

Multifrequency Search & Classification of the Transient Radio Sky



Adam Stewart

Department of Physics

University of Oxford

This dissertation is submitted for the degree of

Doctor of Philosophy

Brasenose College

February 2017

To R.A. & T.P. for putting me on the path.

Declaration

I certify that the work in this thesis is wholly my own, except where explicitly stated.

Parts of this work have been published as

Stewart, A. J., Fender, R. P., Broderick, J. W., et al., *LOFAR MSSS: Detection of a low-frequency radio transient in 400 hrs of monitoring of the North Celestial Pole*, 2016, MNRAS, 456, 2321

Fender R., Stewart A., Macquart J. P., Donnarumma I., Murphy T., Deller A., Paragi Z., Chatterjee S., *The Transient Universe with the Square Kilometre Array*, 2015, AASKA14, p. 51

Parts of this work will soon be published as

Stewart, A. J., Muñoz-Darias, T. & Fender, R. P., *The optical counterparts of radio transients and variables*, 2017, submitted to MNRAS

Adam Stewart

February 2017

Acknowledgements

Firstly, I would like to thank my supervisor Rob Fender for giving me the opportunity to work with him and be part of 4 Pi Sky. While the number of LOFAR transients may have been slightly over estimated at the beginning, the PhD has allowed me to meet and work with brilliant people and experience visiting places I never expected to. For this I am very grateful. Also, thanks for the patience when it came to the 'science stuff'.

Much thanks is also given to all those I worked with over the years, though none more so to Jess Broderick, who managed to get through my constant questioning about radio astronomy, LOFAR and pretty much anything else I had questions on in or outside of astronomy. In return I gave you the stress of becoming a competitive Fantasy Premier League manager which I expect you'll be forever grateful. Thanks to the entire 4 Pi Sky group: Teo, Dan and Aidan for sharing the passion in football (shout-out to the Hurricanes charity football team!), Rene for the discussions on whether the latest Apple product was worth it (usually no but yes), Tim for whom I felt guilty every time I sent one of the three thousand LOFAR machine support emails and to Gemma, Tom, Retha, Tony and Mari who made working in 4 Pi Sky so enjoyable. I'd also like to thank those that were in Southampton in the early years of the PhD, in particular Indi, Maxi and Volker, and also the Southampton astronomy group for developing my interest in astronomy during my undergraduate years.

I would also like to say thank you to all those I have collaborated with from LOFAR, the Transients Key Science project and elsewhere with an extra thank you to George Heald for your general help and making sure that I had safely returned from hiking up Table Mountain. A mention also to Daniel Cunnam, a networking success from a conference in Stellenbosch where we have spent the years after collaborating about FPL on a weekly basis to this day.

This PhD would not have been possible without the huge and constant support from my parents over the years of my academic pursuit, including lending a big hand in the various unexpected moves during the five years spanning three cities. Also to my brother Graeme for the various proof reads and maths checks. Though I don't blame you for passing up this particular proof read. Also a special mention to Tom Bennett from back home who, on the many train journeys to Southampton, really helped me in the early undergraduate days to set me on this course.

Finally, a thank you to Gogo Pietka - from Glasgow to Nottingham, we journeyed through the PhD together and it simply couldn't have been more fun. Got there in the end.

Abstract

The radio sky can exhibit dynamic changes that provide an excellent opportunity to study the most exotic events in the Universe. In recent years the development and completion of advanced radio facilities has enabled the transient radio sky to be probed to unprecedented levels, especially at low frequencies with telescopes such as the Low-Frequency Array (LOFAR) revolutionising observations in this regime.

This work primarily focuses on the detection of radio transients both at low and high frequencies. I describe the analysis of a 60 MHz LOFAR monitoring campaign centred at the North Celestial Pole, and report the discovery of the first LOFAR transient source. I also present the initial transient search of the CHILES survey that was performed with the VLA at 1.4 GHz, supported by simultaneous optical observations. Three strong transient candidates are discussed along with identifying a further 72 possible candidates.

Furthermore, I investigate the relation between the radio and optical flux properties of known transient and variable sources, with the aim of developing a classification method for these types of events. I show that from just radio and optical flux measurements, it is possible to separate certain classes of objects, with additional information provided by analysing how individual classes evolve in the radio and optical parameter space over time.

Lastly, I explore the future prospects of the LOFAR and CHILES transient searches, detailing what techniques could be used to further define the transient parameter space. I also discuss how the radio and optical classification method can be advanced and incorporated into automatic transient pipelines, which may be a necessity in the future Square Kilometre Array era.

Table of contents

List of figures	xi
List of tables	xv
1 Introduction	1
1.1 A Brief History of Radio Astronomy	1
1.2 Interferometry & Aperture Synthesis	5
1.3 Radio Emission Mechanisms	9
1.3.1 Synchrotron	9
1.3.2 Bremsstrahlung	10
1.3.3 Black Body	11
1.3.4 Coherent	12
1.4 Propagation Effects	13
1.4.1 Dispersion	13
1.4.2 Scattering	14
1.4.3 Ionosphere	14
1.5 Radio Transients	17
1.5.1 Slow Transients	19
1.5.2 Fast Transients	25
1.6 Radio Transient Surveys	29
1.7 This Work	35

2	The Low-Frequency Array	36
2.1	LOFAR Design & Capabilities	36
2.1.1	Low Band Antennas	38
2.1.2	High Band Antennas	38
2.1.3	Stations	39
2.1.4	<i>uv</i> coverage & Sensitivity	42
2.1.5	Pointing LOFAR	44
2.1.6	Signal Path	45
2.1.7	The Long Term Archive	46
2.2	The Transients Key Science Project	47
2.2.1	The Radio Sky Monitor	48
2.2.2	The V404 Cygni 2015 Outburst	53
2.3	The Multifrequency Snapshot Sky Survey	56
2.3.1	Survey Design	56
2.3.2	Commissioning Work	57
2.4	LOFAR Commissioning & TKP Publications	60
2.5	LOFAR Repairs	61
3	Searching for Transients at the NCP with LOFAR	63
3.1	Introduction	63
3.1.1	Previous NCP Transient & Variable Studies	64
3.2	LOFAR Observations of the NCP	65
3.2.1	Observation Epochs	66
3.2.2	Calibration & Imaging	67
3.2.3	Pre-Processing Stages	67
3.2.4	Main Processing Stages	68
3.2.5	Quality Control	70

