the predictions of van Leeuwen & Stappers (2010), in a 25-

day all-sky Galactic survey, 900 new Galactic pulsars within a radius of ~ 2 kpc are potentially detectable with LOFAR; and also potentially giant single pulses from greater distances. Previous surveys for radio pulsars have typically been conducted at 450 MHz or 1.4 GHz. In conjunction for searching for new pulsars, LOFAR will also open a new observing window to study the pulse profiles of known pulsars (see Stappers et al. 2011).

A recently discovered class of object known as Rapidly Rotating Transients (RRATs) are considered to be intermittent pulsars (McLaughlin et al., 2006). A number of individual pulses are received followed by up to hours or days of non-detections. By performing single pulse searches for RRAT type sources, using data obtained from the Parkes Pulsar survey, Lorimer et al. (2007) reported one of the most promising and exciting results in transient science.

A 30 Jy, <5 ms single burst was found in the direction of the Small Magellanic Cloud (SMC) with a dispersion measure which indicated cosmological distances (< 1 Gpc). Figure 2.3 shows the intensity of the radio burst as a function of radio frequency versus time (taken from Lorimer et al. 2007). The burst shows a quadratic shift in arrival time as a function of frequency, which is expected from a signal propagating through the interstellar/galactic medium (due to the refractive index). If the source is astrophysical, the energy needed to provide such a large amount of flux, in such a short time duration is vast. Lorimer et al. (2007) derive a brightness temperature of $T_b \sim 10^{34}$ K and energy released of $\delta E \sim 10^{33}$ J. This pre-requisite rules out many astrophysical progenitors (such as synchrotron emitters), and only leaves extreme cataclysmic possibilities, for example, neutron-neutron star mergers and coalescing black holes.

There was doubt cast on the original validity of the result of Lorimer et al. (2007) by Burke-Spolaor et al. (2011), who reported 16 bursts with similar characteristics. The Parkes multibeam receiver (used to detect these bursts) consists of a 13-beam feed system, mounted in the prime focus of the Parkes telescope. Each beam is sensitive to a different focal plane of the parabolic reflector. The 16 bursts reported by Burke-Spolaor et al. (2011) saturated all 13 Parkes beams. This would imply that they originated from a large field of view on the sky, at the same time. Burke-Spolaor et al. (2011) therefore concluded that these 16 new bursts came from a terrestrial origin i.e. the bursts did not arrive at the receivers from reflecting off the parabolic surface, but from an angle directly into the feeds. The Lorimer et al. (2007) burst did saturate one of the Parkes beams, but this is as expected for an astronomical source. Recent work by Keane et al. (2011), who were searching for further RRAT candidates, found a number of single pulses, of which, some had

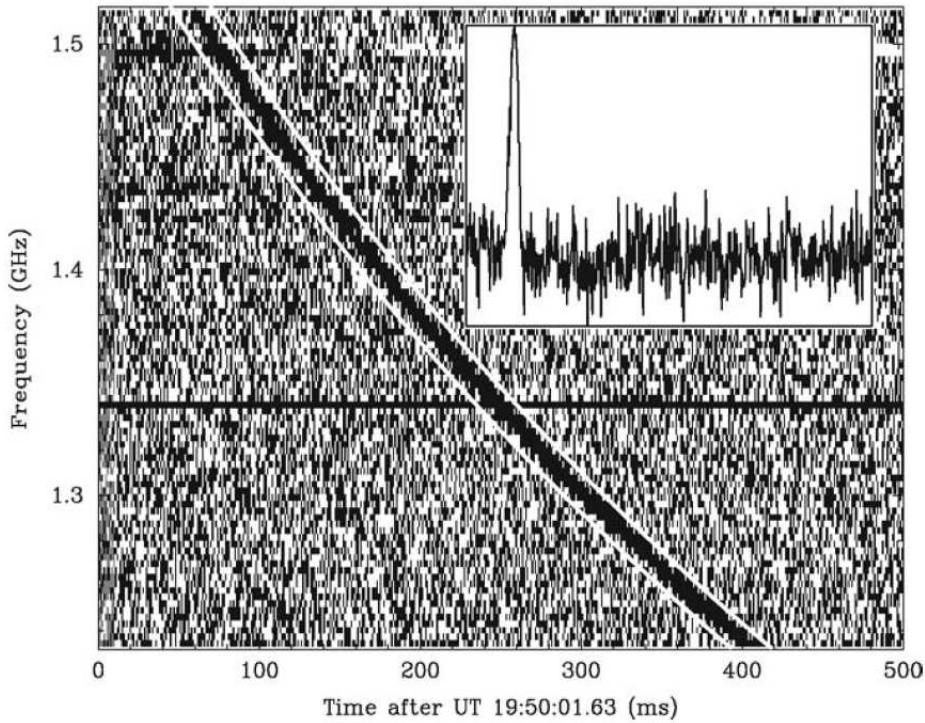


Figure 2.3: The ‘Lorimer’ burst was a 30 Jy burst lasting for just <5 ms, which is postulated to originate from cosmological distance scales (~ 500 Mpc). The figure shows the intensity of this radio burst as a function of radio frequency versus time (figure taken from Lorimer et al. 2007). Inner panel: Integrated pulse profile.

similar properties to those reported by Lorimer et al. (2007).

Regardless of the initial validity, the results of Lorimer et al. (2007) (and potentially Keane et al. 2011) have sparked much interest by both theorists and astrophysics with regard to the possible importance of short duration bursts (e.g. studying the magnetic field and turbulence in the IGM). All of the bursts so far have been detected by the Parkes telescope, so it is of high priority to search for some of these extreme pulses with other instruments, such as LOFAR.

2.2.3 Planets and exoplanets

Jupiter has an intense magnetic field which has been shown to accelerate electrons to produce bright non-thermal coherent low frequency (10-40 MHz) radio bursts (Lecacheux et al. 1998; Zarka 2004a). A major source of charged particles entering the magnetosphere comes from volcanic activity on Io (and other satellites), and also the stellar wind. In conjunction with studying decametric radio bursts from Jupiter, LOFAR will search for bursts from hot Jupiter type exoplanets. A selec-

tion effect in optical exoplanet search techniques typically finds large Jupiter sized planets, orbiting close to a companion star. The proximity to the companion star's stellar wind is predicted to trigger Jupiter type decametric bursts (Lazio et al. 2004; Grießmeier et al. 2007). Detecting these radio bursts would provide the first *direct* detection of an exoplanet, rather than an indirect detection based on the exoplanets influence on its surroundings e.g. it's gravitational influence.

Radio signatures of planetary lightning discharges have been observed by space based and ground based experiments. The Voyager missions have made detections of Saturn (Warwick et al., 1981) and Uranus (Zarka & Pedersen, 1986); the Venera landers detected lightning signatures on Venus (Russell, 1991). Saturn's radio bursts can be bright \sim 0.4 to 1000 Jy, with a typical duration of 30 to 300 ms (Zarka et al., 2004b). Planetary lightning is conceivable for any planet with a strong magnetic field, and should scale approximately with strength of the planetary magnetic field, and the inverse squared distance to the planet. Therefore LOFAR will potentially detect signatures of lightning discharges from a number of planets - potentially ones that have not yet yielded any ground based detections. Planetary lightning is a useful probe of atmospheric dynamics and chemistry, and also the electrification process.

2.2.4 Intraday variability (IDV) and scintillation

Scintillation is the rapid variation in the apparent brightness of a distant luminous object (typically a Quasar) caused by a changing electron density in the interstellar/galactic medium or an intervening scattering screen. There are two types of scintillation: refractive and diffractive. Refractive scintillation is the focusing and defocusing of 'rays'. Diffractive scintillation is the interference amongst different paths of a ray, between a source and receiver (Goodman, 1997). Diffractive scintillation is dependent on source size, therefore only extremely compact sources such as Pulsars are affected. Refractive scintillation commonly affects larger sources such as AGN (Goodman, 1997). Due to the common timescales of these temporal variations ($\sim < 1$ day; Dennett-Thorpe & de Bruyn 2003; Dennett-Thorpe & de Bruyn 2002), and the distances to the sources (which implies a limiting source size), scintillation can only be attributed to propagation effects. Scintillation has been shown to effect pulsars (see Rickett 1977) and γ -ray burst afterglows (Frail et al., 1997) at GHz frequencies. Rickett et al. (1984) also showed that interstellar scintillation would be prominent for flat spectrum extragalactic radio sources at low

frequencies, as limitations in angular resolution usually dictate that a large fraction of sources are unresolved (also see Hunstead 1972; Spangler et al. 1989; Bondi et al. 1996). The SKA will potentially sample a population of sub micro-Jansky radio sources, scintillation will therefore be a prominent effect and persistent cause of variability.

2.2.5 γ -Ray Bursts (GRBs) and Radio Supernovae (RSNe)

γ -ray bursts (GRBs) are extremely bright, distant, short lived (\sim seconds) bursts. A number of possible progenitors are proposed for GRBs; all of them are associated with a unique phase at the end of a stars evolution. One candidate involves the collapse of a star (known as a collapsar), the other is associated with binary mergers of either neutron star binaries, or neutron star - black hole binaries (Mészáros, 2002). The redshifts of GRBs have shown them to originate from cosmological distance scales, which makes them isotropically distributed across the sky (Meegan et al. 1992; Mészáros 2002). GRB afterglows have been detected in the X-ray (Metzger et al., 1997), optical (van Paradijs et al., 1997) and radio (Frail et al., 1997) wavebands. The dominant radio emission mechanism is generally considered to be non-thermal synchrotron radiation emitted by shock accelerated electrons within a jet (van Paradijs et al., 2000). I will present a further description of a ‘jet’, in the AGN and X-ray binaries section, and also Chapter 3. Due to the large inferred distances to GRB (afterglows), scintillation due to the inter-galactic medium can introduce large variability in the flux (Frail et al., 1997). At LOFAR frequencies, the rise time for a typical GRB is predicted to be \sim years, with a peak flux less than ~ 1 mJy (van der Horst et al., 2008). Following up GRB triggers with LOFAR could potentially yield large datasets (of the order of years), which could track GRB afterglows deep into the late stages of evolution.

A further property to GRBs is the ‘orphan’ case, whereby the γ -ray emission is beamed away from our line of sight; the radio emission is emitted isotropically from an expanding shell of material. This yields only an orphan radio afterglow with no high energy counterpart. In the radio, approximately 10-100 GRB orphan afterglows are predicted to be detected isotropically per day, with a lifetime around 1 month and peak observed brightness at GHz frequencies of $100 \mu\text{Jy}$ (Weiler et al., 2004). Blind surveys which probe the (extragalactic) sub-mJy population of radio transients are important to help distinguish between persistent highly variable sources, and ‘one off’ highly variable events such as orphan GRBs.

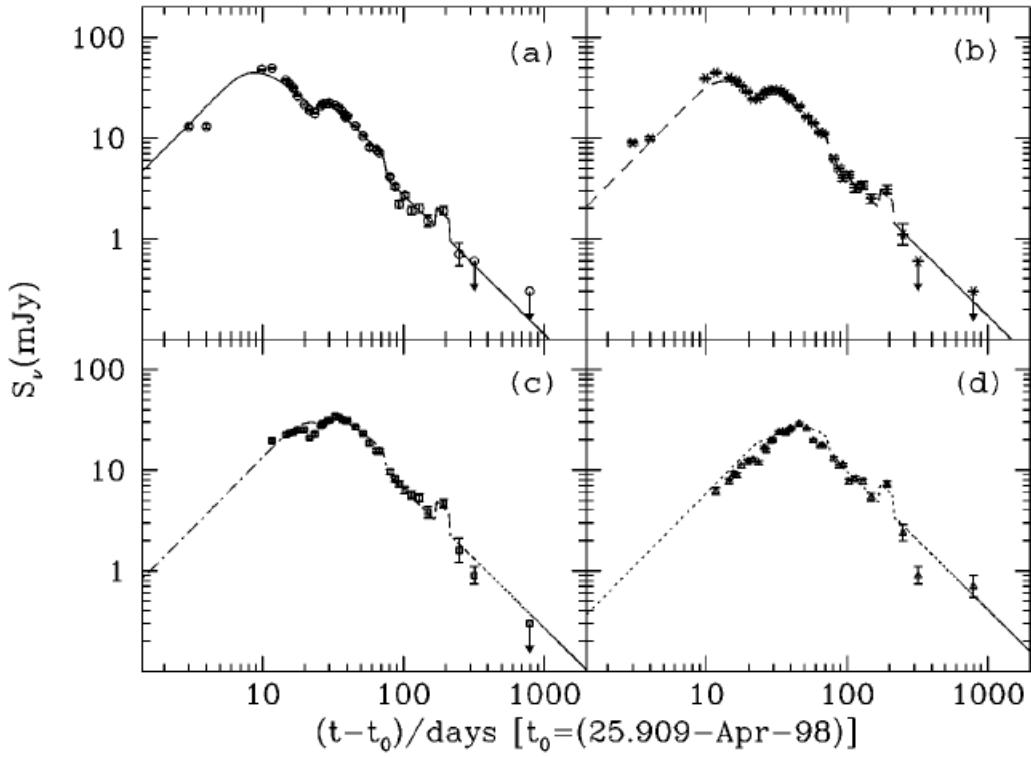


Figure 2.4: Radio lightcurves of Supernovae SN1998bw, taken from Weiler et al. (2002). The frequencies of the observations are (a) 8.6 GHz (b) 4.8 GHz (c) 2.5 GHz and (d) 1.4 GHz. The peak flux of the radio emission arrives later at lower frequencies.

A Supernovae is the explosive death of a massive star; the resulting explosion (and subsequent shock wave) results in an extremely luminous (at all wavelengths), short lived astronomical source. Possible links have been suggested between GRBs and radio supernovae (RSNe), for example, GRB980425 (Soffitta et al., 1998) and SN1998bw (Tinney et al., 1998) are generally considered to be produced by the same progenitor (Weiler et al., 2002). The classification of supernovae is defined by the optical properties: a broad classification of type I and type II divides supernovae based on the detection or non-detection of (optical) hydrogen emission lines; further subdivisions of these classes are provided by the presence of certain emission lines, such as Si and He, and also lightcurve shape (Beswick, 2006). Type Ib/c and II are associated with core-collapse mechanisms and can be radio luminous.

The emission mechanism for core-collapse supernovae is via the non-thermal synchrotron process. Typical radio lightcurves go through an initial fast rise, due to shock accelerated electrons propagating through the circumstellar medium, followed by an expansion phase, which gives a slower decay in the lightcurve. At 8.6

GHz SN1998bw peaked at ~ 400 mJy (at 10 to 12 days, see Figure 2.4), and was detectable up to 200 days post-explosion (Weiler et al., 2002). As the peak radio emission evolves from high to low frequencies, LOFAR will potentially be able to track and explore the later stages of radio supernovae evolution.

2.2.6 Active Galactic Nuclei (AGN) and X-ray binaries

AGN are a class of galaxy where the central nucleus is extremely energetic. The most abundant members of the AGN class are Seyfert galaxies and Quasars. There are however many sub-divisions of the AGN class where the nature of the AGN is described by attributes at a specific frequency, and selection effects such as viewing angle. Radio loud classes of AGN are some of the earliest studied variable radio sources (see Condon & Backer 1975; Dennison et al. 1981; Ryle et al. 1978; Taylor & Gregory 1983; Condon et al. 1979 and references therein). They are bright at all radio frequencies, sometimes problematically bright (\sim kJy) at low frequencies i.e. Cygnus A is so bright its effects are always present in LBA LOFAR observations. AGN harbour some of the largest black holes in our universe and are powerhouses of accretion. Due to accretion and jet production via the super-massive black hole, AGN undergo bright transient outbursts at all frequencies with typical timescales at radio frequencies of months to years (e.g. see Aller et al. 1985 but also references above). The radio emission is produced predominately in large collimated jets via the the synchrotron process, either within shock regions, or through steadily evolving relativistic outflows (e.g. see Figures 2.5 and 2.1 and also Blandford & Konigl 1979).

Low-Mass X-ray Binaries (LMXBs), or microquasars, are the smaller ($\sim 10M_{\odot}$) Galactic versions of AGN. LMXBs consist of a compact object, either a black hole or neutron star, in a binary system with a companion star. Matter is accreted from the companion star via an accretion disk onto the black hole. There is a proposed coupling mechanism in AGN and LMXBs whereby the X-ray emission produced in the accretion disk, tracks the inflow of material onto the black hole (Done et al., 2007); and the radio emission tracks the outflow of material via the jet (Fender, 2006). If the central accretion and ejection processes for AGN and LMXBs are the same, then studying the radio and X-ray variability gives two different windows to probe the coupling mechanism (see Markoff 2010 for a review). The timescales for LMXBs to evolve is very fast when compared with AGN. Complete cycles in the behaviour can be tracked, from the stable, low flux, ‘quiescent’ states (Calvelo et al.,

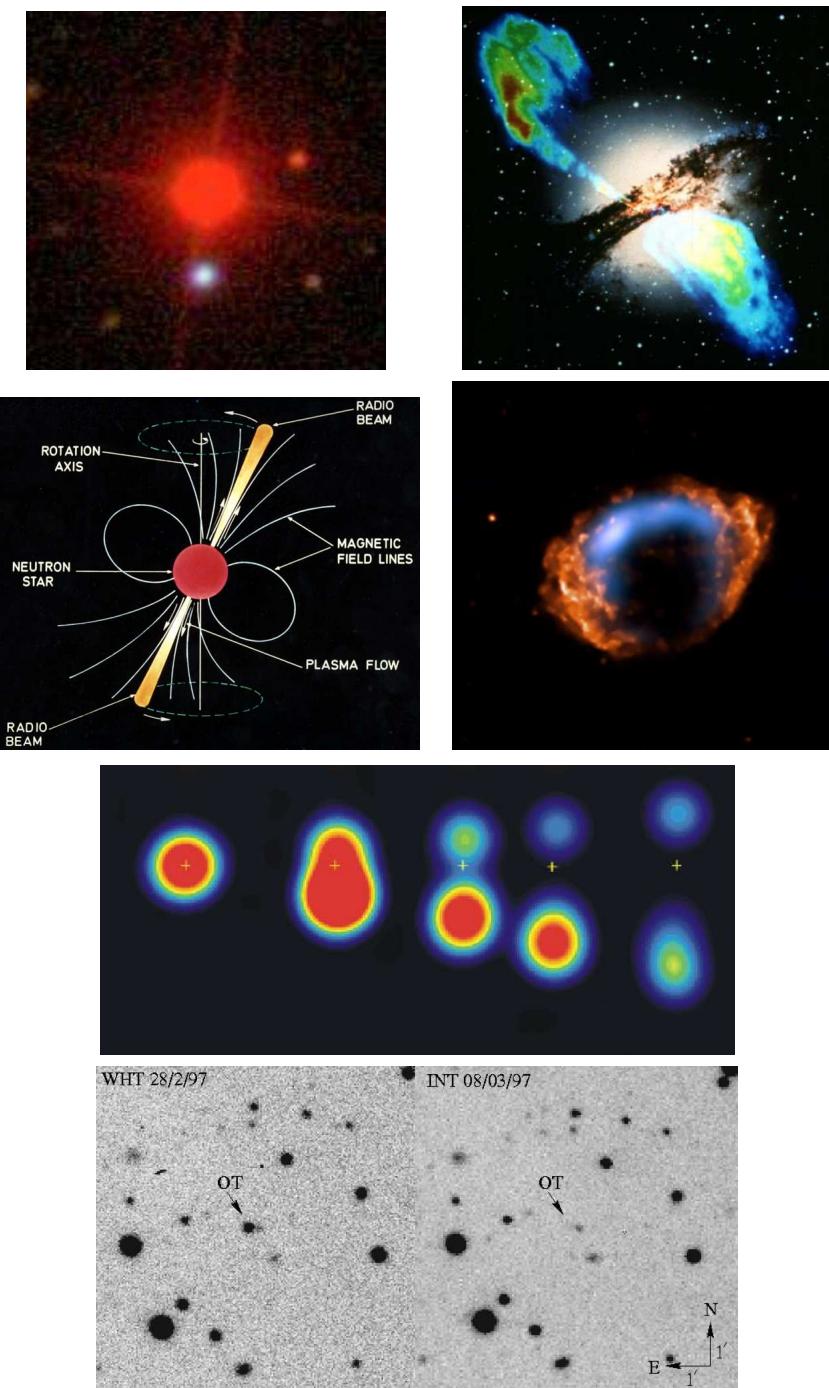


Figure 2.5: The ‘transient zoo’ - examples of transient radio sources. Top left: SDSS image of the flare star GJ 3630 (credit: Galaxy Zoo). Top right: Composite Optical (Hubble) and radio (VLA) image of Centaurus A, showing large scale collimated outflows (blue) from the nucleus (credit: NASA and NRAO). Second row left: Artists impression of a Pulsar (credit: ATNF). Second row right: A combined image of a radio supernovae from Chandra X-ray data (orange) and radio data (blue) from NRAO’s Very Large Array. (X-ray (NASA/CXC/NCSU) 2007; Radio (NSF/NRAO/VLA/Cambridge) 1985). Third row: A radio image of the X-ray binary GRS 1915+105, showing relativistic knots moving away from the core (Mirabel & Rodríguez, 1998). Bottom row: A detection of a γ -ray burst optical afterglow (Groot et al., IAUC 6584).

2010), through to massive accretion and ejection events (high state), which give rise to transient outbursts at both X-ray and radio frequencies (see Figure 2.5 and Fender et al. 1999; Mirabel & Rodríguez 1998). These state changes can last from days to months. In comparison, for AGN these transitions could take many thousands of years (McHardy et al., 2006), therefore what we typically observe in AGN is a constant state. By regularly monitoring a cohort of AGN in their respective states we can probe - what would be considered in LMXB timescales - the fine timescale structure of the disk jet coupling.

LOFAR provides a new frequency window to observe AGN and LMXBs. Potentially through long term, wide-field monitoring of the low frequency sky, coupled with targeted observations of interesting AGN and LMXBs, a deeper understanding of the accretion and ejection mechanisms can be achieved. In Chapter 3 I will present a deeper exploration of the variable radio and X-ray properties of a specific AGN - NGC 7213, and how the behaviour relates to that seen in LMXBs (this will also include a deeper literature review). I will also explore in this chapter the scaling relationship between black hole mass and radio and X-ray luminosity, known as the fundamental plane of black hole activity (see Merloni, Heinz & di Matteo 2003 and Falcke, Körding & Markoff 2004)).

2.2.7 Gravitational wave events

Gravitational waves are so far the undetected legacies of Einsteins general theory of relativity. These gravitational waves are predicted for sources such as binary mergers of neutron stars or black holes (Hansen & Lyutikov, 2001), Type II core collapse supernovae, and also rapidly rotating asymmetric neutron stars (Waldman, 2011). As the gravitational wave sources are predicted to be associated with extreme energy events, finding an electromagnetic counterpart (EM) is a high priority. The Lorimer et al. (2007) burst is potentially one of these EM counterparts, however, until the EM event is correlated with a gravitational wave event, the predictions can not be confirmed. Gravitational wave observatories such as the Laser Interferometer Gravitational-Wave Observatory (LIGO; Waldman 2011) are currently searching for gravitational wave signatures. LOFAR has already accepted triggers from the LIGO to search for electromagnetic counterparts to gravitational wave sources; co-ordinated test observations have also been performed (Abadie et al. in prep). In future operations more gravitational wave triggers will be followed up by LOFAR, and also LOFAR will provide a data stream of EM transient events to LIGO.

2.2.8 Future prospects

The types of *known* transients described above shows the diversity of the transient radio universe. These examples are by no means a definitive list and I have only summarised the typical properties of the sources. Most of the known transient sources are quite common, and typically have multi-wavelength counterparts associated with the radio emission. Potentially small populations of extremely bright or exotic transients exist which are not coupled with multi-wavelength counterparts. Performing blind searches of the radio sky, especially at unexplored frequency windows, such as those at which LOFAR will operate, could yield new classes of objects and physical processes.

Furthermore, there are broader implications to just studying the intrinsic properties of individual transient sources. For example, if a Pulsar was detected that was inspiralling into the gravitational well of a black hole, the regular clock like pulses, coupled with the inward motion would provide an excellent test of General Relativity. Transient studies also have implications for cosmology; γ -ray bursts are some of the most distant and luminous events in the Universe, and can be used as standard candles to probe cosmological distance scales. Also, the propagation of the lower frequency GRB afterglow emission is a useful probe of the electron content in the inter-galactic medium. The event rates (and population studies) of γ -ray bursts and also supernovae (at cosmological distances) give us a further handle on the star formation rate in the early Universe.

2.3 Review of transient surveys to date.

Despite the scientific potential, the transient and time variable radio sky is a relatively unexplored region of parameter space. Radio transient detections have historically been relatively sparse due to the poor survey speed of older radio facilities (Cordes, Lazio, & McLaughlin, 2004; Hessels et al., 2009). Some detections of transients have however been made serendipitously (for examples see Davies et al. 1976, Zhao et al. 1992, van den Oord & de Bruyn 1994, Bower et al. 2003 and Lenc et al. 2008). This limitation will soon be relieved by the next generation of wide-field telescopes and their respective dedicated transient surveys.

A common method to detect radio transients is through multi-wavelength triggered observations from, for example, all sky monitors on X-ray observatories. These have produced radio counterparts to gamma-ray burst (GRB) afterglows and black hole X-ray binary outbursts (for examples see Frail et al. 1997, Gaensler et al.

2005, Eck, Cowan, & Branch 2002). This method relies on having a detectable high frequency counterpart, which may be absent (or difficult to detect) for sources such as X-ray dim isolated neutron stars (XDINs; Ofek et al. 2009) and orphan gamma-ray burst afterglows (Frail et al., 1997), which demonstrates the need for dedicated radio transient programs.

Despite the historical challenges, dedicated or commensal transients surveys have produced a number of interesting results. The Galactic centre (GC) has been the area for some intense observing campaigns; these studies have so far detected a number of radio transients, the most recent being GCRT J1742-3001 (Hyman et al., 2009) and GCRT J1746-2757 (Hyman et al., 2002) – also see Bower et al. (2005). In the high time resolution domain, McLaughlin et al. (2006) discovered short duration transient radio bursts from neutron stars, i.e. RRATs (Rapidly Rotating Transients) – as well as many detections of new pulsars (for example, from the Parkes multi-beam survey).

Nine bursts – named the WJN transients – in excess of 1 Jy have been discovered using drift scan observations with the Waseda Nasu Pulsar Observatory at 1.4 GHz (summarised in Matsumura et al. 2009, but also see Kuniyoshi et al. 2007, Niinuma et al. 2007, Kida et al. 2008, Niinuma et al. 2009 for further details). These are some of the brightest transients reported in the literature and so far remain unexplained. Recently Croft et al. (2010) published results from the ATA Twenty Centimetre Survey (ATATS): no transients were detected and an upper limit on the snapshot rate of events was given. Subsequently the Pi GHz Sky Survey (PiGSS) surveyed the sky with the ATA at 3.1 GHz, providing the deepest static source catalogue to date above 1.4 GHz (Bower et al., 2010). No transients were reported in this survey and an upper limit on snapshot rate was placed (also see Bower et al. 2011).

Searching for highly variable *known* radio sources can also be a useful diagnostic in examining the dynamic radio sky. For example, Carilli et al. (2003) found a number of highly variable ($\Delta S \geq \pm 50\%$) radio sources in a small number of repeated observations of the Lockman Hole at 1.4 GHz. Frail et al. (2003) also found four highly variable radio transient sources from follow-up observations of GRBs at 5 and 8.5 GHz: the rates of these events are consistent with those reported in Carilli et al. (2003). Reporting 39 variable radio sources, Becker et al. (2010) recently characterised the surface density of variables in the direction of the Galactic plane at 4.8 GHz. Most of the variable sources detected had no known multi-wavelength counterparts. This is an important study as the rates of transient and variable sources may differ in the direction of Galactic plane when compared with an extra-Galactic pointing. In particular, large numbers of flare stars are known to produce bright co-

herent bursts (White et al., 1989), and could dominate detections at low frequencies (Bastian et al., 1988). A deeper discussion on the difference between transient and variable processes will follow later.

Radio telescope archives potentially contain many hours of data which currently remains unsearched for radio transients. An archival study comparing the NVSS (NRAO VLA Sky Survey; Condon et al. 1998) and FIRST (Faint Images of the Radio Sky at Twenty-cm; Becker et al. 1995; White et al. 1997) catalogues was conducted by Levinson et al. (2002), with a follow-up study by Gal-Yam et al.(2006): a number of radio transient sources were identified. Thyagarajan et al. (2011) have performed the most thorough analysis of the FIRST survey to date, reporting 1627 transient and variable sources down to mJy levels, over timescales of minutes to years (also see Ofek & Frail 2011). Bower et al. (2007) analysed 944 epochs of archival VLA data at 4.8 and 8.4 GHz spanning a period of 22 years. In this survey ten radio transients were reported, with the host galaxies possibly identified for four out of the ten sources, and the hosts and progenitors of the other six unknown. Bannister et al. (2011) recently published results from a search for transient and variable sources in the Molonglo Observatory Synthesis Telescope (MOST) archive at 843 MHz. 15 transient and 53 highly variable sources were detected over a 22 year period. Bannister et al. (2011) use these detections to place limits on the rates of transient and variable sources. Bower & Saul (2010) have published further archival work examining observations of the calibrator 3C286 at 1.4 GHz. A total of 1852 epochs are examined spanning a time range 23 years: no radio transients were reported.

2.4 This thesis

- In Chapter 3 of this thesis I will present a variability case study of the low luminosity active galactic nuclei NGC 7213.
- In Chapter 4 of this thesis I will present the design, development and testing of the LOFAR transient detection pipeline. I will also present the anticipated broader functionality of the algorithms needed for full LOFAR operations.
- In Chapter 5 of this thesis I will present the testing of the LOFAR transient detection algorithms on archival VLA data. This will include, reproducing the results of the transient detections reported in Bower et al. (2007); the detection of an outburst of the black hole candidate source Swift J1753.3-

0127; and finally a blind survey of archival VLA calibrator fields to search for transient sources.

- In Chapter 6 of this thesis I will address the task of using the LOFAR standard imaging pipeline to produce LOFAR images. In this Chapter I will also present the findings from a transient search conducted on one of the LOFAR zenith monitoring fields.

Beauty is variable, ugliness is constant.

DOUGLAS HORTON (1891 - 1968)

3

X-ray and radio variability in the low-luminosity active galactic nucleus NGC 7213

In this chapter I present the results of a ~ 3 year campaign to monitor the low luminosity active galactic nucleus (LLAGN) NGC 7213 in the radio (4.8 and 8.4 GHz) and X-ray bands (2-10 keV). AGN and variability science is a key component of the broader ‘transient science’ definition. It is a goal for instruments such as LOFAR to survey for unknown transient sources, however, they will also regularly monitor known variable sources – such as NGC 7213. In this chapter I use cross-correlation functions to probe the connection between the X-ray and radio emitting regions. I also explore the previously established scaling relationship between core radio and X-ray luminosities and black hole mass $L_R \propto M^{0.6-0.8} L_X^{0.6}$, known as the ‘fundamental plane of black hole activity’. With a large number of quasi-simultaneous radio and X-ray observations, I explore for the first time the variations of a single AGN with respect to the fundamental plane. The reduction of the ATCA radio data presented in this chapter, and all subsequent analysis was performed by myself. The RXTE X-ray data were reduced by Dr. Phil Uttley. This work has been published in Bell et al. (2010).

3.1 Introduction

The observable properties of active galactic nuclei (AGN) and black hole X-ray binaries (BHXRBS) are consequences of accretion on to a black hole at a variety of rates, in a variety of ‘states’, and within a variety of environments. The major difference between the aforementioned classes of object is the black hole mass. BHXRBS typically have a black hole mass $\sim 10M_{\odot}$ while for AGN it is $10^5M_{\odot} \leq M \leq 10^{10}M_{\odot}$. Theoretically, the central accretion processes should be relatively straightforward to scale with mass, and this is supported by several observed correlations. These include a relation between the X-ray and radio luminosities and the black hole mass (Merloni, Heinz & di Matteo 2003; Falcke, Körding & Markoff 2004), and between X-ray variability timescales, mass accretion rate and mass (McHardy et al. 2006). More qualitative similarities between accretion ‘states’ and radio jet production have also been demonstrated (Körding, Jester & Fender 2006; for the current picture of accretion states in BHXRBS and their relation to radio jets see Fender, Belloni & Gallo 2004).

Studying the delays between different emission regions gives us a further handle on the scalability of black hole accretion, as signals propagate from, for example, the accretion flow to the jet. Variability studies have so far shown that a correlation exists between the X-ray and optical emitting regions of both BHXRBS and AGN, typically reporting small lags, which are consistent with at least some of the optical variations being due to X-ray heating of the disc (Russell et al., 2009; Breedt et al., 2009). A recent study by Casella et al. (2010) has shown that a correlated time lag of ~ 100 ms exists between the X-ray and IR regions (IR lagging X-rays) for the BHXRBS GX339-4, indicating a close coupling between the hot accretion flow and inner regions of the jet. In the case of the BHXRBS GRS 1915+105 a variable X-ray to radio lag of ~ 30 mins (radio lagging X-ray) has been measured (Pooley & Fender, 1997; Fender et al., 1999). Discrete ejection events have also been resolved in both the AGN and BHXRBS, for examples see Marscher et al. (2002) and Pooley & Fender (1997); Fender et al. (1999).

The linear scaling with mass of the characteristic timescale around a black hole means that there are advantages to studying each class of object. In BHXRBS we can track complete outburst cycles, from the onset of disc instabilities through major ejection events, radio-quiet disc-dominated states, and a return to quiescence, on humanly-observable timescales (typically years). For a typical AGN the equivalent cycle may take many millions of years. However, for an AGN we are able to resolve individual variations of the source on time-scales that are comparable to or shorter

than the shortest physical time-scales in the system (e.g. the dynamical time-scale), something which is currently impossible for BHXRBS. In ‘black hole time’ we are able to observe the evolution of sources in fast-forward for BHXRBS and in detailed slow-motion for AGN.

In this chapter I present the results of a long term (~ 3 years) regular monitoring campaign in the X-ray and radio bands of the low luminosity active galactic nucleus (LLAGN) NGC 7213. Previous X-ray studies show that NGC 7213 is accreting at a low rate $\sim 7 \times 10^{-4} L_{Edd}$ (Starling et al., 2005). The hard state in BHXRBS is typically observed at bolometric luminosities below $\sim 1\%$ Eddington, and seems to be ubiquitously associated with a quasi-steady jet. Above $\sim 1\%$, sources can switch to a softer X-ray state, the jets are suppressed (Maccarone, Gallo, & Fender, 2003; Dunn et al., 2009); furthermore transition to this softer state is usually associated with major transient ejection events. As NGC 7213 is considerably below $L_{Edd} \sim 1\%$ I therefore consider it a good candidate for comparison with other BHXRBS in the low/hard state. If I consider AGN to be ‘scaled up’ versions of BHXRBS by exploring the time lag between the X-ray and radio emitting regions I can compare, contrast and hopefully relate the accretion and jet production scenarios for AGN and BHXRBS.

A correlation has been established by Corbel et al. (2003) and Gallo et al. (2003, 2006) relating the radio luminosity (L_R) and X-ray luminosity (L_X) for BHXRBS in the low/hard and quiescent states, where $L_R \propto L_X^{0.6-0.7}$. Merloni, Heinz, & di Matteo (2003) - hereafter MHdM03 and Falcke, Körding, & Markoff (2004) extended the BHXRBS relationship using two samples of AGN to form the ‘fundamental plane of black hole activity’. By accounting for the black hole mass (M) the relationship $L_R \propto L_X^{0.6-0.7}$ has been extended to cover many orders of magnitude in black hole mass and luminosity i.e $L_R = (0.6_{-0.11}^{+0.11}) \log L_X + (0.78_{-0.09}^{+0.11}) \log M + 7.33_{-4.07}^{+4.05}$. Further refinements were made to the fundamental plane by Körding, Falcke, & Corbel (2006) - hereafter KFC06, using an augmented and updated sample to examine the fitting parameters (also see Gültekin et al. 2009 and Broderick & Fender 2011).

Throughout this chapter I define the ‘intrinsic’ behaviour of AGN and BHXRBS as multiple measurements (in the radio and X-ray) of the *same* source. I define the ‘global’ behaviour as single (or average) measurements of *multiple* sources, both with respect to the fundamental plane.

For the BHXRBS in the low/hard state the relationship described above has not only been established globally but in some cases intrinsically, i.e. GX 339-4, V404 Cyg and a small number of other systems have been shown to move up and down the

correlation seen in the fundamental plane (Fender et al., 2007). However, in recent years an increasing number of outliers have been found below the correlation, i.e. less radio-loud than expected (Xue & Cui 2007; Gallo 2007; Soleri et al. 2010; Calvelo et al. 2010; Rushton et al. 2010) as well as some sources which move in the plane with a different slope (e.g. Jonker et al. 2010 and Coriat et al. 2011). To date the correlation found from the fundamental plane has only been measured globally for AGN, not intrinsically. Note, with respect to the global measurements of the AGN population, the specific measurements of the radio and X-ray flux used in the correlation are sometimes taken at different times and thus could be a source of error in the correlation (Körding, Falcke, & Corbel, 2006).

As well as establishing the time lags, another goal of this work, is to establish, through quasi-simultaneous observations the intrinsic relationship between L_R , L_X and M observed in a LLAGN and its relevance to the fundamental plane of black hole activity. I use the MHdM03 and KFC06 samples for comparison both with an updated BHXR sample taken from Fender, Gallo & Russell (2010) - hereafter FGR10. I also explore the possible scatter in AGN data points away from the fundamental plane and place limits on this deviation.

3.2 NGC 7213 Background

NGC 7213 is a face-on Sa galaxy hosting a Seyfert 1.5 nucleus located at a distance of 25 Mpc assuming $H_0=71\text{km}\text{s}^{-1}\text{Mpc}^{-1}$. Narrow low ionisation emission lines are observable in its nuclear spectrum, also making it a member of the LINER class (Hameed et al., 2001).

The radio properties of NGC 7213 are intermediate between those of radio-loud and radio-quiet AGN. Previous radio studies at 8.4 GHz have not resolved any jet emission from the nucleus on scales 3 mas to 1 arcsec (Bransford et al. (1998), Thean et al. (2001), Blank, Harnett, & Jones (2005)). The Long Baseline Array (LBA) is an Australian six station VLBI instrument (Tzioumis, 1997); LBA observations at 1.4 and 0.843 GHz have shown some evidence for larger scale emission that could be due to jet-fed radio lobes (Blank, Harnett, & Jones, 2005). However, as the nucleus remains unresolved it could be that the jet is oriented to some degree in the direction of the observer. As suggested by Blank, Harnett, & Jones (2005), the lower frequency emission could possibly be associated with a ‘kink’ in the jet at a larger distance. This could be consistent with a general model proposed by Falcke, Patnaik, & Sherwood (1996) where radio-intermediate objects are in fact

radio - quiet objects whose jets are to some extent aligned in the observers direction and relativistically boosted.

Previous X-ray studies have so far failed to show significant evidence for a soft X-ray excess, Compton reflection, or a broad Fe K α line in the NGC 7213 spectrum e.g. see Bianchi et al. (2003), Bianchi et al. (2008), Starling et al. (2005). A narrow Fe K α line is observed and Bianchi et al. (2008) show evidence that it is produced in the broad line region (BLR). Starling et al. (2005) also report that the expected UV bump is either absent or extremely weak. The weakness or absence of these signatures suggests that the inner accretion disc is missing, perhaps replaced by an advection dominated accretion flow (ADAF) (as suggested by Starling et al. 2005) or similar hot flow, consistent with the low/hard state interpretation of NGC 7213.

Blank, Harnett, & Jones (2005) estimate the black hole mass of NGC 7213 using the nuclear velocity dispersion calculation of Nelson & Whittle (1995) and the velocity-dispersion versus black-hole-mass relationship of Ferrarese & Merritt (2000) to obtain $9.6^{+6.1}_{-4.1} \times 10^7 M_\odot$.

3.3 Observations

3.3.1 ATCA data analysis

Bi-weekly monitoring was obtained with the Australia Telescope Compact Array (ATCA). The interferometer setup was such that 128 channels of 1 MHz bandwidth were used to form two continuum channels centred at 4.8 and 8.4 GHz respectively. The radio observations have been reduced using the MIRIAD radio reduction package (Sault, Teuben, & Wright, 1995). Flux and bandpass calibration was achieved using (in most cases) PKS J1939-6342 (B1934-638) and for the phase calibration PKS J2218-5038 (B2215-508).

A variety of fitting techniques were then tested to extract the flux density of the source from the observations. These included testing point and Gaussian fitting to the source in the image and *uv* plane. As ATCA is an East-West array and observations were short in duration (\sim 2 hours) an elongated synthesised beam was typically produced. Image-plane fitting often failed to converge and/or produced non-physical results. Therefore fitting was discarded in the image plane for the more (arguably) reliable *uv* plane method. The final radio light curves and spectral indexes (between the radio bands only) are shown in Figure 3.1.

To complement the snapshot observations a full 12 hour integration was performed on 2008 June 29; the purpose of which was to explore the polarisation prop-

erties of the jet, and to detect any faint extended structure. Additional steps were taken to ensure adequate polarisation calibration was performed. This observation was also used to make a deep image of the source; beyond the 0.5 arcsec beam no weak extended radio emission was observed above a noise level $\sigma=78 \mu\text{Jy}$ at 4.8 GHz.

3.3.2 RXTE data analysis

Daily monitoring was obtained using 1 ksec snapshots obtained with the Rossi X-ray Timing Explorer (RXTE) Proportional Counter Array (PCA), allowing a long-term light curve to be obtained in the 2-10 keV band. The data were reduced using FTOOLS v6.8, using standard extraction methods and acceptance criteria. The background was calculated from the most recent background models which corrects for the recent problems with the *RXTE* SAA (South Atlantic Anomaly) history file. The final 2-10 keV fluxes were calculated by fitting a power law to the observed spectra. This allows us to take into account changes in the instrumental gain over the duration of the monitoring. The final un-binned X-ray light curve is shown in Figure 3.1 and is used in all subsequent analysis.

3.4 Results

3.4.1 X-ray/radio light curves and cross-correlation

Figure 3.1 shows X-ray flux (top panel), radio flux (middle panels) and radio spectral index α where $S_\nu \propto \nu^{+\alpha}$ (bottom panel) versus time. The X-ray light curve shows a general decrease in flux until MJD 54600. Two distinct X-ray flares are observed at MJD 53920 and MJD 54390, both appear to be correlated with events in the radio bands.

I use the Discrete Correlation Function (DCF) method of Edelson & Krolik (1988) to calculate the cross-correlation coefficients between the entire X-ray and radio bands to find the lag. To calculate the centroid lag τ_{cent} , I use a weighted mean of the positive lags above an 85% threshold of the peak DCF value. I use τ_{cent} , rather than the peak value τ_{peak} , because the centroid has been shown to better represent the physical lag (Koratkar & Gaskell, 1991). It has been shown that it is difficult to interpret τ_{peak} as a physical quantity, and if it is used to calculate the lag, it usually offers an underestimate when compared with τ_{cent} (Peterson et al., 1998). The centroid width is calculated as

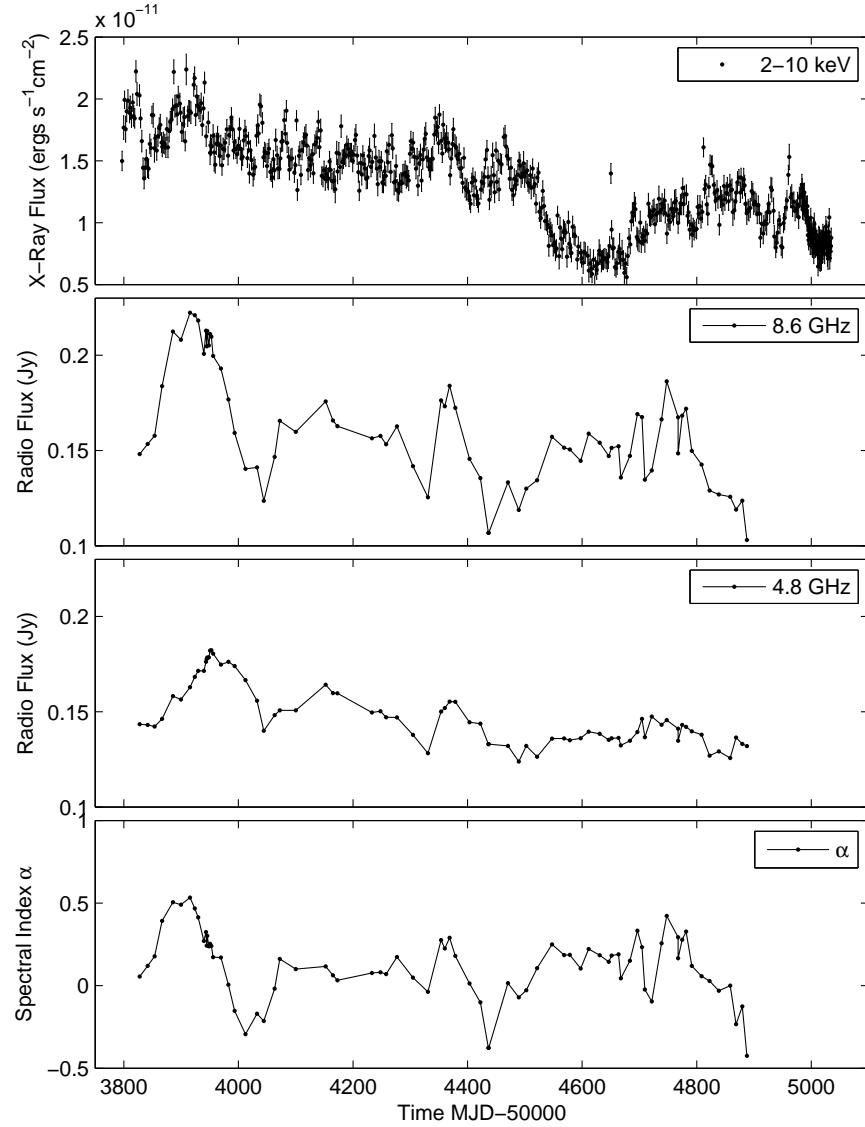


Figure 3.1: Top Panel: X-ray Flux versus Time (MJD-50000). Centre Top: Radio flux $S_{v8.4GHz}$ versus time. Centre Bottom: Radio flux $S_{v4.8GHz}$ versus time. The errors are too small to show on both radio flux figures but are typically ~ 0.4 mJy/Beam. Bottom: Spectral index (8.4-4.8 GHz) α versus time where $S_v \propto v^{+\alpha}$.

Table 3.1: Summary of the lag times which have been calculated using the discrete correlation function for the X-ray-8.4 GHz and X-ray-4.8 GHz light curves. I used the z-transformed discrete correlation function to calculate the lag for the 8.4 GHz-4.8 GHz and flare 1 *only* light curves.

Time (MJD)	X-ray-8.4 GHz Lag (days)	X-ray-4.8 GHz Lag (days)	8.4 GHz-4.8 GHz Lag (days)
All	24 ± 12 ($DCF_{max}=0.56$)	40 ± 13 ($DCF_{max}=0.7$)	20.5 ± 12.9 ($ZDCF_{max}=0.81$)
Flare 1 (53800-54000)	14 ± 11 ($DCF_{max}=0.78$)	35 ± 16 ($DCF_{max}=0.73$)	-

$$\tau_{cent} = \frac{1}{N} \sum_i \frac{c_i \tau_i}{c_i} \quad (3.1)$$

Where c_i and τ_i are the DCF coefficients and lags, τ_{cent} is the centroid lag. As the DCF coefficients are only defined for a given lag, I interpolate to find the peak correlation coefficient DCF_{max} at the centroid lag .