3.3	Transient & Variability Search Method	71
3.3.1	Time-scales Searched	71
3.3.2	The Transients Pipeline	72
3.4	Results	75
3.4.1	Image Quality	75
3.4.2	Variability Search Results	83
3.4.3	Transient Search Results	86
3.5	Transient Candidate ILT J225347+862146	87
3.5.1	A Mirrored Ghost Source	89
3.5.2	Transient Flux Density	95
3.5.3	Testing if a Source Can be Created by the Sky Model	101
3.5.4	Further Validity Testing	101
3.5.5	What is this Transient?	104
3.6	Transient Candidate ILT J130641+865946	112
3.7	Transient Surface Density & Rates	114
3.7.1	Comparison to Other Transient Surveys	116
3.7.2	Comparison of Time-dependent Surface Densities	117
3.8	Conclusions	119
4	The Optical Counterparts of Radio Transients and Variables	121
4.1	Introduction	121
4.2	Data Collection & Methods	124
4.2.1	Quasars	127
4.2.2	Stellar Objects	128
4.2.3	Isolated Neutron Stars	129
4.2.4	Cataclysmic Variables	129
4.2.5	X-ray Binaries	130

4.2.6	Supernovae & Gamma-ray Bursts	130
4.2.7	Dynamic Data	131
4.3	Results	133
4.3.1	Description of the Base Result	133
4.3.2	Dynamic Characteristics	136
4.4	Discussion	140
4.4.1	Extending the Quasar & Stellar Populations	141
4.5	Transient Diagnostic Application	144
4.5.1	Diagnostic Test	147
4.6	Conclusions	150
5	The Initial Transient Search of the CHILES 1.4 GHz VLA Survey	156
5.1	Introduction	156
5.1.1	Predicted Numbers of Transient & Variable Objects	157
5.2	The CHILES 2013–2015 Dataset	158
5.2.1	Calibration & Imaging	158
5.2.2	Radio Data Quality	161
5.2.3	Calibrators	162
5.2.4	Deep Image	164
5.2.5	Optical Data	164
5.3	Transient & Variable Search - Details & Methods	164
5.3.1	Time-scales Probed	167
5.3.2	TraP Settings	167
5.3.3	Interpreting the TraP Results	170
5.3.4	Method I - Experimental Automated Rapid Search	172
5.3.5	Method II - Offline Deep Search	173
5.4	Offline Transient Search Results	174

Table of contents	x
5.4.1 Advantages of the Deep Image	174
5.4.2 V_V Candidates	176
5.4.3 η_V Candidates (Variables)	177
5.4.4 New Sources (Transients)	183
5.5 Transient Surface Density	197
5.6 Conclusions	199
6 Conclusions & Future Work	208
6.1 The First LOFAR Radio Transient	208
6.1.1 Developments Since the Completion of the NCP Project	209
6.2 The Optical Counterparts of Radio Transients	210
6.3 The Initial Transient Search of CHILES	211
6.4 Future Work	212
6.4.1 LOFAR Transients	212
6.4.2 Radio & Optical Flux Transient Diagnostic	215
6.4.3 CHILES Transients	215
6.5 Final Remarks	217
References	219

List of figures

1.1	Radio telescopes through the ages	4
1.2	The uv plane	7
1.3	Examples of uv coverage	8
1.4	The spectrum of synchrotron emission	11
1.5	Comparison of spectral shapes between emission mechanisms	12
1.6	Examples of dispersion	15
1.7	Scattering of a signal	15
1.8	The ionosphere and long baselines	17
1.9	Luminosity and time-scale phase space of radio transients	18
1.10	Example of AGN: Cygus A	20
1.11	Illustration of an X-ray binary and tidal disruption event	21
1.12	Example of GRB: GRB 050709	22
1.13	Example of SN: SN 2013cu	23
1.14	Illustration of a Cataclysmic Variable and a Solar Flare	24
1.15	Jupiter's radiation belts	26
1.16	Pulsar schematic diagram and example of an FRB	27
1.17	Low-frequency radio transient surface density limits from the literature . . .	33
1.18	High-frequency radio transient surface density limits from the literature . .	34
2.1	Schematic map of LOFAR across Europe	37

2.2	LOFAR LBA photograph and spectrum	39
2.3	LOFAR HBA photograph and spectrum	40
2.4	LOFAR stations design	41
2.5	Aerial view of the Superterp	42
2.6	LOFAR uv coverage	43
2.7	How LOFAR observes in a specific direction	45
2.8	Radio Sky Monitor observing technique	49
2.9	Processing steps of the Radio Sky Monitor Pipeline	51
2.10	Examples of Radio Sky Monitor fields	52
2.11	LOFAR detection of V404 Cygni	55
2.12	LOFAR and AMI lightcurve of V404 Cygni	55
2.13	Early example of MSSS-LBA peeling	59
2.14	Example of HBA peeling improvement	60
2.15	Fixing LOFAR-UK photo	62
3.1	Number of NCP observations per day	66
3.2	Examples of the NCP field maps at different time-scales	76
3.3	uv coverage of the NCP observations	77
3.4	Normalised primary beam of the NCP observations	78
3.5	Average rms for each NCP time-scale dataset	79
3.6	NCP sources flux and position compared to VLSS	81
3.7	NCP deep map	82
3.8	Variability parameters distribution for 55 and 297 min time-scales	85
3.9	Transient candidate ILT J225347+862146 and ghost source	90
3.10	Simulated transient and ghost source	93
3.11	Transient extracted flux and sky model flux comparison	97
3.12	ILT J225347+862146 2 min light curve	99

3.13	Time sequence images of ILT J225347+862146	102
3.14	Optical follow up image of ILT J225347+862146 field	106
3.15	Transient candidate ILT J130641+865946	113
3.16	Log N, Log S plot of NCP and previous transient surface densities	117
3.17	Log N, Log S with time-scale information of transient surface densities	118
4.1	Distribution of the sample radio frequencies and optical bands	126
4.2	F_r-F_0 distribution plot	134
4.3	Schematic representation of the F_r-F_0 distribution	135
4.4	Light curves of the dynamic sources	138
4.5	Dynamic sources tracks plotted on F_r-F_0 distribution	139
4.6	Distribution of F_r/F_0 ratio per class	142
4.7	Simulating the extension of the quasar and stellar populations	144
4.8	FIRST transient and variable sources plotted on F_r-F_0 distribution	149
5.1	Histogram of CHILES integration times	159
5.2	CHILES data quality: measured rms	162
5.3	CHILES data quality: rms relation to integration time	163
5.4	CHILES calibrator light curves	163
5.5	Preliminary deep CHILES image	165
5.6	Optical CHILES mosaic image	166
5.7	Distribution time difference between consecutive CHILES observations	169
5.8	Flowchart of rapid transient search pipeline	173
5.9	Distribution of η_v and V_v for the CHILES dataset	175
5.10	V_v candidate #15741 light curve	178
5.11	Variable η_v candidate #15701 light curve	179
5.12	η_v candidate #16037 light curve	180
5.13	η_v candidates postage stamps	182

5.14	η_V candidate #15701 TraP vs. BLOBCAT light curves	183
5.15	Example objects of transient candidate categories	186
5.16	Transient candidate #17230	189
5.17	Transient candidate #17289	190
5.18	Transient candidate #17333	192
5.19	Radial distribution of CHILES transient candidates	196
5.20	Positions and selection of CHILES bright sources	196
5.21	Distribution of candidate distances to nearest bright source	197
5.22	Radio transient Log-N, Log-S plot including CHILES results	200
6.1	LOFAR 2014 NCP data image	214
6.2	CHILES subtraction example	217

List of tables

1.1	Summary of radio transient surveys	32
2.1	LOFAR sensitivities	44
3.1	NCP calibrators	69
3.2	Details of NCP transient search datasets	77
3.3	Results of candidate validity tests	88
3.4	ILT J225347+862146 optical follow up observation details	105
3.5	Estimated pulse width of FRBs at 60 MHz	111
3.6	Previous surveys and NCP transient surface densities	115
4.1	Number of F_r and F_o measurements in sample	126
4.2	Details of the dynamic data sample	132
4.3	Number of sources in FIRST transients and variables sample	150
4.4	Radio pulsar sample information	152
4.5	CV sample information	152
4.6	XRB sample information	152
4.7	SN sample information	153
4.8	GRB sample information	154
5.1	Estimated number of CHILES variable and transient objects	158
5.2	Details of the 2013–2015 CHILES dataset	159

5.3	TraP source extraction settings	168
5.4	Total number of CHILES variable and transient candidates	176
5.5	Transient candidate categories	185
5.6	Time-scale analysis of three transient candidates	194
5.7	CHILES radio transient surface density properties	198
5.8	Details of η_v variable candidates	203
5.9	Details of CHILES 75 transient candidates	205
6.1	Comparing the old and new NCP surveys	213

Chapter 1

Introduction

1.1 A Brief History of Radio Astronomy

In the early 1930s an engineer working for Bell Laboratories, named Karl Jansky, was tasked with determining the levels of radio interference at 20.5 MHz. Jansky found that the background mostly consisted of emission from thunderstorms, however, there was one component, described as a steady hiss, that could not be attributed to any known source. After investigating further he discovered that the interference varied along with the sidereal day, and that the peak of the noise could be found in the direction of the constellation of Sagittarius. Jansky concluded that the origin of the signal must be from sources beyond the Earth - the serendipitous discovery of astronomical radio emission (Jansky, 1933).

The discovery by Jansky did not have an immediate impact on the astronomical community. Instead, it would not be until the post-war years that significant advancements were made in the field of radio astronomy, helped by the development of radio technologies during World War II. During the 1950s and 1960s a number of major discoveries were made through radio astronomy such as the observation of the theorised 21-cm hydrogen line (Ewen & Purcell, 1951; Muller & Oort, 1951; Pawsey, 1951), quasars (Schmidt, 1963), the cosmic microwave background (Penzias & Wilson, 1965) and pulsars (Hewish et al., 1968).

Contrary to the early pioneering work in radio astronomy which concentrated on metre wavelengths, the following decades saw a shift towards observing the radio sky at higher, gigahertz frequencies. This was primarily due to the desire to have high angular resolution, along with rapid developments in areas such as receiver technology, interferometry and aperture synthesis, advancements that will be returned to in this chapter.

During the late 1980s and 1990s, a renewed interest in low-frequency radio astronomy began to emerge, due to a variety of scientific reasons (inverted radio spectra, steep spectrum objects). A major factor however, was the ambition to be able to detect neutral hydrogen at cosmological distances, the one known method of probing the dark ages of the Universe. For example, due to the expansion of the Universe and by using the relation of $\lambda_{\text{obs}} = \lambda_0(1 + z)$, where λ_{obs} and λ_0 are the observed and original wavelengths, respectively, and z is the redshift, the 21-cm signal 550 million years after the big bang ($z = 9$) now appears at 2.1 m - a frequency of 140 MHz. It was suggested that a telescope with a collecting area of one square kilometre would be required to detect the cosmological emission (Wilkinson, 1991), a suggestion that has now formed into the global project of the Square Kilometre Array (SKA; Dewdney et al. 2009; Ekers 2012). The SKA will be the largest radio telescope ever constructed and, in its first phase, will be split into two components: ‘SKA Low’ in Australia covering a frequency range of 50–350 MHz, and ‘SKA Mid’ in South Africa spanning the frequency range of 0.350–14 GHz. The construction of the first phase of the SKA is scheduled to begin in 2018, with the first observations taking place in ~ 2020 . Such an ambitious project meant that the current era of radio astronomy has somewhat been defined by precursor and pathfinder instruments, dedicated to developing and testing technology which will be incorporated into the SKA design. This includes instruments such as MeerKAT (Booth & Jonas, 2012) in South Africa, and telescopes in Australia: The Australian SKA Pathfinder (ASKAP; Johnston et al. 2007, 2008) and the Murchison Widefield Array (MWA; Tingay et al. 2013). This thesis makes extensive use of a European telescope, the Low-Frequency Array (LO-

FAR; van Haarlem et al. 2013), an instrument that has indirectly also become a technology precursor for the SKA, which is detailed in Chapter 2. Jansky's original array and other telescopes through the ages, including an artists impression of the SKA, can be seen in Figure 1.1.