Figure 3.2 show the DCFs for the X-ray-8.4 GHz and X-ray-4.8 GHz bands respectively. I find a lag of $\tau_{cent} = 24 \pm 12$ days at $DCF_{max} = 0.56$ for X-ray-8.4 GHz and $\tau_{cent} = 40 \pm 13$ days at $DCF_{max} = 0.7$ for X-ray-4.8 GHz. The time lags taken from Figure 3.2 are summarised in Table 3.1. The errors in the lag are calculated using the flux randomisation/random subset selection (FR/RSS) method of Peterson et al. (1998). For 10^4 samples, I randomly reduce the number of data points in our original X-ray and radio light curves by 37% and then re-calculate the DCF. I take the centroid lag from each of the 10^4 simulations and use the rms spread in the distribution of centroid lags to obtain the quoted errors. Note that since the radio light curves are strongly correlated with one another, the statistical errors on the lags are not independent between the two bands, so that the difference in lag between the two radio bands with respect to X-rays is likely to be real (this is also highlighted by the radio-radio correlation, see below).

To ascertain a confidence level in the DCF calculation I used the method of Timmer & Koenig (1995) to generate Monte-Carlo simulations of un-correlated red-noise light curves. I assumed a broken power-law-shape power spectrum with low-frequency slope -1, high-frequency slope of -2.3 and break frequency $\nu = 0.012 \text{ d}^{-1}$, based on the best-fitting model to fit the X-ray power spectrum (Summons, private communication). I used this model to generate 10^4 random X-ray light curves and cross-correlated them with the real radio light curves in both bands. I then used these correlations to estimate the distribution of cross-correlation values at *each* lag to obtain a local confidence level, corresponding to the CCF value which exceeds P per-cent of the simulated CCF values at that lag. From the simulations I plot on Figure 3.2 the 95% and 99% local confidence levels: I therefore assign a local confidence of $> 95\%$ for both of the peaks with the 4.8 GHz peak reaching close to 99%.

It is important to note however, that without a priori expectation of what the lag should be, it is more statistically rigorous to estimate the significance of the correlation using the ‘global significance’, which is the fraction of simulated random light curves which show CCFs at better than the observed confidence level for *any* given lag. For example, to estimate the global significance of a peak which appears at the

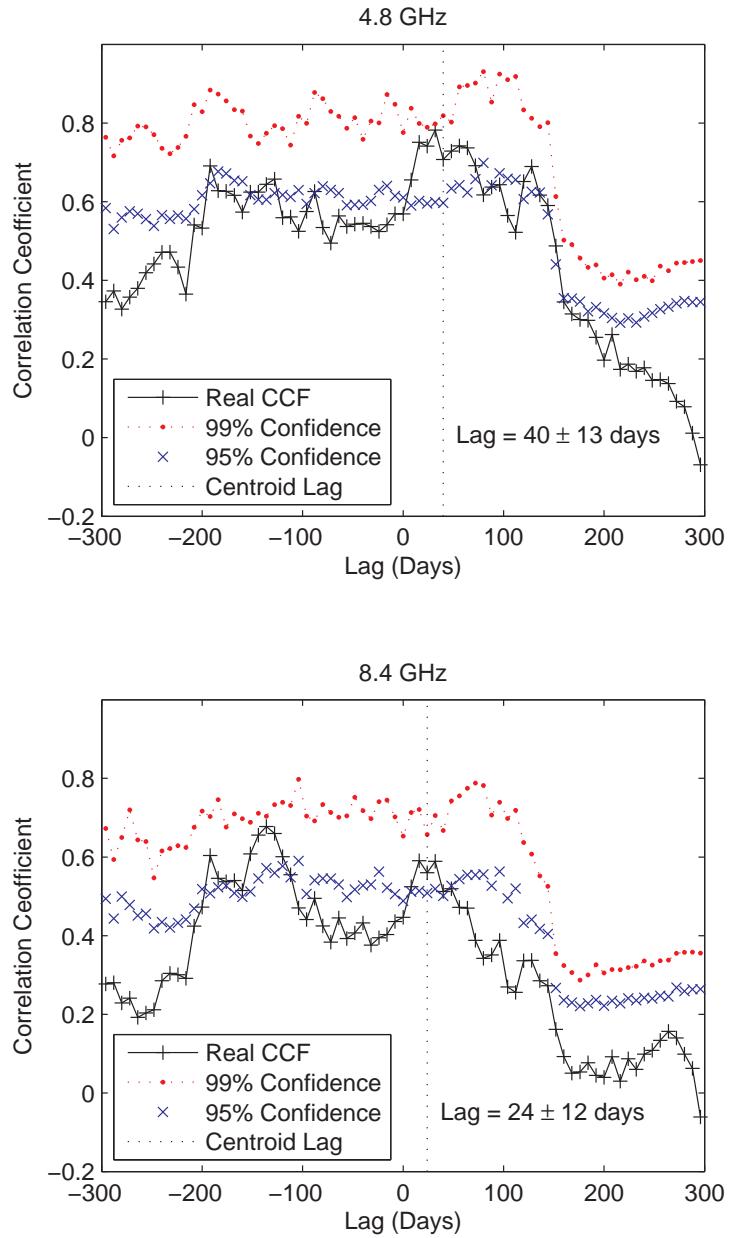


Figure 3.2: Top Panel: Discrete cross-correlation function (DCF) plot showing time lag against correlation coefficient for X-ray to radio (4.8 GHz lagging X-ray) (black crosses). A lag of 40 ± 13 days at $DCF_{max} = 0.7$ is calculated using the weighted mean of the lags above an 85% threshold of the peak DCF value. Bottom Panel: Discrete cross-correlation function (DCF) plot showing time lag against correlation coefficient for X-ray to radio (8.4 GHz lagging X-ray) (black crosses). A lag of 24 ± 12 days at $DCF_{max} = 0.56$ is calculated using the weighted mean of the lags above an 85% threshold of the peak DCF value. On both plots the 99% and 95% local significance confidence levels are plotted. The errors are calculated via the method described in section 3.4.1

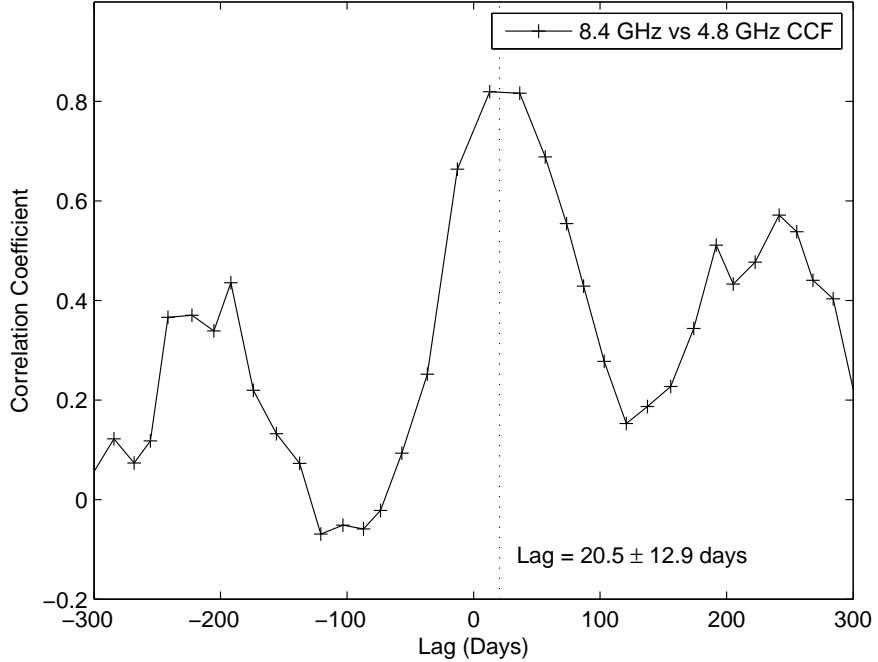


Figure 3.3: The z-transformed discrete cross-correlation (ZDCF) function between the 8.4 GHz and 4.8 GHz light curves (4.8 GHz lagging 8.4 GHz). A centroid peak of 20.5 ± 12.9 days is found at $ZDCF_{max}=0.81$.

99% local confidence level, I search the 10^4 simulated CCFs for any peak values (at any lag) which exceed the 99% local confidence level for that lag. The fraction of CCFs which do not show such a peak defines the global significance level for the correlation observed at the given local confidence level. In this way, I account for the fact that we are effectively searching over a broad range in lags, and so are sampling from a larger population of potential ‘false positives’ than we would expect from searching for correlations at a single lag.

I find for the entire range of computed lags (-300 to +300 days) the global confidence levels at 4.8 GHz are 82.0% (at 99% local significance) and 56.2% (at 95% local significance); at 8.4 GHz, 77.1% (at 99% local significance) and 46.0% (at 95% local significance). If I restrict the range of lags to positive only (> 0 days) I find at 4.8 GHz, 91.0% (at 99% local significance) and 74.7% (at 95% local significance); at 8.4 GHz, 86.6% (at 99% local significance) and 63.3% (at 95% local significance).

Note, that in the case where I do not restrict the lag to lie within any given range, the global significance is low. For example, it is not surprising that the 8.4 GHz versus X-ray CCF shows two peaks at greater than 95% local confidence, since the corresponding global confidence at this level is only 56%, so I expect a spurious peak in nearly half of the CCFs sampled from random light curves. However the

peak corresponding to a negative lag of ~ 140 days (i.e. radio leading X-ray) cannot easily be explained with any physical model. When I presume that there must be a positive lag (radio lagging X-ray) - as I would indeed predict based on our current understanding of the accretion disk/jet - this global significance level increases. Put another way, it is statistically unlikely that a spurious correlation should appear in the range of lags which matches our physical expectations, as seems to be the case here. Therefore, I assume for the remainder of the chapter that the correlation is real. Note that if the correlation is real, the lags can be well-determined: the significance of a lag and the significance of a correlation itself are unrelated quantities, since a lag can be very well determined from just a single well-sampled ‘event’ (flare or dip) in red noise data, if the data are indeed correlated. I will explore possible complexities in the disk/jet connection with respect to our calculated lags in section 3.5.1

3.4.2 Radio/radio cross-correlation

To calculate the cross-correlation between the two radio bands I used the z-transformed discrete correlation function (ZDCF) of Alexander (1997), see figure 3.3. Note that the two radio light curves are sampled at the same time: as noted in Edelson & Krolik (1988); to keep the normalisation correct for the DCF I need to omit the zero lag pairs as the errors between the two bands are correlated. Therefore for convenience (and as the two methods are comparable) I switch to the ZDCF which omits zero lagged pairs. I calculate the errors using the same method as described in section 3.4.1. I thus find $\tau_{cent} = 20.5 \pm 12.9$ days which is consistent with the difference in lags between X-ray-4.8 GHz and X-ray-8.4 GHz.

3.4.3 X-ray/radio cross-correlation when flaring

I split the light curve up into data points surrounding the first flare at MJD 53920, which is summarised in Table 3.1. I then calculated the ZDCF to ascertain an accurate measurement of the lag for this flare *only*. I find $\tau_{cent} = 14 \pm 11$ days for X-ray-8.4 GHz, and $\tau_{cent} = 35 \pm 16$ days for X-ray-4.8 GHz. Both lag times are below the global lag for the entire light curve. I do not perform this analysis for the second flare at MJD 54380 as it was only sparsely sampled in the radio.

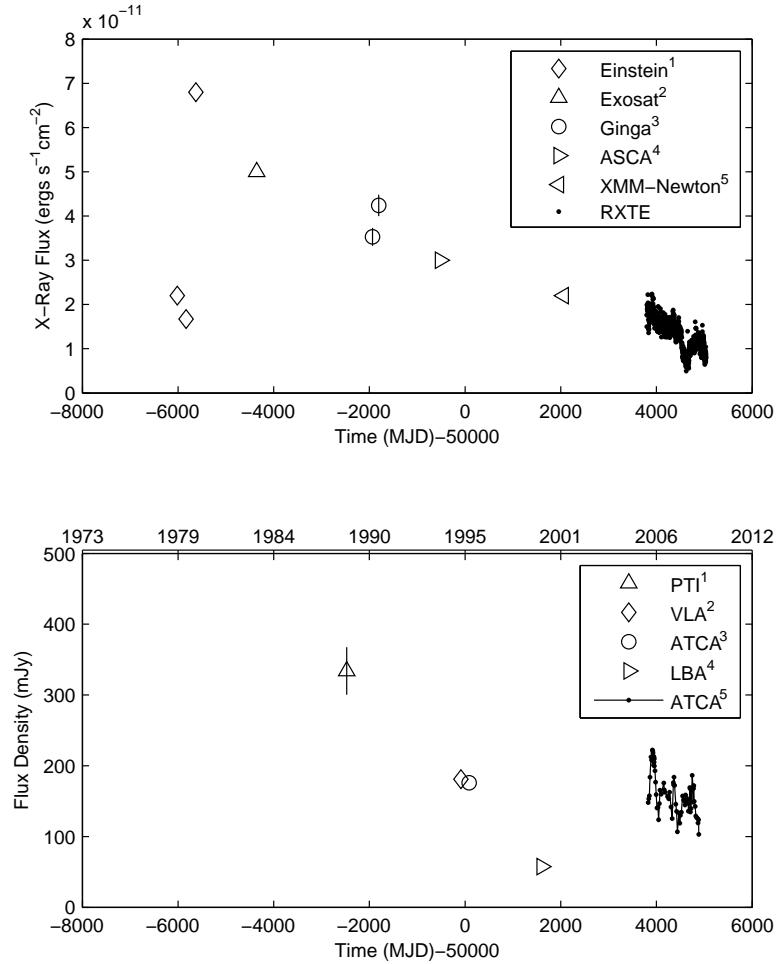


Figure 3.4: Top Panel: Light curve of archival X-ray data points in the 2-10 keV range. Flux values cited in $\diamond 1$ =(Halpern & Filippenko, 1984), $\triangle 2$ =(Turner & Pounds, 1989), $\circ 3$ =(Nandra & Pounds, 1994), $\triangleright 4$ =ASCA using the tarturus database (Turner et al., 2001), $\triangleleft 5$ = (Bianchi et al., 2008). Bottom Panel: Light curve of archival radio data points at 8.4 GHz. Flux values cited in $\triangle 1$ = (Slee et al., 1994), $\diamond 2$ =(Thean et al., 2001), $\circ 3$ = (Bransford et al., 1998) and $\triangleright 4$ = (Blank, Harnett, & Jones, 2005)

3.4.4 Long term variability

The long term X-ray variability, covering a time-range of about 30 years, is plotted in the upper panel of figure 3.4. Archival data were searched for in the literature from the instruments *Einstein* (Halpern & Filippenko, 1984), EXOSAT (Turner & Pounds, 1989), *Ginga* (Nandra & Pounds, 1994), ASCA using the Tartarus database (Turner et al., 2001) and *XMM-Newton* (Bianchi et al., 2008). Around MJD 44400 (~ 1980) a sharp flare which is brighter than the flares visible in our dataset was observed. Since then the flux has been gradually decreasing.

Table 3.2: Summary of archival observations found in the literature. The data points between the horizontal lines indicate where the two fluxes were used to calculate the X-ray-radio correlation. Two X-ray and two radio points were taken at similar times thus they were both used to calculate the correlation, (R) and (X) denotes the time (MJD) of the observation for radio and X-ray respectively. δt gives the time difference between observations (Radio-X-ray). The instrument used to derive the flux is given in brackets (see section 4.4.4 for references) and errors are given only when reported in the literature.

Time (MJD)	Radio Flux S_v 8.4 GHz (mJy)	X-ray Flux 2-10 keV (ergs^{-1})	δt_{meas} (Radio - X-ray) (days)
43985	-	2.2×10^{-11} (Einstein)	-
44168	-	1.7×10^{-11} (Einstein)	-
44374	-	3.8×10^{-11} (Einstein)	-
45644	-	5.0×10^{-11} (Exosat)	-
47527 (R) - 48065 (X)	334 ± 33 (PTI)	$3.5 \times 10^{-11} \pm 0.2 \times 10^{-11}$ (Ginga)	-538
47527 (R) - 48193 (X)	334 ± 33 (PTI)	$4.2 \times 10^{-11} \pm 0.25 \times 10^{-11}$ (Ginga)	-666
49913 (R) - 49479 (X)	181 ± 1 (VLA)	3.0×10^{-11} (ASCA)	434
50083 (R) - 49479 (X)	176 ± 1 (ATCA)	3.0×10^{-11} (ASCA)	604
51604 (R) - 52057 (X)	57.6 ± 1.3 (LBA)	2.2×10^{-11} (XMM-Newton)	-453

The long term radio variability at 8.4 GHz is plotted in the lower panel of figure 3.4. The flux values were taken from a variety of instruments and publications (which are referenced on the plot). In order of increasing time the first point on the plot was taken from Slee et al. (1994) using the Parkes-Tidbindilla interferometer (PTI). Then near simultaneous data points were taken with the VLA (Thean et al., 2001) and then ATCA (Bransford et al., 1998). The flux from the long baseline array (LBA) observation by Blank, Harnett, & Jones (2005) was the last point plotted prior to the ATCA monitoring which is presented in this chapter.

When comparing the flux expected from the start of our ATCA monitoring with the LBA data point, the LBA point seems too low. The core is un-resolved with all of the radio telescopes described above. The PTI observations constitutes the first data point in figure 3.4 and the LBA the last (before ATCA monitoring). Note, the Parkes-Tidbindilla baseline is part of the LBA network. Blank, Harnett, & Jones (2005) comment that in the LBA observations, no decrease in flux was observed with respect to baseline length i.e. it was not resolved; they conclude from this that the decrease in flux with respect to the other observations is accurate.

However, both the PTI and the LBA are not sensitive to the same spatial resolutions as the VLA and ATCA (the largest ATCA baseline is 6km and the shortest LBA baseline used in the archival observation is 113km). It is possible that some flux (on ATCA equivalent baselines) is not accounted for: which could plausibly give a decrease in flux i.e. indicating brighter, larger scale structure. The logical consequence of this is that the PTI data point is also missing some flux as it samples on a singular baseline of 275km. I will discuss these measurements within the context of the fundamental plane and attempt to draw some conclusions in section 3.5.2

3.4.5 Linear polarisation

A full 12 hour synthesis ATCA observation was performed on 2008 June 29 at 8.4 and 4.8 GHz. The purpose of this observation was to perform a reliable polarisation calibration to ascertain an accurate measurement of the percentage linear polarised (LP) flux from the source. Bower, Falcke, & Mellon (2002) showed that from a sample of 11 LLAGN the mean LP was $\sim 0.2\%$. Bignall et al. (2004) from a survey of 22 blazars report a mean LP $\sim 3\%$ with a standard deviation $\sigma = 1.5$. At the time of observation the polarisation of NGC 7213 was $< 0.1\%$ at 4.8 and 8.4 GHz respectively, which is more consistent with the reported value of typical LLAGN and not blazars. The percentage polarisation was calculated at other epochs, however

without a full 12 hour synthesis parallactic angle calibration the values were often suspect and will not be presented here.

3.5 Discussion

3.5.1 X-ray/radio jet connection

I have shown that a weak but statistically significant delay exists between the X-ray and radio emitting regions, with the radio lagging behind the X-ray. A number of models have been proposed to explain and interpret the X-ray to radio lag in both BHXRBS and AGN. The most notable of these are the ‘internal shock model’ (Blandford & Konigl, 1979; Rees, 1978) and a ‘plasmon model’ (van der Laan, 1966). I will briefly explore these models and comment on the relevance - if any - to our data.

I first consider a plasmon model. After the accretion mechanism(s) have pushed matter into the jet, and somehow shock accelerated a non-thermal population of electrons, an adiabatically expanding initially self-absorbed synchrotron emitting plasmon travels at relativistic speeds down the path of the jet. At a given time and frequency the matter becomes optically thin and is ‘detectable’ at that given frequency (in this case either 8.4 and 4.8 GHz). If the jet is fed at a constant mass rate, density and velocity the delay time for material to become optically thin will be constant. If these parameters are not constant the delay between X-ray and radio will be variable (e.g. see van der Laan 1966) assuming that the disk and the jet are indeed coupled (Falcke et al., 2009). Note, the higher frequency radio emission (8.4 GHz) will become optically thin first, thus a delay from 8.4 to 4.8 GHz is always expected.

Another model for explaining the emission seen in jets is the ‘internal shock model’ (Blandford & Konigl, 1979; Rees, 1978). The synchrotron lifetime of an emitting region is too short to adequately explain the scales of jets observed in AGN. These emitting regions - commonly referred to as ‘knots’ - are often displaced from the central core emission. Localised shocks within these regions are needed to explain the time-scales of variation observed (Rees, 1978; Felten, 1967). Jet shock scenarios have also been used to model the common flat spectrum jet observed in BHXRBS and AGN (e.g. see Spada et al. 2001, Jamil, Fender, & Kaiser 2009).

Figure 3.1 shows the evolution of spectral index; indicating that during a flare event α increases ($\alpha > 0$), flattens and steepens ($\alpha < 0$) shortly afterwards. The initial re-energisation given by a shock would push/compress the plasma into the

optically thick regime; subsequently this process moves the material into the optically thin regime. The variable time lag seen in the DCF could be interpreted as the time taken for newly injected matter to ‘catch up’ with older matter expanding adiabatically in the jet to shock and produce a flare. The lag would be dependent on the Lorentz factor of the newly injected material which in turn is related to the accretion rate (Falcke et al., 2009).

Both of these models can adequately explain the variations observed in our light curves, although we cannot completely disentangle them. It is possible that both of these models are in some part responsible for the behaviour. As the jet remains unresolved I cannot identify an area of localised shocks or indeed discrete resolved events that I could use to support a particular model.

To explore - in a very simplistic sense - how the time lags seen in BHXRBS scale up to AGN I offer two comparisons; that of a simple mass scaling and also mass-Eddington ratio scaling. For example, discrete resolved ejection events have been seen in GRS 1915+105 with a variable time lag between X-ray and radio (GHz), with the clearest examples of events having lags of $T_{GRS1915} = 20\text{-}30$ mins (Pooley & Fender, 1997; Mirabel et al., 1998). Taking the characteristic timescale of X-ray-radio lags measured in GRS 1915+105 and scaling up to NGC 7213 with mass only (assuming the mass of GRS 1915+105 as $\sim 10M_{\odot}$ and the mass of NGC 7213 as $9.6 \times 10^7 M_{\odot}$ and using $T_{lag} = T_{GRS1915} \times (M_{NGC7213}/M_{GRS1915})$) I infer $T_{lag} = 2 \times 10^5$ days, much longer than I observe. GRS 1915+105 is however accreting at a much higher rate, therefore including a scaling by the ratio of the Eddington luminosities (using $T_{lag2} = T_{lag} \times (L_{Edd-NGC7213}/L_{Edd-GRS1915})$ and taking $L_{Edd-GRS1915} \sim 1$) I find $T_{lag2} = 140$ days, reasonably comparable to our measured lag.

The difficulty in comparing the actual X-ray-radio time lag measured for NGC 7213, and that of the GRS 1915+105 scaled time lag is that the radio emitting regions being probed are significantly different: in a broader view the spectral energy distributions are different. In fact until the structure of jets and how it scales with accretion rate and mass at a given frequency are better understood, such attempts at quantitative comparison are of marginal value, although the qualitative comparison is valuable.

3.5.2 The fundamental plane of black hole activity

The fundamental plane of black hole activity (Merloni, Heinz, & di Matteo, 2003; Falcke, Körding, & Markoff, 2004) shows that the correlation $L_R \propto M^{0.8} L_X^{0.6}$ holds over many orders of magnitude of X-ray and radio luminosities, and black hole

masses. Figure 3.5 shows; at high luminosities a sample of AGN taken from MHdM03, and at the low luminosities an updated BHXR sample taken from FGR10. The updated BHXR X-ray data were observed in the 0.5-10 keV energy range, therefore careful steps were taken (using WebPIMMS) to correctly convert into the 2-10 keV range to compare with the MHdM03 and NGC 7213 data points. I plot on Figure 3.5 the best fit parameters $L_R = (0.6^{+0.11}_{-0.11}) \log L_X + (0.78^{+0.11}_{-0.09}) \log M + 7.33^{+4.05}_{-4.07}$ as defined by MHdM03. Note, I do not re-fit the best fit line for the updated BHXR sample.

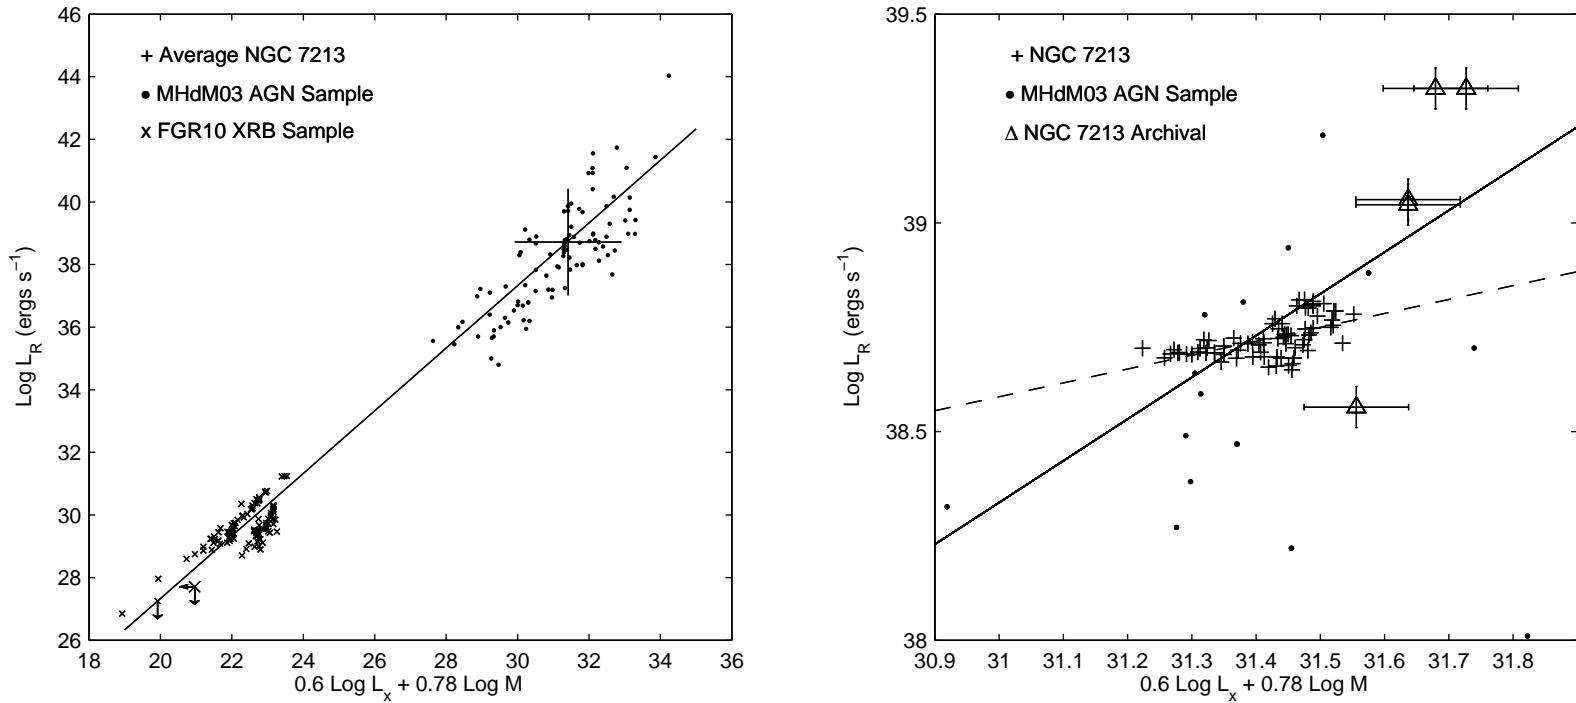


Figure 3.5: Left Panel: The fundamental plane of black hole activity. The large cross indicates the average location of NGC 7213. Black dots are the MHdM03 sample with the black line indicating the best fit. Crosses indicate the updated BHXR sample of FGR10; data points have been converted into the 2-10 keV energy range. Right Panel: A close up of the NGC 7213 data points with least squares best fit indicated with a dashed line. The MHdM03 fit is shown with a solid line. The archival data points are also shown but are not used in the NGC 7213 data best fit. NGC 7213 measurements are taken at 4.8 GHz and the archival are at 8.4 GHz.

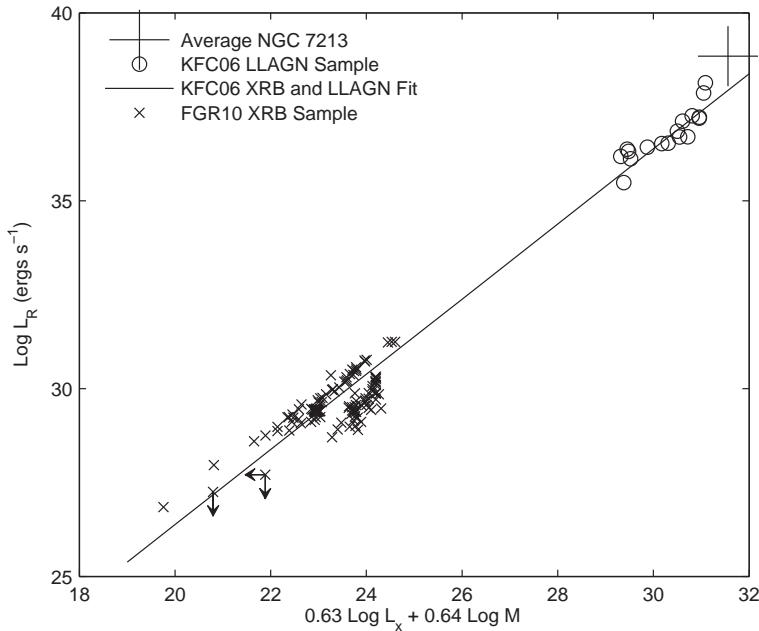


Figure 3.6: The fundamental plane using the KFC06 LLAGN sample only (circles) and updated FGR10 BHXR sample (crosses). The large cross indicates the average location of NGC 7213. The solid line indicates the KFC06 best fit.

I paired the X-ray and radio data points to calculate the correlation by finding the closest X-ray point (in time), to the radio, because the X-ray sampling was more frequent. The average for NGC 7213 sits well on the predicted correlation. The right hand panel of Figure 3.5 shows a close up of the NGC 7213 data. A least squared best fit for the NGC 7213 data points only is shown and parametrised by $L_R = (0.2^{+0.056}_{-0.056}) \log L_X + 28.3^{+1.77}_{-1.77}$ (the constant includes the mass term in the context of MHdM03).

I include in Figure 3.5 the correlation found from the archival X-ray and radio data; however, as they were not simultaneous I do not include them in the best fit. I paired the X-ray and radio data points to calculate the correlation according to the closest in time. If either two X-ray, or two radio points were very close in time, both are included in the correlation. Also note that these data were calculated between 8.4 GHz and 2-10 keV and not 4.8 GHz. Assuming that the same correlation (i.e. \sim flat spectrum) holds between the 8.4 and 4.8 GHz observations the plot seems to show that two of the archival points do not follow the trend predicted by the fundamental plane.

To assess whether the deviation/scatter in the archival points (and indeed the other MHdM03 AGN) away from the fundamental plane can be explained due to a delay in sampling the X-ray and radio fluxes; I take the fraction RMS scatter in the

RXTE/ATCA X-ray and radio data and plot this errorbar with the archival points. I find the fractional X-ray RMS to be $fRMS_x=26\%$ and the fractional Radio RMS $fRMS_r=11\%$. The full light curve presented in this chapter spans ~ 1000 days and the longest change in time between the archival observations is ~ 500 days (see Table 3.2). Although this method only offers an estimate of the error between sampled data points, it does allow us to assess outliers.

Within this framework it appears that the first (PTI correlation) and last (LBA correlation) data points that were used in the correlation are outliers: the ATCA and VLA correlation points sit close to the best fit. Even when taking into account the full RMS from the 1000 day light curve the error bars do not bring the data points close to the MHdM03 best fit. I have also speculated earlier that there could be a certain amount of missing flux associated with these long baseline observations. Moving both the PTI and LBA points up by some set amount will still leave the PTI point away from the correlation. Although applying the intrinsic variation of the source to the archival points cannot bring all points onto or close to the correlation; the scatter in these points is within that permitted by the other MHdM03 points shown on the plot.

By using the fractional RMS of X-ray and radio variability from this study I have placed a constraint in the deviation of the archival data points away from the fundamental plane. These constraints suggest that other forms of scatter (apart from bad sampling and missing flux) could be affecting the data. As summarised in KFC06 other forms of scatter could be attributed to beaming, source peculiarities and spectral energy distribution - but note, these should remain relatively constant for the *same* source.

In Figure 3.6 I plot a LLAGN *only* sample and best fit parameters taken from KFC06 with the updated BHXR sample. Note, the X-ray data in the KFC06 sample are taken in the 1-10 keV band. As the NGC 7213 data points are in the 2-10 keV band I assume a photon index of $\Gamma=1.8$ and add a correction factor to the NGC 7213 data points to make them comparable. I apply a similar correction to the BHXR sample. From this plot it is clear that the NGC 7213 data points are positioned slightly above the best fit line.

Considering the position of NGC 7213 on the fundamental plane with respect to other LLAGN (see figure 6), I calculate its radio loudness parameter to assess the difference. I calculate the radio loudness parameter $R = L_{6cm}/L_B$ (where L_{6cm} and L_B are the radio and optical luminosities); I use a B band magnitude of 16.3, and find the optical flux S_{opt} using $B = -2.5\log S_{opt} - 48.6$ (Halpern & Filippenko, 1984): giving $R = 134.8$. In this scheme radio-loud sources are typically defined as

having an R parameter > 10 , while radio-quiet range between 0.1-1 (Kellermann et al., 1989).

Using the alternative radio loudness parameter of Terashima & Wilson (2003) which utilises the X-ray instead of optical luminosity, $R_X = L_{6cm}/L_{2-10keV}$, I find $\log R_X = -3.28$. Panessa et al. (2007) show that for a sample of low-luminosity Seyfert Galaxies $\log R_X = -3.64 \pm 0.16$ while for a sample of low-luminosity radio Galaxies (LLRGs) $\log R_X = -1.40 \pm 0.11$. Therefore, with respect to the X-ray radio loudness, NGC 7213 is only slightly higher than that of a sample of low-luminosity Seyfert Galaxies; while under the standard definition of radio loudness it is indeed radio loud. These results are consistent with the position of NGC 7213 on Figure 3.6.

As was discussed earlier in this chapter there is an apparent time lag between events in the X-ray and radio. Therefore comparing the fitting parameters found from the NGC 7213 data with the MHdM03 relationship without correcting for the lag might give rise to errors as we are not matching the correct data points. The width in the cross-correlation peaks shows that the time lag is variable. Thus, for example, the two radio flares could have different lag times associated with them. Therefore shifting the entire radio light curve back by a set amount to match the X-ray could still give a scatter. To simplify this problem I separated out the data for the first flare only because I have a more accurate measurement of the lag in this specific case. I then shifted the radio data -35 days which was the time lag measured for this singular flare using the DCF at 4.8 GHz (see Table 3.1). The top panel of Figure 3.7 shows the uncorrected data on the MHdM03 plot while the middle panel shows the corrected data; for completeness the bottom panel shows all radio data points shifted back.

For the first flare correcting the data appears to reduce the scatter and increase the gradient more in line with the MHdM03 best fit. It is now described by $L_R = (0.58_{-0.14}^{+0.14}) \log L_X + 8.5_{-7.6}^{+7.6}$. To check the statistical significance of this I measured the gradient for a variety of shifts. From 0-25 days the gradient gradually steepens until it gets close to 1 (giving a coefficient of ~ 0.6). From 30-50 days it plateaus ~ 1 and from 50 days the gradient decreases towards 0. Thus it does appear that moving the flare back by the amount given from the DCF function does seem to better represent the data with respect to the MHdM03 fit. However, it should be noted that as this is a log - log plot, measuring gradients from such a small range of values should be treated with care. It would, however, be interesting in future studies to assess the importance of this parameter.

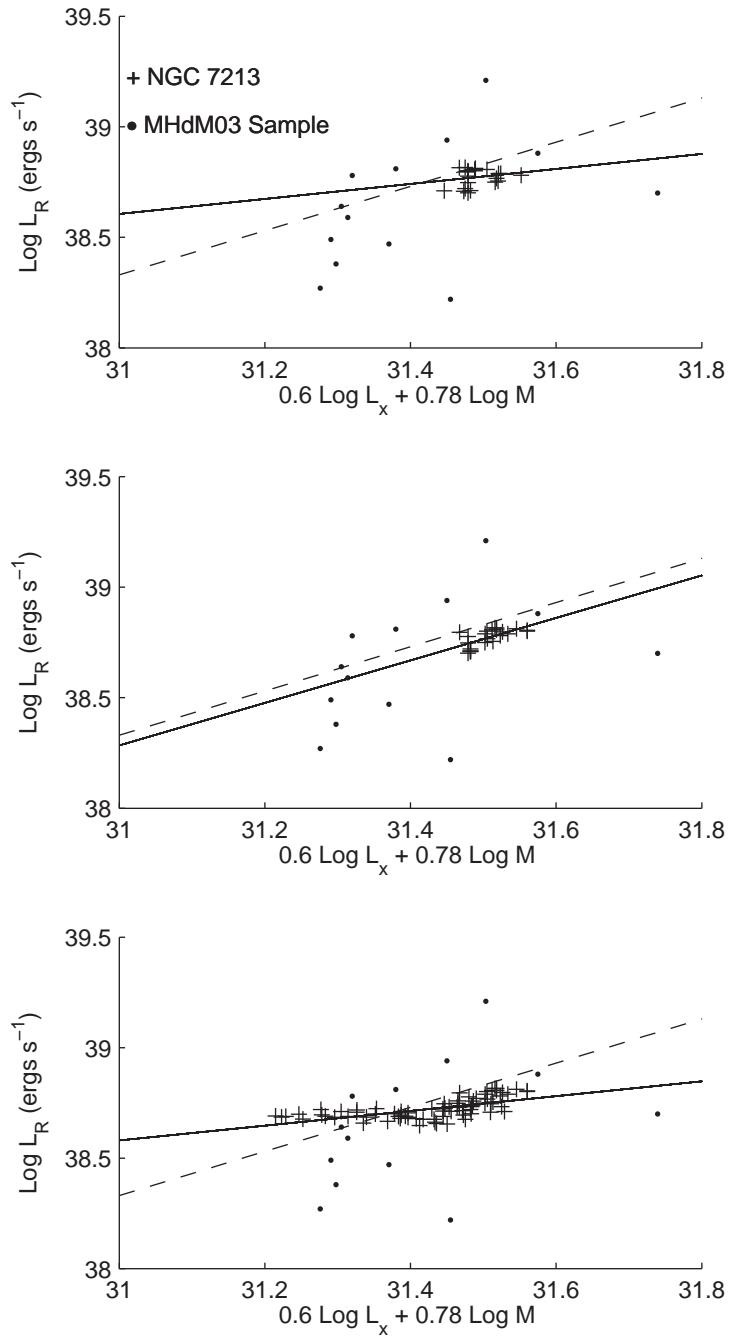


Figure 3.7: Top Panel: Flare 1 radio data points at 4.8 GHz uncorrected for lag. On both graphs the solid line indicates the least squared best fit for the NGC 7213 data only; the dashed line indicates the MHdM03 best fit. Middle Panel: Flare 1 radio data points shifted -35 days to correct for time lag calculated from the cross-correlation function. Bottom Panel: All radio data points shifted back -35 days.

3.6 Conclusion

We have used the Australian Telescope Compact Array and the Rossi X-ray Timing Explorer to conduct a long term study of AGN variability in the LLAGN NGC 7213. I have used the cross-correlation function to show that a complex and weakly significant correlated behaviour exists between the X-ray and radio emitting regions. Although the statistics only show a weak correlation, this study is the first definitive campaign to probe this type of behaviour.

We have shown that NGC 7213 sits well on the predicted fundamental plane of black hole activity plot when compared with the MHdM03 sample. However I have shown that when comparing NGC 7213 with a revised BHXR and LLAGN sample that the data points are above the expected correlation; which is however consistent with the calculated radio and X-ray loudness parameters. I have also shown some support that by correcting for the time lag between events in X-ray and radio the gradient of data points agree better with the best fit derived from the MHdM03 sample on the fundamental plane.

Monitoring of all types of AGN at low frequencies (with LOFAR and other pathfinder instruments) has the potential to deliver long term datasets to examine variability. The low frequency component of the jet emission may be further removed from the X-ray emission, as it is produced further down the jet. Therefore, finding correlations between the X-ray and radio emitting regions might be more complex. It has been difficult to perform long term monitoring campaigns with GHz instruments, as the observations need to be spread over many semesters, for many years - and thus require maintained requests for observing time. Wide field GHz facilities such as ASKAP and the SKA potentially might observe large numbers of AGN serendipitously, which will in turn provide more data to examine the correlations presented in this chapter.

One man's constant is another man's variable.

ALAN PERLIS (1922 - 1990)

4

Testing and refining the LOFAR transients detection pipeline

In this chapter I will present the contribution I have made to the testing and development of the LOFAR transients detection pipeline. I will present a prototype transients pipeline which I have constructed, with the help of LOFAR Transients Key Science Project (TKSP) members, to perform the testing conducted in this thesis. I will refer to this prototype pipeline explicitly as LTraP. I will state clearly when code has been developed by other LOFAR TKSP members - as it is a major project with many components. I will also present some functionality of the fully operational LOFAR transients detection pipeline, for clarity. I will give examples within this chapter of some of the intricacies of transient detection, and offer some suggestions for further improvement with regard to full operations. The total time I have contributed to testing and developing the LTraP is \sim one year, spread over four PhD years.

4.1 Introduction

The (fully operational) LOFAR transients detection pipeline is designed to perform source extraction on incoming LOFAR images (or indeed any acceptable astronomical images); to catalogue the results (from source extraction) in a database; to search the database for transient and variable radio sources, including multiwavelength counterparts; to classify the variability (if possible); and to issue alerts to the wider astronomical community (if appropriate).

4.2 The prototype LOFAR transients detection pipeline

Figure 4.1 shows a flow diagram which describes the prototype LOFAR transients detection pipeline (LTrAP). Some of the components in Figure 4.1 have been added to illustrate the broader functionality of the fully operational pipeline. For example, the design and operation of the Virtual Observatory Event (VOEvent) alert system is not within the scope of this thesis, however, it will be an important component to alert the broader astronomical community to interesting transients. The version of LTrAP described in Figure 4.1 is the latest version I have developed, which has been augmented over the duration of my PhD. In the following sections I will describe in detail the components of the LTrAP, including some of the testing that has been performed.

4.2.1 Input: LOFAR images and metadata.

The input to the LTrAP will be provided by the LOFAR Standard Imaging Pipeline (SIP). In Figure 4.1 I describe the output from the SIP in terms of an image cube and metadata. The image cube can either be a sequence of images in time, or in frequency, or both i.e. a 4D hypercube. Typically we have used the Flexible Image Transport System (FITS; Wells et al. 1981) format to define a given image. The Hierarchical Data Format 5 (HDF5; Anderson et al. 2010) will potentially be used in future LOFAR operations. The LTrAP does not necessarily just work with LOFAR images, any radio image (or multi-wavelength image with well characterised point spread function) can be examined by the LTrAP for variable and transient sources.

As the SIP and the LTrAP are both written in Python, and operate using the same architecture, we can pass metadata and commands between the SIP and the LTrAP. I have added a metadata branch into Figure 4.1 to highlight this property. By designing the LTrAP to receive metadata from the SIP, the LTrAP can interpret

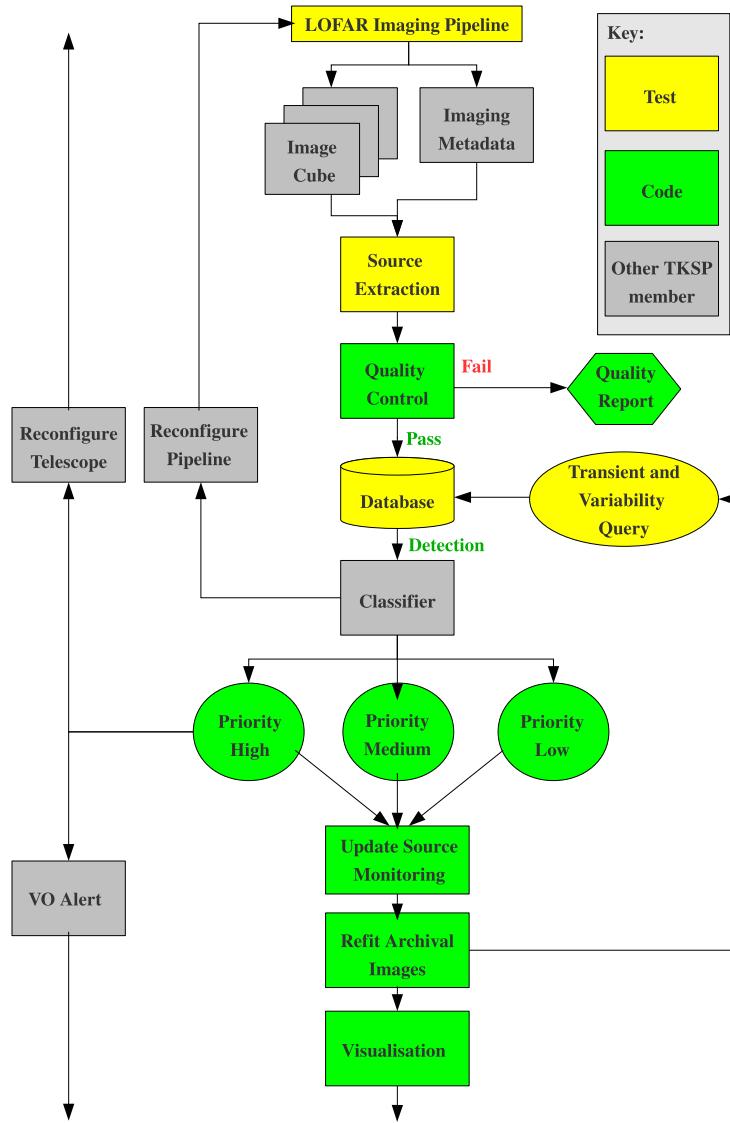


Figure 4.1: A flow diagram which describes the prototype LOFAR transient detection pipeline (LTrAP) that has been constructed for the testing described in this thesis. The functionality coloured in yellow I was responsible for *testing*; the functionality coded in green I was responsible for *coding* (at least in the preliminary sense); the functionality in grey was the responsibility of other TKSP and LOFAR members. Some of the functionality has been added to illustrate the broader functionality of the fully operational LOFAR transients detection pipeline.

what data reduction settings were used to make the images. The LTrAP can then reconfigure (if appropriate) the SIP so that images are reconstructed with a variety of different settings, for example, to achieve better dynamic range.

4.2.2 Source extraction

The source extraction algorithms have been developed by Dr. Hanno Spreuuw (see Spreuuw 2010). The source extraction procedure is designed to isolate islands of pixels, within an image, above a given source extraction threshold (typically the RMS). Once these islands have been identified elliptical Gaussians are fitted, and the results (including all associated metadata and information of the dataset being inspected) are stored in a custom designed database. The inputs to the source extraction algorithm (apart from a sequence of images) are the source extraction and analysis thresholds. Once an island of pixels has been identified above the source extraction threshold, the analysis threshold is then used to fit an elliptical Gaussian. The analysis threshold can be set lower than the source extraction threshold to achieve better source fitting. For example, typical values used to retrieve almost all sources within an image are: 5σ for source extraction and 3σ for the analysis threshold (where σ is the RMS). Alternatively, if the user wanted to only identify bright sources within an image, a source extraction threshold of 20σ can be used, with an analysis threshold of 3σ . Keeping the analysis threshold low insures that the Gaussian fits are accurate.