The design of a radio telescope needs to meet two fundamental requirements: sensitivity and angular resolution. The sensitivity is maximised by having the largest collecting area possible, in addition to minimising the contribution of unwanted noise (e.g. from the electronics). The angular resolution on the other hand is governed by the physical dimensions of the telescope. The latter point presents a unique challenge to the design and construction of radio telescopes because of the wavelengths involved in the radio band. The angular resolution (θ) is defined by $\theta \sim \lambda/D$, where λ is the wavelength and D is the telescope's diameter. For example, the Hubble Space telescope has a 2.4 m mirror allowing for an angular resolution of approximately 0.04 arcsec at a visible wavelength of 500 nm. To replicate this angular resolution at 14.6 m (20.5 MHz), the wavelength at which Jansky observed the sky, a radio telescope of a diameter 70,000 km would be required. Considering the diameter of the Earth is 12,700 km, this telescope would be $5.5 \times$ the size of the Earth - one can appreciate why radio telescopes are large in size! This also demonstrates why higher gigahertz frequencies were more popular in the late 20th Century in a bid to obtain the best angular resolution possible. Examples of single dish telescopes include the 76.2 m diameter Lovell Telescope in the UK; the Green Bank Telescope in the USA that has a dish diameter of 100 m, and the Arecibo Observatory in Puerto Rico - currently the largest operational single dish telescope in the world with a diameter of 305 m¹.

Constructing ever larger single aperture dishes in order to improve the angular resolution quickly becomes impractical in terms of both cost and structural design. However, this problem can be solved by using the techniques of interferometry and aperture synthesis (Ryle

¹Once complete, the Five hundred meter Aperture Spherical Telescope (FAST) in south west China will become the largest physical single dish with a diameter of 500 m (effective diameter of ~ 300 m).




Fig. 1.1 Radio telescopes through the ages. (i) Jansky's 'merry-go-round' antenna used to detect astronomical radio emission for the first time in the early 1930s; (ii) a current photograph of the 76.2 m diameter Lovell telescope in Cheshire, England, completed in 1957; (iii) the world's current largest single aperture radio telescope, Arecibo (305 m diameter) in Puerto Rico, built in 1963; (iv) the Very Large Array in New Mexico, USA, constructed during the 1970s and consisting of 27 antennas; (v) one of the 128 tiles of the Murchison Widefield Array in Western Australia, a low-frequency telescope completed in 2012; (vi) a single dish of the MeerKAT telescope which is currently under construction, the final array will contain 64 such dishes and (vii) an artist impression of the Square Kilometre Array telescope, the South African portion, with first light hoped to be achieved in 2020.

& Hewish, 1960). The single dish approach can be replaced by a number of smaller, cost-effective dishes spread over an area forming an array; the signals received by each antenna are then combined using interferometry effectively mimicking a single aperture telescope. When observing at metre wavelengths, the individual antennas can be formed by arrays of dipoles (the SKA for example will use both dishes and arrays of dipoles). The diameter of the telescope is then determined by the maximum distance between any two antennas, called the ‘baseline length’. This allows for significantly larger ‘diameters’ than those which could be achieved with a single dish, thus greatly improving the angular resolution. While the method allows for more manageable physical structures to be built, the transport and analysis of the vast amount of data produced by an interferometric system presents its own unique challenge, requiring very fast network connections and often supercomputers to perform the processing. The technique was first utilised by Ryle & Vonberg (1946) to find the angular diameter of sunspot radiation, and since then many other interferometer arrays have been constructed such as: the Karl G. Jansky Very Large Array (VLA; Perley et al. 2011) in the USA, the Westerbork Synthesis Radio Telescope (WSRT; Baars & Hooghoudt 1974) in the Netherlands and the Australian Telescope Compact Array (ATCA; Frater et al. 1992) in Australia.

The work presented in this thesis uses two interferometric arrays, LOFAR and the VLA, to observe the radio transient sky. Hence, what follows is a brief description of the aperture synthesis method.

1.2 Interferometry & Aperture Synthesis

In an array, a radio signal is detected by each baseline pair of antennas, in which the signals are combined together by accounting for the geometrical path delay (such a delay is discussed further in the context of LOFAR in Chapter 2). The combined signal forms an interference pattern that arises from the alternately constructive and destructive interference of the signals

from each antenna, also known as the fringe pattern. The length of the baseline corresponds to what source sizes the baseline itself is sensitive to; short baselines produce wide fringe patterns (i.e. low resolution) whereas long baselines result in much more narrow fringe patterns (higher resolution). It is for this reason that an array will commonly consist of a variety of baseline lengths, such that the instrument is sensitive to both large and small scale structures on the sky. The fringe pattern is most often expressed as a complex number of a defined amplitude and phase, named a complex visibility V , which is one component of the Fourier transform of the sky brightness distribution. The precise position of the component depends upon how the baseline is oriented relative to the source of interest in the radio sky, and this is expressed with the use of the u , v and w co-ordinate system as shown in Figure 1.2. In this system, u and v define the projected easterly and northerly directions, and w is directed towards the observed object in the sky. The latter w term can be neglected however if the array antennas can be considered as being coplanar. Each visibility measurement, $V(u,v)$, from a single baseline at a specific time provides a single point on the so called uv plane (also refer to Figure 1.2).

Sampling of the uv plane depends on the amount of baselines that exist in the array and the physical positioning of the antennas, in addition to the elevation of the observed source. Furthermore, observations can be allowed to run for a period of time, during which the projected lengths of the baselines change due to the rotation of the Earth. This effectively ‘fills in’ the uv plane, sampling long arcs or rings of Fourier components as opposed to discrete points. The technique of exploiting the Earth’s rotation is referred to as aperture synthesis. Figure 1.3 shows the uv coverage over different observing times for two telescopes: the WSRT, which is purely an East-West (E-W) array, and the VLA, which is comprised of 27 antennas arranged in a two dimensional ‘Y’ configuration. It can be seen that the instantaneous uv coverage for the WSRT E-W array is quite poor in comparison to the good instantaneous coverage obtained by the VLA. Such short observations are often referred to as




Fig. 1.2 The uv plane. This schematic diagram demonstrates how baseline and source positions relate to the uv plane. The hollow, dashed baseline labelled B , represents the projected baseline when the source being observed is straight overhead, represented by the dashed source. In this case the uv plane would be in the same plane as the telescope plane. It is then shown how the orientation changes when observing an offset source, and how the baseline is now projected differently, labelled B' .

‘snapshots’. The figure then shows how the uv coverage is improved by a full 12 hr synthesis, which is especially true for the WSRT.

Taking a raw Fourier transform of the visibility measurements produces a ‘dirty map’, which is the sky brightness distribution convolved with the response of the interferometric array to the point source, otherwise known as the ‘dirty beam’. The form of the dirty beam depends upon how well the uv plane was sampled during the observation. To obtain the true sky map the process of deconvolution can be applied. There are various different deconvolution methods but the most commonly used is the CLEAN algorithm (Högbom, 1974; Clark, 1980). It works by assuming that the sky can be represented by a number of point sources. The algorithm then locates the brightest point in absolute intensity in the dirty image and subtracts the dirty beam from this location. The position and magnitude of the brightest point is recorded in a model and the step is then repeated in an iterative manner until




Fig. 1.3 Comparison of the uv coverage between the WSRT East-West array (left) and the VLA (right). The top panels show the array design for each respective telescope. The middle panels show the instantaneous uv coverage for each telescope, where it can be seen that the WSRT uv coverage is very sparse when compared to that of the VLA. The bottom panels show the uv coverage of a full 12 hr synthesis, which greatly improves the coverage for the WSRT, highlighting its dependency on the Earth's rotation. *Figure credit: <http://192.96.5.2/synthesisschool/MiodArrayDesign.pdf>.*

the peak is below a user-defined level. The recorded point source model is then convolved with a CLEAN beam to produce the final map.

1.3 Radio Emission Mechanisms

This section provides an overview of the different types of radio emission mechanisms. From here on, the frequency dependence of the flux density F_ν at frequency ν is defined as $F_\nu \propto \nu^\alpha$, where α is the spectral index. For more comprehensive overviews of the discussed subjects I refer the reader to works such as Rybicki & Lightman (1986) and Longair (2011).

1.3.1 Synchrotron

Much of high energy astrophysics is dominated by synchrotron radiation, a process which is responsible for emission such as that from our Galaxy, supernova remnants and extragalactic radio sources. It is produced by the acceleration of relativistic charged particles, commonly electrons, in a magnetic field due to the Lorentz force. The charged particles follow a spiral path along the direction of the magnetic field, with the frequency of the emitted radiation governed by the strength of magnetic field and the energy of the particle. Due to the relativistic nature of the particle, the emission is confined to a narrow cone pointing in the direction of motion; this is known as ‘beaming’. Cyclotron and gyro-synchrotron radiation refers to emission from non-relativistic and mildly relativistic charged particles accelerated in a magnetic field, respectively.

Synchrotron radiation is usually (but not exclusively) a type of non-thermal emission; the emission does not depend on the temperature of the emitting source. In other words, the energy distribution of the emitting particles cannot be described by a Maxwellian distribution. Instead, a population of synchrotron emitting electrons can be described by the power law energy distribution $N(E)dE = N_0 E^{-p} dE$, where p is the energy distribution index. Within

the optically thin regime, the spectrum of synchrotron emission is derived by summing the contribution of the emission spectra from the individual electrons. The electrons emit at a range of frequencies as they accelerate around the magnetic field lines, peaking at the so-called ‘critical frequency’. The spectral index is related to the energy distribution index by $\alpha = (1 - p)/2$, and is usually a value in the range of $-2 \leq \alpha \leq 0$.

As the intensity of synchrotron radiation within a source reaches a significantly high point, the spiralling electrons can begin to reabsorb low energy photons. This is called ‘synchrotron self-absorption’. At sufficiently low frequencies (e.g. ~ 100 MHz) the brightness temperature², T_b , of the source will reach the effective electron temperature. At this point the spectrum has reached the peak value and becomes inverted and is now within the optically thick regime. This point is also called the ‘turnover’. This means that the source in question has become opaque with any radiation detected coming from the surface only. In this regime the spectral index is independent from the energy distribution index and changes as $F_\nu \propto \nu^{2.5}$. Figure 1.4 shows a schematic plot of the spectrum of a synchrotron emitting source.

1.3.2 Bremsstrahlung

Also known as free-free radiation, bremsstrahlung refers to radiation due to the deceleration of a free electron in the electric field of a heavier ion (the electron is unbound before and after the collision hence the name free-free). It is a thermal emission mechanism that is similar to that of synchrotron in that instead of a magnetic field, bremsstrahlung utilises a charged plasma. For example, such radiation is commonly seen from HII regions which are interstellar regions of ionised atomic hydrogen. The free-free spectrum is relatively flat at high frequencies ($\alpha \approx 0$), with a turnover at lower frequencies where the source becomes optically thick and is described by $\alpha = 2$ (i.e. a black body). The spectrum shape of bremsstrahlung compared to other emission mechanisms can be seen in Figure 1.5.

²The brightness temperature is the temperature at which a black body would have to be to produce the observed intensity.




Fig. 1.4 Illustrations of the spectrum of a synchrotron emitting source. *Left*: The spectral distribution of the optically thin stage of the spectrum demonstrating how it is a superposition of the spectral distributions of the individual electrons. *Right*: The full radio spectrum of a partially self-absorbed synchrotron source. The self-absorption occurs at low-frequencies when the emission is optically thick, in this regime the flux density, F , changes with frequency as $F \propto \nu^{2.5}$. A turnover then takes place at higher frequencies and F then depends on the energy distribution of the electrons $F_\nu \propto \nu^{(1-p)/2}$.

1.3.3 Black Body

Black body radiation is observed from optically thick thermal sources. As a type of thermal emission, it is possible to describe both the intensity of the observed radiation, along with the shape of the spectral distribution of the emission, by only knowing a single parameter - the objects temperature, T . The Rayleigh-Jeans law approximation states that the flux density of a black body source increases with ν^2 , and has a spectrum of $\alpha = 2$. At higher frequencies quantum effects become more important and cause the emitted flux density to peak, which is followed by a steep decline as $F_\nu \approx e^{-h\nu/kT}$. The spectrum shape of black body radiation compared to other emission mechanisms can be seen in Figure 1.5.

The most famous example of black body radiation in astronomy is that of the cosmic microwave background, the closest natural spectrum to that of an ideal black body.




Fig. 1.5 Schematically shows the spectral shapes of the three emission mechanisms, synchrotron, bremsstrahlung and black body.