There are however drawbacks to using a simple RMS detection level source extraction system. In an ideal radio image containing no astronomical sources the noise can be described with a Gaussian distribution centred around zero. The commonly quoted value of the ‘noise’ is the RMS - or one standard deviation of the Gaussian distribution. Setting a source extraction threshold of 3σ would ensure that only $\sim 0.3\%$ of the pixels were above the source extraction level. These ‘bright’ pixels could potentially trigger a detection of a source, artificially. The number of false detections is proportional to the image size, for example, a 512×512 image would yield 786 pixels above a 3σ source extraction threshold; a 4096×4096 image would yield 5×10^4 . To account for this the source extraction system contains a False Detection Rate (FDR) algorithm, as described in detail in Hopkins et al. (2002). The FDR sets a limit on the total number of false detection permitted by the user, and adjusts the source extraction threshold accordingly. This method is independent of image size and dynamic range and has been shown to be superior to simple threshold based source extraction. The FDR algorithm was not tested within

the work presented in this thesis, but it will be used in future LOFAR operations.

There are numerous other settings to the source extraction system, the only parameter of major consequence to this thesis is the backsize. The backsize (in X and Y pixel coordinates) defines a box size, for which a given image is divided into when calculating an RMS grid. This RMS grid is subsequently used to calculate if an island of pixels is above the detection threshold. Changing the backsize parameter can influence if certain sources are detected or not, and I will discuss this further in Chapter 5. Using an RMS grid is important for wide field images as it allows for variations in the RMS across the image, which is not always constant. See Spreuuw (2010) for further details on the source extraction algorithms.

The source extraction algorithms have been tested extensively on the data presented in this thesis. In conjunction with the chapters presented later, source extraction has also been tested on reproducing the lightcurves for the NGC 7213 dataset presented in Chapter 3. The NGC 7213 images were fed into the LOFAR source extraction algorithms and the Gaussian fits were compared with those provided by the MIRIAD fitting routines. The LOFAR source extraction Gaussian fits reproduced the results from MIRIAD with excellent accuracy.

4.2.3 Quality control

Quality control will be an important filter to prevent results from unreliable images being catalogued into the database. At the time of writing quality control was not implemented in the transients pipeline, however, some offline quality checks had been performed which could easily be incorporated in the future. For example, in the case of LOFAR the SIP will currently produce one image per subband, for a total of 244 subbands. A subband is defined as a dataset containing a fraction of the total available bandwidth. Because the LOFAR data volumes are large, individual images are produced for each subband, they are then averaged together in the image plane to produce a final image. Note, the full architecture and discussion of the LOFAR SIP will follow in Chapter 6. An *a priori* sky model is currently used to calibrate the LOFAR observations. The Gaussian parameters of a number of bright sources within the FoV of the observation are specified, with a spectral index, and are passed to the calibration routines.

A simple quality check is to measure the flux of the calibrator sources (in the image plane) in each subband, and to compare it with the expected flux from the sky model. Figure 4.2 shows two plots calculated from LOFAR commissioning observations (observation ids L2010_20033 and L2010_09948 see Chapter 6) of

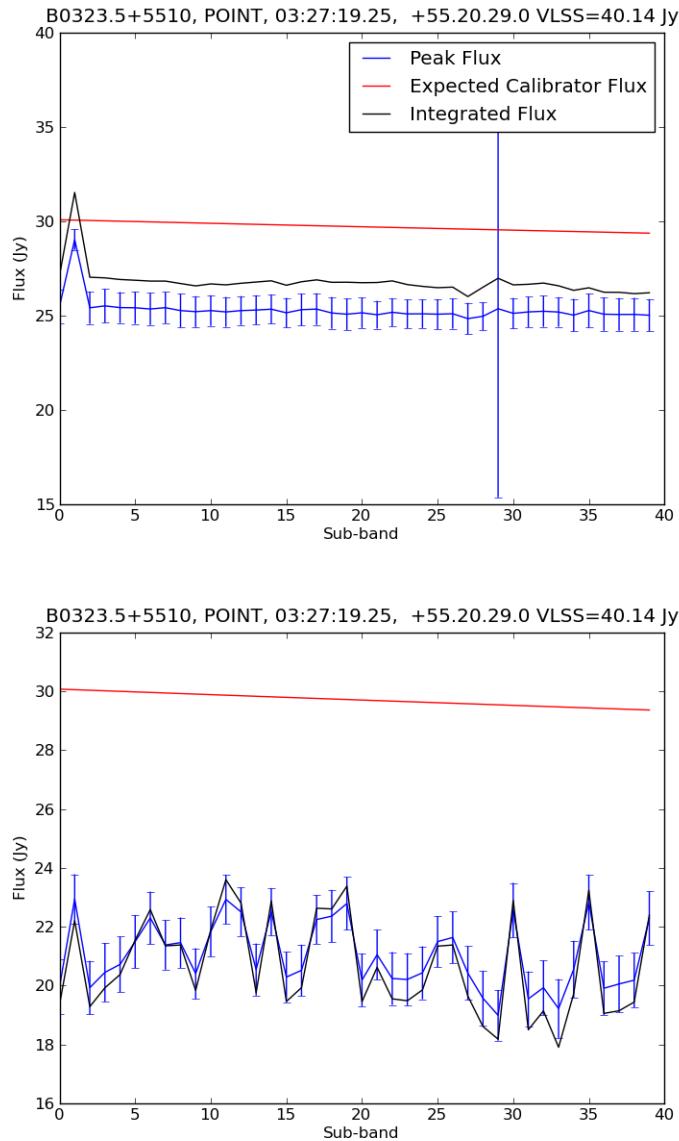


Figure 4.2: Both figures show the fluxes (per subband) of the calibrator source B0323.5+5510, measured in the image plane for two different epochs of LOFAR commissioning data (top panel: LOFAR dataset L2010_20033, bottom panel: LOFAR dataset L2010_09948, see Chapter 6 for further details). The LOFAR imaging pipeline currently produces one image per subband. To check the quality of the calibration I fit an elliptical Gaussian to the position of the calibrator in each subband. The red line shows the expected fluxes based on the spectral index derived from the VLSS (74 MHz) and WENSS (325 MHz) fluxes.

flux against subband. The peak flux is shown in blue, integrated in black, and the expected flux for a given subband (and hence frequency) is shown in red (the first 40 subbands are shown only). The spectral index of the calibrator source B30323.5+5510 was calculated from the VLSS and WENSS fluxes respectively. It can be seen from both plots that the measured flux in all subbands sits well below the expected flux. This was anticipated because at the time of these observations absolute flux calibration was uncertain. However, the bottom panel shows a large scatter from subband to subband. This information could be easily used to reject images from progressing further in the LTrAP.

Other diagnostics of quality that might be used could be examining the amount of data flagged data out by the imaging pipeline. Images with bad fidelity may have lots of visibilities removed before imaging, if the LTrAP can interpret this information it could further reject images from being used in subsequent stages in the pipeline. If the quality control algorithms reject a given image, a report should be issued and the pipeline process terminated for that image. If the quality is adequate then the results from source extraction should be stored in the database. Note, if bad images are extracted into the database, new transients might be associated with older incorrect artifacts, which could have a knock on effect for scientific interpretation.

4.2.4 MonetDB database and transient searches

A MonetDB (Boncz 2002) database has been developed by Dr. Bart Scheers (see Scheers 2010). MonetDB is a MySQL type database structure designed to execute and return query results with speed. It is used to store the results from source extraction and to provide an architecture to search for variable and transient sources. The database consists of a structure of tables and columns which are populated with the metadata and source properties of all images and extracted sources. The database will not only store information about the extracted sources, but it will automatically associate new sources with those from previous images and catalogues. Each new source entering the database is assigned a source id. number. The right ascension and declination (including errors) of new sources entering the database are compared with previously inserted sources, and those found in the stored radio catalogues. Currently we have the NVSS (Condon et al., 1998), WENSS (Rengelink et al., 1997) and VLSS (Cohen et al., 2007) catalogues stored, more can be added later. If two or more sources are found to be associated, the database will assign an associated source id. indicating the association. For further information on the full

architecture of the LOFAR database see Scheers (2010).

A simple query (written by Dr. Scheers) is shown below which was used extensively in the archival transient searches presented in Chapter 5. The query returns a list of unique sources that appear in only one epoch i.e. single epoch transients. It achieves this by searching for sources with a unique source id. and no associated source ids. This query is quite crude and needs some massaging to perform adequately. For example, simply having a unique id. does not completely qualify a source as a true radio transient: many imaging artifacts have unique source ids. This query in reality offers a starting point for further investigation. A large number of single epoch transient detections were triggered with this query when inspecting the archival VLA calibrator fields presented in Chapter 5. I will offer some further discussion surrounding the topic of single epoch transients, and the use of this query later in this chapter, and in Chapter 5.

Example single epoch transient database query:

```
SELECT image_id, xtrsrc_id, det_sigma
      FROM assocxtrsources, extractedsources, images
     WHERE xtrsrc_id = xtrsrcid
       AND image_id = imageid
       AND ds_id = dsid
```

In parallel to searching for single epoch sources the database can be queried for variable sources. I use the measures of variability adopted and derived by Scheers (2010), which are also used by Bower et al. (2011) and Ofek et al. (2011). The index of variability at a given frequency V_ν is calculated using the standard deviation of the flux σ_ν , divided by the mean flux \bar{S}_ν .

$$V_\nu = \frac{\sigma_\nu}{\bar{S}_\nu} = \frac{1}{\bar{S}_\nu} \sqrt{\frac{N}{N-1} (\bar{S}_\nu^2 - \bar{S}_\nu^2)} \quad (4.1)$$

N is the number of data points within a given sample (or lightcurve). For clarity, within the brackets is the mean of the squared fluxes, minus the mean of the fluxes squared. To give the reader an intuitive feel for the index of variability, consider a sample consisting of five measurements, all measuring 2 Jy i.e. no variability. When substituted into equation 4.1 the brackets will cancel yielding $V_\nu = 0$. Substituting one of the five measurements with a flux of 4 Jy i.e. a 100% increase for one measurement will yield $V_\nu = 0.42$. Typically when using this function to search the database for variable sources I set $V_\nu > 0.5$ as a first pass to search for highly

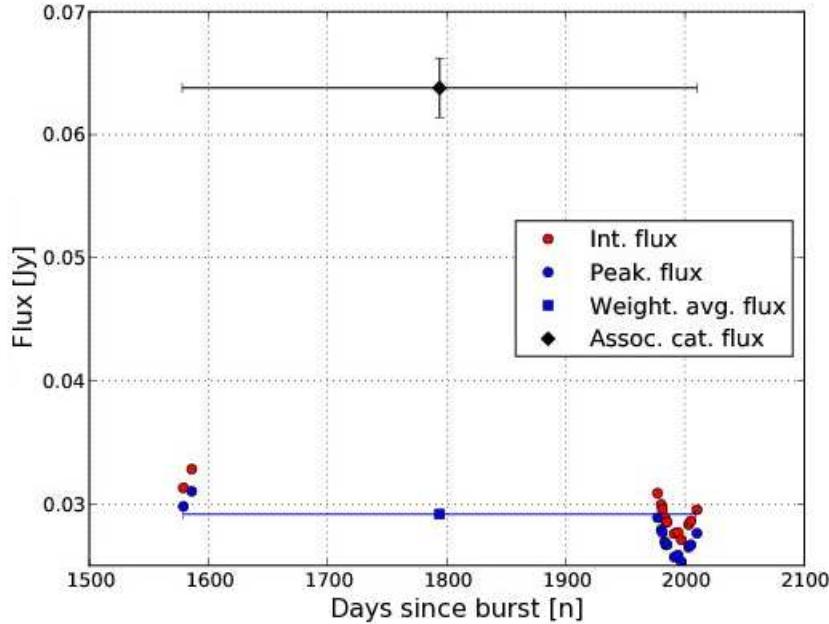


Figure 4.3: Example lightcurve of a variable source detected in a 14 epoch WSRT dataset. The source had an index of variability $V_v = 0.59$. The X-axis shows time in days measured from an arbitrary time in the past. Assoc. cat. flux shows the flux of the NVSS counterpart. The images containing the transient source did not have primary beam correction applied, therefore there is a discrepancy in the NVSS and transient flux density (including intrinsic variability of the transient source).

variable sources. This variability query can be executed via (an adapted version of) the MySQL syntax described above to return the source id. numbers for all highly variable sources. Another variability statistic is described in Scheers (2010) (and coded in the LOFAR algorithms repository) which calculates the significance of the variability using χ^2 statistics: this was not tested as part of this thesis. It should be noted that indeed *any* metric of variability, or search, can be coded in the MySQL syntax to return sources. This flexibility will certainly be important for more difficult data mining exercises in the future.

Once the source id. numbers of interesting objects are retrieved, further information can be obtained from the database, for example, to construct lightcurves. An example lightcurve is shown in Figure 4.3. The source was identified in testing conducted on 14 epochs of WSRT data that is not discussed in this thesis (but will be pursued in future work). I extract the integrated and peak flux (including errors), time and NVSS flux of a candidate variable source (triggered from the MySQL variability search query) from the database, then construct a lightcurve (using Pylab). Note, this was initially coded by Dr. Scheers but was adapted, to retrieve different information in other scenarios, by myself. The source had an index of variability

$V_v = 0.59$. The source was correctly identified by the database to be associated with an NVSS source, and its NVSS flux (labelled Assoc. cat. flux) was added onto the plot in Figure 4.3 automatically. Note, the LTrAP currently has no knowledge of the primary beam correction needed for WSRT observations, thus the peak and integrated fluxes are below the NVSS flux: the source was considered variable based on the *it*'s flux alone. In future operations, queries can, and should be constructed to assess the variability of sources by comparing the catalogued and extracted fluxes, in conjunction with examining the intrinsic variability.

4.2.5 Source classification and alerts

Once the images have been passed through source extraction and the database transient searches, information can be gathered from the database to further characterise the source, and hopefully offer a classification. The character of the lightcurve can be examined and classed, for example, as periodic/pulsating, explosive i.e. fast rise exponential decay, scintillating etc. Other data such as multi-wavelength counterparts can be used to further classify a source and produce VO Event alerts. The fully operational version of the classification code is currently under development by Dr. Evert Rol.

In the testing conducted in this thesis I only used a basic classification of single epoch and variable source, derived from the queries described above. I also incorporated a three tiered ranking based on the detection level ($\overline{S_v}/\overline{RMS}$) of each candidate source. Detections were ranked in significance where 5 to 10σ = low priority, 10 to 15σ = medium priority and 15 to $\geq 20\sigma$ = high priority. This was a very simplistic interpretation of classification, however, it was important to assess at which stage of the LTrAP the transient search should be implemented and how the pipeline might react. Waiting for a sequence of images to finish processing through the pipeline and then performing the transient queries, is not a time effective way to detect transients in real time. In terms of the alerts that I coded in LTrAP; I only implemented a print-line which outputted to the screen when a high priority transient candidate had been identified in pseudo real time. By 'pseudo' I refer to the fact that the testing was conducted after the observations had been taken, but the pipeline could issue an acknowledgement that a transient had been detected as soon as source extraction and the database transient had been completed.

A broader view of how the alert system may operate in the future is shown below.

Case 1

- Good quality control metric (minimal RFI exclusion and stable calibration).
- Fast rise lightcurve with detections in multiple epochs.
- Multi-wavelength association with a black hole X-ray binary.
- High priority $> 20\sigma$ detection and increasing.

Action: Trigger VO Event (X-ray and optical follow-up); over-ride LOFAR and continue monitoring (e.g. bringing in international stations, making higher bandwidth observations etc); trigger other radio facilities.

Case 2

- Poor quality control metric (large amounts of RFI removed but stable calibration).
- Single epoch transient, no lightcurve classification.
- No known multi-wavelength association.
- Medium priority 12σ detection.

Action: Continue monitoring but do *not* trigger follow-up.

4.2.6 Monitoring

An instantaneous classification is an extremely important function to catch bright and rare transients. It is however anticipated that sources may change classification once more information has been collected. In the flow diagram shown in Figure 4.1 I have included two steps prior to visualisation - ‘Update Source Monitoring’ and ‘Re-fit Archival Images’. Once the positions of candidate transients have been identified, elliptical Gaussians should be fitted to that position in all further images that are processed and potentially all previous images, if available. Furthermore, if a transient falls – in flux – below the RMS, the RMS should be submitted to the database and subsequently used in the variability search. In the early testing I conducted, the database was not yet capable of correctly handling the RMS measurements. Elliptical Gaussians had to be fitted manually to the region of interest to derive the complete lightcurve.

I therefore experienced a few examples of interesting transient behaviour slipping through the net of the transient detection algorithms, with regards to RMS measurements. In one case a test dataset was created using a VLA image of the source 3C48. The field contains a bright point source at the centre; this source was modelled and removed from the image, leaving a relatively blank field. A sequence of images were then created (by updating the time stamps in the FITS headers) using the image with 3C48 present as well as the image with 3C48 removed. Three images of 3C48 removed were followed (in time) by one image of 3C48 present (unremoved) and continued for a number of cycles. The transient pipeline identified the source when it was ‘on’, however, it did not submit the RMS to the database when it was ‘off’. As the source had the same flux (when it was on) the variability query incorrectly classified it as having $V_v = 0$. After the data had been refitted i.e. utilising the RMS when the source was off, the true variability could be quantified. In the top left panel of Figure 4.4, I show a lightcurve depicting this behaviour for clarity. Both the processes of refitting and monitoring are represented in the flow diagram in Figure 4.1 with a loop that cycles back to the transient search and database. In future versions of the database this loop, including the re-fitting/monitoring will be fully implemented.

A further example is shown in Figure 4.5 (note this figure shows the actual lightcurves extracted from the database). An early 12 hour LOFAR commissioning observation (L2010_06928, see Chapter 6) was calibrated and imaged by Dr. Jess Broderick. The dataset was then divided into 12×1 hour chunks and re-imaged. The final FITS images were then processed through the LTrAP. The full 12 hour image was used as the first image (in the sequence of images being searched), as it had the largest number of sources and best image fidelity. In this particular example, the consequence of using an image with higher dynamic range as the first image, is that many sources in the following images sit below the source extraction threshold, due to the increase in RMS (see the top right panel of Figure 4.4 for a further example). In accordance with the flow diagram the source was refitted/monitored. See the bottom panel of Figure 4.5 for the fully fitted lightcurve. We can see that the source is not particularly variable; there was a slight variability attributed to uncertainty in absolute flux calibration. This was a commonly observed false transient trigger which is largely solved by monitoring sources. It is obviously anticipated that LOFAR will on occasion perform a very deep observation of a given field. This will again potentially trigger a number of single epoch transient detections, if not handled correctly. In parallel to monitoring and re-fitting, the single epoch transient search described above should be improved by using the RMS values, to assess if

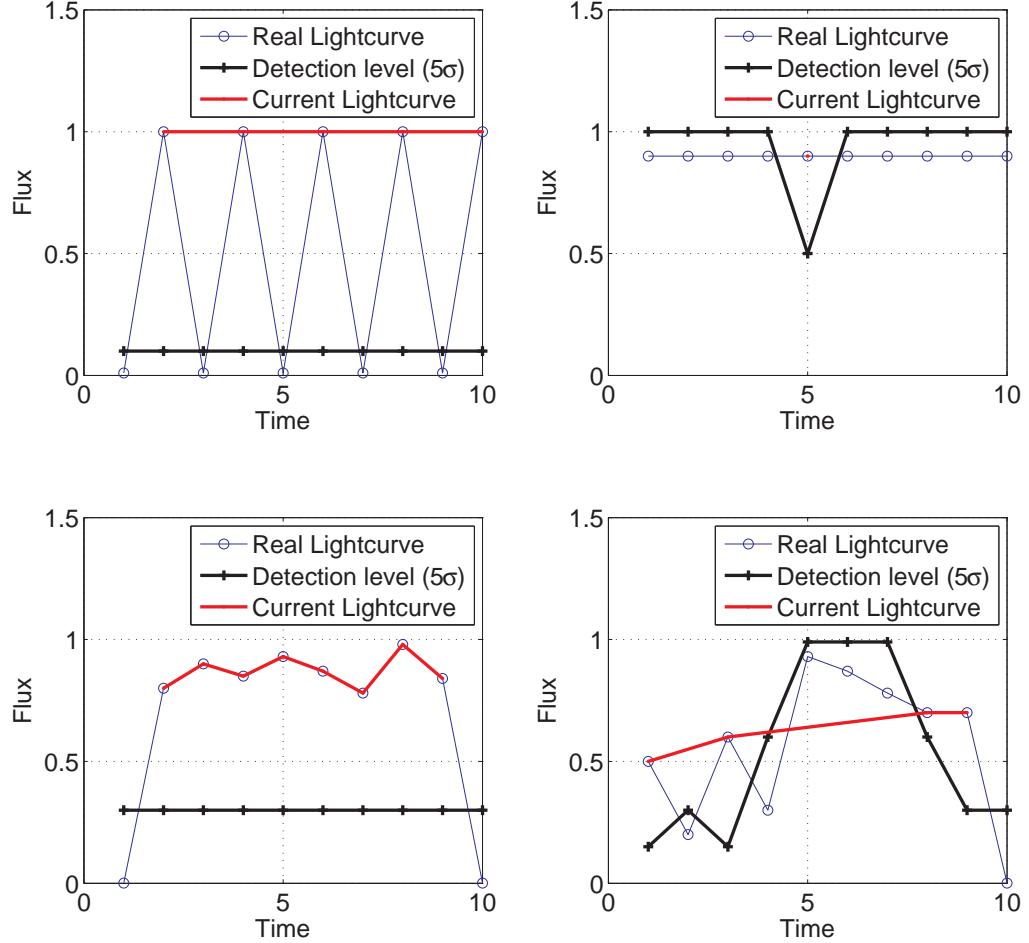


Figure 4.4: Example (simulated) lightcurves of transient behaviour that LTrAP currently struggles to correctly identify. The flux and time are plotted in arbitrary units. The red lines show the current lightcurves that are extracted using the LTrAP. The blue line shows the correct lightcurve, which utilises the RMS measurements in the transient query.

particular sources would be expected to be detected in different images.

Two further examples are shown in Figure 4.4. On the bottom left panel the red line shows the source lightcurve, without the RMS measurements. The blue line shows the true lightcurve of the source. By using the RMS measurements we can only set an upper limit on the variability of the (simulated) source, with respect to the detection threshold. The bottom righthand panel shows a mix of variability and RMS measurements. In this case, from the time 1 to 3, by using the RMS measurements we may define this source as highly variable. From the time 4 to 7 the detection threshold is above the transient flux. It is difficult to use the RMS measurement in this case, because we cannot be sure what changes in flux have occurred. The pipeline will need to understand the difference in these two cases

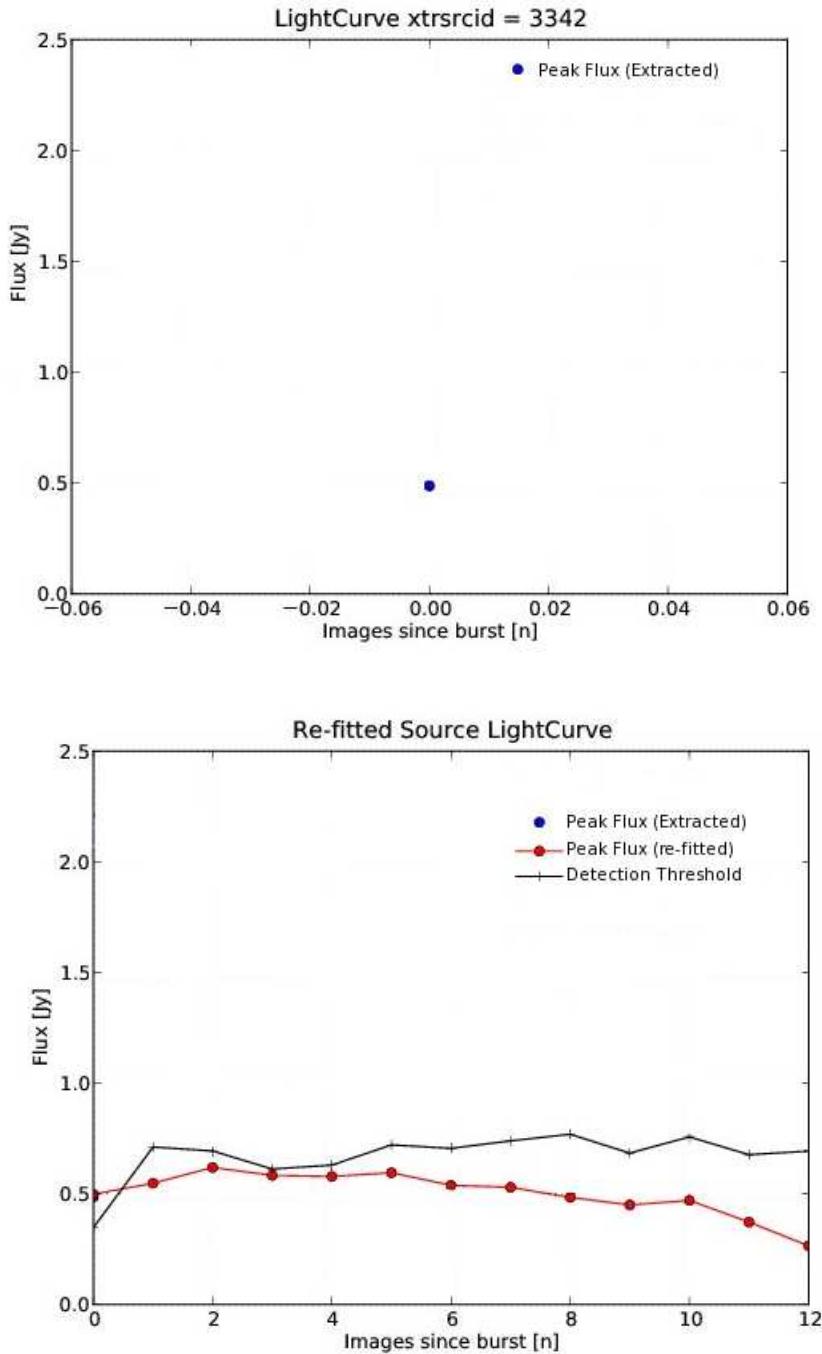


Figure 4.5: Top panel: A lightcurve retrieved from the database of a single epoch transient automatically identified by the pipeline in LOFAR commissioning data. Bottom panel: The full lightcurve of the *same* source when refitted/monitored with elliptical Gaussians. In this example the sequence of images consisted of one 12 hour LOFAR observation, followed by re-imaged one hour consecutive chunks of the original full 12 hour dataset. The first 12 hour image yields one unique source, because the source sits below the detection threshold in the subsequent images. This error is rectified by re-fitting/monitoring unique source positions and then reassessing the variability.

and respond accordingly.

4.3 Visualisation

Subsequent to the transient database search (including all re-fitting, monitoring and re-assessment of variability), an essential part of the pipeline procedure is to inspect the large quantity of information generated. I developed two methods to do this. One method was an interactive tool to browse through data using the FITS viewing package DS9. The second made use of HTML to produce large numbers of webpages which contained all lightcurves and information of potential transients. These were both exploited in the testing conducted throughout this thesis and I discuss them in more detail in the following sections.

4.3.1 Interactive tools

There are two interactive query toolbars shown in Figure 4.6 (written in Python Tkinter) that communicate between the LOFAR database and DS9. In the top toolbar clicking ‘Show Transient Candidates’ will send – via Python – the MySQL query to the database which will return the locations of *all* the transient candidates within a given sequence of images. It will then pipe the locations to DS9 and overlay them on a master image. Each source is labelled with its unique source id. number, and the relevant lightcurve can be retrieved using the toolbar. The lightcurve is generated by querying the database for the correct flux (integrated and peak, both with errors) and time; the results are then plotted with Pylab. The toolbar can also show the locations of the persistent sources within the field, which is shown in Figure 4.6.

The bottom toolbar uses a similar system but instead overlays only the transient candidates in the image which they are detected. The location of each FITS image is stored in the database (as metadata) such that each transient identification is linked to an image location. This information can be pulled from the database to view the transient locations specific to a given image. The toolbar can then be used (at speed) to look through a large sequence of images using the scroll buttons. Both these toolbars we used extensively in the VLA archival transient search presented in Chapter 5, inspecting over 5000 images.

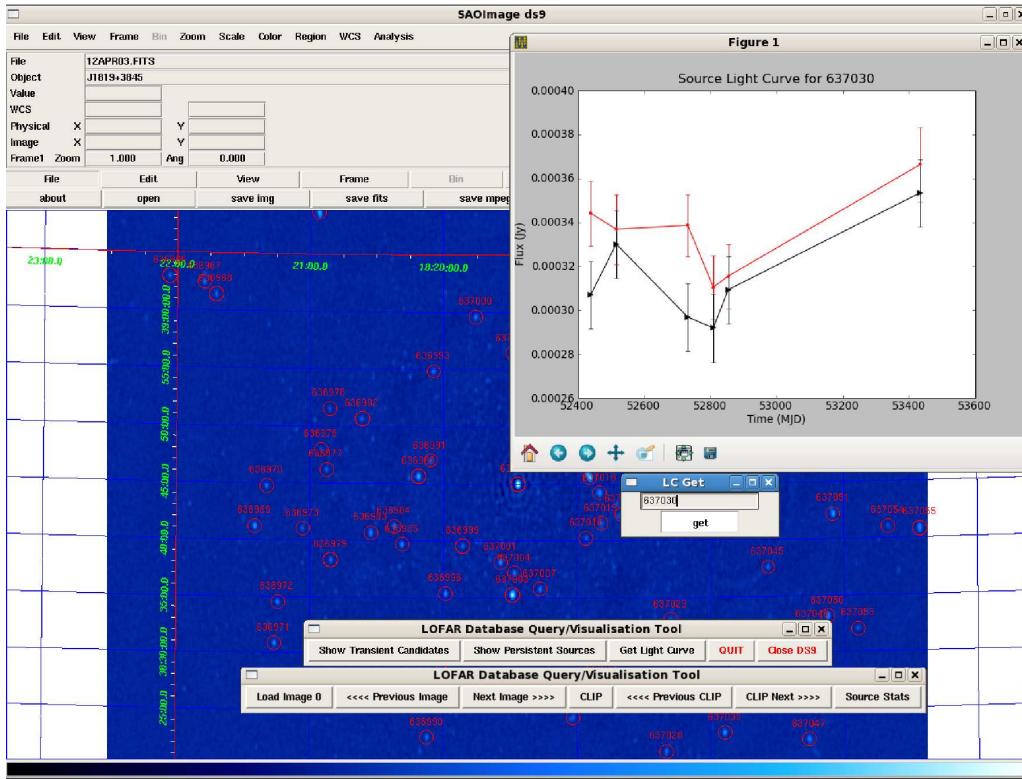


Figure 4.6: Example visualisation GUI showing the use of a Python coded toolbar to interact with the database, in conjunction with DS9. Large numbers of images can be inspected – at speed – and the positions and source id. numbers of transient candidates can be viewed. Lightcurves can also be generated for any interesting source(s). These tools were used extensively in the VLA transient search presented in Chapter 5.

4.3.2 Webpages

An alternative to using the interactive tool was to generate HTML webpages (one per source) containing the all information relevant to the transient candidates. This algorithm was coded within the pipeline to post-process any transient detection (after source extraction and the database search). The sources id. numbers of candidates were piped to a visualisation function which performed the following tasks (for each source independently).

- The locations of the stored FITS images was queried and returned from the database; the pipeline then used the FITSCUT¹ tool to extract a JPEG thumbnail of 50×50 pixels around the transient location (for each image). The locations of the JPEG images were then added to the source webpage. It is possibly anticipated that data storage capabilities might be such that even

¹FITSCUT is a tool to produce a cutout in the form of FITS, PNG or JPEG of a master FITS image. <http://acs.pha.jhu.edu/general/software/fitscut/>

storing FITS images with LOFAR might be impractical, therefore extracting thumbnails may be required.

- The lightcurve was extracted from the database, plotted and added onto the webpage.
- Metadata for each source was added (in text) to the webpage. This metadata included RA and DEC, detection level and timescale of variability (if detected multiple times). The peak and integrated fluxes were compared to assess if the source was a point source – a classification of ‘point’ or ‘resolved’ was then added to the page accordingly.
- If the source was a single epoch detection, all images were refitted with elliptical Gaussian and the lightcurve was added to the page.
- If there was an associated catalogued source (NVSS, WENSS or VLSS) the associated source id. number and distance to the source were also added to the webpage. This was used to check that sources had been correctly associated with the NVSS, WENSS and VLSS counterparts.

An example webpage is shown in Figure 4.7. The infrastructure described above could be modified to retrieve any information from the database (using MySQL), which can subsequently be used to generate outputs on the webpages.

All sources that were deemed variable or transient were processed by using the procedure described above. The final list of transient sources and the links to the webpages generated, were ordered based on either detection threshold (single epoch) or index of variability for variable sources. After the pipeline had finished the post-processing, a firefox browser was automatically loaded to let the user browse the results.

The tools described above have been useful to heavily test various LTrAP settings, to expose bugs within the pipeline code and to proceed with commissioning. They have also been important to test the possible future full scale visualisation needed for LOFAR. The tools developed for the work conducted within this thesis did however have their pitfalls. The toolbars were good with large numbers of images, with a small number of sources. The webpages were useful for large number of sources with a small number of images i.e. it is impractical to view large numbers (>100) of thumbnail images on a webpage. The final LOFAR visualisation toolkit may indeed include implementations of both webpages and software. One choice would be to use the PHP dynamic internet language, which can execute embedded MySQL

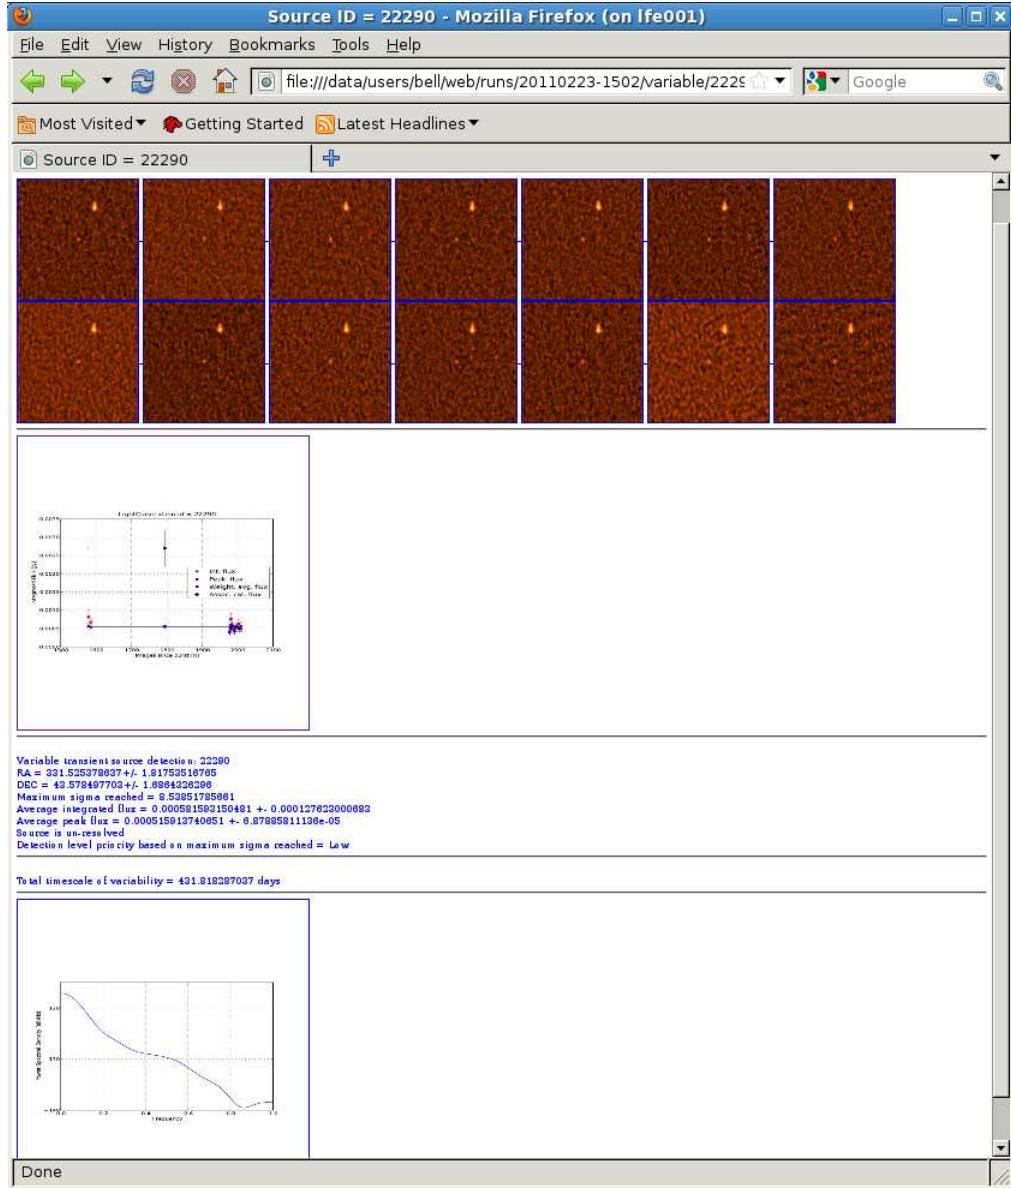


Figure 4.7: Example webpage for a transient source generated from the visualisation library. Included on the webpage are: thumbnail cutouts of a region around the transient position for each epoch; an automatically generated lightcurve which included the fluxes of multi-wavelength counterparts; statistical information; and other diagnostics plots relevant to the detection.

commands; this could be used to generate source webpages, on-the-fly. However, it is still cumbersome when viewing large numbers of thumbnail images, which may therefore require specifically designed software.

4.4 Considerations of single epoch transients

Single epoch (non-repeating) transients will be more prominent in images where the characteristic timescale of transient behaviour is less than the cadence of the observations. In future LOFAR operations it is anticipated that images will be produced on one second timescales; this is advantageous for producing well sampled lightcurves for bright transients. In this thesis (Chapters 5 and 6) I have typically dealt with observations that are spaced a few days apart. I will therefore add some statistical considerations of how we define the properties of single epoch transients, without adequate information on time duration.

The luxury of examining the visibilities long after an observation is not something that is possible with LOFAR (in full operational mode), as we are limited by the data storage restrictions (N.B. It may be possible to store time and frequency averaged visibilities). Therefore, it is important to understand how single epoch transients/bursts manifests themselves in our images and also how we define a statistic governing them - without any knowledge of the visibilities. It is possible that a single bright burst could be recorded in the visibilities that only lasts a fraction of the total observing time, which is then subsequently diluted with sequential measurements of noise. It is of course true that if the event *did* last for the entire observation (or longer) then it was indeed the true flux of the source. This is comparable with an effect observed in Pulsar astronomy. When observing Pulsars the duty cycle of the pulses can be much less than the pulse period. The time averaged flux of the pulse, over an entire rotation period, is therefore much less than the peak flux. This can result in time averaged fluxes as low as μJy , while the peak flux is much higher.

To account for this, if LOFAR detects a single epoch transient we can define a transient flux and duration function. The flux of a single epoch transient in an image S_{Im} can simply be defined as: the fraction of time (T) the source was on, with respect to the integration time (τ), multiplied by the burst flux S_B , plus or minus the noise.

$$S_{Im} = \frac{TS_B}{\tau} \pm \frac{\sigma_{1\text{sec}}}{\sqrt{\tau}} \quad (4.2)$$

In this example I make the assumption that the transient is either on or off, which is a crude assumption, however, using a more complicated burst function i.e. fast

rise exponential decay is possible. I can further rearrange equation 4.2 to find the minimum time the source was on to produce a 3σ detection – if we know the burst flux S_B .

$$T_{min} = \frac{\sqrt{\tau}3\sigma_{1sec}}{S_B} \quad (4.3)$$

To test the equations above I constructed a simulation; the motivation was to check if the inverse Fourier transform and deconvolution stages affected the resulting transient image flux, as predicted by equation 4.2. A *uv* dataset was created using the MIRIAD task UVGEN, which used the antenna positions for the VLA in D-configuration at 1.4 GHz, with a bandwidth of 10 MHz. A system temperature of 60 Kelvin was used and an addition Gaussian noise component at the level of 1% was included. 12×12 hour datasets were then created which contained a transient point source at the phase centre. The transient source was inserted in each of the 12 datasets and had a duration that increased from 0 to 12 hours (in one hour increments). Once the *uv* data had been created, it was imaged and lightly CLEANed with 50 iterations. The datasets had an RMS of ~ 0.4 mJy, I therefore inserted (in two separate simulations) a transient of $S_B=40$ mJy (100σ), and a transient of $S_B=4$ mJy (10σ).

In Figure 4.8 (in both plots) I show the duration the transient was on i.e. the Burst Time T (x-axis) against Flux (Jy), or S_{Im} (y-axis). The black line shows the measured flux of the transient source (using the MIRIAD IMFIT routine), downward triangles denote upper limits. The blue line shows the model fit calculated using equation 4.2. I also include a horizontal red line giving the 3σ detection threshold, and a vertical line showing the minimum burst time T_{min} needed for a 3σ detection, calculated using equation 4.3. The top panel, which used a 100σ detection, shows good agreement with the model. The only image the transient is not detected in is when the transient is completely off. The bottom panel shows when the transient is closer to the noise (10σ); the source is undetected in the first three images and marginally detected ($\sim 2.5\sigma$) in the fourth. The model still fits the measured flux within the errors, but is slightly less accurate.

Furthermore I can solve for the burst flux needed to produce a given image flux (as a function of time) by rearranging equation 4.2.

$$S_B = \frac{S_{Im}\tau}{T} \quad (4.4)$$

I can therefore take the flux of the transient measured in an image and calculate how bright it would have had to be - on specific time scales - to produce the final

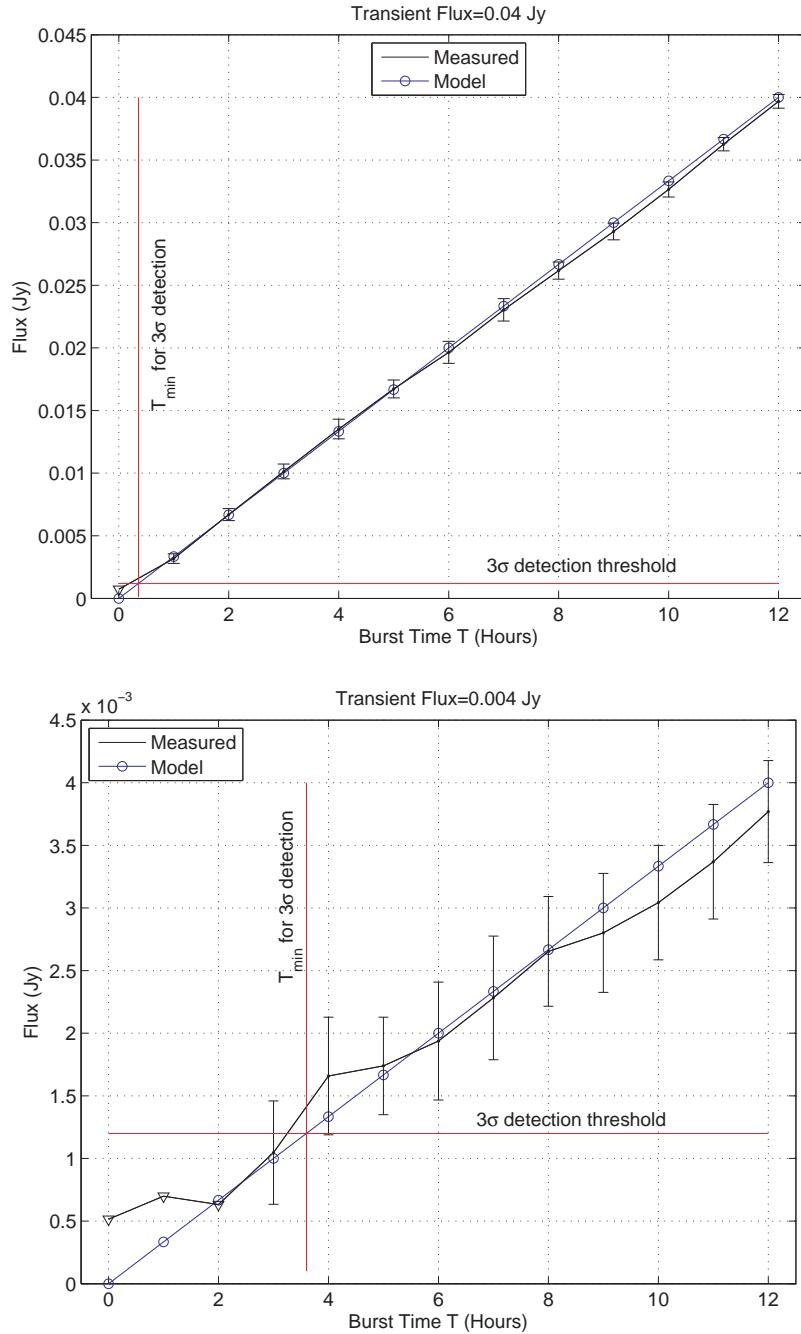


Figure 4.8: Both plots show, on the X-axis, the flux (either measured in the image, or modelled) of simulated datasets that have had a transient source inserted of different durations (Y-axis). The datasets were created using the MIRIAD *uv* simulation tool based on a VLA D-configuration telescope setup at 1.4 GHz, with 10 MHz bandwidth. All 12 datasets that were created had an integration time of 12 hours; a transient source was inserted into each dataset that lasted from 0 to 12 hours. On both plots the horizontal red line shows the 3σ detection limit; the vertical shows the theoretical minimum burst duration to provide a 3σ detection. The top panel shows the simulation with a 40 mJy source inserted, the bottom a 4 mJy.

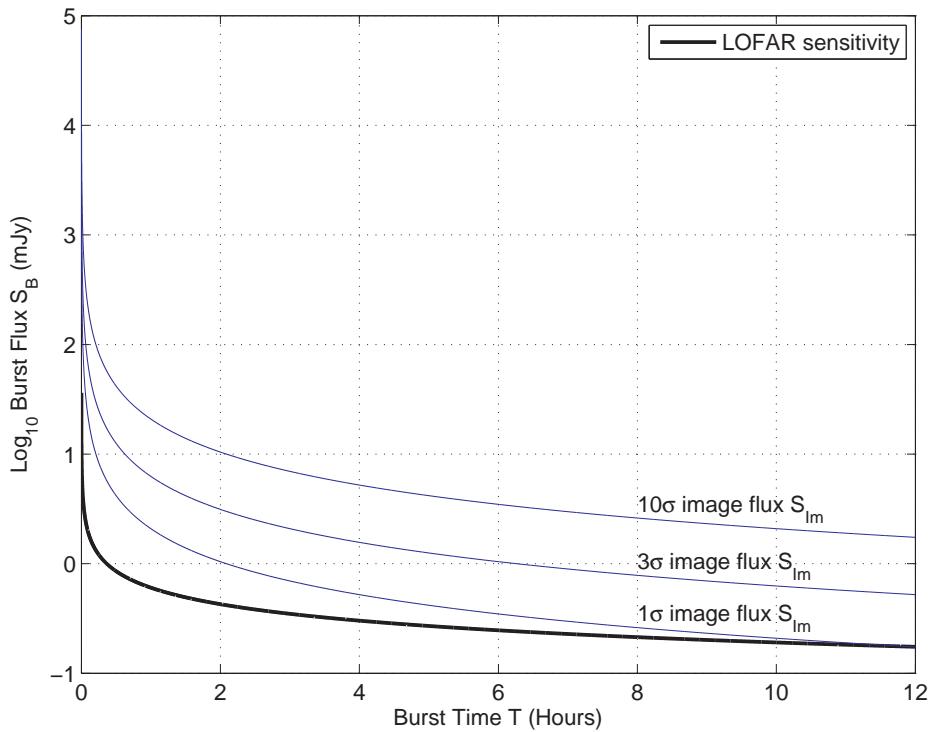


Figure 4.9: Graph showing the burst flux S_B as a function of burst time T needed to produce a variety of final image fluxes S_{Im} . This example is parameterised in terms of LOFAR RMS levels, using 13 core and seven remote stations.

image flux. The motivation for this calculation is that when reporting the flux and duration of a transient to the wider multi-wavelength community, it is important that we do not exclude possible counterparts, by having under-estimated the flux, and over-estimated the duration. Figure 4.9 shows the transient burst flux needed to produce a variety of image fluxes, based on a 12 hour LOFAR observation. The LOFAR noise was calculated via equation 1.6 (Chapter 1). In Figure 4.9 I use the example of 13 core + seven remote stations, with two polarisations and an operating frequency of 150 MHz, with a bandwidth of 4 MHz. The trend in Figure 4.9 shows a typical $1/T$ relationship (see equation 4.4). Note, as the observing parameters change from observation to observation, the LTrAP should be capable of performing the above calculation automatically. This is now coded as a standard data product within the LTrAP for single epoch transients.

4.5 Conclusion

The algorithms presented in this chapter form the basis of the full LOFAR transient detection system. This full system will monitor the LOFAR transient and variable

radio sky and alert the community, in real time, to interesting sources and in certain scenarios, trigger multi-wavelength follow-up. There are a lot of settings within the transient algorithms, some of which have been tested (as part of this thesis) and others which will only be tested in future commissioning. Optimised settings, for example, within the source extractor and the MySQL transients searches will need to be refined and augmented before full operations. I will elaborate on this subject in Chapter 6 when I apply the algorithms to LOFAR commissioning data. In Chapter 5 I will also present how these algorithms have been used to interrogate archival VLA images.

In principle the algorithms that have been developed for LOFAR will potentially be deployed to other instruments. Applying the algorithms to different waveband images may however require different treatment. For example, at GHz frequencies extended sources will be more abundant; image subtraction may need to be applied to remove non-variable extended structure before source extraction. This will mean however that the source extraction algorithms will need the functionality to fit negative peaked elliptical Gaussians within the images. At MHz frequencies, bandwidth smearing, ionospheric effects and wide field imaging artifacts may complicate transient searches further. I will return to some of these concepts in the remainder of this thesis.

In this bright future you can't forget your past.

BOB MARLEY (1945 - 1981)

5

Archival radio transients

In this chapter I present the early stage science which has been produced with the prototype LOFAR Transients Pipeline (LTrAP). Firstly, in section 5.1 I present an automated VLA reduction pipeline which was used to recreate the images presented in the archival radio transient search of Bower et al. (2007) - these images contained a number of transient sources. I then compare the results from source extraction and transient searching using the LTrAP, with those found in Bower et al. (2007). In section 5.2 I present the application of the automated VLA reduction pipeline to a previously unreported X-ray binary outburst of the source Swift J1753.3-0127. This work has been published as part of Soleri et al. (2010). Finally I present a blind radio survey for transients using archival VLA calibrator fields. I discuss the results of the transient search and place upper limits on the snapshot rate of transient events (at GHz frequencies). This work has been published in Bell et al. (2011).

5.1 Testing the LOFAR transient detection pipeline on the Bower field

To test the LTrAP prior to adequate LOFAR images being available, it was decided to use data from a previously published study. The VLA dataset used in Bower et

al. (2007) was selected because a number of radio transients were reported and the data were publicly available in the VLA archive. The aim was to independently reduce the raw data to produce images, then to push the images through the prototype LTrAP. Note, at the time this work was conducted the LTrAP was fairly rudimentary. The source extraction algorithms were operational, however, the database and visualisation infrastructure had not been completed yet. Therefore the testing was more focused on verifying that the source extraction algorithms identified the correct sources, with the correct fluxes, rather than a full algorithmic transient search (using the database etc.).

The field of interest was a relatively blank piece of sky, which had been routinely observed from 1983 to 2005 for the purpose of system calibration checks. Typically three sources were observed per observation; the phase/flux calibrator J1803+784; the North Celestial Pole; and a blank field named ‘OFF784’ - which I will refer to as the Bower field hereafter. The Bower field was observed approximately weekly for around 20 minutes; the field is centred at $\alpha = 15^h 02^m 20.53^s$ and $\delta = 78^\circ 16' 14.905''$ (J2000). The 4.8 GHz (C-band) datasets were chosen for this testing, because seven out of the ten transients were reported at this frequency (see Bower et al. 2007). One transient was reported at 8.4 GHz, the remaining two transients were found in the two month and year averaged images (of the 4.8 GHz data). A total of 626 observations were retrieved from the VLA archive (4.8 GHz) for processing.

5.1.1 Automated VLA reduction script

The ParselTongue package was used to design an automated VLA data reduction script. ParselTongue (Kettenis et al. 2006) is a Python interface to the Astronomical Image Processing System (AIPS; Greisen 2003). It is extremely convenient for this type of processing as it allows two way communication between Python and AIPS. In conjunction with triggering AIPS commands (with specific settings), it can read data tables within AIPS and return the results back to Python. For example, this can be used to read the maximum baseline length of a given observation; Python can then subsequently calculate the optimum pixel/cell size and then trigger imaging with those settings to achieve the best resolution.

A ParselTongue script was constructed to perform the following AIPS tasks, and is summarised in Figure 5.1:

- The dataset was loaded into AIPS.
- The automated flagging routines were applied.

- The visibilities were calibrated using the flux and phase corrections derived from J1803+784 (phase calibrator).
- The calibrated data were imaged and de-convolved with 500 CLEAN iterations and a gain loop of 0.1.
- The primary beam correction was applied to the data.
- The images were then exported from AIPS to disk in FITS file format.
- Each calibrated dataset was concatenated (per-configuration) and imaged once the pipeline had finished processing all the data files.

The procedure was constructed within a Python ‘for’ loop to cycle around multiple datasets. For each VLA configuration I applied the same cell (pixel) size and image size to each observation (per-configuration), so that twice the half-power radius was imaged (see Figure 5.1 for details). In each VLA configuration the maximum resolution (i.e. the cell/pixel size) is set by the maximum baseline length. This is dependent on the available antennas in the outer configuration, which can change. By specifying the same cell and image size for each VLA configuration I could make each image (in that configuration) explicitly comparable.

The automated VLA pipeline, although not being as advanced as that used in later sections, did an excellent job of creating the Bower images with good image fidelity. The typical RMS was $\sim 50 \mu\text{Jy}$: comparable with that quoted in Bower et al. (2007). A number of persistent sources were retrieved in each image, and were consistent in flux with those reported in Bower et al. (2007). To give the reader an intuitive feel for the Bower field, the top panel of Figure 5.2 shows an image constructed from 40 images in D-configuration using the pipeline presented in this chapter. This image has an RMS of $13.8 \mu\text{Jy}$. On the bottom panel I show the deep image (constructed using all VLA configurations) presented in Bower et al. (2007), which has an RMS of $2.6 \mu\text{Jy}$. On both images the half-power (small circle) and twice the half-power (large circle) radii are shown.

After the images were created they were processed through the source extraction algorithms. A detection level of 5σ and an analysis threshold of 3σ was set. The Bower et al. (2007) study set a detection threshold of 5 to 6σ depending on array configuration.

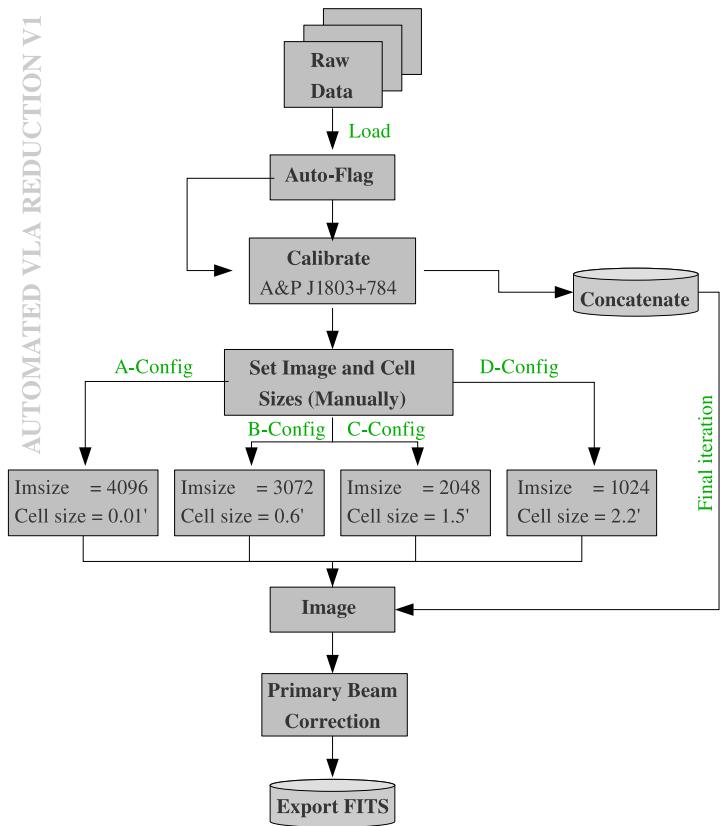


Figure 5.1: Flow diagram showing the automated VLA data reduction pipeline (version 1) used to calibrate and image the archival Bower et al (2007) field.

5.1.2 Results and discussion

The results from source extraction were compared with those reported in Bower et al. (2007) and are summarised in Table 5.1 (the LTrap results have been corrected for primary beam attenuation, in concordance with the Bower results). When examining the results there are two aspects to consider. Firstly, did the ParselTongue pipeline produce the correctly calibrated images? Secondly, did the LTrap correctly fit the transient sources? In the initial testing, the LTrap struggled to pick out the published transients in the field. After further refinement it was found that the RMS background map size had been set too small - approximately 30×30 pixels. The RMS background map size grids the input image into boxes, for which the RMS is measured. If the grid size is set too small, in this case 30×30 pixels, and the source size is around five pixels in diameter, then the noise in that region is large (when compared with that away from point sources) and the transient does not exceed a 5σ detection threshold. A number of the transient sources reported by Bower et al.

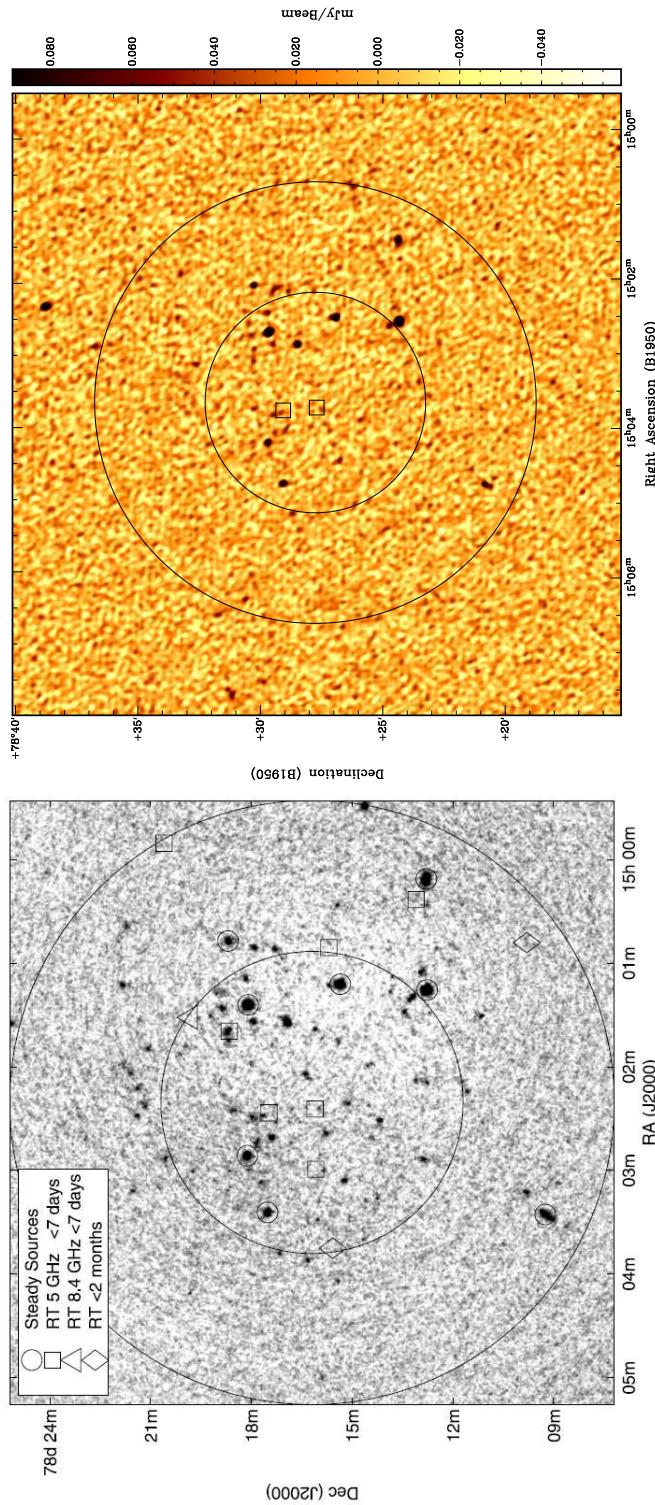


Figure 5.2: Top panel: Averaged image of 40 epochs (at 4.8 GHz) produced using the automated VLA reduction pipeline. The two squares indicate the locations of the transients RT 1986-01-15 (Northern square) and RT 1984-05-02 (Southern square), which are discussed further in this chapter. Bottom panel: The concatenated deep image of all 4.8 GHz data presented in Bower et al. (2007). On both plots the small circle represents the half-power radius, the larger circle represents the twice the half-power radius.

Table 5.1: Summary of the Gaussian fits produced from the LTrAP compared with those reported in Bower et al. (2007).

Transient epoch	VLA Configuration	Bower flux (μ Jy)	LTrAP flux (μ Jy)
RT 1984-05-02	C	448 \pm 74	445 \pm 82
RT 1984-06-13	C	566 \pm 81	480 \pm 61
RT 1986-01-15	D	370 \pm 67	370 \pm 79
RT 1986-01-22	D	1586 \pm 248	Image fidelity bad
RT 1992-08-26	D	642 \pm 101	510 \pm 85
RT 1997-05-28	CnB	1731 \pm 232	1744 \pm 394
RT 1999-05-04	D	7042 \pm 963	See discussion

(2007) were only just above 5σ . Bower et al. (2007) reported that they had not used a RMS background map in their search, therefore I then set the RMS map size to the entire image and a number of the transients were then correctly identified.

In Table 5.1, five out of the seven transients are fitted with reasonable agreement when compared with Bower et al. (2007). In Figure 5.3 I show an image within the region of the transient source RT 1984-05-02 produced by the automated pipeline; on the bottom panel I show the lightcurve constructed using the LTrAP for all images in C configuration. In Figure 5.4 I give a further example of RT 1986-01-15, the lightcurve is constructed using the LTrAP with all images from D-configuration. Both the sources are in good agreement with Bower et al. (2007) and the image fidelity is good.

The observation containing RT 1986-01-22 did not produce a good quality image with the automated pipeline. The data were rereduced by hand, but the image quality was still too poor to detect the source. The data were thoroughly flagged and the calibration and imager settings were consistent with those used in previous pipeline runs. It should be noted that at the time this work was conducted, it was the author's first experience with VLA data. The settings used to make the Bower images were not described in Bower et al. (2007), so potentially a different parameter in the reduction procedure could yield a detection.

The image of the observation containing RT 1999-05-04 did produce a reasonable image, however, the transient was only marginally detected. The source was searched for in the primary beam corrected image of the field, however, it was beyond the correctable beam radius. A Gaussian was fitted in the uncorrected image yielding a flux of $242\pm61\ \mu$ Jy: which is a 4σ detection (see Figure 5.5). As the source was beyond the correctable beam radius, by using the AIPS task PBCOR, I can calculate, by hand, the primary beam correction needed for the VLA using eq.

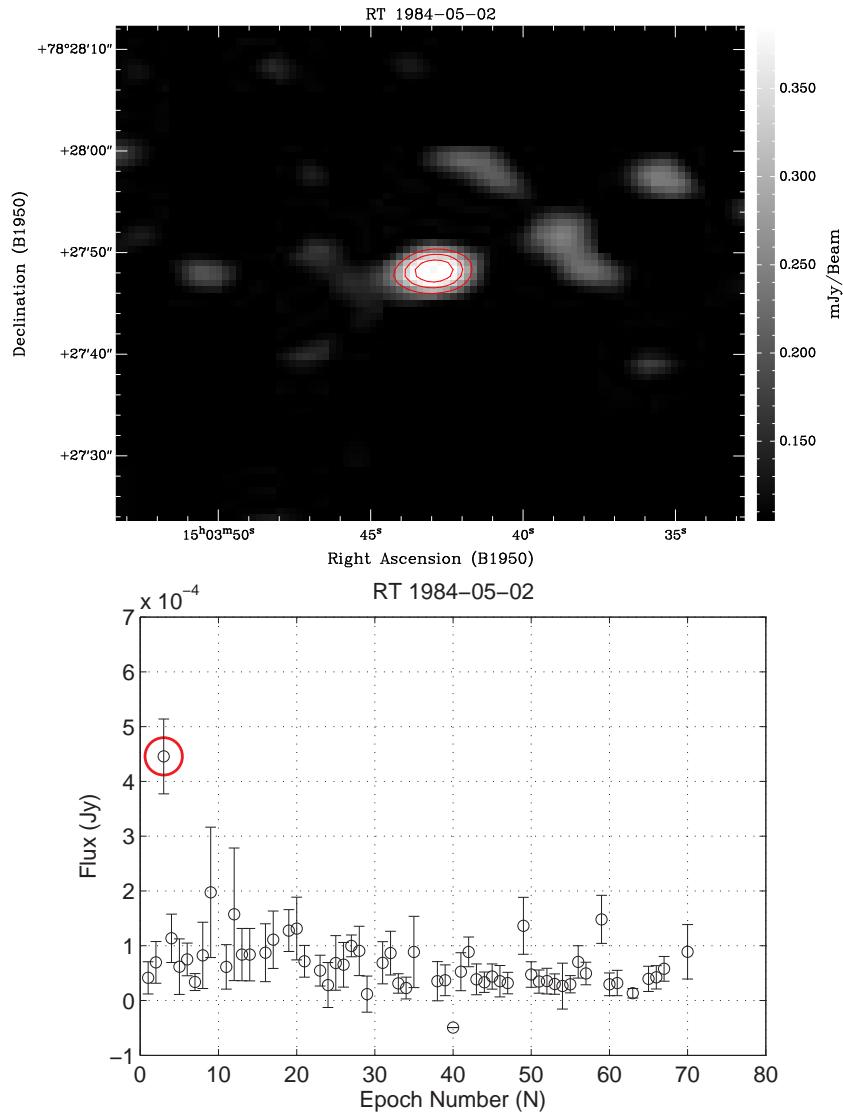


Figure 5.3: Top panel: Image of the transient detection RT 1984-05-02 produced by the automated VLA reduction pipeline (not primary beam corrected). Contours are at 3, 4, 5, 6 and 10×RMS. Bottom panel: Lightcurve of transient produced using the LTrAP source extraction algorithms.

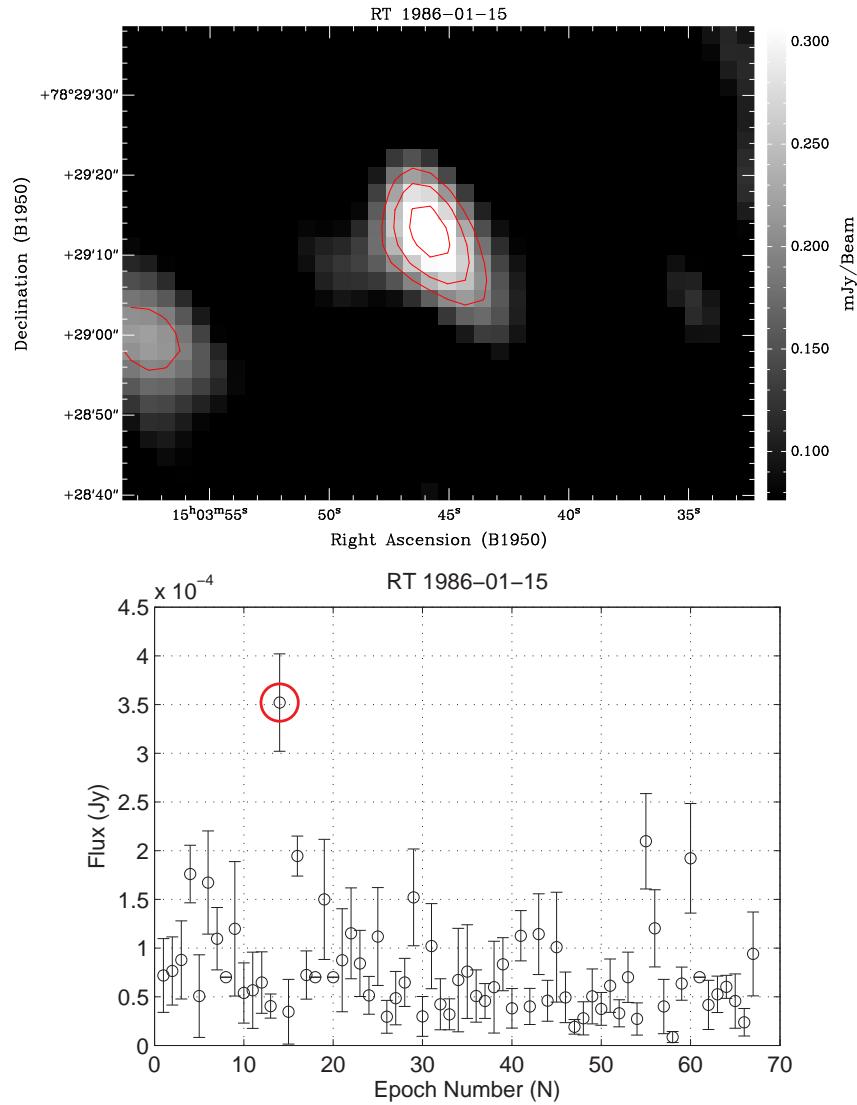


Figure 5.4: Top panel: Image of the transient detection RT 1986-01-15 produced by the automated VLA reduction pipeline (not primary beam corrected). Contours are at 3, 4, 5, 6 and 10×RMS. Bottom panel: Lightcurve of transient RT produced using the LTrAP source extraction algorithms.

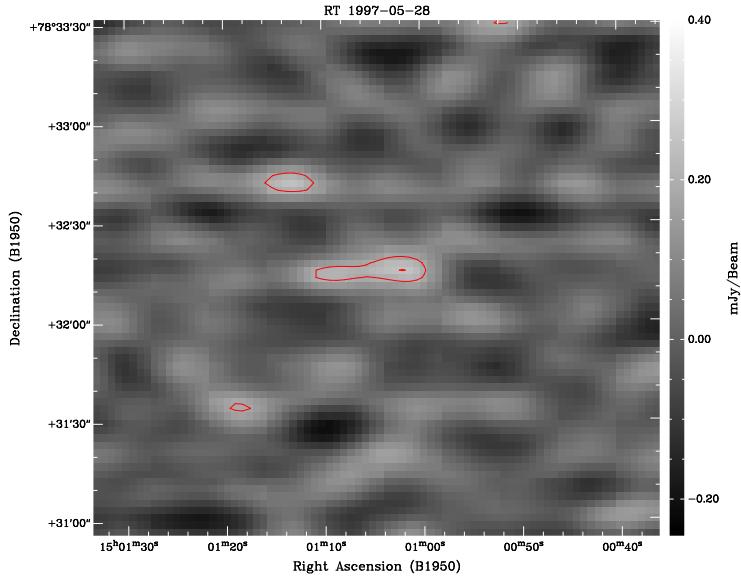


Figure 5.5: Image of the field containing RT 1999-05-04 made using the automated VLA pipeline. Contours are at 3 and 4 σ . This image has not been primary beam corrected.

5.1 and eq. 5.2 below (see footnote¹ for PBCOR reference).

$$f_{PBCOR} \simeq 1 + \frac{X \cdot C_1}{10^3} + \frac{X^2 \cdot C_2}{10^7} + \frac{X^3 \cdot C_3}{10^{10}} \quad (5.1)$$

$$X = \left(\frac{\Delta r}{(arcmins)} \cdot \frac{v}{(GHz)} \right)^2 \quad (5.2)$$

where X is the distance from the source to phase centre (in arc-minutes), v is the frequency in GHz, the constants are defined as $C_1 = -1.372$, $C_2 = 6.940$ and $C_3 = -1.309$ at C-band. The corrected flux can then be calculated using $S_{PBCOR} = S_v / f_{PBCOR}$. Applying this correction to the flux measured in the uncorrected image yields $S_{PBCOR} = -23$ mJy. Applying the correction at the maximum correctable beam radius (measured in the VLA primary beam corrected images), I find $S_{PBCOR} = 32$ mJy which uses a correction factor $f_{PBCOR} = 0.0074$. Applying this correction factor to the flux quoted in Bower et al. (2007) i.e. $S_{PBCOR} = 7042$ μ Jy, I find the transient flux in the Bower uncorrected image would be $S_v = 52.1$ μ Jy - assuming it is at the correctable beam radius. A detection of $S_v = 52.1$ μ Jy would be close to, or below the noise, therefore making it undetectable (typical RMS values are ~ 50 μ Jy). In Bower et al. (2007) they do comment that wide field images were made of all the fields, possibly using a mosaicing mode. Potentially the source was detected in the wide field image and thus the analysis differs from mine. If not, the Bower et al. (2007) RT

¹<http://www.aips.nrao.edu/cgi-bin/ZXHLP2.PL?PBCOR>

1999-05-04 result is incorrect.

5.2 The X-ray binary outburst of *Swift* J1753.3-0127

The X-ray binary candidate source *Swift* J1753.3-0127 was discovered by the hard X-ray Burst Alert Telescope (BAT) on board the *Swift* observatory. The outburst was reported on 2005 May 30 (see Palmer et al. 2005) and a multi-wavelength campaign proceeded on the source. A previously unpublished radio outburst of the source – which was performed in conjunction with *Swift* – was identified by Dr. Paolo Soleri in the VLA archive. Prior to VLA observations the Multi-Element Linked Interferometer Network (MERLIN) reported a detection in the intial fast rise period of the outburst (Fender, Garrington & Muxlow 2005). This VLA dataset was flagged as a good test for both the automated VLA reduction pipeline and the LTraP.

All the publicly available VLA data surrounding the outburst at 1.4, 4.8 and 8.4 GHz were retrieved from the VLA archive (proposal codes AR570, AR572, AR603, AM986, AT320 and S7810). I used a modified version of the automated VLA reduction script presented in the previous section to calibrate and image the data. As the observations were taken in various VLA configurations, and at different frequencies, a variety of different flux and phase calibrators were selected for the observations respectively. To account for this I adapted the automated reduction script to read the sources observed within an observation. I then gave the pipeline a list of the common flux and phase calibrators and allowed it to select them accordingly to perform calibration. Note, the previous version of the script designed for the Bower field only used one flux/phase calibrator, thus the procedure was trivial. After calibration, the images were made using a natural weighting scheme (for best point source sensitivity) and lightly CLEANed with 50 iterations and a gain loop of 0.1. Image self-calibration was not incorporated as the dynamic range of the target source was insufficient. I added a further update to the data reduction procedure to ascertain the maximum baseline length (in each observation), from which, the image and cell sizes were set accordingly. See Figure 5.6 for a flow diagram of the data reduction procedure.

The images were then processed through the LTraP. The outburst was easily detected and the algorithms automatically performed Gaussian fits at the region of interest for all detections above 3σ . Below 3σ the map RMS was used to place an upper limit on the flux densities. As the outburst was unresolved on the longest VLA baseline the peak flux density from the Gaussian fit was used in Figure 5.7

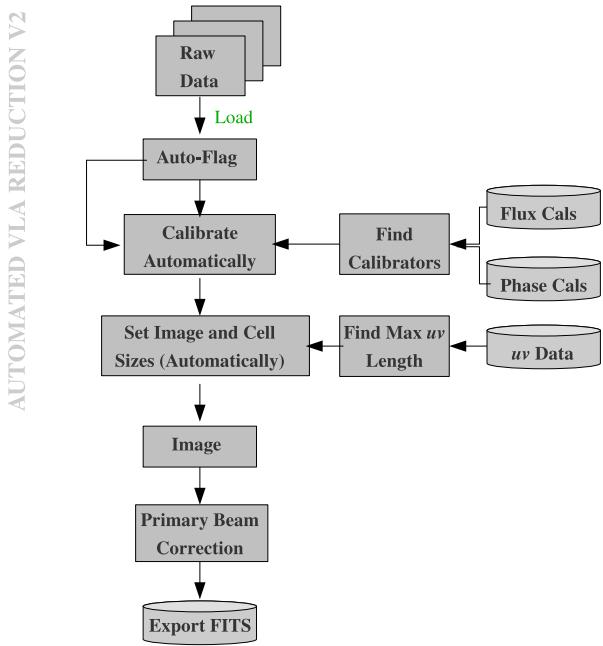


Figure 5.6: Flow diagram showing the automated VLA data reduction pipeline (version 2) used to calibrate the archival VLA observations of *Swift* J1753.3-0127.

(and is shown as it appears in Soleri et al. 2010) to plot the evolution of flux with time. In Figure 5.7, on the top panel, I show all radio data including the VLA points produced through this work. On the bottom panel I show the *Swift*/BAT X-ray lightcurve produced by Dr. P. Soleri and collaborators.

The radio and X-ray data points were used to calculate the position of *Swift* J1753.5-0127 on the fundamental plane, using the scaling relationship presented in Chapter 3. One of the conclusions presented in Soleri et al. (2010) – which is relevant to this thesis – was that the compact jet produced by the source was fainter than expected. Therefore *Swift* J1753.5-0127 was an outlier to the scaling relationship when compared with other X-ray binaries.

5.3 VLA transient survey using the calibrator fields

Bower et al. (2007) with its high yield of transients provided the motivation for this work. The aim was to push archival radio transient studies further, within the framework of testing and refining the transient detection algorithms that will operate

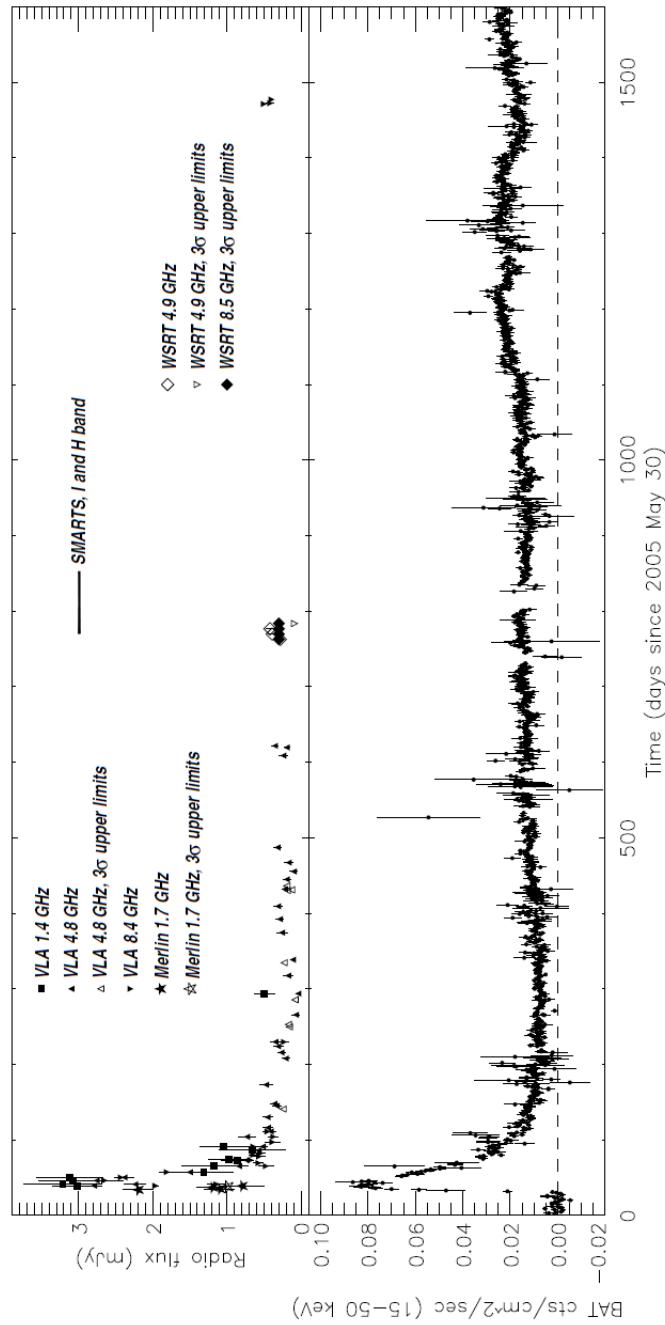


Figure 5.7: Top panel: Radio outburst of *Swift* J1753.5-0127 for the period 2005 May 30 - 2009 October 16. The VLA data points (squares and upward triangles) produced using the automated VLA reduction pipeline and LTrAP track the peak of the outburst from 2005 July 08 until 2009 June 09. Bottom panel: *Swift*/BAT hard X-ray lightcurve of the outburst. This figure was taken from Soleri et al. (2010).

on the LOFAR radio telescope (Fender et al., 2008). I present the findings of an 8.4, 4.8 and 1.4 GHz study of the repeatedly observed flux and phase calibrator fields found in the VLA archive. I explore from some of the publications discussed in the introduction the reported snapshot rates of either detections of radio transients, or upper limits based on non-detections: I place my own upper limit on these values and discuss the implications.

5.3.1 VLA calibrator fields

Extending the Bower et al. (2007) approach to more observations, I have searched a larger part of the VLA archive for transients. In order to optimise the chances of success I searched for the most repeatedly observed fields in the VLA archive. The flux and phase calibrator fields were chosen as the backbone of this new transient study. These calibrators are observed routinely and are a standard and necessary calibration technique in radio interferometry. The calibrator fields usually contain a relatively bright compact object, typically a quasar. The calibrators were selected to fulfil one or more of the following criteria: (A) they were observed frequently; (B) they should be unresolved on the longest A-configuration VLA baseline; (C) they had a relatively large integration time per observation. The chosen VLA fields are summarised in Table 5.2. Note that I initially focused my efforts on the flux calibrators, specifically 3C48; however, because of the bright source in the field and the typically short integration time (~ 1 minute) the images often suffered from artifacts. Therefore I quickly switched my attention to the phase calibrators – which proved, due to longer integration times (~ 5 mins) to have better image fidelity.

Table 5.2: Number of images reduced and searched with respect to observing frequency. Note, 4.8 and 8.4 GHz were the primary frequencies of interest, 1.4 GHz images were only produced for one field. The calibrators are referred to by their J2000 epoch name. $\langle \delta T_{next} \rangle$ is the average change in time between sequential observations (including observation on the same day – see section 5.1). $\langle \tau \rangle$ is the average integration time spent on the calibrator.

Field	R.A. (J2000)	Dec. (J2000)	ℓ	b	# Obs. (8.4 GHz)	# Obs. (4.8 GHz)	# Obs. (1.4 GHz)	$\langle \delta T_{next} \rangle$ (days)	$\langle \tau \rangle$ (mins)
1800+784	$18^h 00^m 45^s.7$	$+78^{\circ} 28^m 04^s.0$	110.0	+29.1	992	908	151	4.3	5.6
0508+845	$05^h 08^m 42^s.4$	$+84^{\circ} 32^m 04^s.5$	128.4	+24.7	205	413	-	13.6	5.1
1927+739	$19^h 27^m 48^s.5$	$+73^{\circ} 58^m 01^s.6$	105.6	+23.5	171	19	-	45.3	9.0
1549+506	$15^h 49^m 17^s.5$	$+50^{\circ} 38^m 05^s.8$	80.2	+49.1	480	558	-	8.3	5.0
0555+398	$05^h 55^m 30^s.8$	$+39^{\circ} 48^m 49^s.2$	171.6	+7.2	123	190	-	27.3	3.5
2355+498	$23^h 55^m 09^s.5$	$+49^{\circ} 50^m 08^s.3$	113.7	-12.0	183	99	-	29.3	$7.6/15.3^a$
3C48	$01^h 37^m 41^s.3$	$+33^{\circ} 09^m 35^s.1$	25.1	+33.4	-	545	-	16.2	4.6
Total					2154	2732	151	$1.9/3.0^b$	5.2

^a 7.6 mins at 8.4 GHz and 15.3 mins at 4.8 GHz.

^b 1.9 days including observations on the same day / 3.0 days disregarding them.

A total of 5037 flux and phase calibrator images have been reduced and searched, with a total observing time of 435 hours. The average integration time spent on all sources was $\langle \tau \rangle = 5.2$ mins. A full statistical description of the measured noise per image is given in section 5.6.

The mean separation between observations regardless of pointing and frequency was $\langle \delta T_{next1} \rangle = 1.9$ days. When calculating this average I have included observations that occurred on the same day. Note that it is quite common for an observation of the same source to be performed at a number of different frequencies (i.e. sequentially within the same observation). This obviously produces a bias that reduces the average time between observations. When producing the images for this survey I only logged the date of the observations, not the exact start and stop time. Extracting the start and stop time was not easily executed within the imaging pipeline framework. I therefore add in an arbitrary time delay of 4 hours in calculating the averages for observations that have the same date. Ignoring observations on the same date yields an average $\langle \delta T_{next2} \rangle = 3.0$ days (regardless of pointing and frequency).

See the top panel of Figure 5.8 for a histogram of the time differences between sequential observations for *all* pointings and frequencies. The bottom panel of Figure 5.8 shows a (summed) histogram of the time differences between observations *per pointing*. Table 5.2 summarises the average time difference between observations and the average integration time per pointing. The lowest average time difference between observations achieved for one single pointing was 4.3 days (1803+784); the highest was 45.3 days (1927+739). As I have sampled a number of different fields at different cadences I state the range 4.3 to 45.3 days as the timescale of transient behaviour that we would be sensitive to. In quoting these numbers I do make the assumption that each image has an equal chance of making a detection – i.e. that there is an isotropic distribution of radio transients and there is no frequency dependence (between 1.4 and 8.4 GHz) for detection.

5.3.2 VLA data analysis

The automated VLA reduction pipeline presented in the previous sections was further modified to calibrate and image the VLA calibration fields. The pipeline was constructed to perform the following tasks (for a flow diagram see Figure 5.9):

- The data were loaded into AIPS.
- The antenna table was searched for antennas that were ‘out’ or designated

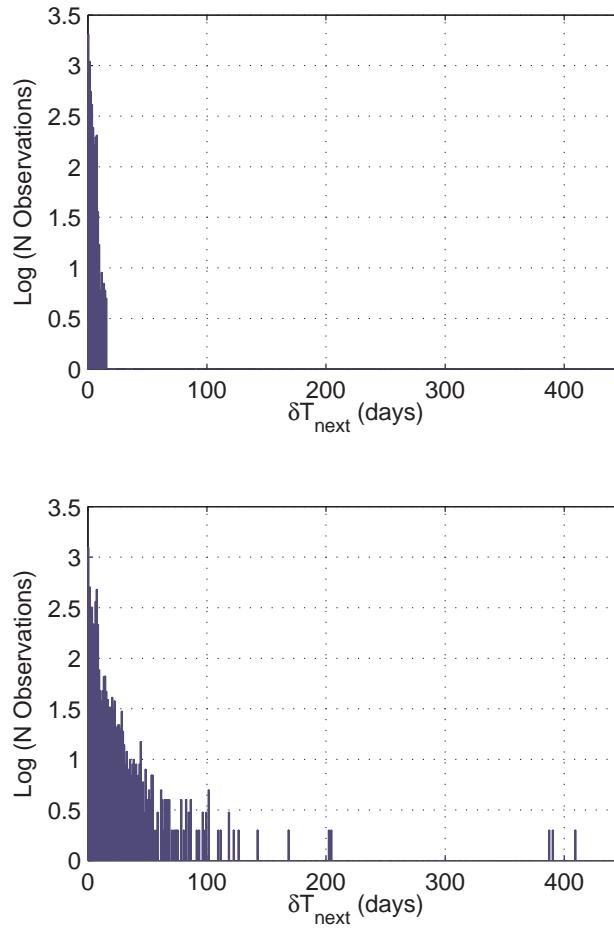


Figure 5.8: Top Panel: Histogram showing the time difference between observations – for *all* pointings and frequencies within this survey. Bottom Panel: A (summed) histogram of the time differences between observations *per pointing* including all frequencies. For a description of the average cadence per pointing see table 1.

‘EVLA’ (only relevant 2006 onwards) during the observations. These antennas were then flagged from being used in subsequent calibration and imaging tasks.

- The antenna table was searched for a suitable reference antenna that was close to the central ‘Y’ hub.
- The maximum baseline length in the observation was ascertained; this was then utilised to optimise cell and image sizes respectively in the subsequent imaging steps.
- The automated flagging procedures were applied to flag unwanted and erroneous visibilities.

- The observation log was searched for *any* known flux or phase calibrator used at the VLA within an entire observation. A comprehensive list of all the flux and phase calibrators was download from the VLA website². This list could be searched with Python and compared with the observation log. For any identified calibrator the pipeline proceeded with the reduction.
- For *any* identified calibrator, standard calibration was performed.
- The calibrator was subsequently imaged and deconvolved, followed by three iterations of phase self-calibration only, and then one iteration of amplitude and phase self calibration. The time interval for phase calibration was set appropriately with respect to the integration time on source.
- The calibrator was boxed off, modelled and removed from the image using UV component subtraction.
- Finally the subtracted image was lightly CLEANed with 150 iterations (Högbom, 1974).

The script was designed to be run on large volumes of data without interruption. Python exception handling was used to catch potential errors and remove bad data from further processing. An example image of 3C48 produced by the pipeline before and after source subtraction is shown in Figure 5.10. Both images show contours and Grey scale to give the reader an intuitive feel of the image quality. The source to the North of 3C48 is persistent with a flux ~ 30 mJy and will be discussed further in section 5.4.

5.4 Transient Search

The images produced using the automated imaging pipeline were then passed to the LTraP. This version of LTraP utilised the full database architecture and toolbar visualisation discussed in the previous chapter. For each image a background RMS map was calculated over the entire image. For any island of pixels above 8σ i.e. eight times the noise measured from the RMS map, source extraction was performed by fitting elliptical Gaussians.

A conservative threshold of 8σ was chosen after initial tests and pipeline refinements indicated that a badly calibrated and reduced image – considering the typical

²Found at <http://www.aoc.nrao.edu/~gtaylor/csource.html>

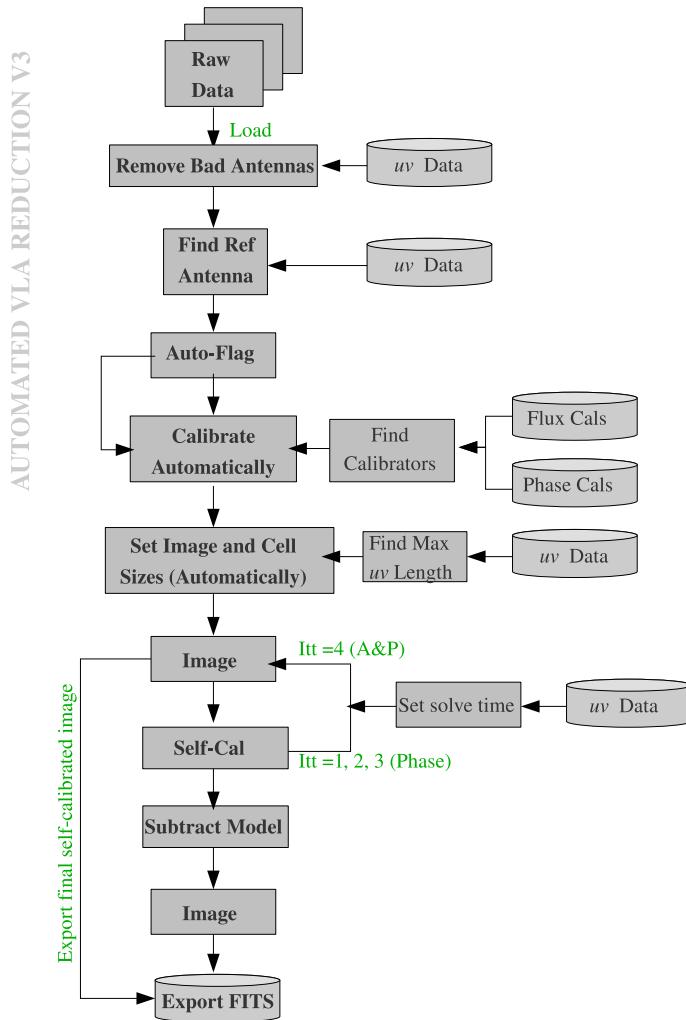


Figure 5.9: Flow diagram showing the automated VLA data reduction pipeline (version 3) used to calibrate the archival calibrator observations.

strength of the calibrator flux – could produce a large number of artifacts and thus false source detections. All images that were processed by the imaging pipeline were processed through the LTrAP – errors included. The rationale for this was to explore the effects of badly calibrated images with respect to source extraction and transient detection. LOFAR will incorporate a false detection rate (FDR) algorithm in the source extraction system, where the global detection threshold for source extraction is set to minimise the number of false positives and is governed by the individual image statistics (Hopkins et al., 2002).

After source extraction the LOFAR transients database (see Scheers 2010) was then populated with the measured properties and associated data of the extracted sources. The source properties include: position and associated errors, all Stokes parameters of peak and integrated flux including the Gaussian fitting parameters.