1.3.4 Coherent

The previous examples of emission mechanisms have all been incoherent emission, where the emission is generated by the electrons (or other particles) acting independently from one another. This means that the resultant photons have no phase relationship. Coherent emission on the other hand refers to the process when electrons work collectively to emit radiation. In this case all the particles undergo acceleration in phase, therefore producing photons which are also in phase and at the same frequency and direction. It is a non-thermal emission process and uniquely, the mechanism allows for observed T_b values to far exceed the actual temperature or kinetic temperature of the electrons.

An example of coherent emission is that of an astrophysical maser (Microwave Amplification by Stimulated Emission of Radiation). In thermal equilibrium, a gas, for example, will usually have a greater number of atoms in a low energy state than those in a high energy state. This results in the net absorption of any radiation passing through the gas. However, an external energy source (known as the pump) can act upon the gas to cause this ratio to invert, otherwise known as a population inversion. This causes stimulated emission to dominate over absorption, and hence any incident radiation is amplified. In this situation, an incident photon

can cause the stimulated emission of an additional photon at the same phase and frequency. Both of these photons continue to propagate initiating a chain reaction of stimulated emission, creating a coherent signal. Pulsars are also a famous example of coherent emission, which are discussed further in Section 1.5.2. Star forming regions consisting of young, dusty stellar nurseries, commonly produce coherent emission.

1.4 Propagation Effects

Between the emission from the source, to reaching the telescopes on Earth, the radio signal can be distorted as it passes through the inhomogeneous interstellar medium (ISM) which contains gas and dust. This section briefly describes the effects which are relevant to the work presented in this thesis, along with the effect of the Earth's ionosphere on incoming radio waves. A more comprehensive overview of the following effects can be found in works such as Spitzer (1998) and Lorimer & Kramer (2012).

1.4.1 Dispersion

Charged particles in the ionised regions of the ISM give the medium a refractive index which is a function of frequency. In turn, the light travel speed through the ISM depends upon the refractive index. This causes the light to be 'dispersed', where the lower frequency component of the signal travels at a lower speed through the ISM than that of the higher frequency. Therefore, on Earth the higher frequency emission arrives at an earlier time than that of the lower frequency. This also has the effect of reducing the overall intensity of the received signal, especially in the case of short-time duration events where the time delay, Δt_{DM} , caused by dispersion can be greater than the intrinsic duration of the event itself. Δt_{DM} is characterised by the dispersion measure (DM), which is a measure of the number of free electrons along the line of sight. Examples of how dispersion affects a signal can be seen in

Figure 1.6 which shows both the shift in arrival time of different frequencies of a radio burst discovered by Lorimer et al. (2007), along with a simulated example of how the amplitude of a signal is reduced.

Any signal which is dispersed can be recovered by using the de-dispersion technique, which involves shifting the individual frequency channels relative to each other by Δt_{DM} before being combined. This process is repeated for a large sample of DM values in order to search for signals across the whole range of possible dispersion.

1.4.2 Scattering

As a radio signal travels through the ISM it can encounter fluctuations in the electron density. Such regions are usually modelled as thin ‘scattering screens’ at some distance between the source and the observer. The fluctuation causes different components of the incoming signal to be scattered in different directions. The scattered signal is then received by the observer on Earth with a delay compared to the portion of the signal that was not scattered. This causes the signal as a whole to broaden in time with an exponential decay, along with becoming fainter. The observed width of a signal with an intrinsic width of τ_0 due to a scattering time τ_s is expressed as $\sqrt{\tau_0^2 + \tau_s^2}$. The scattering time is also dependent on the observing frequency as $\tau_s \propto \nu^{-4}$, meaning that the broadening of a signal is much more significant at lower frequencies, and in turn, the amplitude is also more significantly reduced. Figure 1.7 demonstrates the effect of scattering on a signal.

1.4.3 Ionosphere

Before cosmic radio emission reaches radio telescopes on the ground, the signal can be influenced by a layer in the Earth’s atmosphere called the ionosphere. It is a layer of partially ionised gas at a height of $\sim 50\text{--}1000$ km above the surface, primarily ionised by solar radiation, though cosmic radiation can also have an ionising effect. It is a very dynamic




Fig. 1.6 Examples of dispersion. *Left*: This shows how dispersion by the ISM causes a shift in arrival time of the different frequency components of a short duration radio burst. The inset plot shows the profile of the recovered pulse. Figure from Lorimer et al. (2007). *Right*: Demonstrates the effect of dispersion on the amplitude of a signal. In grey is how the signal would appear with no dispersion, followed by how the signal then appears when affected by a DM of 50 and 150 kpc cm^{-3} in blue and red, respectively. The amplitude is significantly reduced with dispersion greatly hindering the chances of detection without de-dispersion. Simulation produced by M. Piętko.




Fig. 1.7 An initial boxcar pulse with a flux density of S_{v0} (black line) is broadened by scattering from the ISM. The peak flux density is reduced to S_v and undergoes a change of profile (grey line), increasing the pulse width to $\sim \sqrt{\tau_0^2 + \tau_s^2}$. Figure from Hassall et al. (2013).

medium, meaning that the electron density can fluctuate in both time (order of 10 s) and space (order of 10 km) (van der Tol et al., 2007).

Low-frequency radio signals (< 300 MHz) can be heavily influenced by the ionosphere, primarily due to the effect of refraction. The ionised gas, if it is assumed as being cold and without a magnetic field, has a refractive index, n , which is dependent upon the free electron density and is expressed as

$$n = 1 - \frac{v_p^2}{v^2}, \quad (1.1)$$

where v is the radio frequency and v_p is the plasma frequency³. Any radio frequency below v_p is reflected by the ionosphere. This effect allowed for long distance radio communications in the early 20th Century by using the ionosphere to reflect the radio transmission. Typically the value of v_p is in the range of 1–10 MHz but can rise to values of up to 200 MHz. In terms of radio interferometry imaging, ionospheric refraction changes the path difference of the signal received at different array elements, which can cause phase errors. Hence, further calibration techniques are required to correct for such effects. It should be noted that effects can be mild when an array is observing through a relatively small patch of ionosphere as then all signals will be affected by the same path length changes.

Ionospheric effects on low-frequency observations were usually ignored due to the poor resolution and sensitivity of instruments and the lack of computing power available to accurately correct for them. However, with new powerful low-frequency telescopes such as MWA, LOFAR and the future SKA, being able to solve for ionospheric effects is crucial in reaching theoretical sensitivity limits. This is especially true for LOFAR where - unlike compact arrays - different elements of long baselines will be viewing the sky through different patches of the ionosphere (see Figure 1.8). Hence, special care is required when calibrating such long baselines at low-frequencies and is both a technical and computational challenge.