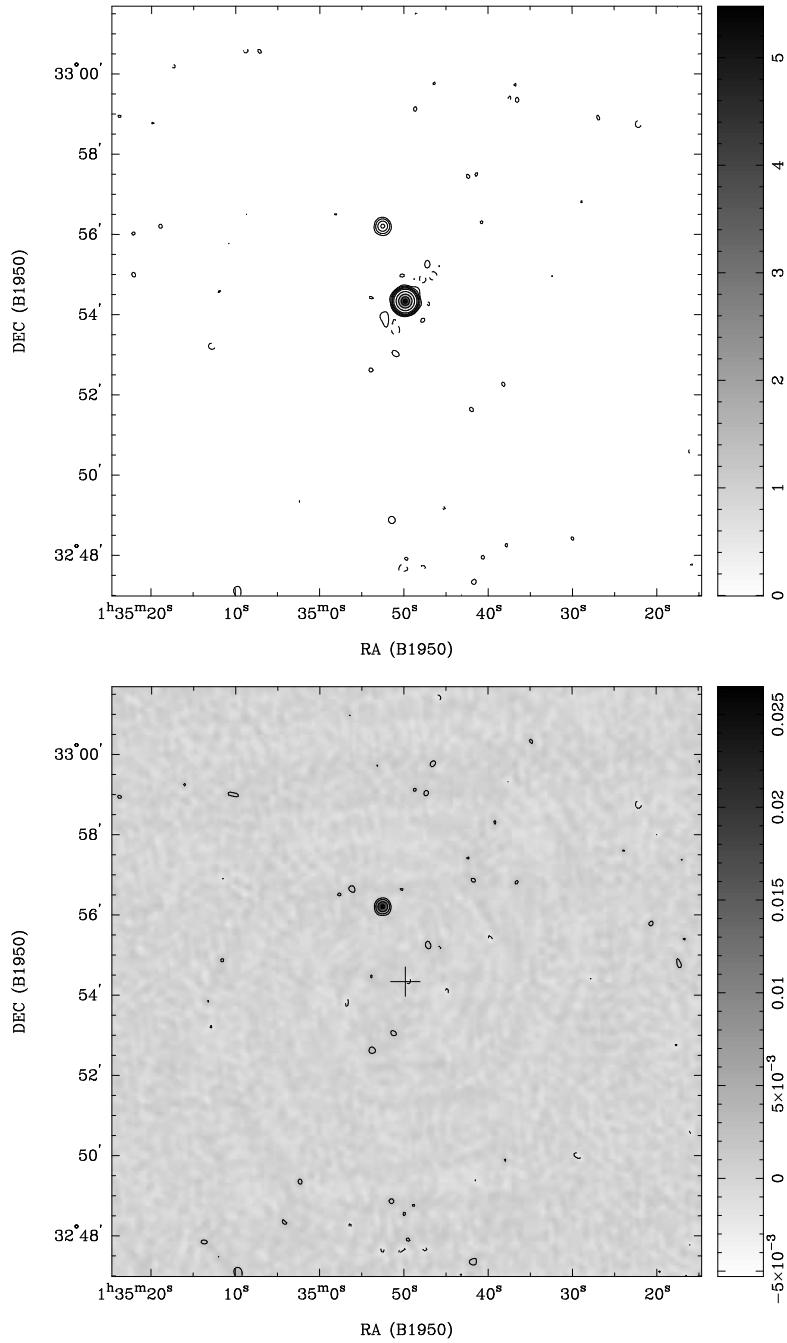


Figure 5.10: An example image of 3C48 at 4.8 GHz produced by the pipeline (observation date 1984-10-05). The top panel shows a CLEANed image before source subtraction. The bottom image shows the same epoch after subtraction of 3C48, a cross denotes the original position of 3C48. The wedge on the right of each image shows the intensity of the Grey scale in Janskys. The integration time was ~ 4 mins, yielding an RMS of 0.6 mJy. Contour levels are $-3, 3, 8, 20, 40, 100, 500, 2000, 5000, 8000 \times \text{RMS}$ for both images.

The associated data included, for example, time of observation, operating frequency and beam properties. The database was searched for either unique sources that had no previously known counterparts (in previous images), or a known source whose flux had varied by a significant amount. I adopt the same metric used in Carilli et al. (2003) to define significant variability as $\Delta S \geq \pm 50\%$, which approximately corresponds to $V_\nu=0.5$ (see Chapter 4).

For the majority of fields in this survey, after the calibrator source was subtracted, the fields were left almost devoid of sources: thus the transient search was relatively trivial. A few \sim mJy radio sources were present in some of the fields and only some of the time, due to changing sensitivity. Therefore, I concentrated my efforts on locating unique sources, rather than characterising the variability of known sources. Once the final list of candidate transient sources had been produced, lightcurves were automatically generated and the images of interest were checked for calibration errors and image fidelity.

5.5 Results

A total of 5037 images at various pointings and frequencies have been searched at a detection level of 8σ . Nine candidate transients that were detected in images with adequate image fidelity were scrutinised further. Four of these candidates were found to lie consistently on the dirty beam. After re-reduction they were shown to be calibration errors which had been cleaned to a point source.

One candidate was detected 76 times at 4.8 and 8.4 GHz and was found to be significantly variable. After careful consideration and review of the literature it was concluded that this source was created by a bug in the VLA recording system, whereby the pointing of telescopes was changed without updating the header information; this error had affected a previously reported transient VLA J172059.9+3852 26.6 (see Ofek et al. 2009 for further details). This error was *not* detected in any of the other calibrator images taken around the same time.

Three of the candidates were found to be associated with known weaker radio sources. These sources were detected in the deepest observations (~ 30 mins) of the respective fields. As observations at this depth were very sparse, these radio sources were considered – by the database algorithms – as transient. These candidates were quickly removed when cross referenced with the radio catalogues. The last transient candidate although surviving some of the re-reduction tests was discarded after it was re-calibrated and the significance level dropped below adequate levels.

The 3C48 field did contain a persistent source at $\alpha = 01^h37^m44^s.2$ and $\delta = +33^\circ11'26''$ (J2000) with $S_v \sim 30$ mJy. This source could *not* be identified in the NVSS catalogue, due to insufficient resolution to separate it from 3C48. The FIRST survey did not cover the position needed to catalogue this source. The source is however previously identified in high dynamic range studies of 3C48 (see Briggs 1995). This source was searched for significant variability but none was found.

5.6 Snapshot rate upper limit

As I have detected no radio transients with this survey, I use the area surveyed per observation, coupled with the typical sensitivity to constrain the snapshot rate of transient events. To calculate the 2σ upper limit of the snapshot rate of transients from our survey I assume a Poisson distribution; for zero detections ($n=0$) I use:

$$P(n) = e^{-\rho N} \quad (5.3)$$

where ρ is the snapshot rate of transients; and the 2σ confidence interval is defined as $P(n) = 0.05$ at the 95% confidence level. N is the sum of the number of images multiplied by the field of view (Ω) at that given frequency i.e.

$$N = (\Omega_{8.4} \times N_{8.4}) + (\Omega_{4.8} \times N_{4.8}) + (\Omega_{1.4} \times N_{1.4}) \quad (5.4)$$

Note, I only consider a search area within the half-power radius per image. Evaluating equations 5.3 and 5.4 yields a snapshot rate of $\rho \leq 0.032 \text{ deg}^{-2}$.

To evaluate the flux density limit that we are sensitive to when searching for transients, we must statistically consider the noise measured in all the images. Figure 5.11 presents this information in two different ways. Firstly, on the top panel I show a histogram of the measured noise σ_m in all the images; included on the plot is an indicator of the theoretical noise in a 5 minute observation (0.16 mJy at 4.8 GHz), and also ten times this value. Secondly – in the bottom panel – I show a histogram of the measured noise divided by the theoretical noise σ_t (i.e. σ_m/σ_t) for all images: thereby taking into consideration that not all observations were ~ 5 minutes in duration.

It can be seen that the bulk of the images achieved an image noise less than $10\sigma_t$. Possible deviations away from the theoretical noise could be attributed to, for example, unremoved RFI, bad calibration solutions, effects of a bright source in the field and, in general, settings and assumptions within the imaging pipeline

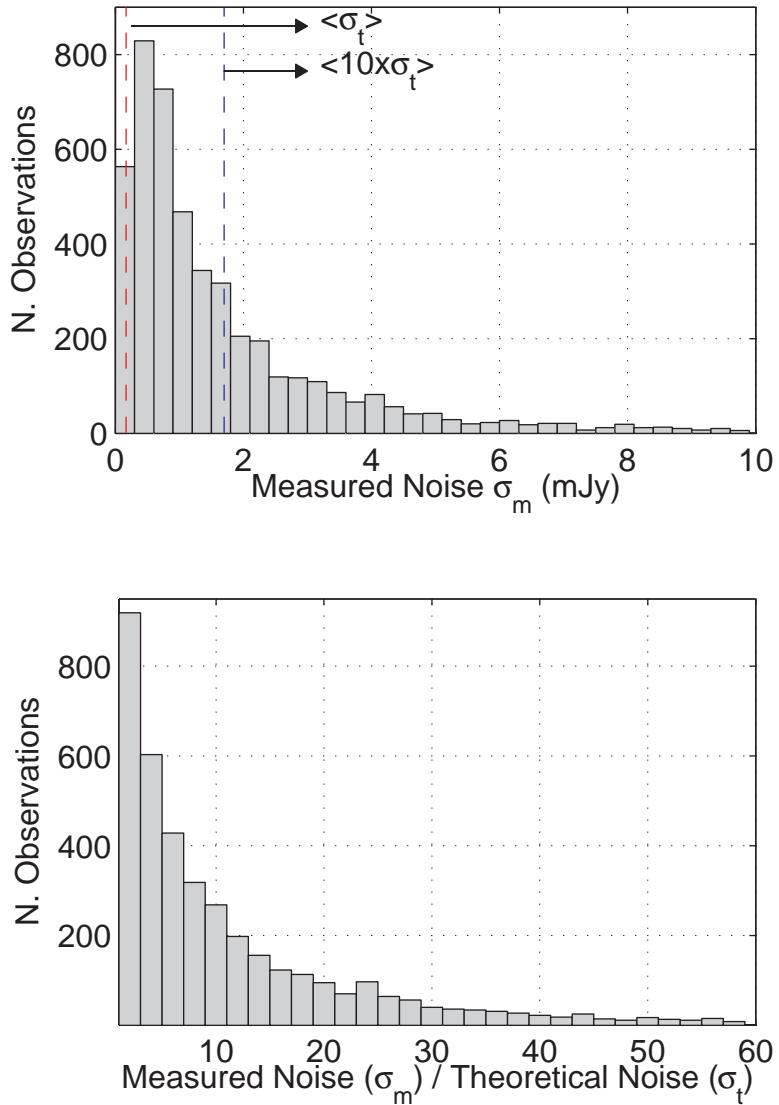


Figure 5.11: Top panel: Histogram showing the measured noise – calculated in the same region – for all images. $\langle \sigma_t \rangle$ gives the theoretical noise calculated from the average integration time for all observations; $\langle 10 \times \sigma_t \rangle$ gives this limit multiplied by 10. Bottom panel: Histogram showing the measured noise divided by the theoretical per observation – which accounts for different integration times.

that do not lend themselves to a given observation. Taking the median value of all the measured noises I find $\sigma_{median} = 1$ mJy (or $6.25 < \sigma_t >$). Using $8\sigma_{median}$ as the global detection threshold for the entire survey, I find that I would be sensitive to transients >8 mJy, with typical timescales 4.3 to 45.3 days.

In calculating this upper limit I have included *all* images reduced by the pipeline; however, not all images were reduced successfully. Although some images contained artifacts, they rarely contaminated the entire image, thus some area could still be searched effectively. If a unique transient point source was detected in an image with errors, the image was re-reduced by hand and checked for reproducibility.

In Figure 5.12 I compare the limit imposed on the snapshot rate of sources from this study, with those found in the literature – this figure is derived from Fig. 9 of Bower et al. (2007) and Fig. 20 of Croft et al. (2010). I do not include a typical timescale of the transient duration in the snapshot rate calculation as it is not well constrained, or easy to represent in a 2D plot. This information is summarised and referenced in Table 5.3 and described further in the following paragraphs.

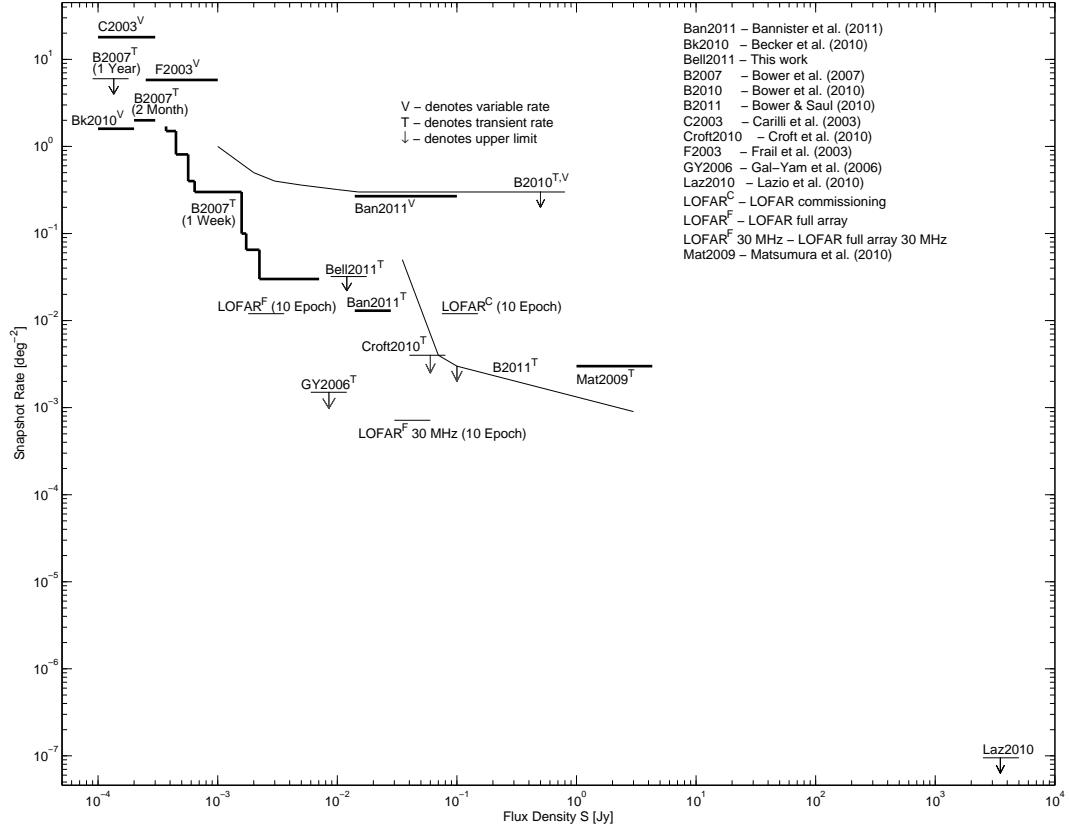


Figure 5.12: Snapshot rate (deg^{-2}) against flux density (Jy) of detections of transients (labelled ‘T’), detections of variable sources (labelled ‘V’) and upper limits based on non-detections (labelled with downward arrows). The thick black line denotes detections; the thin line denotes upper limits. The Bower et al. (2007) 1 week, year and two month limits are indicated as B2007^T with the appropriate time-scale. Ban2011^V and Ban2011^T indicates the separate rates derived for variables and transients reported in Bannister et al. (2011). Bell2011^T indicates the 2σ upper limits derived from this study. ‘LOFAR^F’ indicates the theoretical constraint that LOFAR could provide with zero detections from 10 epochs of 12 hour observations, each of 25 deg^2 fields (using 18 core and 18 remote stations at 150 MHz); ‘LOFAR^C’ indicates the current commissioning capabilities at 150 MHz. ‘LOFAR^F 30 MHz’ shows the rate calculated for a 30 MHz field of view (assuming the final theoretical noise is reached with 18 core and 18 remote stations). Note, that this plot does not contain any information on characteristic time duration and recurrence of transient behaviour as both are currently poorly constrained.

Table 5.3: Summary of snapshot rates reported in the literature. The results are separated out according to upper limits based on non-detections (top); transient detections (middle); and detections of highly variable radio sources (bottom). The flux min column designates the detection threshold of the observations reported in the literature or the minimum flux of detections (indicated as such); the maximum flux is only indicated for transient detections. The Bower et al. (2007) results have been stated three times depending on the characteristic time scale sampled. I do not give the number of epochs for the WJN transients as it is not stated clearly in the literature. Bower et al. (2010) and Bower & Saul (2010) state two different rates depending on flux density, I quote these separately as (A) and (B).

Survey/Paper	Flux Min (μ Jy)	Flux Max (μ Jy)	ρ (deg $^{-2}$)	t_{char}	ν (GHz)	Epochs (N)
This work	>8000 (8σ)	-	<0.032	4.3 - 45.3 days	8.4, 4.8 and 1.4	5037
FIRST-NVSS/Gal-Yam et al. (2006)	$>6000^a$	-	$<1.5 \times 10^{-3}$	-	1.4	2^b
ATATS/Croft et al. (2010)	>40000	-	<0.004	81 days - ~ 15 years	1.4	12^b
Bower et al. (2007)	>90	-	<6	1 year	4.8 and 8.4	17
PiGSS-I/Bower et al. (2010)(A)	>1000	-	<1	1 month	3.1	75
PiGSS-I/Bower et al. (2010)(B)	>10000	-	<0.3	1 month	3.1	75
Bower & Saul (2010)(A)	>70000	-	$<3 \times 10^{-3}$	1 day	1.4	1852
Bower & Saul (2010)(B)	$>3 \times 10^6$	-	$>9 \times 10^{-4}$	1 day	1.4	1852
Lazio et al. (2010)	$>2.5 \times 10^9$ (5σ)	-	$<9.5 \times 10^{-8}$	5 mins	0.0738	~ 1272
Bannister et al. (2011)	14000 (5σ)	6.5×10^6	1.3×10^{-2}	days - years	0.843	3011^b
Bower et al. (2007)	370	7042	1.5 ± 0.4	20 mins - 7 days	4.8 and 8.4	944
Bower et al. (2007)	200	697	2	2 months	4.8 and 8.4	96
WJN/Matsumura et al. (2009)	1×10^6	4.3×10^6	3×10^{-3}	~ 1 day	1.4	-
Bannister et al. (2011)	>14000	-	0.268	days - years	0.843	3011^b
Carilli et al. (2003)	>100	-	<18	19 days and 17 months	1.4	5
Becker et al. (2010)	>100	-	1.6	~ 15 years	4.8	3^b
Frail et al. (2003)	>250	-	5.8	~ 1 day	5 and 8.5	-

^a Different noise values were found in each survey map thus global threshold taken above 6 mJy

^b Combined mosaic

The Bower et al. (2007) survey reported the snapshot rate of transients to be $\rho=1.5\pm0.4 \text{ deg}^{-2}$ (labelled ‘B2007^T 1 Week’ in Figure 5.12) from eight detections, with characteristic timescale 20 minutes $< t_{\text{char}} < 7$ days, above a flux density 370 μJy (with typical image noise $\sim 50 \mu\text{Jy}$ at the pointing centre). Two transients were detected in the 2 month averaged images above a flux density of 200 μJy , giving a 2σ limit on the snapshot rate $\rho \sim 2 \text{ deg}^{-2}$ (labelled ‘B2007^T 2 Month’ in Figure 5.12). No transients were detected in the year long averages above 90 μJy : limiting the 2σ snapshot rate to $\rho < 6 \text{ deg}^{-2}$ (labelled ‘B2007^T 1 Year’ in Figure 5.12).

The PiGSS-I survey using the ATA at 3.1 GHz sets an upper limit on the snapshot rate of transients to $\rho < 1 \text{ deg}^{-2}$ at 1 mJy, and $\rho < 0.3 \text{ deg}^{-2}$ at 10 mJy (labelled ‘B2010^{T,V}’ in Figure 5.12), with characteristic timescale one month (Bower et al., 2010). A recent study by Bower & Saul (2010) - using archived VLA observations of the flux calibrator 3C286 at 1.4 GHz - set an upper limit of $\rho < 3 \times 10^{-3} \text{ deg}^{-2}$ at 70 mJy, and $\rho < 9 \times 10^{-4} \text{ deg}^{-2}$ at 3 Jy (labelled ‘B2011^T’). The work of Bower & Saul (2010) is very comparable to the work presented in this chapter, however, in this study by using predominantly the phase calibrator fields, I slightly push the mean sensitivity down.

The Carilli et al. (2003) study set an upper limit on the rate of highly variable radio sources $\geq 100 \mu\text{Jy}$ to $< 18 \text{ deg}^{-2}$ with characteristic timescales of 19 days and 17 months (labelled ‘C2003^V’ in Figure 4). Note that these were detections of variable radio sources. Frail et al. (2003) derived a comparable quantity to Carilli et al. (2003) of $\rho \sim 5.8 \text{ deg}^{-2}$ with four highly variable sources with characteristic timescale ~ 1 day, above a flux density 250 μJy (labelled ‘F2003^V’ in Figure 5.12). Similar to these surveys but in the direction of the Galactic plane Becker et al. (2010) found 39 variable radio sources between a flux density range 1 to 100 mJy, varying on timescales of years: they derived $\rho \sim 1.6 \text{ Galactic sources deg}^{-2}$ (labelled ‘Bk2010^V’ in Figure 5.12).

In the context of this survey, the difference between a variable and a transient – from a purely observational sense – is a matter of detectability: transients sit, most of the time, below the detection capabilities of the instrument; while variables sit above or close to it. However, the underlying astronomical processes associated with transient and variable sources may differ, and should require a different treatment when considering the rates of events. For example, we might expect the rates of variable sources to differ from ‘one off’ explosive transients such as GRB afterglows – which have a finite lifetime and will be undetectable beforehand. For future surveys, a spectrum of transient and variable behaviour will be observed depending on the cadence and sensitivity. The boundaries between the definitions will become

more blurred as the cadence and sensitivity is increased.

The WJN transients summarised in Matsumura et al. (2009) range in flux density from 1 to 4.3 Jy with characteristic timescale ~ 1 day, yielding a snapshot rate $\rho \sim 3 \times 10^{-3} \text{ deg}^{-2}$ (labelled ‘Mat2009^T’ in Figure 5.12). In comparison the Croft et al. (2010) survey set a 2σ upper limit on the snapshot rate of events > 40 mJy to be $\rho < 0.004 \text{ deg}^{-2}$; by comparing their source fluxes with those in the NVSS catalogues the characteristic timescale is ~ 15 years (labelled ‘Croft2010^T’ in Figure 5.12). The most stringent limit placed on the snapshot rate of sources is set by Gal-Yam et al. (2006) to be $\rho < 1.5 \times 10^{-3} \text{ deg}^{-2}$ for flux densities > 6 mJy (labelled ‘GY2006^T’ in Figure 5.12). Note, the FIRST survey has improved angular resolution (5'') when compared with NVSS (45''), therefore correct source association effects transient identification. I do not state a characteristic time scale for the FIRST-NVSS comparison as both individual surveys took a number of years; specific timescales can only be considered on a source by source basis.

Bannister et al. (2011) set a limit on the snapshot rate of transient sources (calculated from detections) at 0.848 GHz to be $\rho < 1.3 \times 10^{-2} \text{ deg}^{-2}$ above 14 mJy at a variety of timescales (labelled ‘Ban2011^T’ in Figure 5.12). For variable radio sources a rate of $\rho < 0.268 \text{ deg}^{-2}$ is expected between a flux density 14 to 100 mJy (labelled ‘Ban2011^V’ in Figure 5.12).

The upper limit derived from this study is consistent with the detections reported by Bower et al. (2007) – I might have expected of order one detection at our flux density thresholds, assuming that the transient population sampled is isotropically distributed. Bower et al. (2007) did note that an overdensity of galaxies was found within their field. The rate derived from this work is also broadly consistent with that of Bower et al. (2010), Bannister et al. (2011) and Bower & Saul (2010). If the Bower et al. (2007) and Bannister et al. (2011) detections are of a similar nature, then some of the surveys that report upper limits sit very close to the ‘real’ Log N - Log S. This is the best benchmark to date to predict the parameter space a given survey should probe to find transients. However, measurements such as frequency dependence and characteristic timescale of transient behaviour still need to be constrained.

5.7 Predictions for LOFAR

Commissioning observations are currently underway with LOFAR that are probing the parameter space described in this chapter. I include in Figure 5.12 a currently

theoretical upper limit of $\rho < 0.012 \text{ deg}^{-2}$ based on zero detections from 10 epochs of 12 hour observations, each of 25 deg^2 fields at 150 MHz (labelled LOFAR^C). Early observations with LOFAR around August - September 2010 yield a typical RMS of 15 mJy (75 mJy for a 5σ detection which is plotted) with a bandwidth of 48 MHz spread over 247 sub-bands (16 channels per subband). However as more baselines have come on-line, and next-generation data reduction strategies have been implemented, improvements upon this value have been made (and will be discussed further in the next chapter). A more realistic final theoretical noise of 0.36 mJy based on 18 core and 18 remote stations is indicated in Figure 5.12 (labelled LOFAR^F). I also include a theoretical prediction based on 10 epochs of 419 deg^2 fields at 30 MHz (labelled ‘LOFAR^F 30 MHz’). This sets an upper limit on the snapshot rate of transients to $\rho < 0.0018 \text{ deg}^{-2}$ above a detection threshold 30 mJy (6 mJy RMS noise). Of course if the MHz transient population is similar to the GHz population then there will be detections, not upper limits.

Recent work by Lazio et al. (2010) using the Long Wavelength Demonstrator Array (LWDA) – a 16 dipole phased array with all-sky imaging capabilities – have performed an all-sky blind transient search. A total of 106 hours of data were searched for radio transients at 73.8 MHz - the largest survey yet (in imaging mode) at low frequencies. With no detections of radio transients outside of the solar system above a flux density 500 Jy, an upper limit of $10^{-2} \text{ yr}^{-1} \text{ deg}^{-2}$ is placed on the rate of events. With a typical integration time of 5 minutes, this converts to $\rho < 9.5 \times 10^{-8} \text{ deg}^{-2}$; I plot this limit in Figure 5.12 (labelled ‘Laz2010’) assuming a 5σ (2500 Jy) detection is needed. The Lazio et al. (2010) results tell us that extremely bright radio transients with characteristic timescale ~ 5 minutes are very rare, especially at low frequencies. Observations with LOFAR at 30-78 MHz (with 48 MHz bandwidth) would be complementary to the Lazio et al. (2010) survey. Approximately 20 tiled pointings could offer the same solid angle coverage as the LWDA i.e. the whole sky, with increased sensitivity. Pushing into this parameter space on a logarithmically spaced range of timescales is a goal for LOFAR, as well as an all-sky monitoring functionality to catch the brightest and rarest exotica.

We can see from Figure 5.12 that if LOFAR observations were separated \sim weekly, we would be able – via sampling similar parameter space to the Bower et al. (2007) detections – to test the differences between a GHz and MHz population of radio transients. If the emission mechanism for the GHz population is predominantly via the synchrotron process, then many sources will be initially optically thick within the LOFAR band. The rise time for a distant, luminous, population of radio transients – such as GRB afterglows – could be months or years; with lower

peak fluxes (e.g. see van der Horst et al. 2008). Therefore a steep spectrum population of coherent emitters, such as the various flavours of isolated neutron stars, might dominate detections in the LOFAR band. This coherent population will not be limited by the brightness temperature limit of the synchrotron sources and they could also have more erratic cadences (i.e. switch on and off) which should in turn affect the snapshot rate of events. Over ten observations of the same field, with approximately a weekly cadence have already been obtained with LOFAR at 150 MHz - I will discuss the data reduction and results in the next chapter.

5.8 Conclusion

In this chapter I have presented an automated pipeline for reducing VLA data. With the images generated from this pipeline I have used the LTrAP to produce some early stage science. With regard to the VLA transient search, it should be noted that in general for future surveys our transient detection algorithms should be capable of recognising (and flagging) common errors associated with interferometry imaging, hence reducing the number of false detections produced. For example, quality control measurements should be included in the LTrAP that assess and extract measurements from the observation to remove images from further transient searching.

Even in a very simplistic implementation, these could include, measuring the flux of the calibrator source(s) and removing images where the flux had deviated away from the correct flux, or ignoring any image where more than > 100 sources (or any number more than expected) are extracted. In this survey, a number of candidate transient sources were found to lie on the dirty beam structure. Source extraction could be performed on the dirty beam and cross checked for associations in the CLEANed image. Both false and real transient sources can be expected to lie on the dirty beam, however, this information could be used to lower the significance of a given detection in further analysis.

If false detections find their way into the database, they should be systematically removed to avoid chance source associations in future observations. More complex, and computationally intensive interrogations of the database should also be performed to find lower significance detections. Greater consideration should also be given to the automated flagging algorithms. For this survey flagged data was very minimal. For LOFAR, however, large amounts of data may be removed, due to the nature of low frequency RFI.

This VLA calibrator survey did not detect any radio transients and I have placed a constraint on the snapshot rate of radio transients. I have compared this constraint with results from other surveys and although I did not detect any transients it is clear that large volumes of parameter space still remain unexplored. As new surveys push into the sub-mJy regime with various cadences and at different frequencies, definitive transient source populations should become apparent. In general, these populations could potentially populate distinct regions of the Log N - Log S - Log t. By exploring the available parameter space at predefined cadences, new surveys can explore how different populations of transients behave as a function of timescale, flux density and abundance. Therefore with the next generation of radio telescopes such as LOFAR becoming available soon, the problem of inadequate sampling of rare transient phenomena will be alleviated. Due also to the triggering of multi-wavelength followup the classification and interpretation of these events will come of age.

It seems clear that a large population of bright ($>$ mJy), GHz, frequent radio transients does not exist. Therefore the sub-mJy GHz regime is clearly an important part of parameter space to probe for radio transients. The high dynamic range, and wide-field capabilities, of GHz instruments such as APERTIF (APERture Tiles In Focus, wide field upgrade to the Westerbork Synthesis Radio Telescope; Verheijen et al. 2008) and ASKAP make them attractive, potentially high yield, discovery instruments.

Experience without theory is blind, but theory without experience is mere intellectual play.

IMMANUEL KANT (1724 - 1804)

6

Blind searches for transients in LOFAR commissioning data

6.1 Introduction

In this chapter I will present the LOFAR Standard Imaging Pipeline (SIP) which is currently used to flag (removal of radio frequency interference), calibrate and image LOFAR data. I will also present the findings from a transient search conducted using the LOFAR Transients detections Pipeline (LTraP, see Chapter 4), on 14 observations obtained for LOFAR commissioning.

6.2 Test field B0329+54

The Transients Key Science Project has obtained 14 observations, spread over a year in time, of a field centred on the pulsar PSR B0329+54 (phase centre: $\alpha = 03^h32^m59^s.4$ and $\delta = 54^\circ34'43''$ (J2000); Oster & Sieber 1978) using the High Band Array (HBA) component of LOFAR. Repeat observations of the pulsar PSR B0329+54 were the original focus of the LOFAR pulsar commissioning effort; but because high time resolution and image (visibility) data can be obtained simultane-

ously, it was proposed as an interesting field to commission image plane transient searches. Furthermore, the field containing PSR B0329+54 is within the LOFAR zenith monitoring track (which will form part of the routine radio sky monitoring, see Figure 1.6 in Chapter 1), and it also contains a bright source 3C84, which is useful to generate adequate calibration solutions. Additional motivation for the transient search was provided by examining the Log N - Log S produced and reviewed in the VLA calibrator transient study (Chapter 5). By exploring the low frequency component of the Log N - Log S, with a small number of wide field LOFAR observations, we could potentially distinguish between a GHz and MHz population of transient and variables sources. Table 6.1 summarises the observations obtained with LOFAR.

For each observation a total of 247 subbands, spanning a bandwidth of 48 MHz, centred at 150 MHz, are split into either 256, 64, 32 or 16 channels per subband. The volume of data collected is proportional to the number of channels each subband is divided into, therefore using 256 channels per subband significantly increases the data volumes. In contrast, using 16 channels per subband reduces data volumes. The rationale for observing with different combinations of channels per subband, was to test the impact of reducing the volume of data (i.e. less channels), with respect to the science quality of the images. Radio Frequency Interference (RFI) is typically very narrow band, thus with only 16 channels - and a broader bandwidth (per-channel) - more channels must be removed to mitigate RFI. This, in turn, affects the final achievable dynamic range in the images. It is now generally accepted by the LOFAR imaging team that using between 32 to 64 channels per subband, gives the best balance between data volumes and impact on the science quality of the data.

The integration time was typically 12 hours per observation; three observations only had a six hour integration. The array configuration incorporated Dutch core and Dutch remote stations only, no international stations were used. Dutch remote stations are intermediate between the core and international baseline lengths i.e. typically $2 \text{ km} < B_L < 60 \text{ km}$. For example, the observation taken on 2010-Dec-03 used a total of 25 stations, of which 19 were Dutch core and 6 were Dutch remote stations (which yielded 300 baseline pairs). A uv plot and a map of the station locations are shown in Figure 6.1. The uv plot is dominated by the 19 central core stations and the impact of the Dutch remote baselines can be seen. A maximum baseline length of $\sim 24 \text{ km}$ is achieved in the ‘ u ’ direction, and $\sim 30 \text{ km}$ in the ‘ v ’ direction. We can therefore achieve a maximum resolution of $\theta = 8.6 \text{ arc-seconds}$ at 150 MHz, with a field of view of $\sim 5^\circ \times 5^\circ$.

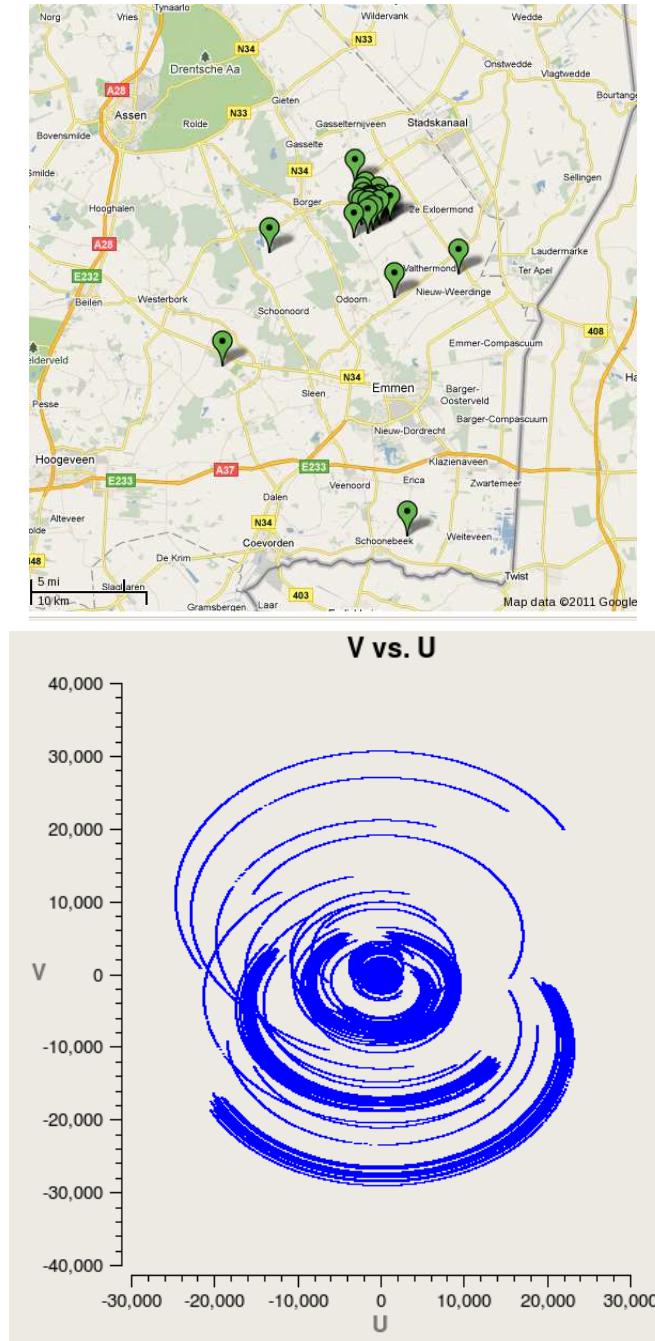


Figure 6.1: Top panel: A map of the station locations used in the 2010-Dec-03 HBA observation. There are five remote stations to the South of the core, and one on the Northern tip. Bottom panel: A uv plot of the 2010-Dec-03 HBA observation (note axes are in metres not λ). The Dutch core-core baselines dominate below 2 km. The Dutch core-remote baselines dominate above 2 km. Credit: The station map was taken from <http://www.astron.nl/~heald/lofarStatusMap.html>; the uv plot was generated using CASA.

Table 6.1: Transient commissioning observations obtained with the LOFAR High Band Array. The # Channels column gives the number of channels per subband. The Image RMS column gives the noise calculated in the same region for each image.

Obs. Date	Obs. ID	Int. Time (Hours)	# Stations	# Channels	Success	Image RMS (mJy)	Comments
2010-04-09	L2010_06928	12	17	256	Yes	21.5	See comment A
2010-08-13	L2010_09851	12	16	16	Yes	20.5	
2010-08-14	L2010_09852	12	16	16	No	—	Image reduction failed
2010-08-15	L2010_09870	5	14	16	No	—	Image fidelity is bad
2010-08-20	L2010_09936	12	14	16	Yes	16.5	
2010-08-21	L2010_09948	12	14	16	Yes	22.9	
2010-08-27	L2010_20033	12	15	16	Yes	16.1	
2010-09-19	L2010_20473	12	14	16	No	—	Image reduction failed
2010-09-25	L2010_20613	12	13	16	Yes	20.4	
2010-10-02	L2010_20673	12	20	16	Yes	7.2	
2010-10-10	L2010_20851	12	15	16	Yes	12.3	
2010-12-03	L2010_21641	12	25	64	Yes	3.5	
2011-03-19	L2011_24435	6	38	32	Yes	7.9	See comment B

^A This observation had the pointing centre shifted to the brightest calibrator source 3C84 (labelled ‘A’ in Figure 6.4).

^B This observation used a three second interval in the integration time, to reduce data volumes. All other observations used one second intervals.

6.2.1 The LOFAR Standard Imaging Pipeline (SIP)

The LOFAR Standard Imaging Pipeline (SIP) is a data reduction procedure that combines the necessary stages to produce images from raw LOFAR data. It is implemented in parallel across the LOFAR offline processing cluster. Note that as LOFAR is in a rapid stage of commissioning, the SIP has, and will continue to evolve to reflect the needs of the telescope and its users. At the time of writing, some standard techniques that are typically used in radio data reduction, such as (phase and amplitude) self calibration loops, were only partially implemented within the SIP. Furthermore, next generation calibration techniques will certainly be needed to correctly image low frequency data.

In particular, direction dependent gains which correct more accurately for amplitude and phase, for all sources (or regions) within the field, must be used. In the analysis presented in this chapter we apply a correction based only on a four source sky model (used in calibration), we then interpolate this correction to sources that are large distances away from the calibrators. This is a typical 2nd generation calibration strategy that is implemented on interferometers such as the VLA. However, the LOFAR images presented in this chapter encompass a much larger field of view (25 deg² compared to 0.018 deg² for the VLA at 4.8 GHz).

The ionosphere is a non-homogeneous time-varying partially ionised layer of gas (Intema et al., 2009). Density fluctuations in the total electron content cause refraction, propagation delay and Faraday rotation of in-coming radio waves (and is more prominent < 300 MHz; Thompson et al. 2001). These effects are a function of antenna location, viewing direction and time, thus the same astronomical source viewed from different Dutch LOFAR stations (and indeed international LOFAR stations) will encounter phase errors. Phase calibration can calculate the corrections needed for any given source, but the solutions will degrade as a function of distance moving away from the chosen phase calibrators phase centre (which is also potentially non-linear). Next generation calibration strategies based on the full implementation of the Radio Interferometer Measurement Equation (RIME; Hamaker et al. 1996) are in progress. Initial tests of the direction dependent gains have been conducted on the datasets described in this chapter. For a further discussion on direction dependent effects see Chapter 1 and also Hamaker et al. (1996) and Smirnov (2011,a,b,c).

The set of observations described in Table 6.1 – even with early stage data reduction techniques – can provide a valuable first look at the low frequency sky, as well as potentially providing early stage science. These observations are also important

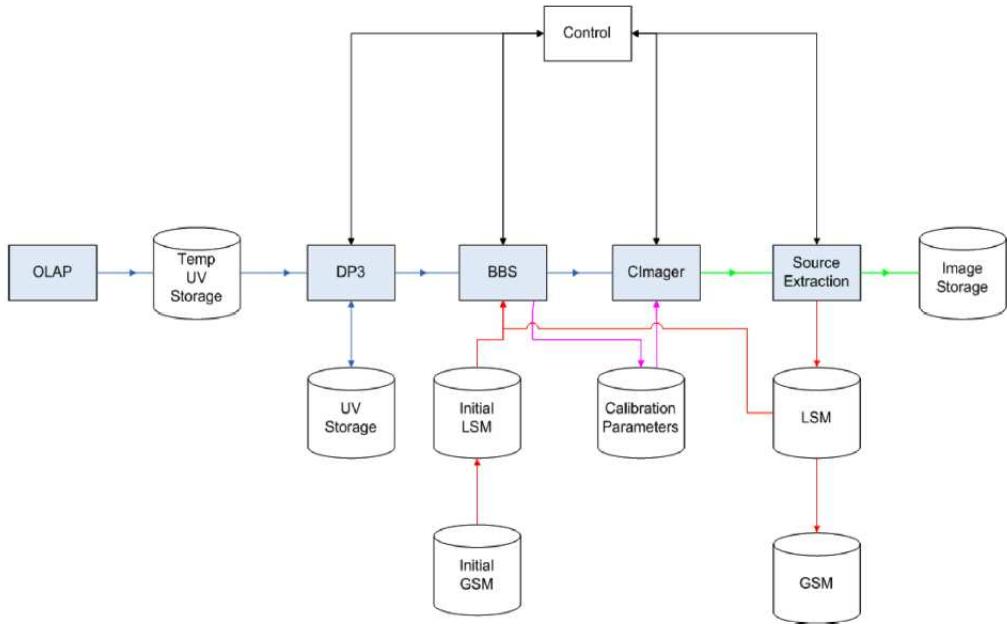


Figure 6.2: A schematic of the steps involved in the SIP taken from the LOFAR imaging cookbook. OLAP is the BlueGene/P central correlator processor. LSM is the Local Sky Model used for calibration. GSM is the Global Sky Model from which a Local Sky Model (which is explicitly used for calibration) can be extracted. DP3 (or New Default Pre-Processing Pipeline NDPPP) is the automated flagging routines. BBS (Black Board Selfcal) is the main calibration stage.

for testing and commissioning the imaging and transients pipeline i.e. software issues are reported back to team leaders. The steps involved in the data reduction using the SIP are described in the following sub-sections, and also in the LOFAR imaging cookbook¹.

To run the SIP a number of parssets, or parameter-set-files, must be defined, these files contain all the necessary settings to run the flagging, calibration and imaging stages. Example parset files, for the data reduction procedure used in this chapter, are given in Appendix A. A schematic showing the procedure is also given for clarity in Figure 6.2 (this figure is taken from the LOFAR imaging cookbook). In the following sections I will take a 64 channel per subband dataset as an example to describe each stage of the SIP explicitly. This exact procedure was also used to produce the image shown in Figure 6.4.

Step 1. Initial flagging and compression

The raw visibilities were flagged using the New Default Pre-Processing Pipeline (NDPPP or NDP³) automated flagging routines; a time window of five seconds,

¹http://www.astron.nl/sites/astron.nl/files/cms/Lofar_Imaging_Cookbook_v5_1.pdf

and a frequency window of five channels were processed in the first pass. The data were then collapsed in frequency from 64 channels down to one channel, the first and last three channels were discarded due to bandpass roll off. The data were then flagged again with a time window of 51 seconds. The visibilities were then collapsed in time by a factor of five, and reflagged with a time window of 20 minutes. This flagging strategy was designed to effectively remove RFI that lasted for different durations in time. NDPPP examines the amplitudes of the visibilities as a function of time and frequency in steps defined by the time and frequency window. If the amplitude of a given visibility is above the median of the amplitudes within that time and frequency window respectively, then they are flagged out. RFI is typically very narrow band, therefore using a frequency window of five channels will remove large amounts of RFI. Flagging the visibilities with different time windows is effective in removing RFI of different durations in time. For example, the flaggers may struggle to remove RFI that is persistently high for 10 mins, when a time window of five seconds is applied, because no visibilities are above the median on five second timescales. See Figure 6.3 for an example time versus amplitude plot before and after flagging. For more information about the NDPPP flagging routines see Offringa et al. (2010).

Step 2.1 Black Board Selfcal (BBS) calibration

The four brightest sources within the B0329+54 field were identified in the NVSS and VLSS catalogues and were used as a Local Sky Model (LSM) to calibrate the observations. These sources are indicated with crosses in Figure 6.4. Note, in Figure 6.4 there is a bright source to the West of the most Northerly calibrator source; this was omitted from the sky model as it sits within the side lobes of the Northern calibrator source 3C84 (which is also the brightest in the field $\sim 40\text{Jy}$ and labelled ‘A’). The fluxes of these calibrators, at 150 MHz, were calculated based on the spectral index of the sources derived from the 74 MHz (VLSS) and 1.4 GHz (NVSS) fluxes (the sky model was generated by Dr. Jess Broderick for this analysis). The spectral indices of the four calibrator sources are: calibrator A = -0.453, calibrator 1 = -0.684, calibrator 2 = -0.830 and calibrator 3 = -0.999. The RA, DEC and spectral index of the four sources are used by the BBS routines to calculate calibration solutions before imaging. In the future, low frequency sky models will be derived from the LOFAR Million Source Sky Survey (MS³), rather than interpolating from other radio catalogues. The MS³ catalogue will define the Global Sky Model (GSM) in Figure 6.2. Furthermore, after the first imaging stage, the local sky model can be updated with CLEAN components from the CLEANed imaged, and this can be used

for further self-calibration.

Step 2.2 Subtraction

A further option in BBS is to subtract a number of bright sources to increase the dynamic range of the final image. Calibrator source A was subtracted in all of the observations.

Step 3. Post calibration flagging

After calibration the data were reflagged to remove calibration errors and previously unremoved erroneous visibilities. A preflag condition was set on all baselines so that any flux greater than 250 Jy would be flagged out. A level of 250 Jy was chosen because no sources within the field could combine through constructive interference to be above this flux. When testing this routine it was found that the later stages of flagging performed better if this preflag condition was set. After this preflag condition was applied, the data were then flagged using timescales of 25 seconds and 2, 4, 6 and 20 minutes. In general, this flagging strategy achieved good RFI mitigation. See Figure 6.3 for an example flagged (bottom panel) and unflagged (top panel) amplitude vs. time plot of one baseline taken from the observation L2010_20673.

Step 4. Imaging

The calibrated visibilities (both subtracted and unsubtracted) were imaged and deconvolved using the Cimager algorithm. A cell size of 40 arc-seconds and an image size of 512 x 512 pixels were chosen to speed up image processing time. A W-projection term is needed in the deconvolution to account for curvature over the large field of view, which slows image processing down considerably. Currently the SIP produces one image per subband; for each image the beam is fitted by the Cimager, with a robust parameter set to zero (see Briggs 1995 for a discussion of weighing schemes). We also tested manually specifying the beam within the imaging steps for each subband. This created unrealistic point sources which proved problematic for the source extraction algorithms, therefore we continued with allowing the beam to be fitted by the Cimager.

Each image was CLEANed using 200 iterations with a loop gain of 0.1 (using the Högbom algorithm; Högbom 1974). I estimated the number of CLEAN components needed using $\frac{S_{max}}{RMS} = (1 - G)^{N_{CL}}$ per source; where G is the loop gain, N_{CL} is the number of CLEAN components needed, S_{max} is the peak flux of the source, and the RMS is the noise in the image (Taylor et al., 1999). To CLEAN the bright-

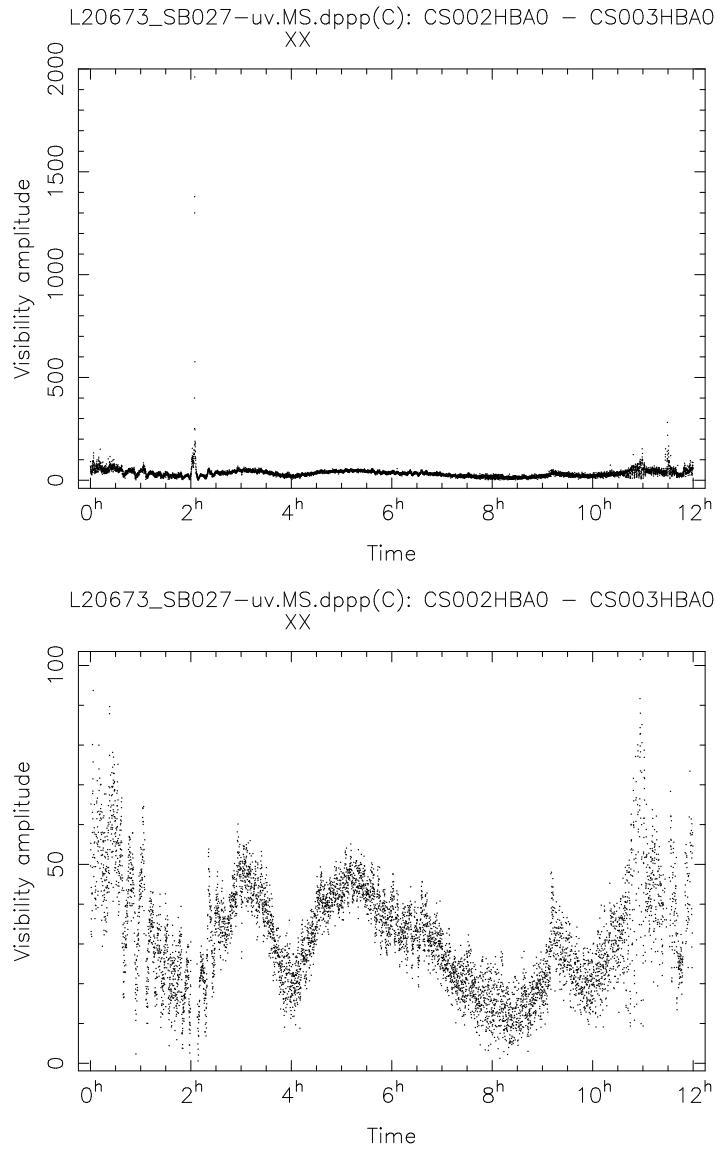


Figure 6.3: Top panel: Post calibration unflagged amplitude vs. time plot for the baseline CS002HBA0 - CS003HBA0 (y-axis is in Janskys), XX polarisation is plotted only. Bottom panel: Flagged amplitude vs. time plot. Around two hours into the observation a bright < 2000 Jy RFI spike can be seen, lasting around 10 minutes. This is correctly flagged out by firstly applying the preflag condition that removes all visibilities above 250 Jy; and secondly by applying a time window that is sensitive to RFI of 10 minutes in duration (and other timescales). Note, the brightest source in the field is ~ 40 Jy. Credit: These plots made use of `uvplot.py` written by Dr. George Heald.

est source (~ 40 Jy) in one subband (RMS ~ 50 mJy) with a loop gain of 0.1, 73 CLEAN iterations would be needed. There are typically four sources per subband per image (three of which have lower peak fluxes); if they are all at 40 Jy, 292 CLEAN components would be needed: therefore 200 CLEAN iterations is a reasonable amount.

The final images of each subband produced by the SIP are averaged together in the image plane, using a weighted average based on the RMS of each image. Note, the beam size varies for each subband across the entire bandwidth. Ideally all UV data would be combined together before imaging and multi-frequency synthesis imaging (Sault & Wieringa, 1994) would be used. However, the data volumes are currently too large to implement this. When LOFAR transitions to producing images on timescales less than \sim ten seconds, it may be possible to perform multi-frequency synthesis imaging as the data volumes will be greatly reduced. The consequence of having to estimate the beam properties in the final averaged image, has a knock on effect for effective transient searching. It makes the Gaussian fits often erroneous, and yields uncertainties in the positional accuracy of the fits. Also, artifacts and sidelobes that would typically be removed - in an image combined in the *uv* plane - by the CLEAN algorithm, are only marginally removed, and become more prominent when the images are averaged together (in the image plane).

6.3 Imaging results and transient search

For all observations reasonable image fidelity was realised, and the sky remained relatively ‘static’ from image to image i.e. the majority of bright sources detected in one image, matched the sources in the next (and so on). This was reassuring from the point of view of calibration and imaging. However, some bright artifacts were typically found around brighter sources, perhaps due to simplifications in the calibration and imaging applied. A typical noise of ~ 15 mJy was achieved (with all subbands) from observations between August to October. Due to more baselines coming on-line, improved station calibration and 64 channels per subband a noise of 3.5 mJy was achieved from the December observation. The station level calibration adds extra stability to the flux scale, and is independent from the calibration applied in the SIP. The RMS values for all images are summarised in Table 6.1. An example image of the December observation (observation id. L2010_21641) using all subbands is shown in the top panel of Figure 6.4. The four sources used in the BBS calibration procedure are indicated with crosses (and labelled 1, 2, 3 and A).

On the bottom panel I show an image of one subband only (subband 20): this image has an RMS of 83.9 mJy

To assess the stability of the flux calibration from epoch to epoch the flux of the calibrator sources (in the final combined images) were measured. Note, in the data reduction procedure the brightest source (3C84, labelled ‘A’) is subtracted from the *uv* data, I therefore only measure the fluxes of the three remaining calibrators. These three sources are labelled 1, 2 and 3 in Figure 6.4. In Figure 6.5 I show a lightcurve of these sources, the grey dotted line denotes when station calibration was applied to the visibilities. Before station calibration is applied, the calibrator sources show a large amount of variability; afterwards there is still some variation but it is slightly more prominent for calibrator 3. The index of variability V_v can be calculated for the three calibrator sources which yields: calibrator 1 = 0.11, calibrator 2 = 0.19 and calibrator 3 = 0.14. Note, I only calculate the index of variability for the fields which have the same pointing centre. The first observation of this field (L2010_06928) had the pointing centre placed on calibrator source A. The primary beam attenuation of LOFAR is currently not well characterised, therefore a change in flux would be expected with respect to a shift in the phase centre, thus this observation was omitted from the analysis. The instability in the absolute flux scale has implications for effective transient searches. The change in flux of the calibrators will introduce a variability to all sources in the field, and must be considered when searching for transient and variable sources.

The images produced by the SIP (which had the same pointing centre) were then processed through the LTrAP. A source extraction threshold of 5, 10 and 15σ was used (in different runs) to locate sources within the images - see Chapter 4 for discussion of the full LTrAP procedure. A number of difficulties were found with source extraction as the beam properties were not recorded in the final combined FITS headers. I approximated the beam properties by using an average semi-major axis, semi-minor axis and position angle over the total bandwidth, for each observation respectively. The Gaussian fits derived from source extraction occasionally had unphysically large errors, and integrated fluxes as high as 10^6 Jy. This is potentially a consequence of making approximations of the beam properties, and combining data in the image plane. It is an ongoing area of active research to find the optimum beam properties, and in general, the best data reduction strategy that complements the source extraction routines.

After processing the images through LTrAP, they were first inspected by eye, and then algorithmically, with respect to the archival WENSS image of the field. The WENSS survey had an RMS of ~ 4 mJy at 325 MHz, therefore comparing the LO-

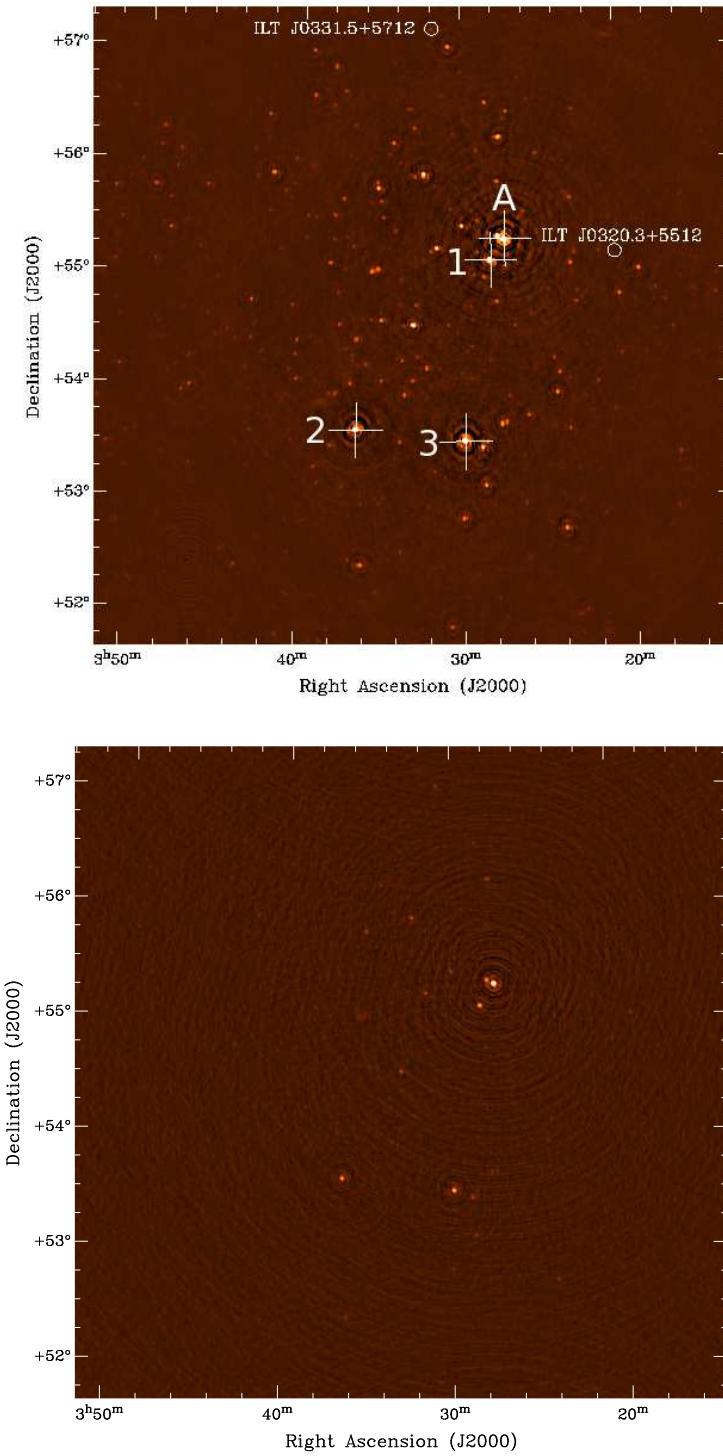


Figure 6.4: Top panel: LOFAR observation of the field 0329+54 taken on December 3rd 2010. The image is centred on the pulsar PSR 0329+54. The four bright sources used in the calibration sky model are indicated with crosses. The transient candidates J0320.3+5512 and J0331.5+5712 which are discussed in this chapter are indicated as such. An RMS of ~ 3.5 mJy is achieved in this image using 19 Dutch core and 6 Dutch remote stations, and all subbands. The calibrator source labelled ‘A’ is typically removed in the imaging stages. Bottom panel: An image of the same field with one subband only (subband 30), this image has an RMS of 83.9 mJy

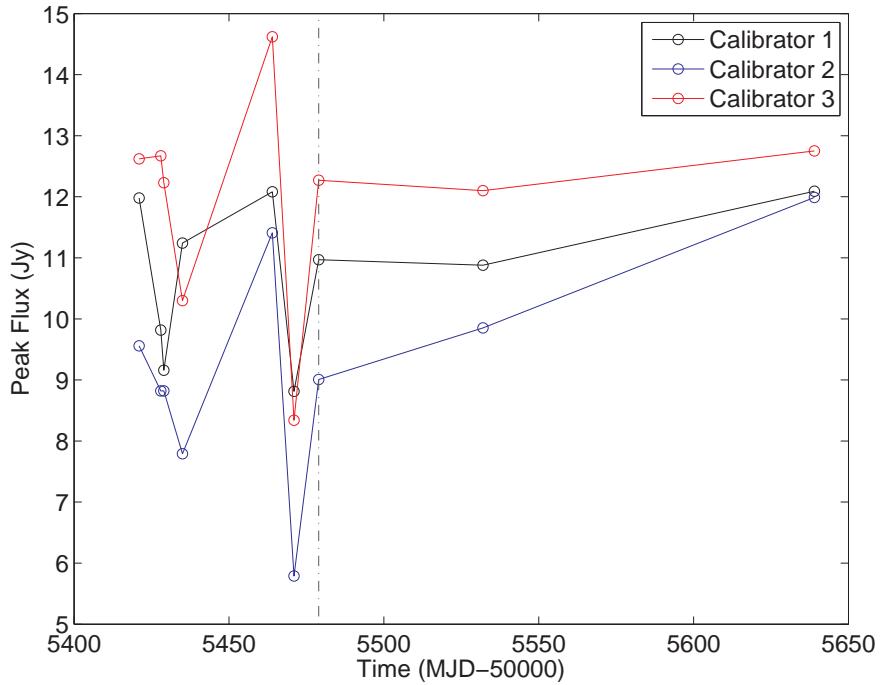


Figure 6.5: Lightcurve of the three unsubtracted calibrators (labelled 1, 2 and 3 in Figure 6.4). The black dashed line indicates when station calibration was applied, which adds extra stability into the flux scale. The indexes of variability for each source respectively are: calibrator 1 = 0.11, calibrator 2 = 0.19 and calibrator 3 = 0.14. Before station calibration is applied the indexes of variability for each source respectively are: calibrator 1 = 0.14, calibrator 2 = 0.21 and calibrator 3 = 0.18. After station calibration is applied they are: calibrator 1 = 0.06, calibrator 2 = 0.14 and calibrator 3 = 0.03

FAR and WENSS images should yield very similar sources (above $\sim 10\sigma$): some steep spectrum sources may only be detected in the LOFAR images close to the RMS. In Figure 6.6 I show a comparison between the WENSS and LOFAR sources. The red circles indicate the locations of the LOFAR sources which were detected in more than two epochs, above 5σ . The green circles indicate the WENSS associations. I have overlayed this comparison on the LOFAR 2010-Dec-03 observation (L2010_21641); this observation had the greatest dynamic range when compared with the other observations. Some of the weaker sources in this image were *only* detected in this image (and not the previous images), due to the change in sensitivity, so they are not circled in red and green (I set the query to exclusively return source that were detected in two or more epochs). See Figure 6.7 for the same comparison but with the 2010-Oct-10 LOFAR observation. The comparison between the 2010-Dec-03 observation and WENSS image by eye does show good agreement. Using the LTrAP we can produce this type of figure for each image processed through the system for an accurate comparison, and to identify sources which are not in the ma-

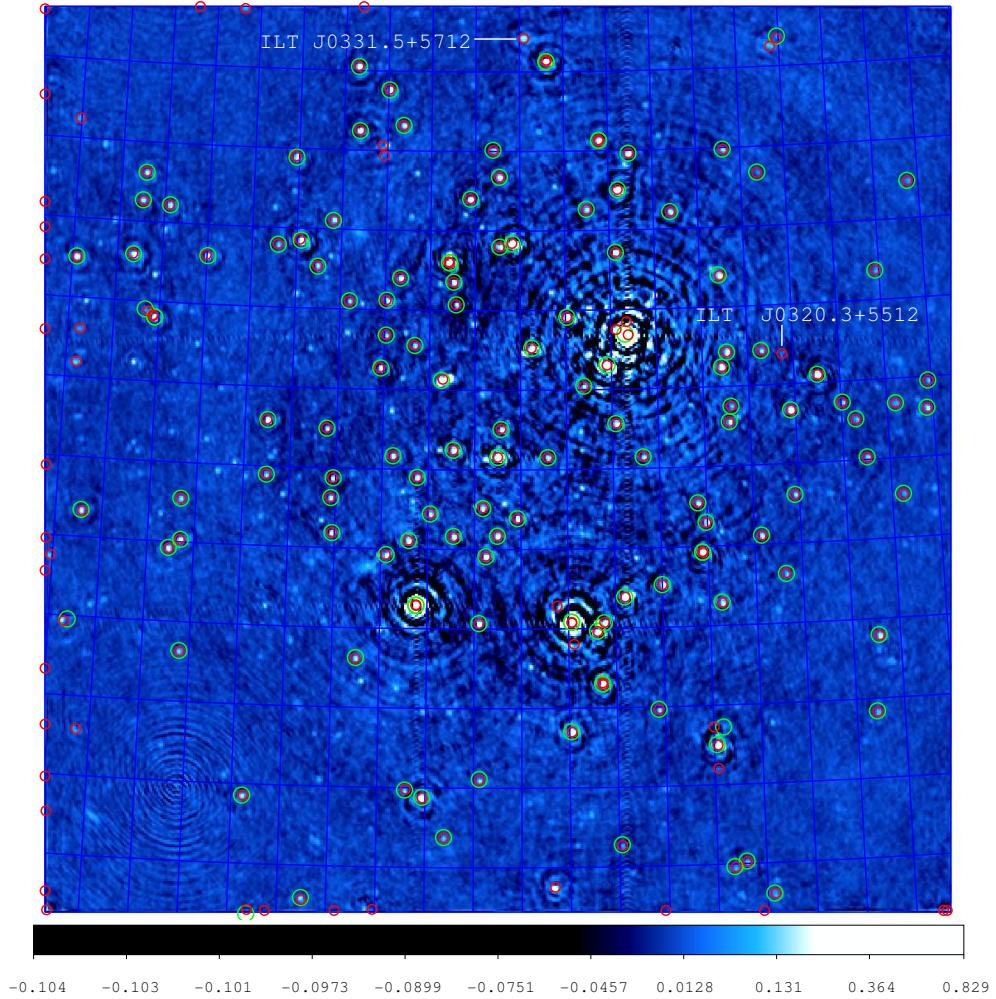


Figure 6.6: An image of the 2010-Dec-03 LOFAR observation (L2010_21641). The red circles show the sources detected in two or more LOFAR images, above 5σ . Green circles denote the WENSS associations. Note, the 2010-Dec-03 LOFAR observation had the best dynamic range of all the images, therefore some of the sources within this image, are only detected in this image only (see Figure 6.7 for the same comparison but for a different observation). The sources detected in this image only are not identified on this plot. It is possible to compare the catalogued and LOFAR sources per image, which will be implemented in the future. Two transient candidates are indicated in this image and will be discussed in later sections. Further candidates (indicated with red circles only) are currently being investigated.

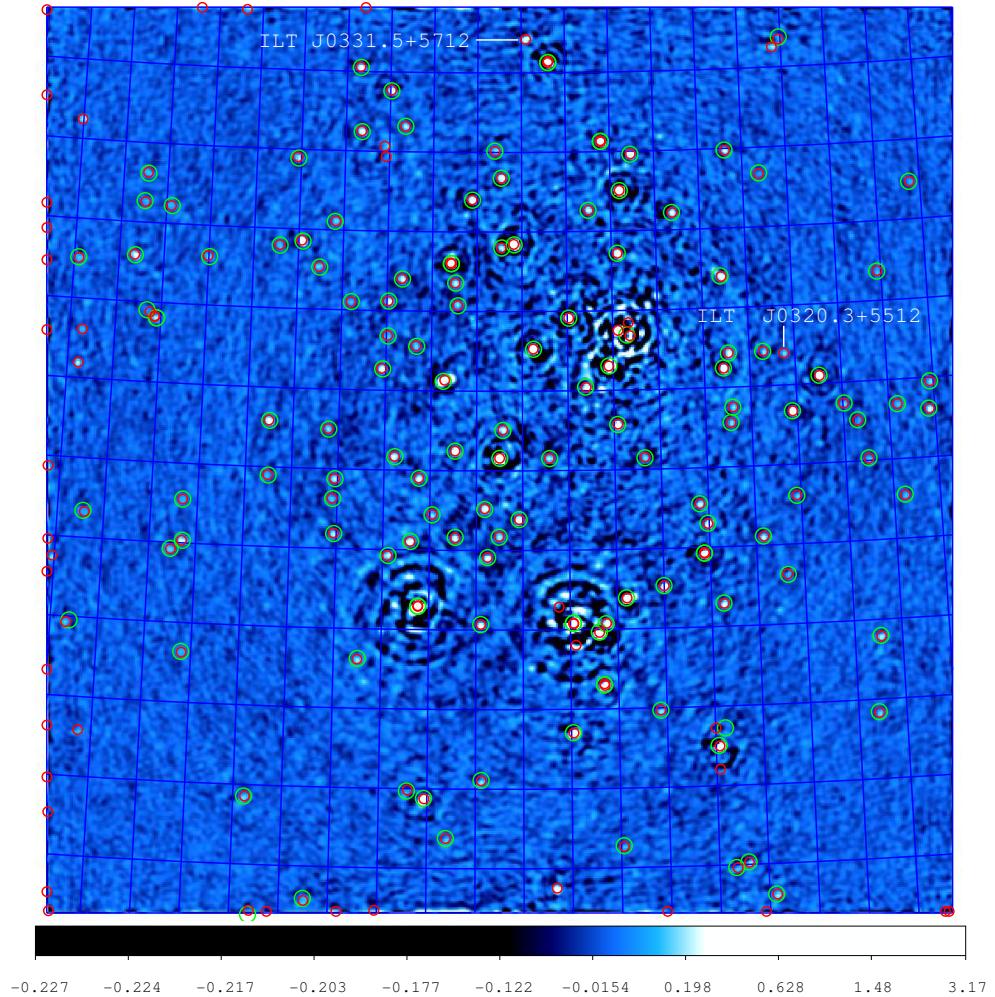


Figure 6.7: An image of the 2010-Oct-10 LOFAR observation (L2010_20851). The red circles show the sources detected in two or more LOFAR images, above 5σ . Green circles denote the WENSS associations. Less sources are detected in this image because a smaller number of stations were available for the the observation.

major radio catalogues (including NVSS, VLSS etc). A number of sources in Figure 6.6 are only circled in red, indicating that they were only detected in the LOFAR images and not the WENSS catalogue. A large number of these sources sit on the edge of the image, these are easily removed using a margin in the LTrAP, so that 50 pixels around the edge of the entire image are ignored. The LTrAP struggles to fit elliptical Gaussians at the edge of images because it cannot directly derive some of the fitting parameters, without the relevant pixels (i.e. off the edge of the image).

In Figure 6.8 I show the associated WENSS flux against the averaged extracted LOFAR flux. The red line shows the trend that would be expected from a one-to-one agreement between the LOFAR and WENSS fluxes. The dotted blue line shows the

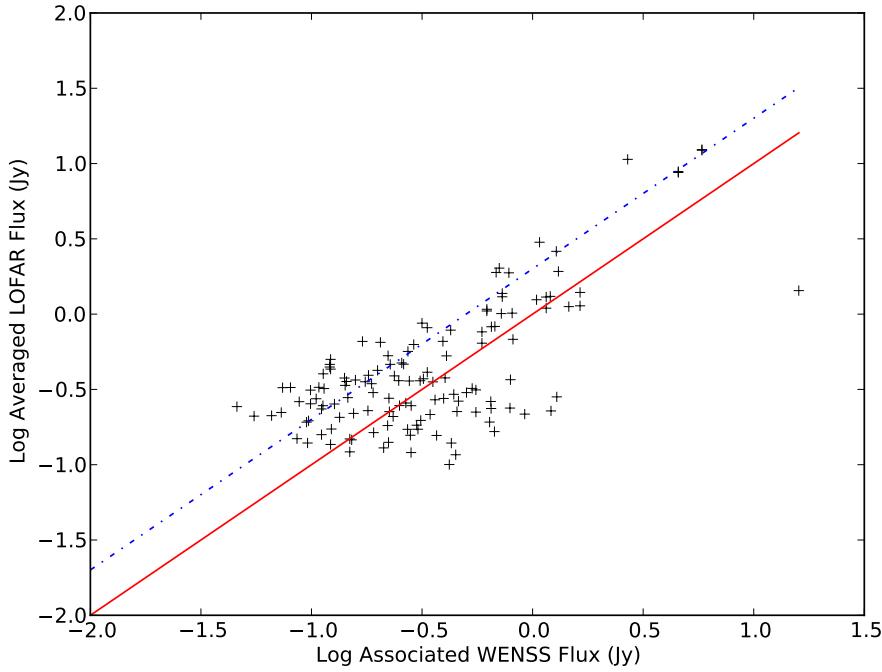


Figure 6.8: Comparison of the catalogued WENSS fluxes with the associated averaged LOFAR source fluxes. The red line shows the trend that would be expected from a one-to-one agreement between the LOFAR and WENSS fluxes. The dotted blue line shows the expected trend from a two-to-one LOFAR-to-WENSS flux relation. The expected increase in flux from 325 to 150 MHz for a source with $\alpha = -0.7$ is \sim two. A large amount of scatter exists because the LOFAR fluxes have not been corrected for primary beam attenuation, or a change in spectral index between 325 MHz (WENSS) and 150 MHz (LOFAR). In future LOFAR operations this type of diagnostic plots can be generated automatically to search for variable sources. The sample of sources used to make this plot are shown in Figure 6.6 with red and green circles.

expected trend from a two-to-one, LOFAR-to-WENSS flux relation. The LOFAR flux has not been corrected for primary beam attenuation or the change in flux due to the spectral index. Both these effects - as well as the uncertainty in absolute flux calibration - will introduce a scatter in this plot. Once the corrections above have been implemented within the LTrAP, using these plots we can identify variable sources with respect to the WENSS and other major radio catalogues.

It was noticed that the positional accuracy of the LOFAR sources, with respect to the WENSS sources, degraded moving away from the pointing centre. This was more prominent in the early observations, but improved in the observations from December (L2010_21641) onwards. In Figure 6.9 I show a region to the edge of the December (L2010_21641) LOFAR field, with the red crosses indicating the lo-

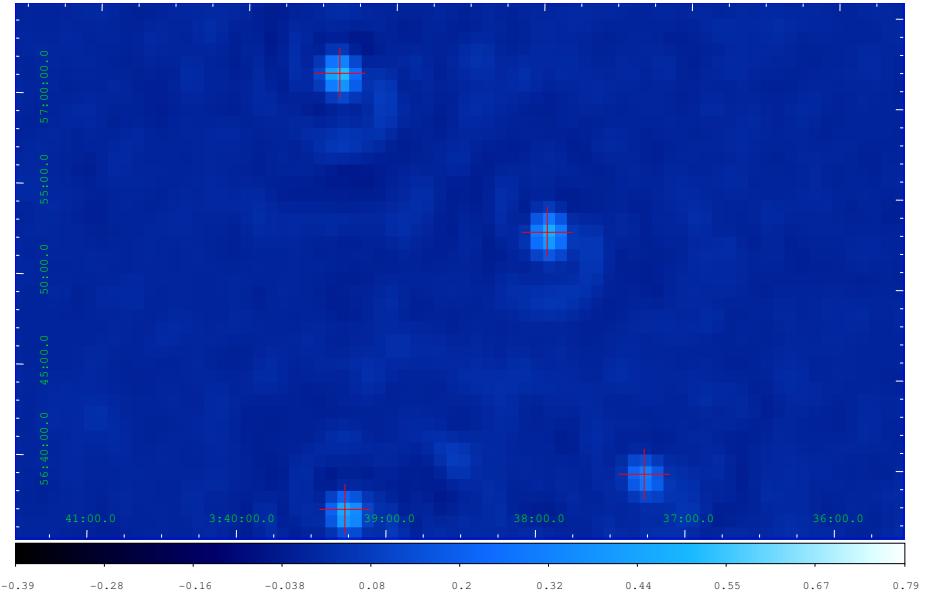


Figure 6.9: A ‘zoom in’ to the edge of the LOFAR 2010-Dec-03 observation (L2010_21641). The red crosses denote the locations of the WENSS sources. Good positional accuracy was found in this case, however, in earlier observations the positional accuracy degrades moving away from the pointing centre.

cations of the WENSS sources. Good positional accuracy is needed for the source association routines to perform effectively, and relies on good direction dependent calibration.

To assess the overall variability of all sources within the field, I perform a transient database query which retrieves all sources that have a index of variability $V_v > 0$. This query will retrieve all sources that are detected in more than one epoch i.e. two measurements are needed to calculate the index of variability. A total of 153 sources were returned with this query. In Figure 6.10, in the top three panels, I show the index of variability V_v as a function of distance from the pointing centre, source id. number and mean flux, for each extracted source using a 5σ source extraction level. The index of variability shows a slight broadening with respect to distance from the pointing centre. The index of variability does seem to have a dependence on the mean flux of the source, however, this might be biased because there are many more weaker sources. Figure 6.10 does indicate that a large number of sources within the field are highly variable. Typically for the science presented in this thesis I consider $V_v > 0.5$ to be highly variable (also see Carilli et al. 2003). On all of the plots in Figure 6.10 I indicate an index of variability of $V_v = 0.5$ with a blue line. The index of variability is theoretically dependent on the detection level of the source. The index of variability is defined as:

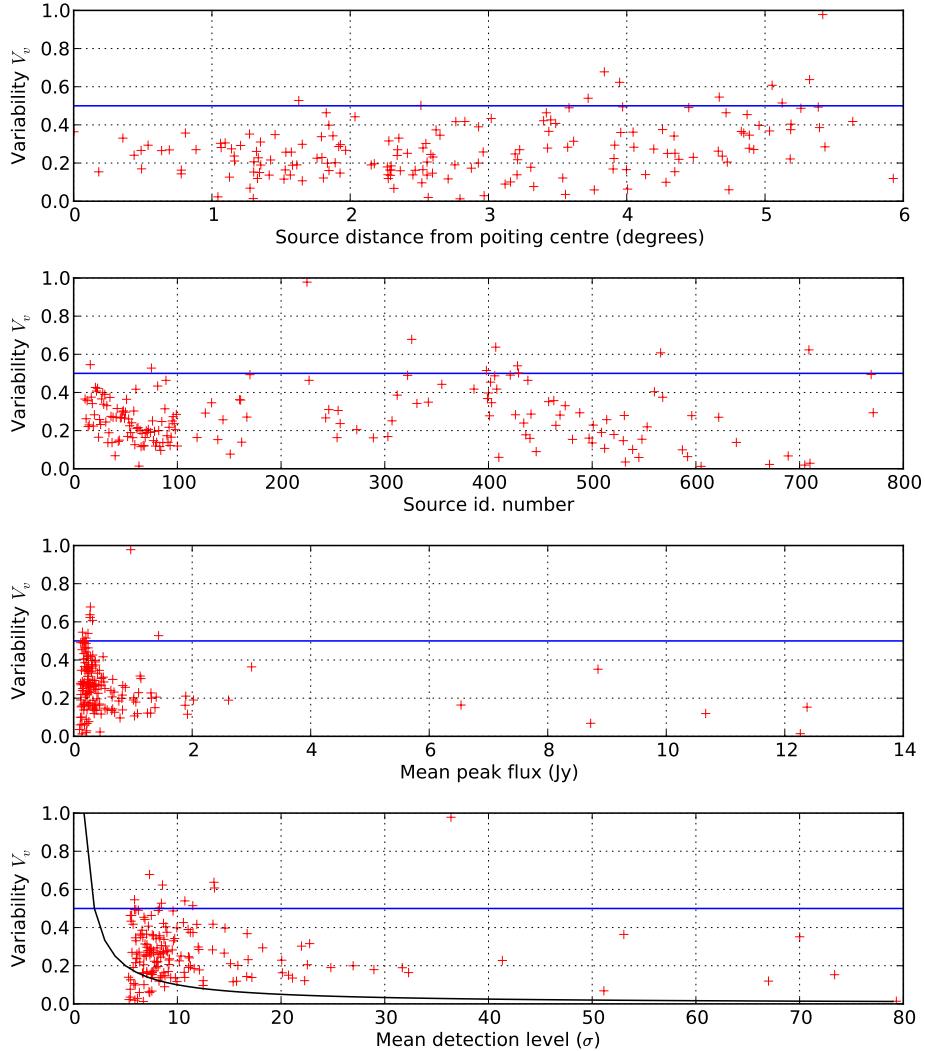


Figure 6.10: Top panel: Distance of extracted source(s) from the pointing centre against index of variability. The variability of sources appears to get slightly broader moving away from the pointing centre. Upper middle panel: Source id. number against variability. The source id. number is the unique number assigned to each source entering the database. Some sources do not have variability measures because they have already been associated with sources within the database from previous images, or they are unique sources for which an index of variability cannot be calculated. Lower middle panel: Mean peak flux of the extracted sources against index of variability. Sources with lower fluxes appear to show a larger amount of variability. Bottom panel: Mean source detection level against index of variability. The black lines show the theoretical index of variability based on the contribution from the noise exclusively.

$$V_v = \frac{\sigma_v}{\bar{S}_v} \quad (6.1)$$

Therefore for a $\bar{S}_v = 5\sigma$ detection, the index of variability $V_v = 0.2$. The noise therefore starts to dominate the index of variability at low detection levels. On the bottom panel of Figure 6.10 I plot the mean detection level of all the extracted sources against the index variability. I include on this plot the expected variability as a function of detection level (black line).

In all cases the large amount of sources with a high index of variability is a result of instability in the absolute flux scale, and hinders effectively detecting highly variable sources. One source was detected to have a significantly higher index of variability $V_v = 0.96$, than other sources (see Figure 6.10), but it was found to be associated with an artifact close to the edge of the image.

Further investigation into the variability of the sources within this field is ongoing. Further work may include refinements in the flagging, calibration and imaging stages, and also the transient detection algorithms. Currently the instability in absolute flux calibration and difficulty in source fitting impedes an effective algorithmic transient search. However, this work can be used as a benchmark and pilot study for future observations and testing. Despite the technical challenges of the algorithmic transient search, two promising candidate transient sources were easily detected, by eye. All LOFAR images from August to October were averaged together, and compared with the April and December LOFAR observations, and also the archival WENSS image. This gives the best sensitivity to long duration changes within the field. These candidates were not initially identified by the LTrAP, because the functionality that deals with non-detections was being commissioned (see Chapter 4 for discussion of non-detections), but they were subsequently detected when the index of variability query was dropped below $V_v = 0.5$. I will discuss these candidates in detail in the following sections.

6.4 Transient Candidate ILT J0320.3+5512

The transient candidate ILT J0320.3+5512 was detected seven times in the observations from 2010-Aug-13 to 2010-Oct-10, and also once on 2011-03-19; it was *not* detected in observations on 2010-Apr-13 nor 2010-Dec-03. See the top panel of Figure 6.11 for the August to October averaged image (with the transient on), see the middle panel for the December observation (transient off), and see the bottom panel for a cut out of the WENSS image with the transient candidate indicated with

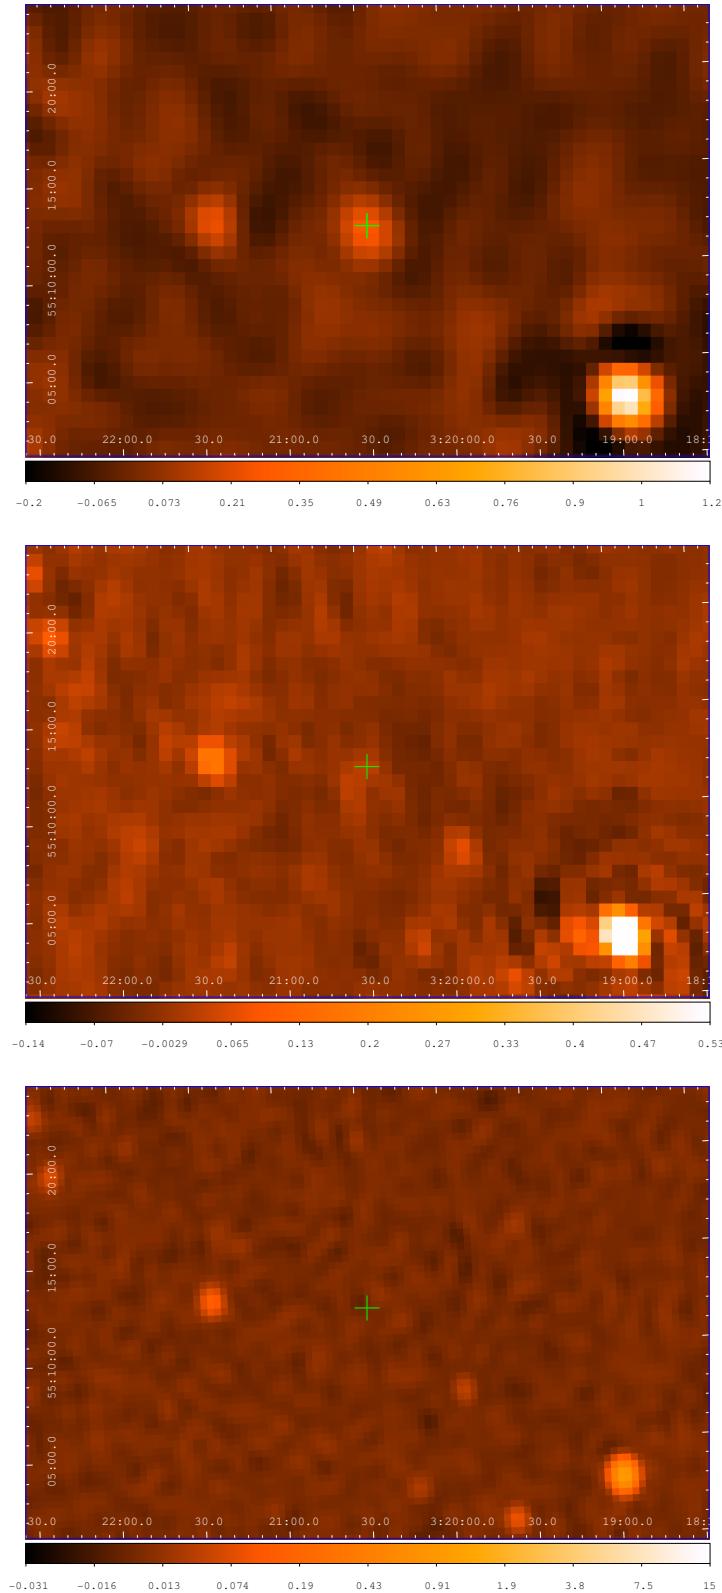


Figure 6.11: Top panel: A zoom in of the transient candidate **ILT J0320.3+5512** taken from a combined image of all observations from August to October 2010. Middle panel: The same region taken from the 2010-Dec-03 observation. Bottom panel: WENSS image of the same region with the transient position indicated with a green cross. The WENSS survey took place between 1991 to 1996.

a green cross. From epoch to epoch the fluxes of all sources within the field vary, due to instabilities in the absolute flux calibration - and the variations of all sources seem to be correlated. I correct for this uncertainty and trend in flux change (from epoch to epoch) by assuming a source near to the transient position is constant; I then apply the fractional change in the flux of this source, to the candidate transient (as well as another persistent source). See the top panel of Figure 6.12 for a lightcurve with the corrections applied, see the bottom panel for the uncorrected lightcurve.

The LTrAP does not currently use upper limits to calculate the index of variability. Thus for the time range August to October - between the two upper limits on Figure 6.12 - the index of variability calculated by the LTrAP is $V_v = 0.36$. The index of variability calculated by hand for the entire lightcurve - including upper limits - is $V_v = 0.61$. Hence in previous LTrAP searches, where the RMS values were not correctly handled, the transient nature of the object was not correctly identified. Also, as was discussed previously I set a cut off $V_v > 0.5$ to define highly variable sources, which without the RMS measurements this source is below. The index of variability for the two nearby persistent sources are: persistent source 1 = 0.09 and persistent source 2 = 0.16.

The candidate transient is located at $\alpha = 03^h 20^m 29^s.115$ and $\delta = +55^\circ 12' 08''.92$ (J2000). In all the observations I set a pixel size of $40''$, and approximately three pixels per beam. I therefore set a conservative upper limit on the positional uncertainty of $120''$. Fitting a Gaussian to the transient position in each epoch (when it is on) yields smaller positional errors. In Figure 6.13 I plot the position of the Gaussian fits with respect to the mean position of *all* the Gaussian fits, including errors. Figure 6.13 shows that there is a greater systematic error in the position, than the derived error from the individual Gaussian fits. From Figure 6.13 I therefore consider a refined positional uncertainty of $< 30''$. I will use both $30''$ and $120''$ in searching for multi-wavelength counterparts.

6.4.1 Multi-wavelength followup and catalogue search

Optical data have been obtained via the Liverpool Telescope at the position of the candidate transient source, in the *i* and *r* filters (*i* filter = 6930 - 8670Å, *r* filter = 5560 - 6890Å). The observations were obtained on 2011-Jan-14 with an exposure time of 25 minutes with each filter (data reduction courtesy of Dr. D. Bersier). See Figure 6.14 for the *r* filter image, with the $120''$ and $30''$ error circles, the transient position is indicated with a cross. The field of view of the Liverpool telescope (LT)

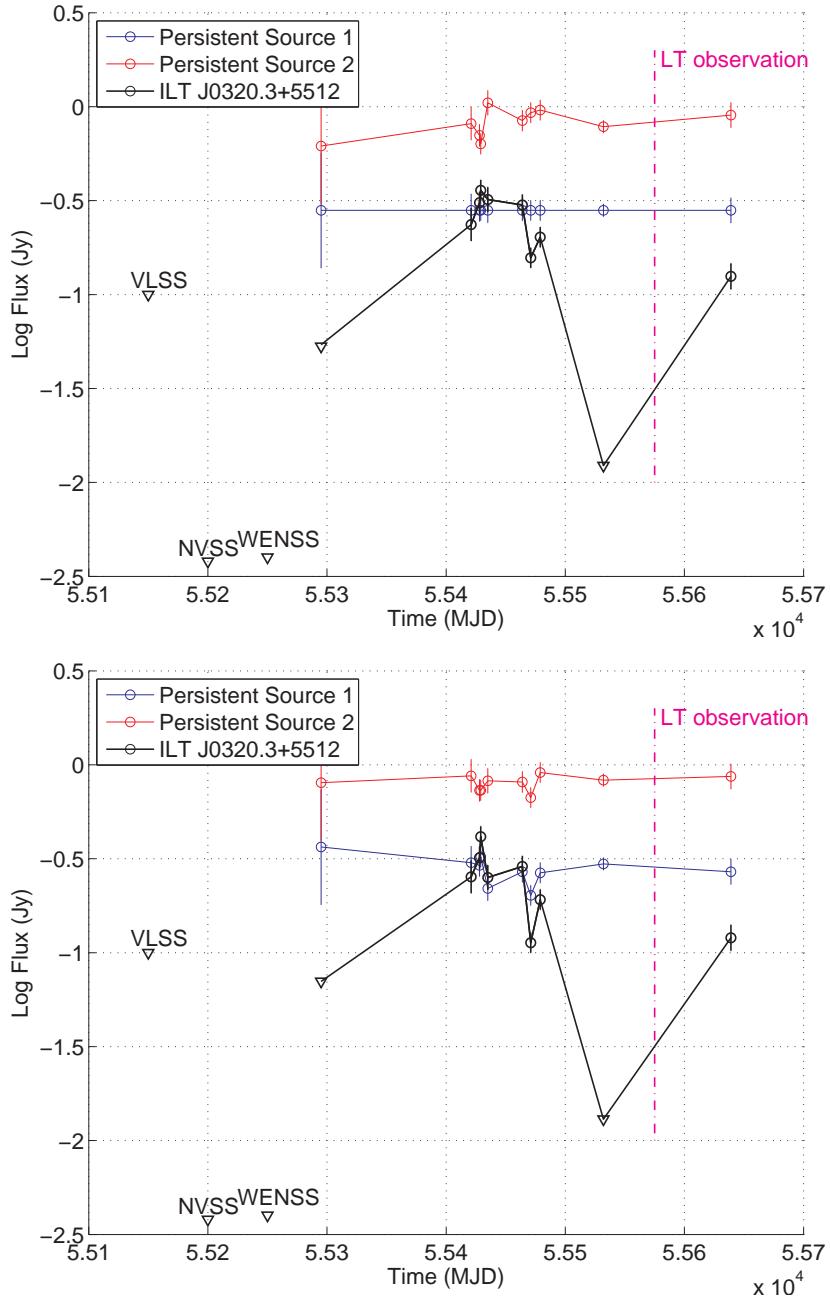


Figure 6.12: Top panel: Lightcurve of **ILT J0320.3+5512** (black line). A correction has been applied to both the candidate transient and persistent source 2 to account for uncertainty in flux. The peak flux has been used on the plot and the errors are derived from the local RMS. The WENSS, NVSS and VLSS upper limits are plotted at arbitrary times; they all in fact to place more than 5 years beforehand. Downward arrows denote upper limits from non-detections. The dashed purple line shows the time of the Liverpool telescope observation. Bottom panel: Lightcurve of **ILT J0320.3+5512** without the corrections applied.

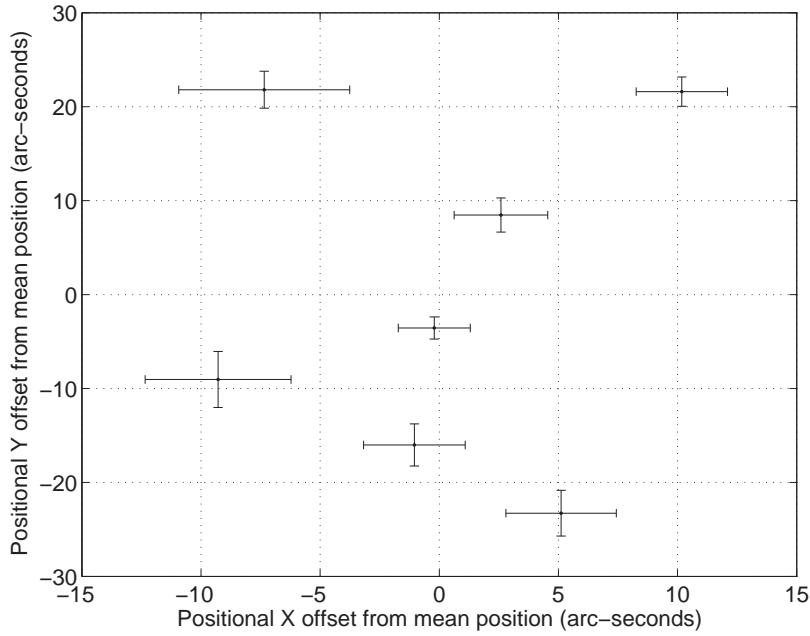


Figure 6.13: Plot showing the positional offset of the Gaussian fits (including errors), to the transient candidate **ILT J0320.3+5512**, with respect to mean position of *all* the Gaussian fits.

is ~ 2.5 arc-minutes; a number of sources lie within the $120''$ error circle of the transient location; and approximately eight sources lie within the $30''$ error circle.

It should be noted that the LT observation was obtained on 2011-Jan-14, which lies between the two LOFAR detections on 2010-Oct-10 and 2011-Mar-19. Furthermore, the LOFAR observation on 2010-Dec-03 (which is close in time to the LT observation) finds the source at its lowest flux. The LOFAR multi-wavelength follow-up team have compared the LT image with an archival optical image of the field, from the Isaac Newton Telescope (INT). No new sources are detected in the field, and none of the sources show significant variability. However, considering the timing of the observation the optical emission could have been in a state of ‘low flux’ and was potentially undetectable at the time - assuming that the radio and optical emission are linked. The four persistent optical sources within the $30''$ error circle are currently being investigated further, to ascertain if they are possible transient progenitors. If we assume that no new optical sources have been detected within the error circle of the transient position, and that the progenitor transient source is not a persistent source in the field: we can set an upper limit of $m_R=22.1$ and $m_I=21.4$ on the optical emission (D. Bersier, private communication).