³In an electrically neutral plasma containing a gas of positively charged ions and negatively charged electrons, if the electrons are displaced with respect to the ions then the Coulomb force acts to pull the electrons back. Hence the Coulomb force acts as a restoring force, and the frequency of this oscillation is known as the plasma frequency.




Fig. 1.8 A simplistic example of how a long baseline instrument can be affected by the ionosphere. When treated individually the two arrays can maintain correlation due to looking through the same local ionospheric conditions, respectively. However, when combining the two arrays, the array on the right is observing through a different and more active part of the ionosphere. Hence, the correlation between the two arrays is destroyed because of the different ionospheric conditions at their respective locations.

1.5 Radio Transients

The work detailed in this thesis is focussed upon the search, detection and classification of radio transient events, at both megahertz and gigahertz frequencies. Various astrophysical events, from many types of sources contribute to the dynamic radio sky, and that is demonstrated by Figure 1.9 (from Pietka et al., 2015). It shows the phase space of known transient (and variable) sources in terms of luminosity and time-scale of the event, and covers a wide range of phenomena, from flares associated with nearby flare stars, to emission from radio luminous galaxies. The transient types are generally split into those which emit through an incoherent emission mechanism, usually synchrotron, and those which are coherent (the regions marked blue and white on the figure, respectively). Coherent events are typically of very short duration (< 1 s) and are observed through time-series techniques, whereas incoherent events evolve over long time-scales and are mainly observed in the image plane. These




Fig. 1.9 The phase space of known radio transient and variable populations, where radio luminosity is plotted against the time-scale of the included events. The events plotted originate from two types of emission mechanisms: incoherent and coherent, or slow and fast transients, respectively. They are separated by the constant brightness temperature value of 10^{12} K, which is the theoretical upper limit for incoherent emission, with the blue region highlighting the incoherent area of the plot. Figure from Pietka et al. (2015).

are known as fast and slow transient populations respectively, with the border defined by the maximum T_b of 10^{12} K that incoherent events can achieve (Kellermann & Pauliny-Toth, 1969). However, it should be noted that coherent events can also be detected in the image plane in certain circumstances.

What follows is a brief overview of the different types of transient sources both slow and fast. Quoted time-scales are from data compiled by Pietka et al. (2015) which also contains typical luminosity information for the different types of transient objects.

1.5.1 Slow Transients

Active Galactic Nuclei (AGN)

These are extremely energetic galaxies which are powered by a supermassive black hole (typically $10^5 - \sim 10^9 M_{\odot}$; Hughes, 2005) at the centre. There are many subclasses of AGN, such as Quasars and Seyfert Galaxies, which are defined by the viewing angle or certain features at specific frequencies. In the radio, the majority of the emission is observed from optically thin synchrotron emission from large jets emanating from the central black hole, which also drive into the surrounding medium producing bright radio lobes (Blandford & Rees, 1974). They are observable at all radio frequencies, and at low frequencies can become incredibly bright objects in the sky. AGN can vary in brightness over long time-scales (Condon et al., 1979; Taylor & Gregory, 1983) but due to the powerful jets and accretion processes they can also undergo outbursts of emission on shorter time-scales; likely due to the propagation of shocks within the jets (Aller et al., 1985; Lainela, 1994). Note that in Figure 1.9 some of these sources are seemingly above the 10^{12} K limit. This is due to jets being relativistic and hence the radiation can be beamed towards Earth giving an apparent brightness temperature above 10^{12} K. An example of an AGN observation is shown in Figure 1.10.

X-ray Binaries (XRB)

Similar to AGN but on a much smaller scale, XRBs are systems that consist of a compact object such as a neutron star or black hole, orbited by a companion star. The black holes in X-ray binaries are known as ‘stellar black holes’ and are roughly 5 – several tens of solar masses. X-ray emission is produced in the accretion disk formed by the matter falling from the companion star, whereas radio emission is observed from the jets (Fender, 2006) - hence the likeness to AGN (XRBs are also known as ‘microquasars’). The time-scales of radio




Fig. 1.10 A VLA 5 GHz image of the radio galaxy Cygnus A spanning ~ 2 arcmin. The jets and resulting radio lobes are clearly seen emanating from the AGN core of the galaxy, with red showing the areas of the brightest radio emission and blue the fainter emission. *Image courtesy of NRAO/AUI.*

outbursts from XRBs are typically on the order of days or months which is very fast in comparison to AGN where similar behaviour can take many years. One of the most famous XRBs detected in the radio is Cygnus X-3 where large radio flares have been seen with a peak brightness reaching $10\text{--}100\times$ the quiescent level (Waltman et al., 1995; Fender et al., 1997). Figure 1.11a shows an artists impression of a black hole XRB.

Tidal Disruption Events (TDE)

A TDE refers to the event where a star strays too close to a black hole. Due to the gravitational forces the star is tidally pulled apart, with roughly half the matter accreted by the black hole and the other half escaping the system. During the process jets are formed which are detected in the radio (Zauderer et al., 2011; van Velzen et al., 2016). Only a few of these events have so far been observed, but it is predicted that the SKA could observe TDEs at a rate of one per week. An artists impression of a TDE is shown in Figure 1.11b.




Fig. 1.11 (a): an artists impression of a black hole X-ray binary. Matter from the companion star is shown flowing into an accretion disk around the black hole, along with the production of jets from the centre. *Image credit: ESO/L.Calçada.* (b): an artists impression of how a tidal disruption event may occur. The stars path is intersected by a black hole which pulls and rips the star apart. Some of the matter is then accreted onto the black hole (the whiter gas in the illustration) with the remaining gas ejected into space. *Image credit: NASA/CXC/M.Weiss.*

Gamma-ray Bursts & Radio Supernovae

Gamma-ray bursts (GRB) are very bright and intense short duration bursts that occur at cosmological distances. The accepted theory as to their origin is the collapsar model (MacFadyen et al., 2001) where an extremely massive star collapses at the end of its evolution, creating two very powerful relativistic jets. What can be seen in the radio is the afterglow of the event, in which the ejected matter is colliding with the surrounding environment, and forms a spherically expanding shell emitting via the synchrotron process. In cases where the initial beamed explosion of the GRB is not in the direction of Earth, it is still possible to detect the afterglow - these are known as ‘orphan GRB afterglows’. There appears to be two types of GRBs - short and long - for which the burst lasts for less or more than ~ 2 s respectively. The collapsar model describes that of the long GRB, the short GRBs are postulated to originate from the merger of two compact objects such as neutron stars (Eichler et al., 1989; Narayan et al., 1992). Typically GRB afterglows can decay over a period of weeks or months. An artists illustration of a GRB event is displayed in Figure 1.12.