In the radio waveband, a proposal has been submitted for two 12 hour continuum observations at 1.4 GHz and 350 MHz with WSRT to search for these candidates with an independent instrument. Note, I have searched the WENSS, VLSS and

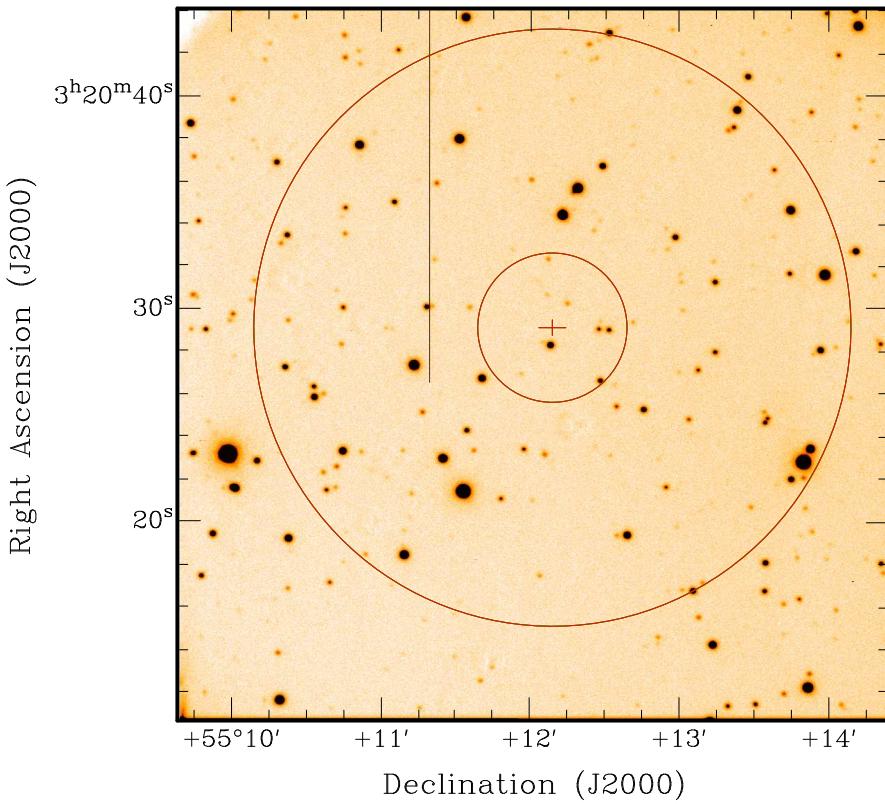


Figure 6.14: Optical image of the **ILT J0320.3+5512** field taken with the Liverpool Telescope in the r band filter. The red circles indicate the 120'' and 30'' error circle on the location of the transient.

NVSS catalogues and images for a counterpart: none was found. The NASA/IPAC Extragalactic Database (NED) and SIMBAD databases have also been searched for a counterpart: no sources were found within 2 arc-minutes of the candidate transient position. Therefore the following possible outcomes are expected from the WSRT observations A.) A radio source is observed within the positional errors of the transient *above* the WENSS and NVSS flux limit, providing the verification of the first LOFAR transient via an independent instrument. B.) A radio source is observed *below* the WENSS and NVSS flux limit, providing a progenitor radio source (produced via an independent instrument). C.) No radio source is observed and upper limits on the flux at 1.4 GHz or 350 MHz (or both) can be placed. In all cases we regard the 1.4 GHz observations as higher priority because we want to achieve the best dynamic range possible. If the source is synchrotron in nature, the increase in dynamic range at 1.4 GHz might be needed considering the peak of the spectral energy distribution may have shifted to lower frequencies. Having an additional observation at 350 MHz will allow us to search for a steep spectrum

(possibly coherent) source, and will also offer a direct comparison with the WENSS catalogue.

In parallel to requesting observations of ILT J0320.3+5512, we also requested observations of a second candidate LOFAR transient ILT J0331.5+5712 (which I will discuss in the next section). At the time the observations were requested, around February this year, ILT J0320.3+5512 was at a low flux state (based on our latest observation at that time, 2010-Dec-03). The second candidate was still bright and active in our latest LOFAR observation of the time (2010-Dec-03), therefore, the TAC decided to prioritise the observations of this source. With the latest LOFAR observations of ILT J0320.3+5512 we will re-request target of opportunity observations of the candidate, and also further LOFAR observations to confirm the result.

6.4.2 ILT J0320.3+5512 Discussion

As LOFAR is in a rapid and critical stage of commissioning we treat this result tentatively. However, as we have multiple detections (and non-detections) at a variety of hour angles, with different combinations of baselines, we believe this result merited further investigation and followup. If we assume the source is real and not a consequence of imaging and calibration errors, then we can rule out certain progenitor sources.

We can rule out a Gamma Ray Burst (GRB) afterglow (including an orphaned event) due to the flux of the candidate transient. One of the brightest and nearest GRBs detected was GRB 030329 at ~ 10 mJy (4.8 GHz; Berger et al. 2003, van der Horst et al. 2005), and $Z = 0.1685$ (Greiner et al. 2003). By modelling the higher frequency light curves of GRB 030329, van der Horst et al. (2008) have shown that a peak flux of < 1 mJy would be expected at 325 MHz approximately 600 - 800 days after the initial burst. Therefore at a flux of $S_{150MHz} \sim 0.2$ Jy, ILT J0320.3+5512 appears to be too bright for a GRB afterglow. The NASA High Energy Astrophysics Archive Research Centre (HEASARC) database was also searched for high energy counterparts and none were found.

If we consider a Galactic X-ray binary as the possible progenitor, we can scale the average radio flux (~ 0.2 Jy), using assumed values in the scaling relationship between the X-ray and radio luminosity, and mass presented in Chapter 3, to find a predicted X-ray flux. Using $L_R = 0.6 \log L_X + 0.78 \log M$ (see Merloni, Heinz, & di Matteo 2003) with a distance of 10 kpc and a mass of $10M_\odot$, I find $L_X = 1.6 \times 10^{38}$ ergs s^{-1} . This is reasonable with respect to the Eddington limit, however, such a

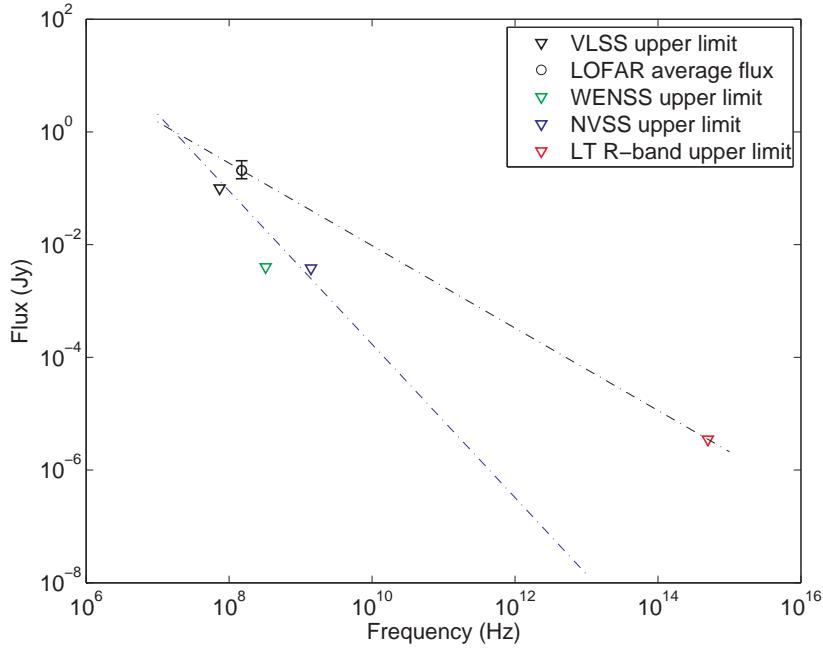


Figure 6.15: Spectral energy distribution of the candidate source **ILT J0320.3+5512**. The black dotted line shows the best fit to the LOFAR detection and Liverpool telescope optical R-band upper limit: which yields a spectral index of $\alpha = -0.71$. The blue dotted line shows the best fit to the radio data only, which yields a spectral index of $\alpha = -1.4$.

bright flux in X-rays should have been easily picked up by X-ray all sky monitors, over such a long duration: therefore an X-ray binary seems unlikely.

Figure 6.15 shows a preliminary spectral energy distribution of this candidate. The black dotted line shows the best fit to the LOFAR detection and Liverpool telescope optical R-band upper limit: which yields a spectral index of $\alpha = -0.71$. In Chapter 3 I used the ratio of the radio flux (at 6 cm), to the optical flux in the *B* band, to evaluate the radio loudness parameter of NGC 7213. Radio loud AGN are typically defined as having an *R* parameter > 10 , while radio quiet range between 0.1-1 (Kellermann et al., 1989). Using an *R*-band magnitude of $m_R = -22.1$ (upper limit), and a average LOFAR flux $S_{150MHz} = 0.2$ Jy, I find the ratio of the radio to optical flux to be $R \geq 57000$. In Sikora et al. (2007) and also Broderick & Fender (2011) AGN have been shown to have an *R* value of up to $\sim 10^7$ (using total radio power). I have calculated the radio to optical ratio above using a frequency of 150 MHz; assuming a spectral index of $\alpha = -0.71$, at 6 cm, the radio flux of the candidate would be $S_v = 17$ mJy, which in turn would yield $R = 5050$. Note, I have not converted from an R-band flux to a B-band flux so this is only an estimate. Both these *R* parameter estimates are within the permitted range for AGN: so we cannot completely rule this scenario out based on the *R* parameter alone.

It is however difficult to speculate on an AGN origin when considering the upper limits from the radio data points. The blue dotted line in Figure 6.15 shows the best fit to the radio data only, which yields a spectral index of $\alpha = -1.4$. A flat spectrum would typically be predicted from the jet of an unresolved AGN (see Markoff 2010), therefore this spectral index appears too steep. This spectral index should however be treated with care because the VLSS, WENSS and NVSS measurements were taken at different times. Confirming a higher frequency radio counterpart using another facility such as WSRT remains a high priority. If a counterpart was detected, obtaining follow-up X-ray observations would be a natural progression, especially if the radio spectral indexes was in a permitted range for a synchrotron type source. If the radio spectral index was steep, then we may consider a neutron star origin. Ofek et al. (2009) have shown that long duration radio transient outbursts are predicted for X-ray dim isolated neutron stars: which have low optical and X-ray luminosities, and also steep radio spectral indices. If this is the case high-time resolution follow-up would allow us to probe for short RRAT like bursts, or even weak radio pulses.

A coherent radio emission would typically produce steep spectrum, high brightness temperature sources. The brightness temperature of an object is defined as:

$$T_B = \frac{S_\nu \lambda^2}{2K_B \Omega_s} \quad (6.2)$$

Where T_B is the brightness temperature, S_ν is flux density, λ is the wavelength, K_B is the Boltzmann constant, and Ω_s is the solid angle of the source. The solid angle can be approximated using:

$$\Omega_s \approx \theta^2 \approx \frac{r^2}{D^2} \quad (6.3)$$

Where r is the size of the source and D is the distance to the source. The size of the source must be limited by the light travel time i.e. $r = c \times \tau$, in a given timescale of variability τ . This would therefore yield a minimum brightness temperature of:

$$T_B \geq \frac{S_\nu \lambda^2 D^2}{2K_B \Omega_s c^2 \tau^2} \quad (6.4)$$

Substituting in a timescale of variability of 107 days i.e. the time difference between the penultimate and final datapoint on Figure 6.12; which results in a change in flux from 12.3 mJy to 125 mJy: a brightness temperature for a source at 10 kpc would be $T_B \geq 2.0 \times 10^{12} K$. The inverse Compton catastrophe limit restricts the brightness temperature of a synchrotron emitting source to $T_B < 10^{12} K$, providing there is no relativistic boosting (Miller-Jones et al., 2008). At 10 kpc this source is at

that limit, increasing the distance rapidly moves the brightness temperature beyond the inverse Compton catastrophe limit. Therefore, if the source is synchrotron in nature, then we can place a limit of ~ 10 kpc on the distance - which further rules out an AGN and GRB origin (unless scintillation is a contributing factor to the variability). Further observations and analysis will be needed (and is ongoing) to check the robustness of this result, and to rule out other possible progenitors.

6.5 Transient Candidate ILT J0331.5+5712

This transient candidate was detected in all LOFAR observations from 2010-Aug-13 to 2011-Mar-19, but *not* in the observation from 2010-Apr-09. The total time the transient has been *observed* on is ~ 218 days. See the top panel of Figure 6.16 for the 2010-Apr-09 image of the transient location, see the bottom panel for the 2010-Dec-03 LOFAR observation (transient on). The candidate has no counterpart in WENSS (see the bottom panel of Figure 6.16), VLSS and NVSS catalogues. The candidate transient is located at $\alpha = 03^h31^m45^s.567$ and $\delta = +57^\circ12^m45^s.56$ (J2000). The candidate reached a peak flux of $S_v \sim 0.5$ Jy; note, this value has not been primary beam corrected. See Figure 6.17 for a light curve of the candidate. A correction has also been applied to this light curve to account for the uncertainty in absolute flux calibration (see previous candidate for details). The index of variability for the transient candidate is $V_v = 0.44$ for the entire lightcurve; the index of variability for a nearby persistent source is $V_v = 0.29$. The transient nature of this object is better defined by the upper limit in the 2010-04-09 observation, and the upper limits from the WENSS, VLSS and NVSS catalogues and images.

6.5.1 Multi-wavelength followup and catalogue search

Optical data have been obtained from the Liverpool Telescope at the position of the transient in the *i* and *r* filters. See Figure 6.19 for optical image with $120''$ and $30''$ error circles on the transient position (data reduction courtesy of Dr. D Bersier). Similar positional errors were found for this candidate as were found for the previous candidate i.e. ~ 30 arc-seconds (see Figure 6.18). The optical images were compared with an archival ILT image of the field (by the LOFAR multi-wavelength team), for unique sources, and sources that showed significant variability: none was found.

As of March 2011 this candidate had been observed consistently in nine consecutive observations; therefore we requested a WSRT Target of Opportunity (ToO)

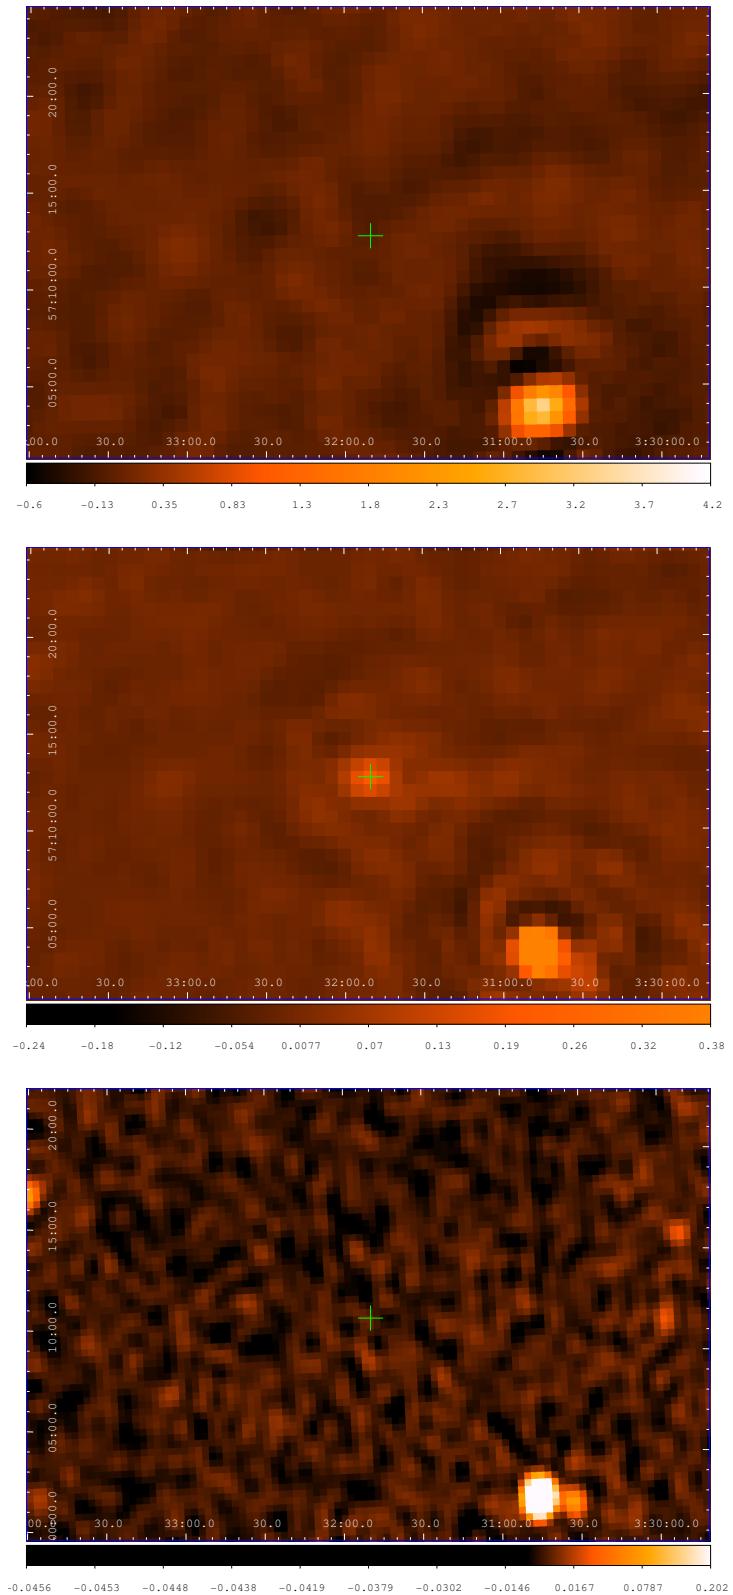


Figure 6.16: Top panel: A zoom in of the transient candidate **ILT J0331.5+5712** taken from the 2010-Apr-09 observation. Middle panel: The same region taken from the 2010-Dec-03 observation. Bottom panel: WENSS image of the same region with the transient position indicated with a green cross.

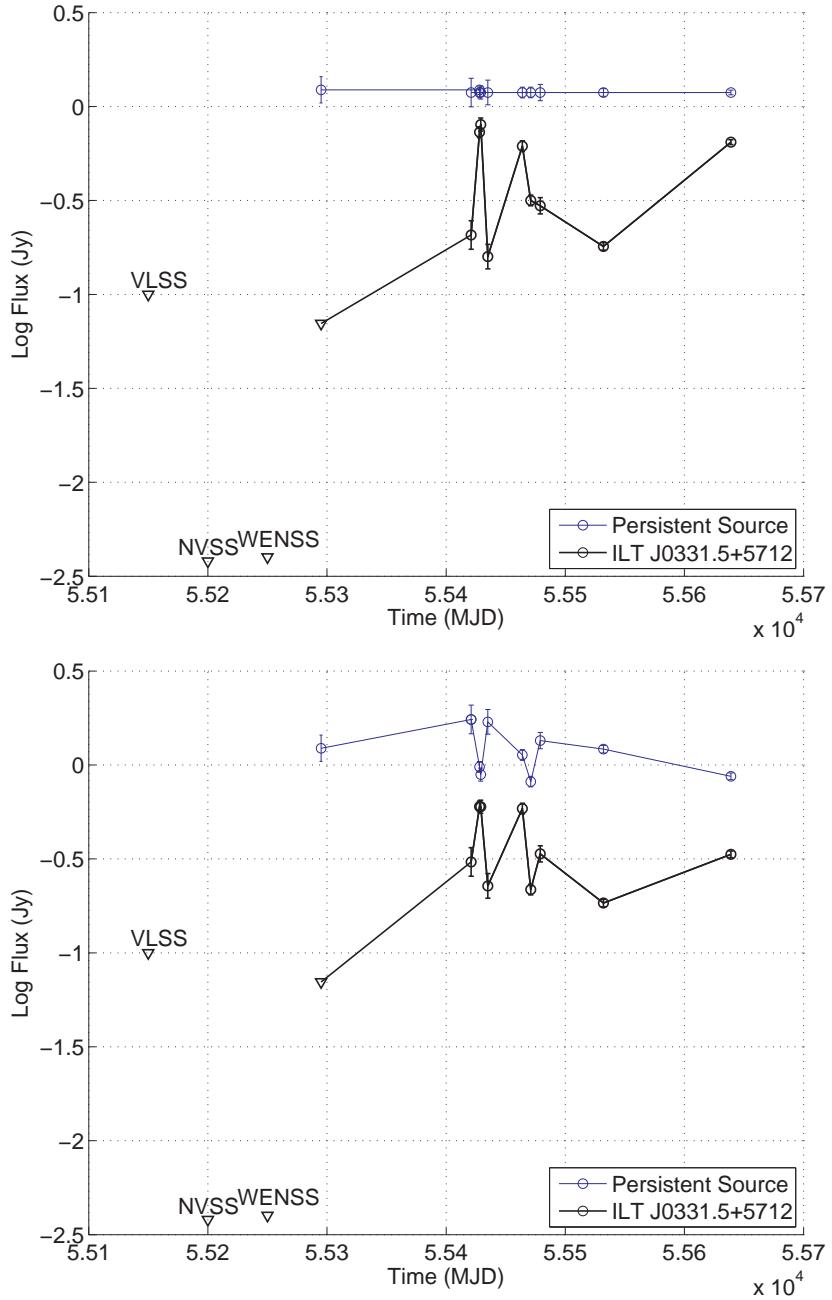


Figure 6.17: Top panel: Lightcurve of **ILT J0331.5+5712** (black line). A correction has been applied to the candidate transient to account for uncertainty in flux. The peak flux has been used on the plot and the errors are derived from the local RMS. The WENSS, NVSS and VLSS upper limits are plotted at arbitrary times. Downward arrows denote upper limits from non-detections. Bottom panel: Lightcurve of **ILT J0331.5+5712** without the corrections applied.

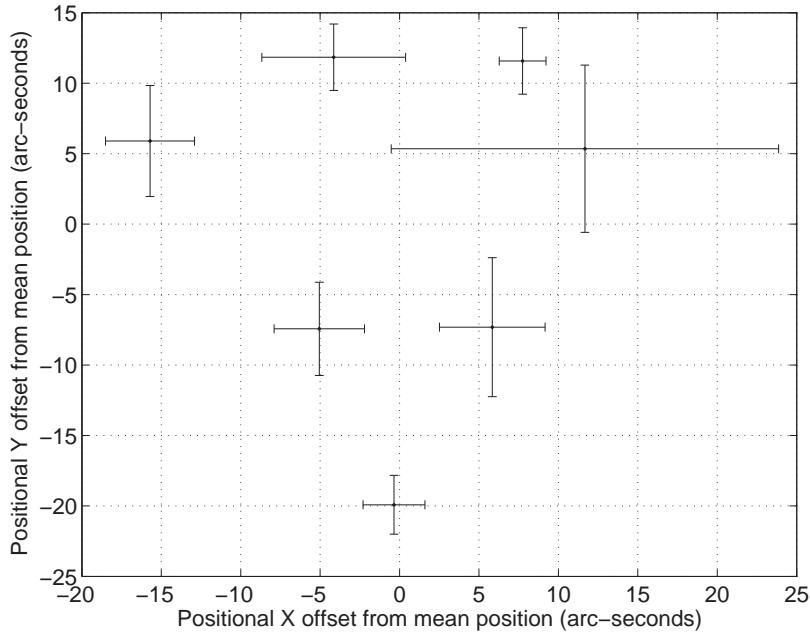


Figure 6.18: Plot showing the positional offset of the Gaussian fits (including errors), to the transient candidate **ILT J0331.5+5712**, with respect to mean position of *all* the Gaussian fits. Note, two of the fits failed to produce positional errors, which is potentially due to estimating the beam volume (attempting to fit within MIRIAD produced the error ‘Failed to determine covariance matrix’), this is under further investigation.

observation. A six hour LOFAR observation was taken on 2011-Mar-19 and showed the transient candidate still on. We obtained one 9.75 hour 350 MHz observation on 2011-Apr-02, and one 8.5 hour observation on 2011-Apr-03 at 1.4 GHz, both with WSRT. Figure 6.20 shows the WSRT 1.4 GHz image with the $120''$ and $30''$ error circles, with the transient location indicated with a cross (data reduction courtesy of Dr. V. Tudose). There is a source within the $120''$ error circle, but outside the $30''$ error circle. This source has a peak flux $S_v = 0.76$ mJy, which is below the 4 mJy RMS limit of the WENSS survey, we can not therefore completely rule this source out as the transient counterpart, however, it is outside the $30''$ error circle so it seems unlikely and the spectral index would be extremely steep.

This transient candidate was also observed with the Lovell telescope on both 2011-Apr-02 and 2011-Apr-03 in high time resolution mode. A LOFAR tied array beam was also used to observed the candidate on 2011-Apr-03. In both cases the rationale was to search for either single pulses, or sequences of pulses. The 350 MHz WSRT observations and the high time resolution observations are still yet to be reduced, however, this work has been suspended for reasons that I will discuss in the next section.

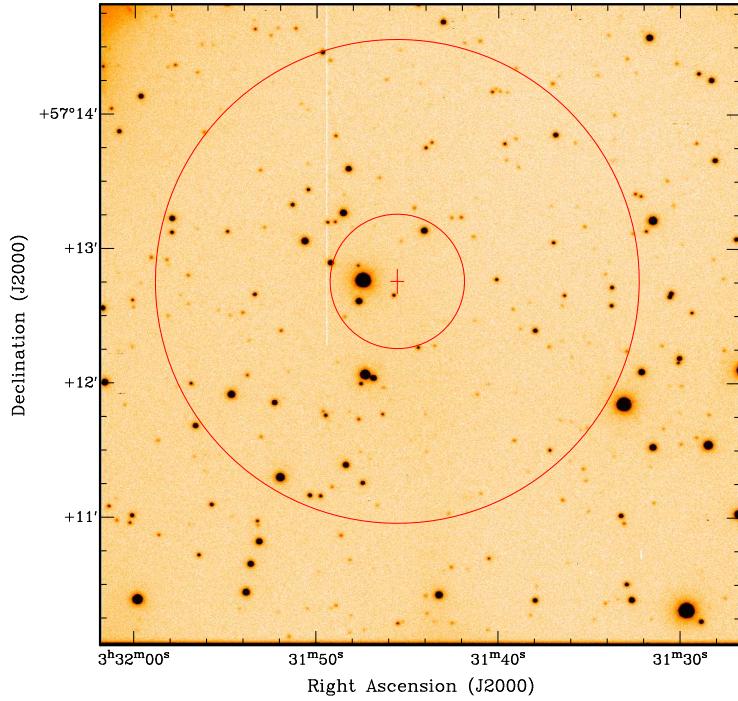


Figure 6.19: Optical image of the **ILT J0331.5+5712** field taken with the Liverpool Telescope in the r band filter. The red circles indicate the 120'' and 30'' error circle on the location of the transient.

6.5.2 ILT J0331.5+5712 Discussion

Around February this year it was considered that ILT J0331.5+5712 was of higher priority than ILT J0320.3+5512, because in the latest observation - at that time (2010-Dec-03) - ILT J0320.3+5512 was in a low flux state. Therefore the candidate ILT J0331.5+5712 was the focus of intense follow-up. Since that time further observations have been carried out, the data reduction procedures have been refined, and further algorithms have become available, or been improved upon. ILT J0331.5+5712 is quite close to the edge of the field ~ 18 pixels in a 512×512 image. The only observation which does not show the transient on, is the first 2010-Apr-09 LOFAR observation. This was the only image in which the phase centre was different from the other observations: it was centred on calibrator A (3C84). To test that shifting the phase centre did not affect the detection, we requested a further LOFAR observation - taken on 2011-04-03 - which had the phase centre shifted to the bright source to the South West of J0331.5+5712, see Figure 6.16. The image from this observation did not detect ILT J0331.5+5712 down to the RMS. So either the transient had significantly dropped in flux since the latest LOFAR observation - 2011-Mar-19 - when the source was detected, or, it was potentially an artifact. It should be noted that ILT J0331.5+5712 had been detected in nine different ob-

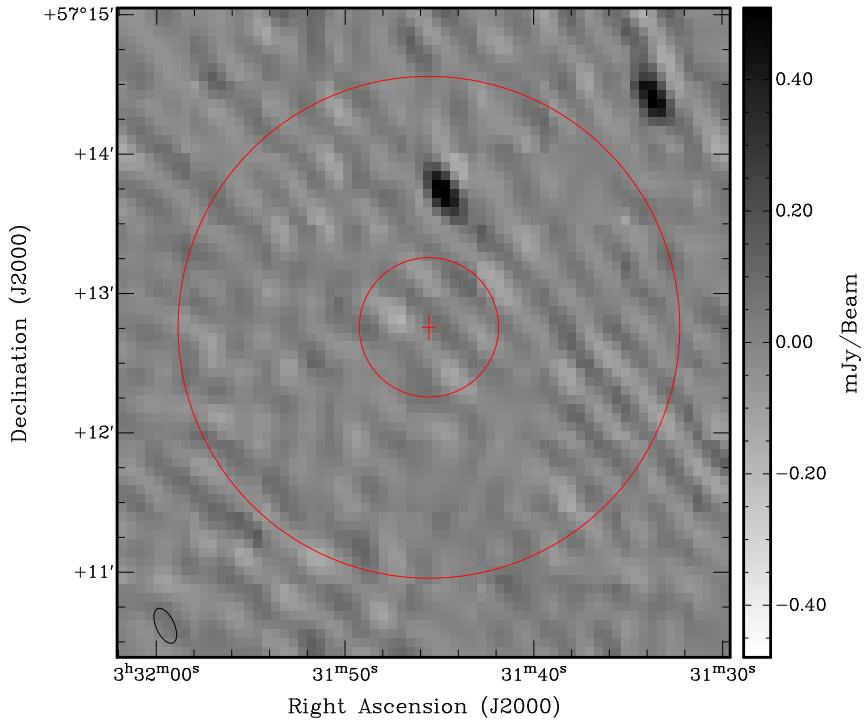


Figure 6.20: WSRT 1.4 GHz observation of the candidate transient **ILT J0331.5+5712**, with $120''$ and $30''$ error circle on the position. One source is detected within the $120''$ error circle at $S_v = 0.76$ mJy, however, it is below the detection threshold of the WENSS survey: therefore it is difficult to conclude if it is a counterpart.

servations with different baseline configurations, different numbers of channels per subband, and different hour angles.

Furthermore, the only imaging algorithm that was implemented in the pipeline architecture at the time, was the Cimager. The Cimager has since been replaced by a new imager based on the CASA data reduction package. The calibrated observations from 2010-Dec-03, which had a detection of ILT J0331.5+5712, were re-imaged with the CASA imager, and the transient source was not detected in the image, see Figure 6.21. However, a few sources close to the field edge that were in the Cimager image had either declined in flux, or were not detectable. One of the sources that was not detected, had a counterpart in the WENSS catalogue. Clearly both imagers are struggling to perform adequately at the edge of the field. Currently we are comparing different calibration and imaging strategies with both the Cimager and CASA imager, to test the robustness of detecting sources to the edge of the image. By differencing two images made with each algorithm, and comparing with the WENSS catalogue, we can more easily detect imaging artifacts. More research is needed to ascertain which imager is performing correctly, and also if ILT J0331.5+5712 is real, or just a spurious artifact. It is difficult to explain how an

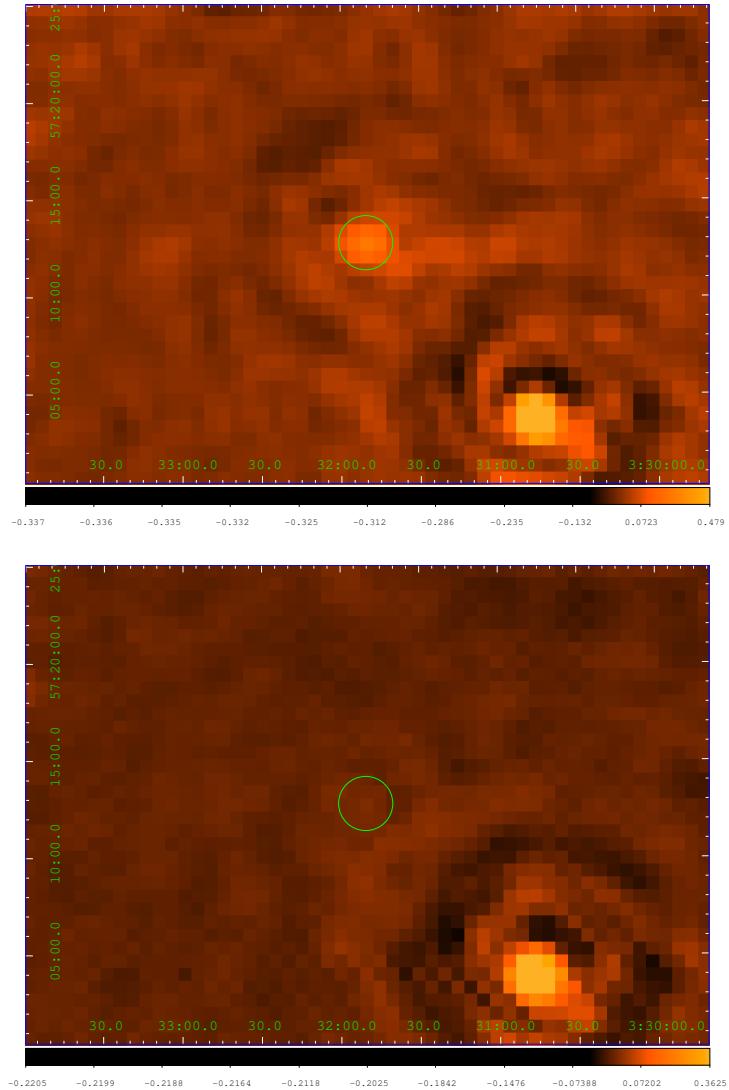


Figure 6.21: Top panel: A zoom in of the transient candidate **ILT J0331.5+5712** taken from the 2010-Dec-03 observation produced with the Cimager. Bottom panel: The same epoch imaged with the CASA imager. The intensity scaling of each image is the same (logarithmic units)

artifact (or error) can manifest itself in the visibilities, or the image production procedure, in the same way for so many different observations. However, LOFAR is in its infancy so we should *expect* to find unexpected results. This process of comparing different data reduction strategies, and obtaining observations from other radio telescopes, is therefore extremely important for future LOFAR transient searches and commissioning.

6.6 Conclusion and Future Work

In this chapter I have presented the data reduction procedure to create LOFAR images. I have also presented the first algorithmic transient search on wide field LOFAR images. There are still many issues to resolve in both these procedures. For the calibration and imaging algorithms absolute flux calibration must be stable, and wide field next generation calibration strategies such as direction dependent gains must be applied. The beam properties also need to be better characterised to produce good Gaussian fits. Other issues such as band-pass calibration and primary attenuation are still in a state of commissioning and impact on the quality of the final images.

I have shown in this chapter this it is possible to run the LTrAP algorithm over a sequence of LOFAR images, which produces a large number of transient candidates and statistics. The main bottle-neck is the current quality of the images, a large majority of sources within the field show a large index of variability, and this impedes detecting truly variable sources. However, by comparing the images by eye with archival images, some interesting candidates have been identified. The LTrAP will soon be capable of automatically detecting (and responding) to the type of attributes that these candidates have i.e. non-LOFAR and archival catalogue detections. The type of logic, testing and re-reduction of images conducted on the LOFAR commissioning images may have to become more algorithmic. For example, if a particular setting in the calibration or imaging stages is prone to producing artifacts (in certain situations), the SIP and LTrAP must be capable of altering these setting to produce comparison images, to automatically assess the robustness of the result. Simply producing one image and assuming it is always accurate may result in lots of false detections. The flexibility of the LOFAR data reduction system is unique in this sense; all algorithms within the pipeline operate under the same architecture. If the source extraction algorithm detects an interesting source, it can spawn new imaging tasks to produce subsequent images with different settings, or re-task the calibration and flagging etc.

The B0329+54 field will form one of the zenith monitoring fields, therefore developing accurate sky models, and data reduction techniques is essential to commission the radio sky monitor. In future observations we can expand the B0329+54 field of view, using multiple LOFAR beams, and using the detailed sky model for this field (plus further sky model sources), we can commission the radio sky monitor. We can also use these early datasets and analysis as a benchmark to assess the long term ($>$ year) variability and transient nature of the field. In this chapter I

have only focussed on producing one combined image from 12 hours of data; it is a goal to produce images on much shorter timescales (i.e. one second). Dividing the datasets described in this chapter into smaller chunks and re-imaging could yield further interesting results, because we would be more sensitive to transients with that characteristic timescale. A single 12 hour dataset would produce some 43200 one second images, which in turn would place stringent limits on the rates of short durations transients: if none were detected.

So far ILT J0320.3+5512 remains the most promising candidate and may well turn out to be the first transient detected with LOFAR, however, it must stand up to rigorous testing with different algorithms and data reduction procedures. Potentially many more observations of the B0329+54 field will be obtained, and assuming that each image holds an equal chance of detection, it should hopefully not be long before we have the first confirmed LOFAR transient to report. In the final chapter to this thesis I will comment further on the parameter space we have probed with this survey, and I will relate this to the potential parameter space other SKA pathfinder instruments might explore.

I like the dreams of the future better than the history of the past.

THOMAS JEFFERSON (1743 - 1826)

7

Conclusion

In this chapter I will present the main findings of this thesis. I will also discuss future work and broader issues pertaining to each chapter, that were not addressed in the individual chapter conclusions. I will place the broader implications of the transient searches presented in this thesis into context, with an emphasis on the future parameter space SKA pathfinder instruments may explore.

7.1 Summary of thesis and future work

In Chapter 3 of this thesis I presented a detailed variability study of the low luminosity active galactic nuclei NGC 7213. Through this analysis I explored the radio and X-ray correlation using the cross correlation function. I showed that a weak but statistically significant delay existed between the X-ray and radio emitting regions. I related this behaviour to that seen in other active galactic nuclei and black hole X-ray binaries. I also used the X-ray and radio luminosities, together with the black hole mass, to show that NGC 7213 was in good agreement with the predicted correlation found on the fundamental plane of black hole activity. Around three years of X-ray and radio data was collected to study NGC 7213 - this is one of the longest campaigns of its type to date. Requesting radio observations over so

many semesters is difficult and time consuming. X-ray data were much more easily obtained with the Rossi X-ray Timing Explorer (RXTE), due to the observatory's flexible scheduling, and willingness to collect long term datasets. Potentially with 'all sky monitor' type properties of the next generation of radio telescope, obtaining large datasets of this type of source (and others) will become easier. It is a goal for these wide field facilities to frequently monitor known (bright) variable sources.

In Chapter 4 of this thesis I presented an overview of the prototype LOFAR transients detection pipeline (LTrAP). I demonstrated and discussed the main components that were used to analyse the datasets presented in Chapters 5 and 6. I discussed the broader functionality of the pipeline that will be needed for future operations, and I also presented pitfalls with current approaches to tackling transient searches. Future work which is needed (and is currently underway) to commission the full pipeline. This includes: development of quality control metrics, automated classification algorithms, an automated (internal and external) response system, more rigorous transient database queries, and an efficient visualisation system to view results.

In Chapter 5 of this thesis I presented an automated VLA data reduction pipeline that was used to flag, calibrate and image VLA data. This imaging pipeline was used to produce images for three different datasets. Firstly the datasets containing the transient detections reported in Bower et al. (2007) were imaged. The properties of the transient sources, and the light curves, derived from the LTrAP were compared with those reported in Bower et al. (2007). Five out of the seven transient results were in good agreement with those reported in Bower et al. (2007) and this work was an excellent test of the LTrAP. Secondly an archival VLA dataset containing the previously unreported outburst of the X-ray binary *Swift* J1753.3-0127, was imaged and processed through the LTrAP. The results were used to explore the X-ray to radio correlation; it was found that the source was less radio luminous than predicted by the relationship derived from the fundamental plane of black hole activity. Finally a blind survey for radio transients was conducted using both the automated VLA reduction script and the LTrAP, on archival calibrator fields. This work placed upper limits on the snapshot rate of radio transients at GHz frequencies; it also explored, and made predictions regarding the Log N- Log S of transient radio sources. In conjunction with producing early stage science, this work contributed to commissioning the LOFAR transient detection system. The VLA archive potentially contains many hours unsearched data, which could contain radio transients. The imaging algorithms developed for this study could be used to further interrogate the archive in the future.

In Chapter 6 of this thesis I presented the procedure to create images from raw LOFAR data, and also a blind transient search using the LTraP. Many stages of the image production procedure are in a state of commissioning, however, it has been shown that reasonable image fidelity can be achieved, and promising results are expected in the future. In Chapter 6 I only concentrated on producing images with an integration time of 12 (or 6) hour timescales; future work will include reducing the integration time to one second, and also developing and implementing 3rd generation calibration techniques that are specifically tuned for transient studies. There are two key issues which have been addressed in Chapter 6. Firstly, what are the ideal parameters to produce *one* adequate LOFAR image? Secondly, what are the ideal parameters (and compromises that should be made) to produce a *sequence* of LOFAR images? So far efforts by other LOFAR imaging commissioners have mainly focused on producing one image, whereas producing multiple images (on short timescales) is clearly a top priority for the Transients Key Project. The transient search using LTraP reflects the current quality of LOFAR data. Until adequate flux calibration is achieved, analysis of the variability of sources within this field is restricted. In this chapter I have however demonstrated that the infrastructure and algorithms are available to do this. Although the flux calibration is unstable, this survey is still sensitive to unique sources not observed in previous images or radio catalogues. I have presented two LOFAR transient candidates, which meet this criteria. In the next section, as a conclusion to this thesis, I will explore the parameter space we have probed with this prototype LOFAR survey, and what can potentially be achieved in the future. I will also present the parameter space that other SKA pathfinder facilities may explore.

7.2 Exploring SKA pathfinder parameter space

In this thesis I have presented two transient surveys, one using archival VLA data, and the second using LOFAR commissioning data. Both of these surveys have started to probe the parameter space where large transient source populations may be abundant. In Figure 7.1 I show a reduced version of the flux density versus snapshot rate plot (see Chapter 5, Figure 5.12). In this figure I plot only the blind transient surveys which have detected transient or variable sources; for clarity I include the Lazio et al. (2010) and Bell et al. (2011) upper limits. I also include the derived least-squares best fit to the Bower et al. (2007) detections of 1 week duration (thin black line). The best fit yields a logarithmic gradient of $m = -1.51$,

which is consistent with a uniformly distributed source population in a Euclidean volume of space (i.e. gradient of -3/2).

In Figure 7.1 the vertical dashed lines represent the RMS noise (in a 12 hour observation) for a number of SKA pathfinder instruments. The corresponding horizontal lines represent the 2σ upper limit on the snapshot rate, derived from 10 epochs of data with each instrument respectively. For the SKA Phase 2 I adopt the values quoted in Carilli & Rawlings (2004); for a frequency of 1.4 GHz, an RMS of 37 nJy, and a Field of View (FoV) of 1 deg 2 is realised. For a frequency of 200 MHz (c.f. LOFAR), an RMS of 0.4 μ Jy, and a FoV of \sim 200 deg 2 is realised. The APERTIF upgrade to WSRT and the ASKAP telescope achieve a similar sensitivity of \sim 10 μ Jy (at 1.4 GHz) in a 12 hour observation, however ASKAP will image a FoV of 30 deg 2 , compared to 8 deg 2 for APERTIF (see Oosterloo et al. 2010 and Johnston et al. 2007 for further details). The MeerKAT telescope is designed for high dynamic range imaging, and will achieve an RMS of \sim 1.8 μ Jy in 12 hours, with a FoV of \sim 1 deg 2 at 1.4 GHz.

Figure 7.1 also shows the current sensitivity (\sim 10 mJy) of the LOFAR commissioning observations (labelled ‘LOFAR (HBA) 2011’), and the potential full array sensitivity (31 μ Jy, see Table 1.1), both values assume a 12 hour integration at 120 MHz and 10 repeat visits to the same field. A greater FoV can be achieved with LOFAR by using multiple beams to tile up a larger FoV, however, this comes with a trade off of sensitivity. Assuming that the transient population detected by Bower et al. (2007) are in some way representative of the global population; by either improving the sensitivity (which pushes the detection limit on Figure 7.1 to the left); or by imaging larger FoVs (which pushes the horizontal snapshot rate limit down), the transient population can be efficiently explored. Note, in Figure 7.1 I do not correct for any spectral index effects, with respect to the GHz and MHz surveys and predictions. As discussed in Chapter 1, an order magnitude increase in flux from 1.4 GHz to 150 MHz would be expected for a source with $\alpha = -0.7$. This in turn would shift the Log N - Log S of the Bower et al. (2007) GHz detections to the right. This is probably an over-simplification, however, it should be considered when plotting fluxes and snapshot rates derived at two different frequencies.

In Figure 7.1 I have also included the snapshot rate derived from *one* LOFAR candidate transient (ILT J0320.3+5512) in the 10 epoch survey. More work remains to ascertain the robustness of this result, it does however indicate the region of parameter space we have probed. The current sensitivity of the commissioning observations needs to improve to adequately explore the predicted population. In Chapter 5 and 6 I discussed the possibility of detecting a difference between a GHz

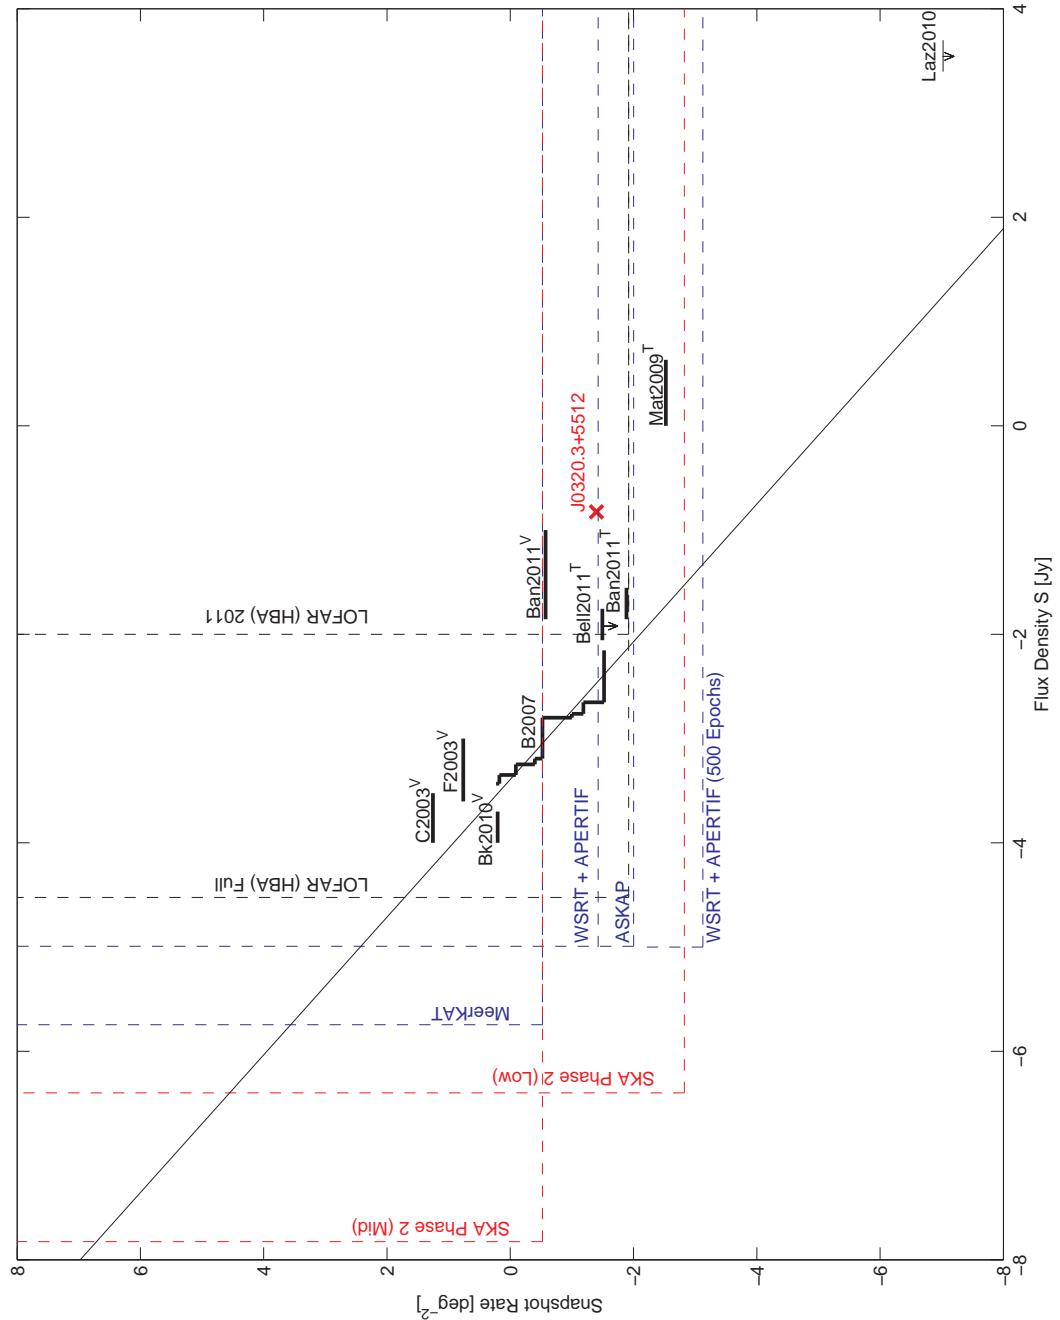


Figure 7.1: Log of the Snapshot rate (deg^{-2}) against log of the flux density (Jy) of detections of transients (labelled ‘T’), detections of variable sources (labelled ‘V’) and upper limits based on non-detections (labelled with downward arrows). The Bower et al. (2007) 1 week limits are indicated as B2007^T . The least squared best fit to the Bower et al. (2007) 1 week detections are indicated with a thin black line. Ban2011^V and Ban2011^T indicates the separate rates derived for variables and transients reported in Bannister et al. (2011). Bell2011^T indicates the 2σ upper limits derived from the VLA transient study. See Chapter 5 for details and description of other surveys. The dashed lines represent the sensitivity (vertical), and 2σ upper limits on the snapshot rate (horizontal), derived for a number of SKA pathfinder instruments, assuming zero detections in 10 epochs of data.

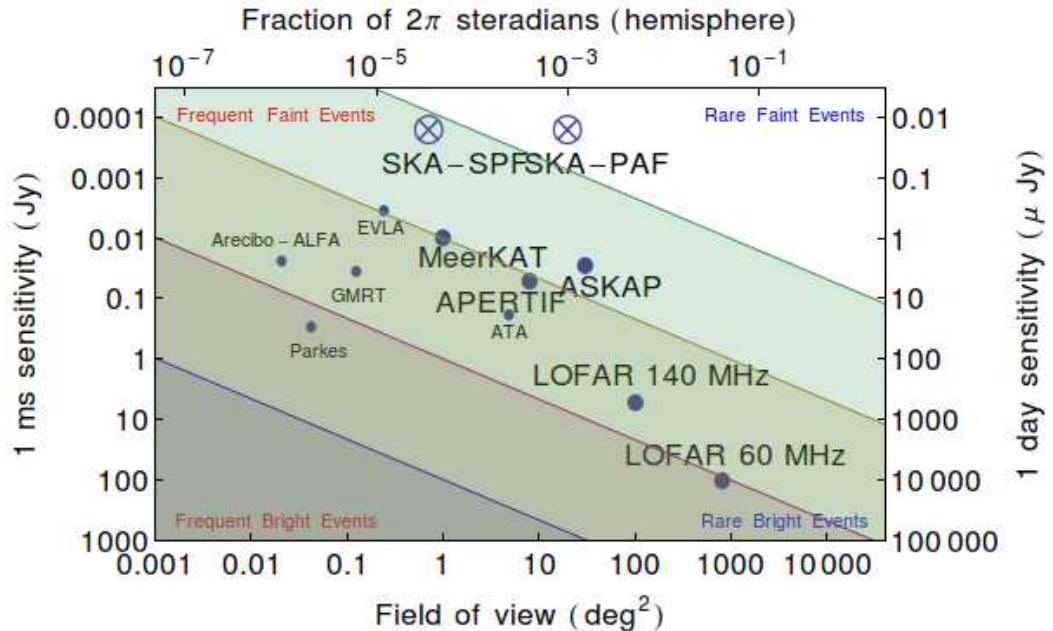


Figure 7.2: Field of view (deg^2) against 1 day sensitivity (μJy - right axis) for current facilities (small circles) and new facilities (large circles). The solid lines corresponds to constant numbers of detected transients, assuming a homogeneous Euclidean population (i.e. $\sigma \propto \text{FoV}^{-2/3}$). Figure courtesy of R. Fender.

and MHz population of transients. The limits provided by this survey currently do not adequately explore the region of parameter space needed to form solid conclusions. The predicted full sensitivity of LOFAR (labelled ‘LOFAR (HBA) Full’), with only 10 epochs, separated by a week in time, should provide a definitive answer.

The field of view and sensitivity of current and future radio telescopes is summarised in Figure 7.2 (courtesy of R. Fender). Considering the FoV advantage of ASKAP, when compared with MeerKAT, it might be predicted that it would yield greater transient detections. Based on the current Log N - Log S, at each of the telescopes sensitivity limits respectively, MeerKAT will in fact detect greater numbers of transient sources (although this prediction is based on the current performance parameters of the telescopes, which could change). The total yield is however dependent on the integral under the Log N - Log S slope, for a range of source flux densities (see Figure 7.1). Included in Figure 7.2 are lines which correspond to constant numbers of detected transients assuming a homogeneous Euclidean population i.e. $\sigma \propto \text{FoV}^{-2/3}$. From this figure it can be seen that ASKAP will yield a larger total number of transient sources (MeerKAT will however observe deeper which will potentially sample a different populations of sources).

Taking the Bower et al. (2007) best fit and extrapolating from the SKA Phase 2 (Mid) sensitivity limit, assuming a 3σ detection is needed, a snapshot rate of $9.7 \times 10^5 \text{ deg}^{-2}$ with a total yield of $9.7 \times 10^7 \text{ yr}^{-1}$ is predicted. I also summarise the predicted snapshot rates, rates per year and the total yield for other instruments in Table 7.1. The predictions are extremely speculative about the true nature of the transient population, however, even if the Log N - Log S is less steep, large numbers of transient detections are still expected.

The examples above have assumed that only ten repeat visits are obtained of the same field. Using the proposed APERTIF surveys¹ as an example, I can calculate how many repeat visits we could expect from a commensal survey. The APERTIF surveys will last for a period of \sim five years and will tackle a number of scientific goals. The proposed surveys fall into the following three categories:

- **Shallow**

$(1 \times 12 \text{ hours}) \times 1250 \text{ fields.}$

Total: $10,000 \text{ deg}^2$ revisited once.

- **Medium deep**

$(10 \times 12 \text{ hours}) \times 125 \text{ fields.}$

$(12 \times 12 \text{ hours}) \times 5 \text{ fields.}$

Total: 1000 deg^2 revisited 10 times.

- **Deep**

$(100 \times 12 \text{ hours}) \times 5 \text{ fields.}$

Total: 40 deg^2 revisited 100 times.

$(500 \times 12 \text{ hours}) \times 1 \text{ fields.}$

Total: 8 deg^2 revisited 500 times.

The shallow continuum survey represents (near) all sky Northern hemisphere coverage to a sensitivity of $\sim 10\mu\text{Jy}$. This survey will provide one of the deepest radio catalogues of the Northern sky at 1.4 GHz and can be compared with existing radio catalogues to identify transient radio sources. Note, a number of other projects are proposing all sky surveys, such as pulsars and HI, however these will be conducted commensally (i.e. Pulsar and HI data can be obtained simultaneously with continuum data). The deep surveys offer the largest number of repeat observations of the same field, potentially 500 revisits. In Figure 7.1, I include the limit on the snapshot rate that would be found from 500 repeat visits to the same field

¹<http://www.astron.nl/radio-observatory/apertif-eoi-abstracts-and-contact-information>

Table 7.1: Predicted rates of radio transients for SKA pathfinder instruments (assuming a 12 hour synthesis for each instrument respectively). The calculations are based on the Log N - Log S of the Bower et al. (2007) detections at 4.8 GHz. The snapshot rate column shows the rate extrapolated to the 3σ detection limit of each instrument respectively (NB, per deg^2). The Bower et al. (2007) survey had a cadence of \sim one week, therefore the rate per year column shows the expected rate of transients in 52 weeks of observing. Note, this is a lower limit because the characteristic timescale of the Bower et al. (2007) transients were between 20 minutes and one week, I also do not consider any recurrence in transient activity. The yield columns give the number of transients expected for each instrument over the instrument’s FoV. The LOFAR rates do not take into account any spectral index considerations. Therefore, assuming a spectral index of $\alpha=-0.7$ (synchrotron) and $\alpha=-2$ (coherent) I show the corrected rates based on these spectral indexes.

Instrument	Snapshot rate (deg^{-2})	Rate per year ($\text{deg}^{-2} \text{yr}^{-1}$)	Yield (yr^{-1})
SKA Phase 2 (Mid)	9.7×10^5	5.0×10^7	5.0×10^7
SKA Phase 2 (Low)	6.7×10^3	3.5×10^5	7.0×10^7
MeerKAT	6.9×10^2	3.6×10^4	3.6×10^4
WSRT + APERTIF	5.0×10^1	2.6×10^3	2.0×10^4
ASKAP	5.0×10^1	2.6×10^3	7.8×10^4
LOFAR (HBA) Full	9.8	5.0×10^2	1.2×10^4
ATCA	1.5×10^2	7.6×10^3	5.3×10^2
EVLA	5.9×10^2	3.0×10^4	2.1×10^3
GMRT	2.7×10^1	1.4×10^3	1.4×10^4
ATA	4.8	2.3×10^2	5.7×10^2
LOFAR (HBA) Full ($\alpha=-0.7$)	1.0×10^2	5.3×10^3	1.3×10^5
LOFAR (HBA) Full ($\alpha=-2$)	8.9×10^3	4.6×10^5	1.2×10^7
SKA Phase 2 (Low)($\alpha=-0.7$)	5.5×10^4	2.8×10^6	5.7×10^8
SKA Phase 2 (Low)($\alpha=-2$)	2.4×10^6	1.2×10^8	2.5×10^{10}

(labelled ‘WSRT + APERTIF (500 Epochs)'). This offers approximately an order of magnitude improvement on the snapshot rate, assuming no transient are detected. Different images could also be averaged together from the deep and medium surveys to improve the sensitivity.

More work is needed to fully characterise the Log N - Log S. A top priority should be cadence, the Log N - Log S should evolve to incorporate the characteristic timescale of transient behaviour i.e. Log N - Log S - Log t. Furthermore, this parameter space (including cadence) should be characterised at a number of different radio frequencies, and also in different directions. For example, in the direction of the Galactic plane, variable radio sources may be more abundant. In a given extragalactic direction, interstellar scintillation may be more prominent (due

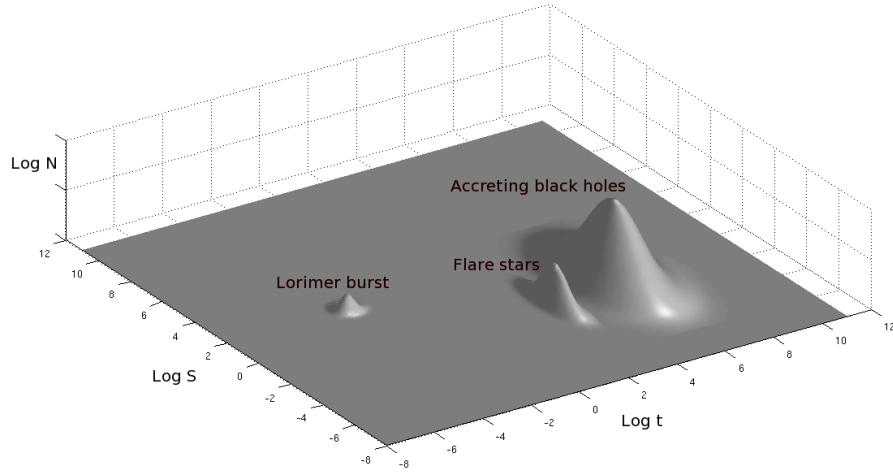


Figure 7.3: An example Log N - Log S - Log t showing a simplification of three known transient source populations. By probing the Log N - Log S - Log t we can examine the abundance of transient source populations and whether they exist within discrete regions within this parameter space. At present we are collapsing and merging together different populations.

to increased electron content), and will also potentially increase short timescale variability at low frequencies (e.g. see Hunstead 1972). It may also be interesting to compare (or even normalise) the transient source Log N - Log S, with the known radio source Log N - Log S. For example, Windhorst et al. (1985) have shown that above ~ 1 mJy source populations are dominated by active galactic nuclei, below this star-forming galaxies are thought to dominate (also see Georgakakis et al. 1999 and Carilli et al. 2003).

In Figure 7.3 I show a toy model of the Log N - Log S - Log t which includes three examples of known transient sources. In this example, accreting black holes populate a region of moderate to long duration transient activity (\sim months to years) with a broad range of fluxes, spanning eight orders of magnitude. Flare stars will potentially populate a region much more constrained in time i.e. the typical flare times are minutes only, with fluxes from mJy to Jy (see Chapter 2). Flare stars are probably less abundant than AGN, which is represented on the plot. Extremely rare events such as the Lorimer et al. (2007) burst would populate a region of timescales \sim milli-seconds, with extremely bright fluxes ~ 30 Jy.

In reality the boundaries between the different known transient source populations may not be so discrete and a variety of transient behaviour, both in timescale and flux, may blur them together (i.e. this figure is just for conceptual purposes only). I have also not included any considerations of a homogeneous extragalactic

of galactic normalisation to this plot, which will potentially impact further on the appearance of the populations.

Large amounts of parameter space clearly remain uncharted, it is possible that there are possibly rare, or even fairly abundant, transient source populations which populate discrete regions of the $\log N$ - $\log S$ - $\log t$ (c.f. type Ia supernovae in the optical), which have so far gone undetected. With the dedicated and commensal transient surveys of the next generation of radio telescope becoming available soon, we are on the cusp of characterising the $\log N$ - $\log S$ - $\log t$. LOFAR will soon start to impact on the exploration of this parameter space at low frequencies, and potentially other SKA pathfinder instruments will soon start to explore different facets of the $\log N$ - $\log S$ - $\log t$.

7.3 A final remark

The process of aperture synthesis is heavily - or even completely - dependent on computers and algorithms to produce images. The human eye has historically been a key tool in interpreting astronomical images. The defining characteristic of the next generation of radio telescope, and transient radio science in the future, will predominantly be the *complete* removal of the astronomer from the imaging, detection, classification and alert system. The algorithms that create this ‘system’ will however have to mimic some of the complex capabilities of the human brain and eye. The brains ability to logically interpret the scientific significance of pixels in an image, whilst considering a broad wealth of astronomical information is currently unsurpassed. At least for the not too distant future anyway.

APPENDICES

Research is what I'm doing when I don't know what I'm doing.

WERNHER VON BRAUN (1835 – 1910)

A

Example Parsets

Initial Flagging and Compression Parset

```
uselogger = True
msin.startchan = 3
msin.nchan = 58
msin.datacolumn = DATA # is the default
msout.datacolumn = DATA # is the default
steps = [preflag,flag1,avg1,flag2,avg2,flag3,count]
preflag.type = preflagger # This step will flag the autocorrelations.
preflag.corrtype = auto
```

```
flag1.type = madflagger
flag1.threshold = 4
flag1.freqwindow = 5
flag1.timewindow = 5
flag1.correlations = [0,3] # only flag on XX and YY
```

```
avg1.type = squash
avg1.freqstep = 58 # it compresses the data by a factor of 64 in frequency
avg1.timestep = 1 # is the default
```

```
flag2.type = madflagger
flag2.threshold = 3
flag2.timewindow = 51
```

```
avg2.type = squash
avg2.timestep = 5 # it compresses the data by a factor of 5 in time
```

```
flag3.type = madflagger
flag3.threshold = 3
flag3.timewindow = 241
```

```
count.type = counter
count.showfullyflagged = True
```

Four Source Sky Model

```
# (Name, Type, Ra, Dec, I, Q, U, V, ReferenceFrequency='74.000e6',
SpectralIndexDegree='0', SpectralIndex:0='0.0', SpectralIndex:1='0.0', MajorAxis,
MinorAxis, Orientation) = format
B0323.5+5510, POINT, 03:27:19.25, +55.20.29.0, 40.14, 0, 0, 0, , 1, -0.453, 0
B0326.1+5322, POINT, 03:29:52.16, +53.32.34.4, 14.82, 0, 0, 0, , 1, -0.684, 0
B0332.6+5328, POINT, 03:36:28.04, +53.38.51.6, 12.60, 0, 0, 0, , 1, -0.830, 0
J0328+5509, POINT, 03:28:11.74, +55.09.12.9, 13.27, 0, 0, 0, , 1, -0.999, 0
```

BBS Calibration Parslet

```
Strategy.InputColumn = DATA
Strategy.TimeRange = []
Strategy.Baselines = *&
Strategy.ChunkSize = 0
Strategy.UseSolver = F
Strategy.Steps = [solve, subtract, correct]
```

```
Step.solve.Operation = SOLVE
Step.solve.Model.Sources = [B0323.5+5510, B0326.1+5322, B0332.6+5328]
Step.solve.Model.Bandpass.Enable = F
Step.solve.Model.Gain.Enable = T
Step.solve.Model.DirectionGain.Enable = F
```

```

Step.solve.Model.Beam.Enable = F
Step.solve.Model.Ionosphere.Enable = F
Step.solve.Model.Cache.Enable = T
Step.solve.Solve.Parms = [”Gain:0:0:*,”Gain:1:1:”]
Step.solve.Solve.ExclParms = []
Step.solve.Solve.CalibrationGroups = []
Step.solve.Solve.CellSize.Freq = 0
Step.solve.Solve.CellSize.Time = 1
Step.solve.Solve.CellChunkSize = 10
Step.solve.Solve.Options.MaxIter = 20
Step.solve.Solve.Options.EpsValue = 1e-9
Step.solve.Solve.Options.EpsDerivative = 1e-9
Step.solve.Solve.Options.ColFactor = 1e-9
Step.solve.Solve.Options.LMFactor = 1.0
Step.solve.Solve.Options.BalancedEqs = F
Step.solve.Solve.Options.UseSVD = T

Step.subtract.Operation = SUBTRACT
Step.subtract.Model.Sources = [B0323.5+5510]
Step.subtract.Model.Gain.Enable = T
Step.correct.Operation = CORRECT
Step.correct.Model.Sources = [B0323.5+5510, B0326.1+5322, B0332.6+5328]
Step.correct.Model.Gain.Enable = T
Step.correct.Output.Column = CORRECTED_DATA

```

Post-BBS Flagging Parset

```

msin.startchan = 0
msin.nchan = 1
msin.datacolumn = CORRECTED_DATA
msout.datacolumn = CORRECTED_DATA
steps = [preflag,preflag2,preflag3,flag3,flag4,flag5,flag6,flag7]

```

```

preflag.type=preflagger
preflag.corrtype=cross
preflag.amplmax=250
preflag.baseline=[ [CS*,CS*] ]
preflag2.type=preflagger

```

```

preflag2.corrtype=cross
preflag2.amplmax=250
preflag2.baseline=[ [CS*,RS*] ]
preflag3.type=preflagger
preflag3.corrtype=cross
preflag3.amplmax=250
preflag3.baseline=[ [RS*,RS*] ]

```

```

flag3.type=madflagger
flag3.threshold=4
flag3.timewindow=5
flag3.correlations=[0,3]
flag4.type=madflagger
flag4.threshold=3
flag4.timewindow=25
flag5.type=madflagger
flag5.threshold=3
flag5.timewindow=51
flag6.type=madflagger
flag6.threshold=3
flag6.timewindow=71
flag7.type=madflagger
flag7.threshold=3
flag7.timewindow=241

```

Imaging Parset

```

Cimager.Images.Names=image_Clean
Cimager.Images.cellsize=[40arcsec,40arcsec]
Cimager.Images.image_Clean.nchan=1
Cimager.Images.image_Clean.polarisation='I'
Cimager.Images.shape=[512,512]
Cimager.datacolumn=CORRECTED_DATA
Cimager.gridder=WProject
Cimager.gridder.WProject.cutoff=0.001
Cimager.gridder.WProject.limitsupport=0
Cimager.gridder.WProject.maxsupport=800
Cimager.gridder.WProject.nfacets=1

```

```
Cimager.gridder.WProject.nwplanes=257
Cimager.gridder.WProject.oversample=1
Cimager.gridder.WProject.wmax=15000
Cimager.gridder.padding=1
Cimager.ncycles=1
Cimager.restore=True
Cimager.restore.beam=fit
Cimager.solver=Clean
Cimager.solver.Clean.algorithm=Hogbom
Cimager.solver.Clean.gain=0.1
Cimager.solver.Clean.niter=200
Cimager.solver.Clean.scales=[0, 3]
Cimager.solver.Clean.verbose=True
Cimager.preconditioner.Names=Robust
Cimager.preconditioner.Robust.robustness=0.0
```


BIBLIOGRAPHY

Abada-Simon, M., & Aubier, M. 1997, Å

Alexander, T. 1997, Astronomical Time Series, 218, 163

Aller, H. D., Aller, M. F., Latimer, G. E., & Hodge, P. E. 1985, *Astrophys. J. Suppl. Ser.*, 59, 513

Anderson, K. R., Alexov, A., Bähren, L., Grießmeier, J.-M., Wise, M., & Renting, G. A. 2010, ISKAF2010 Science Meeting,

Baars, J. W. M., & Hooghoudt, B. G. 1974, *Astron. Astrophys.*, 31, 323

Bannister, K. W., Murphy, T., Gaensler, B. M., Hunstead, R. W., Chatterjee, S. 2011, *MNRAS*, 31

Bastian, T. S., Dulk, G. A., & Slee, O. B. 1988, *AJ*, 95, 794

Bastian, T. S. 1990, *Solar Physics*, 130, 265

Becker, R. H., White, R. L., Helfand, D. J. 1995, *ApJ*, 450, 559

Becker, R. H., Helfand, D. J., White, R. L., Proctor, D. D. 2010, *AJ*, 140, 157

Bell, M. E., et al. 2010, *MNRAS*, 1825

Bell, M. E., et al. 2011, *Mon. Not. R. Astron. Soc.*, 415, 2

Belloni, T., et al. 1999, *Astrophys. J.*, 527, 345

Belloni, T. 2010, Lecture Notes in Physics, Berlin Springer Verlag, 794,

Berger, E., et al. 2001, *Nature*, 410, 338

Berger, E. 2002, *Astrophys. J.*, 572, 503

Berger, E., et al. 2005, *Astrophys. J.*, 627, 960

Beswick, R. 2006, Proceedings of the 8th European VLBI Network Symposium,

Berger, E., et al. 2003, *Nature*, 426, 154

Bhat, N. D. R., Tingay, S. J., & Knight, H. S. 2008, *ApJ*, 676, 1200

Bianchi S., Matt G., Balestra I., Perola G. C., 2003, *A&A*, 407, L21

Bianchi S., La Franca F., Matt G., Guainazzi M., Jimenez Bailón E., Longinotti A. L., Nicastro F., Pentericci L., 2008, *MNRAS*, 389, L52

Bignall H. E., Tzioumis A. K., Jauncey D. L., Venturi T., Clay R. W., 2004, *NuPhS*, 132, 149

Blandford R. D., Konigl A., 1979, *ApJ*, 232, 34

Blank D. L., Harnett J. I., Jones P. A., 2005, *MNRAS*, 356, 734

Bondi, M., Padielli, L., Fanti, R., Ficarra, A., Gregorini, L., & Mantovani, F. 1996, *Astron. Astrophys. Suppl. Ser.*, 120, 89

P.A. Boncz, *Monet: A Next-Generation DBMS Kernel For Query-Intensive Applications*, PhD Thesis, University of Amsterdam, 2002

Booth R. S., de Blok W. J. G., Jonas J. L., Fanaroff B., 2009, arXiv,

arXiv:0910.2935

Bower G. C., Falcke H., Mellon R. R., 2002, ApJ, 578, L103

Bower G. C., Plambeck R. L., Bolatto A., McCrady N., Graham J. R., de Pater I., Liu M. C., Baganoff F. K., 2003, ApJ, 598, 1140

Bower G. C., Roberts D. A., Yusef-Zadeh F., Backer D. C., Cotton W. D., Goss W. M., Lang C. C., Lithwick Y., 2005, ApJ, 633, 218

Bower G. C., Saul D., Bloom J. S., Bolatto A., Filippenko A. V., Foley R. J., Perley D., 2007, ApJ, 666, 346

Bower, G. C., et al. 2010, ApJ, 725, 1792

Bower, G. C., & Saul, D. 2010, arXiv:1101.0121

Bower, G. C., Whysong, D., Blair, S., Croft, S., Keating, G., Law, C., Williams, P. K. G., & Wright, M. C. H. 2011, arXiv:1107.1517

Bransford M. A., Appleton P. N., Heisler C. A., Norris R. P., Marston A. P., 1998, ApJ, 497, 133

Breedt, E., et al. 2009, MNRAS, 394, 427

Briggs, B. S. *High Fidelity Deconvolution of Moderately Resolved Sources*, PhD Thesis, The New Mexico Institute of Mining and Technology , 1995

Broderick, J. W., & Fender, R. P. 2011, arXiv:1105.3769

Broderick, J. W., & Fender, R. P. 2011, arXiv:1105.3769

Burgasser, A. J., & Putman, M. E. 2005, Astrophys. J., 626, 486

Burke-Spolaor, S., Bailes, M., Ekers, R., Macquart, J.-P., & Crawford, F., III 2011, Astrophys. J., 727, 18

Calvelo, D. E., et al. 2010, arXiv:1007.2313

Calvelo, D. E., et al. 2010, Mon. Not. R. Astron. Soc., 409, 839

Carilli, C. L., Ivison, R. J., Frail, D. A. 2003, ApJ, 590, 192

Carilli, C. L., & Rawlings, S. 2004, NewA&R, 48, 979

Casella, P., et al. 2010, arXiv:1002.1233

Cohen, A. S., Lane, W. M., Cotton, W. D., Kassim, N. E., Lazio, T. J. W., Perley, R. A., Condon, J. J., Erickson, W. C. 2007, AJ, 134, 1245

Corbel S., Nowak M. A., Fender R. P., Tzioumis A. K., Markoff S., 2003, A&A, 400, 1007

Condon, J. J., & Backer, D. C. 1975, Astrophys. J., 197, 31

Condon, J. J., Ledden, J. E., Odell, S. L., & Dennison, B. 1979, Astron. J., 84, 1

Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., Broderick, J. J. 1998, AJ, 115, 1693

Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., Broderick, J. J. 1998, AJ, 115, 1693

Cordes J. M., Lazio T. J. W., McLaughlin M. A., 2004, *NewAR*, 48, 1459

Coriat, M., et al. 2011, *Mon. Not. R. Astron. Soc.*, 414, 677

Clark, B. G. 1980, *Astron. Astrophys.*, 89, 377

Croft, S., et al. 2010, arXiv:1006.2003

Davies R. D., Walsh D., Browne I. W. A., Edwards M. R., Noble R. G., 1976, *Nature*, 261, 476

Dennett-Thorpe, J., & de Bruyn, A. G. 2002, *Nature*, 415, 57

Dennett-Thorpe, J., & de Bruyn, A. G. 2003, *Astron. Astrophys.*, 404, 113

Dennison, B., Broderick, J. J., Ledden, J. E., Odell, S. L., & Condon, J. J. 1981, *Astron. J.*, 86, 1604

Done, C., Gierliński, M., & Kubota, A. 2007, *Astron. Astrophys. Rev.*, 15, 1

Dunn, R., Fender, R., Koerding, E., Belloni, T., & Cabanac, C. 2009, arXiv:0912.0142

Dulk, G. A. 1985, *Ann. Rev. Astron. Astrophys.*, 23, 169

Eck C. R., Cowan J. J., Branch D., 2002, *ApJ*, 573, 306

Edelson, R. A., Krolik, J. H. 1988, *ApJ*, 333, 646

Ellingson, S. W., Clarke, T. E., Cohen, A., Craig, J., Kassim, N. E., Pihlstrom, Y., Rickard, L. J., & Taylor, G. B. 2009, *IEEE Proceedings*, 97, 1421

Falcke H., Patnaik A. R., Sherwood W., 1996, *ApJ*, 473, L13

Falcke H., Körding E., Markoff S., 2004, *A&A*, 414, 895

Falcke, H., & Key Science Project, L. C. R. 2007, *Astronomische Nachrichten*, 328, 593

Falcke, H., Markoff, S., Bower, G. C. 2009, *AAP*, 496, 77

Felten J. E., 1967, *AJ*, 72, 796

Fender, R. P., Garrington, S. T., McKay, D. J., Muxlow, T. W. B., Pooley, G. G., Spencer, R. E., Stirling, A. M., Waltman, E. B. 1999, *MNRAS*, 304, 865

Fender, R. P., Garrington, S. T., McKay, D. J., Muxlow, T. W. B., Pooley, G. G., Spencer, R. E., Stirling, A. M., & Waltman, E. B. 1999, *Mon. Not. R. Astron. Soc.*, 304, 865

Fender, R., Garrington, S., & Muxlow, T. 2005, *The Astronomer's Telegram*, 558, 1

Fender, R. P., et al. 2006, VI Microquasar Workshop: Microquasars and Beyond,

Fender, R. 2006, *Compact stellar X-ray sources*, 381

Fender, R., Koerding, E., Belloni, T., Uttley, P., McHardy, I., & Tzioumis, T. 2007, arXiv:0706.3838

Fender R., Wijers R., Stappers B., LOFAR Transients Key Science Project t., 2008, arXiv, arXiv:0805.4349

Ferrarese L., Merritt D., 2000, *ApJ*, 539, L9

Fix, J. D., Claussen, M. J., & Doiron, D. J. 1980, *Astron. J.*, 85, 1238

Frail D. A., Kulkarni S. R., Nicastro L., Feroci M., Taylor G. B., 1997, *Nature*, 389, 261

Frail, D. A., Kulkarni, S. R., Berger, E., Wieringa, M. H. 2003, *AJ*, 125, 2299

Frater, R. H., Brooks, J. W., & Whiteoak, J. B. 1992, *Journal of Electrical and Electronics Engineering Australia*, 12, 103

Gaensler, B. M., & Hunstead, R. W. 2000, *Publ. Astron. Soc. Aust.*, 17, 72

Gaensler B. M., et al., 2005, *Nature*, 434, 1104

Gallo E., Fender R. P., Pooley G. G., 2003, *MNRAS*, 344, 60

Gallo E., Fender R. P., Miller-Jones J. C. A., Mervloni A., Jonker P. G., Heinz S., Maccarone T. J., van der Klis M., 2006, *MNRAS*, 370, 1351

Gallo, E. 2007, *The Multicolored Landscape of Compact Objects and Their Explosive Origins*, 924, 715

Gal-Yam, A., et al. 2006, *ApJ*, 639, 331

Goodman, J. 1997, *NewAst*, 2, 449

Greiner, J., Peimbert, M., Estaban, C., Kaufer, A., Jaunsen, A., Smoke, J., Klose, S., & Reimer, O. 2003, *GRB Coordinates Network*, 2020, 1

Greisen, E. W. 2003, *Astrophysics and Space Science Library*, 285, 109

Grießmeier, J.-M., Zarka, P., & Spreeuw, H. 2007, *Astron. Astrophys.*, 475, 359

Gudel, M., Benz, A. O., Bastian, T. S., Furst, E., Simnett, G. M., & Davis, R. J. 1989, *Astron. Astrophys.*, 220, L5

Gültekin, K., Cackett, E. M., Miller, J. M., Di Matteo, T., Markoff, S., & Richstone, D. O. 2009, *Astrophys. J.*, 706, 404

Hallinan, G., et al. 2007, *Astrophys. J. Lett.*, 663, L25

Hallinan, G., Antonova, A., Doyle, J. G., Bourke, S., Lane, C., & Golden, A. 2008, *Astrophys. J.*, 684, 644

Halpern J. P., Filippenko A. V., 1984, *ApJ*, 285, 475

Hankins, T. H., Kern, J. S., Weatherall, J. C., & Eilek, J. A. 2003, *Nature*, 422, 141

Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, *Astron. Astrophys. Suppl. Ser.*, 117, 137

Hansen, B. M. S., & Lyutikov, M. 2001, *Mon. Not. R. Astron. Soc.*, 322, 695

Hameed S., Blank D. L., Young L. M., Devereux N., 2001, *ApJ*, 546, L97

Hessels J. W. T., Stappers B. W., van Leeuwen J., Transients Key Science Project L., 2009, arXiv, arXiv:0903.1447

Hodapp K. W., et al., 2004, *AN*, 325, 636

Högbom, J. A. 1974, *aaps*, 15, 417

Hopkins, A. M., Miller, C. J., Connolly, A. J., Genovese, C., Nichol, R. C., & Wasserman, L. 2002, AJ, 123, 1086

Hyman S. D., Lazio T. J. W., Kassim N. E., Bartleson A. L., 2002, AJ, 123, 1497

Hyman S. D., Wijnands R., Lazio T. J. W., Pal S., Starling R., Kassim N. E., Ray P. S., 2009, ApJ, 696, 280

Hjellming, R. M., & Bignell, R. C. 1982, Science, 216, 1279

Hunstead, R. W. 1972, ApJLet, 12, 193

Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Nature, 217, 709

Intema, H. T., van der Tol, S., Cotton, W. D., Cohen, A. S., van Bemmel, I. M., Rottgering, H. J. A. 2009, Astron. Astrophys., 501, 1185

Jackson, P. D., Kundu, M. R., & White, S. M. 1989, Astron. Astrophys., 210, 284

Jamil O., Fender R., Kaiser C., 2009, arXiv, arXiv:0909.1309

Jones, S., McHardy, I., Moss, D., Seymour, N., Breedt, E., Uttley, P., Kording, E., & Tudose, V. 2010, arXiv:1011.6633

Jonker, P. G., et al. 2010, MNRAS, 401, 1255

Johnston, S., et al. 2007, Publ. Astron. Soc. Aust., 24, 174

Johnston S., et al., 2008, ExA, 22, 151

Keane, E. F., Kramer, M., Lyne, A. G., Stappers, B. W., & McLaughlin, M. A. 2011, Mon. Not. R. Astron. Soc., 838

Kellermann, K. I., Sramek, R., Schmidt, M., Shaffer, D. B., & Green, R. 1989, ApJ, 98, 1195

Kettenis, M., van Langevelde, H. J., Reynolds, C., Cotton, B. 2006, Astronomical Data Analysis Software and Systems XV, 351, 497

Kirkpatrick, J. D., Henry, T. J., & Irwin, M. J. 1997, Astron. J., 113, 1421

Kida S., et al., 2008, NewA, 13, 519

Koratkar, A. P., Gaskell, C. M. 1991, APJS, 75, 719

Körding E., Falcke H., Corbel S., 2006, A&A, 456, 439

Kundu, M. R., White, S. M., Jackson, P. D., & Pallavicini, R. 1988, Astron. Astrophys., 195, 159

Kuniyoshi, M., et al. 2007, PASP, 119, 122

Lang, K. R., Bookbinder, J., Golub, L., & Davis, M. M. 1983, Astrophys. J. Lett., 272, L15

Lazio, T. J. W., Farrell, W. M., Dietrick, J., Greenlees, E., Hogan, E., Jones, C., & Hennig, L. A. 2004, Astrophys. J., 612, 511

Lazio, T. J. W., et al. 2010, AJ, 140, 1995

Lecacheux, A., Boudjada, M. Y., Rucker, H. O., Bougeret, J. L., Manning, R., &

Kaiser, M. L. 1998, *Astron. Astrophys.*, 329, 776

Lenc E., Garrett M. A., Wucknitz O., Anderson J. M., Tingay S. J., 2008, *ApJ*, 673, 78

Levinson, A., Ofek, E. O., Waxman, E., & Gal-Yam, A. 2002, *ApJ*, 576, 923

Lonsdale C. J., et al., 2009, arXiv, arXiv:0903.1828

Lorimer, D. R., et al. 2006, *Mon. Not. R. Astron. Soc.*, 372, 777

Lovell, B. 1969, *Nature*, 222, 1126

Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, *Science*, 318, 777

Maccarone T. J., Gallo E., Fender R., 2003, *MNRAS*, 345, L19

Maron, O., Kijak, J., Kramer, M., & Wielebinski, R. 2000, *Astron. Astrophys. Suppl. Ser.*, 147, 195

Marscher, A. P., Jorstad, S. G., Gómez, J.-L., Aller, M. F., Teräsranta, H., Lister, M. L., & Stirling, A. M. 2002, *nat*, 417, 625

Markoff S., et al., 2008, *ApJ*, 681, 905

Markoff, S. 2010, *Lecture Notes in Physics*, Berlin Springer Verlag, 794, 143

Matsumura, N., et al. 2009, *Astron. J.*, 138, 787

Miller-Jones J. C. A., Gallo E., Rupen M. P., Mioduszewski A. J., Brisken W., Fender R. P., Jonker P. G., Maccarone T. J., 2008, *MNRAS*, 388, 1751

McHardy I. M., Koerding E., Knigge C., Uttley P., Fender R. P., 2006, *Natur*, 444, 730

McLaughlin M. A., et al., 2006, *Nature*, 439, 817

McKean, J., et al. 2011, arXiv:1106.1041

Meegan, C. A., Fishman, G. J., Wilson, R. B., Horack, J. M., Brock, M. N., Paciesas, W. S., Pendleton, G. N., & Kouveliotou, C. 1992, *Nature*, 355, 143

Merloni A., Heinz S., di Matteo T., 2003, *MNRAS*, 345, 1057

Mészáros, P. 2002, *Ann. Rev. Astron. Astrophys.*, 40, 137

Metzger, M. R., Djorgovski, S. G., Kulkarni, S. R., Steidel, C. C., Adelberger, K. L., Frail, D. A., Costa, E., & Frontera, F. 1997, *Nature*, 387, 878

Mirabel, I. F., Dhawan, V., Chaty, S., Rodriguez, L. F., Marti, J., Robinson, C. R., Swank, J., & Geballe, T. 1998, *A&P*, 330, L9

Mirabel, I. F., & Rodríguez, L. F. 1998, *Nature*, 392, 673

Mutel, R. L., & Weisberg, J. M. 1978, *Astron. J.*, 83, 1499

Nandra K., Pounds K. A., 1994, *MNRAS*, 268, 405

Nelson C. H., Whittle M., 1995, *ApJS*, 99, 67

Niinuma, K., et al. 2007, *Astrophys. J. Lett.*, 657, L37

Niinuma, K., et al. 2009, *Astrophys. J.*, 704, 652

Noordam, J. E., & Smirnov, O. M. 2010, *Astron. Astrophys.*, 524, A61

Ofek E. O., Breslauer B., Gal-Yam A., Frail D., Kasliwal M. M., Kulkarni S. R., Waxman E., 2009, arXiv, arXiv:0910.3676

Ofek, E. O., Frail, D. A., Breslauer, B., Kulkarni, S. R., Chandra, P., Gal-Yam, A., Kasliwal, M. M., & Gehrels, N. 2011, arXiv:1103.3010

Ofek, E. O., & Frail, D. A. 2011, *Astrophys. J.*, 737, 45

Offringa, A. R., de Bruyn, A. G., Zaroubi, S., & Biehl, M. 2010, arXiv:1007.2089

Oosterloo, T., Verheijen, M., & van Cappellen, W. 2010, ISKAF2010 Science Meeting,

Oster, L., & Sieber, W. 1978, *Astron. Astrophys.*, 65, 173

Palmer, D. M., Barthelmy, S. D., Cummings, J. R., Gehrels, N., Krimm, H. A., Markwardt, C. B., Sakamoto, T., & Tueller, J. 2005, *The Astronomer's Telegram*, 546, 1

Panessa, F., Barcons, X., Bassani, L., Cappi, M., Carrera, F. J., Ho, L. C., & Pellegrini, S. 2007, *A&A*, 467, 519

Penzias, A. A., & Wilson, R. W. 1965, *Astrophys. J.*, 142, 419

Peterson, B. M., Wanders, I., Horne, K., Collier, S., Alexander, T., Kaspi, S., Maoz, D. 1998, *PASP*, 110, 660

Phan-Bao, N., Osten, R. A., Lim, J., Martín, E. L., & Ho, P. T. P. 2007, *Astrophys. J.*, 658, 553

Pooley, G. G., Fender, R. P. 1997, *MNRAS*, 292, 925

Ptak, A. 2001, X-ray Astronomy: Stellar Endpoints, AGN, and the Diffuse X-ray Background, 599, 326

Rau A., et al., 2009, arXiv, arXiv:0906.5355

Rawlings, S., & Schilizzi, R. 2011, arXiv:1105.5953

Rees M. J., 1978, *MNRAS*, 184, 61P

Rengelink, R. B., Tang, Y., de Bruyn, A. G., Miley, G. K., Bremer, M. N., Roettgering, H. J. A., & Bremer, M. A. R. 1997, *A&A*, 124, 259

Rickett, B. J., Coles, W. A., & Bourgois, G. 1984, *Astron. Astrophys.*, 134, 390

Rickett, B. J. 1986, *Astrophys. J.*, 307, 564

Rickett, B. J. 1977, *Ann. Rev. Astron. Astrophys.*, 15, 479

Rushton, A., Spencer, R., Fender, R., & Pooley, G. 2010, *Astron. Astrophys.*, 524, A29

Russell, C. T. 1991, *Space Science Review*, 55, 317

Russell, D. M., Lewis, F., Roche, P., Clark, J. S., Breedt, E., & Fender, R. P. 2009, arXiv:0911.4501

Ryle, M., & Hewish, A. 1960, *Mon. Not. R. Astron. Soc.*, 120, 220

Ryle, M., Hine, G., Shakeshaft, J., & Caswell, J. L. 1978, *Nature*, 276, 571

Sault, R. J., & Wieringa, M. H. 1994, *Astron. Astrophys. Suppl. Ser.*, 108, 585

Sault R. J., Teuben P. J., Wright M. C. H., 1995, *ASPC*, 77, 433

Scheers, B. *Transient and Variable Radio Sources in the LOFAR Sky. An Architecture for a Detection Framework.*, PhD Thesis, University of Amsterdam, 2010

Sikora, M., Stawarz, L., & Lasota, J.-P. 2007, *Astrophys. J.*, 658, 815

Simonetti, J. H., & Cordes, J. M. 1990, *Astrophys. J.*, 349, 97

Singh, K., Bähren, L., Falcke, H., & et al. 2008, International Cosmic Ray Conference, 4, 429

Slee, O. B., Sadler, E. M., Reynolds, J. E., Ekers, R. D. 1994, *MNRAS*, 269, 928

Slee, O. B., Willes, A. J., & Robinson, R. D. 2003, *Publ. Astron. Soc. Aust.*, 20, 257

Spada M., Ghisellini G., Lazzati D., Celotti A., 2001, *MNRAS*, 325, 1559

Spreeuw, J. N. *Search and detection of low frequency radio transients*, PhD Thesis, University of Amsterdam, 2010

Spangler, S. R., Shawhan, S. D., & Rankin, J. M. 1974, *Astrophys. J. Lett.*, 190, L129

Spangler, S. R., & Moffett, T. J. 1976, *Astrophys. J.*, 203, 497

Spangler, S., Fanti, R., Gregorini, L., & Padrielli, L. 1989, *Astron. Astrophys.*, 209, 315

Smirnov, O. M. 2011, *Astron. Astrophys.*, 527, A106

Smirnov, O. M. 2011, *Astron. Astrophys.*, 527, A107

Smirnov, O. M. 2011, *Astron. Astrophys.*, 527, A108

Stappers, B. W., et al. 2011, *Astron. Astrophys.*, 530, A80

Starling R. L. C., Page M. J., Branduardi-Raymont G., Breeveld A. A., Soria R., Wu K., 2005, *Ap&SS*, 300, 81

Strassmeier, K. G., Hall, D. S., Fekel, F. C., & Scheck, M. 1993, *Astron. Astrophys. Suppl. Ser.*, 100, 173

Soffitta, P., et al. 1998, *IAU Circular*, 6884, 1

Soleri, P., et al. 2010, *Mon. Not. R. Astron. Soc.*, 406, 1471

Swarup, G. 1990, *Indian Journal of Radio and Space Physics*, 19, 493

Swinbank J., 2007, in Tzioumis T., ed., *Proc. Bursts, Pulses and Flickering: Wide-Field Monitoring of the Dynamic Radio Sky*. p. 44

Taylor, A. R., & Gregory, P. C. 1983, *Astron. J.*, 88, 1784

Taylor, G. B., Carilli, C. L., & Perley, R. A. 1999, *Synthesis Imaging in Radio Astronomy II*, 180,

Taylor, J. H., Fowler, L. A., & McCulloch, P. M. 1979, *Nature*, 277, 437

Timmer J., Koenig M., 1995, A&A, 300, 707

Terashima, Y., & Wilson, A. S. 2003, Active Galactic Nuclei: From Central Engine to Host Galaxy, 290, 403

Tinney, C., Stathakis, R., Cannon, R., & Galama, T. 1998, IAU Circular, 6896, 3

Thean, A., Pedlar, A., Kukula, M. J., Baum, S. A., O'Dea, C. P. 2001, MNRAS, 325, 737

Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 2001, ISRN Astronomy and Astrophysics,

Thyagarajan, N., Helfand, D. J., White, R. L., & Becker, R. H. 2011, arXiv:1107.5901

Turner T. J., Pounds K. A., 1989, MNRAS, 240, 833

Turner, T. J., Nandra, K., Turcan, D., & George, I. M. 2001, X-ray Astronomy: Stellar Endpoints, AGN, and the Diffuse X-ray Background, 599, 991

Tzioumis, A. K. 1997, Vistas in Astronomy, 41, 311

van den Oord G. H. J., de Bruyn A. G., 1994, A&A, 286, 181

van der Horst, A. J., et al. 2008, Astron. Astrophys., 480, 35

van der Horst, A. J., Rol, E., Wijers, R. A. M. J., Strom, R., Kaper, L., & Kouveliotou, C. 2005, Astrophys. J., 634, 1166

van der Laan, H. 1966, NAT, 211, 1131

van Paradijs, J., et al. 1997, Nature, 386, 686

van Paradijs, J., Kouveliotou, C., & Wijers, R. A. M. J. 2000, Ann. Rev. Astron. Astrophys., 38, 379

van Leeuwen, J., & Stappers, B. W. 2010, Astron. Astrophys., 509, A7

Verheijen, M. A. W., Oosterloo, T. A., van Cappellen, W. A., Bakker, L., Ivashina, M. V., & van der Hulst, J. M. 2008, The Evolution of Galaxies Through the Neutral Hydrogen Window, 1035, 265

de Vries, W. H., Becker, R. H., White, R. L., & Helfand, D. J. 2004, AJ, 127, 2565

de Vos, M., Gunst, A. W., & Nijboer, R. 2009, IEEE Proceedings, 97, 1431

Waldman, S. J. 2011, arXiv:1103.2728

Warwick, J. W., et al. 1981, Science, 212, 239

Weiler, K. W., Panagia, N., Montes, M. J., & Sramek, R. A. 2002, Ann. Rev. Astron. Astrophys., 40, 387

Weiler, K. W., van Dyk, S. D., Sramek, R. A., & Panagia, N. 2004, NewAR, 48, 1377

Welch J., et al., 2009, arXiv, arXiv:0904.0762

Wilson-Hodge, C. A., et al. 2011, Astrophys. J. Lett., 727, L40

Welch J., et al., 2009, arXiv, arXiv:0904.0762

Wells, D. C., Greisen, E. W., & Harten, R. H. 1981, *Astron. Astrophys. Suppl. Ser.*, 44, 363

White, S. M., Jackson, P. D., & Kundu, M. R. 1989, *ApJS*, 71, 895

White, R. L., Becker, R. H., Helfand, D. J., Gregg, M. D. 1997, *Astrophys. J.*, 475, 479

Xue, Y. Q., & Cui, W. 2007, *AAp*, 466, 1053

Zarka, P., & Pedersen, B. M. 1986, *Nature*, 323, 605

Zarka, P. 2004, *Planetary Science*, 52, 1455

Zarka, P., Farrell, W. M., Kaiser, M. L., Blanc, E., & Kurth, W. S. 2004, *Planetary Science*, 52, 1435

Zhao J.-H., et al., 1992, *Sci*, 255, 1538