Fig. 1.12 An artists illustration of GRB 050709 - thought to have been caused by a collision of two neutron stars. The inset image shows the X-ray afterglow of the event as detected by *Chandra* (Fox et al., 2005). *Image credit: NASA/CXC/Caltech/D.Fox/D.Berry.*

An event which also results in isotropic emission from an expanding cloud interacting with the circumstellar medium is that of supernovae (SNe; Weiler et al. 2002). These arise from the death of a massive star and are categorised broadly into type I and type II, defined by the detection or non-detection of (optical) hydrogen emission lines. Within these classifications types Ib/c and II are associated with the core-collapse mechanism and are detectable in the radio (thus far type Ia supernovae have not been detected in the radio, e.g. Horesh et al. 2012). The duration of a supernova is typically on the time-scales of months to years. Previous works have also explored a possible link between GRBs and SNe with events such as SN 1998bw and GRB 980425 believed to have originated from the same progenitor (Kulkarni et al., 1998). The before and after optical images of a SN can be seen in Figure 1.13.

Cataclysmic Variables (CV)


CVs are close binary systems normally consisting of a white dwarf (WD) which accretes matter from a companion star via Roche-lobe overflow. There are various types of transient




Fig. 1.13 Optical images of before and after the explosion of SN 2013cu located in the galaxy UGC 9379 which is at a distance of 108 Mpc ($z = 0.03$). It was a Type IIb SN and the detection image is approximately 6 hr after the explosion (Gal-Yam et al., 2014). *Image credit: Avishay Gal-Yam, Weizmann Institute of Science.*

phenomena that arise from these systems, two of the best known examples are that of the classical novae and dwarf novae. Classical novae result from a thermonuclear runaway explosion on the surface of the WD which lasts on the order of days before falling back to a quiescent level over weeks or years, and are thought to repeat on the order of thousands of years. Dwarf novae on the other hand are not quite as dramatic or bright and are believed to be caused by changes in the mass transfer rate of the accretion disk (O’Brien, 2009). These outbursts are repeated quite often, on days to years time-scales. Other categories include recurrent novae, essentially novae that repeat more often, and nova-like CVs which are similar to novae but in constant outburst. Finally there are certain systems where a magnetic field can be detected. When the field is strong enough an accretion disk cannot form and instead the gas from the star is channelled along the magnetic field lines onto the poles of the WD.

In the radio only a handful of CVs have been detected. For classical novae, such as V1494 Aql which was detected in the radio during outburst (Eyres et al., 2005), the radio



(a) Cataclysmic Variable



(b) Solar flare

Fig. 1.14 (a): An artists illustration depicting a nova outburst where the white dwarf (left) has been accreting matter from the companion red giant star (right) and undergoes a thermonuclear explosion. *Image credit: David A. Hardy.* (b): the UV detection of a solar flare on 2015 January 12 as captured by the NASA *Solar Dynamics Observatory*, seen as the bright flash on the right side of the Sun. The process that causes such flares can be responsible for the flares detected from distant stars. *Image credit: NASA/SDO.*

emission is believed to arise from free-free radiation from a photoionised ejected shell. Dwarf novae have also been seen in the radio - observations of radio flares from SS Cyg imply that jets are formed during the outburst (Körding et al., 2008). Recently, Coppejans et al. (2015) demonstrated that it is possible to detect nova-like CVs when utilising the better sensitivity of modern radio telescopes. The authors suggest that the observed radio emission arises from multiple mechanisms such as synchrotron, gyro-synchrotron and cyclotron maser emission. An artists impression of a nova outburst can be seen in Figure 1.14a.

Flare Stars & Active Binaries

Flare stars are commonly late spectral type K or M stars, known as red dwarfs or M-dwarfs; cool and small main sequence stars that constitute a large portion of the Milky Way's stellar population. Much like flares from the Sun, radio emission from flare stars is generated by gyro-synchrotron radiation as electrons are accelerated in the stars magnetic field (Bastian, 1990; Osten et al., 2005). Such flares can range from seconds to minutes, though longer

bursts have been detected (Nelson et al., 1979). However, some flares are detected with very high brightness temperatures along with other characteristics that suggest there are also forms of coherent emission from flare stars (Slee et al., 2003), likely from the stellar atmosphere. Protostars have also shown variability in the radio, for example Bower et al. (2003) and Rivilla et al. (2015), with the latter study observing flares at a rate of 0.14 flares per day in the Orion Nebula Cluster. An example of a solar flare, similar to what might be seen from distant flare stars, is shown in Figure 1.14b.

Various types of stellar binary systems also display transient behaviour. Those include Algol binaries, where the donor star is more evolved yet less massive than the accreting star, and RS CVn - stars with active chromospheres orbiting each other. RS CVn produce more frequent flares, about ~ 1 per day, and have fast rise and decay times. Both these systems emit in radio via gyro-synchrotron radiation (Mutel & Weisberg, 1978), however, coherent emission can also be detected from these sources. A Hertzsprung-Russell diagram for stars detected as radio sources was produced by Stephen White (White, 2000)⁴, which presents the radio luminosity and emission type (thermal or non-thermal) for a wide range of stellar objects and systems.

1.5.2 Fast Transients

Planets & Exoplanets

At decimetre wavelengths (10–40 MHz) Jupiter produces bright, fast, non-thermal radio bursts which can be as intense as Solar bursts (Zarka et al., 2001). They are caused by Jupiter’s strong magnetic field accelerating electrons, which emit via the cyclotron maser process. The electrons are supplied by the stellar wind as well as the volcanic activity on Jupiter’s satellite, Io. All the gas giant planets in the Solar System emit at radio wavelengths from such aurorae interactions as indeed does the Earth. Hence, this form of radiation could

⁴Figure available at http://www.astro.umd.edu/~white/images/radio_hr_diag_full.html




Fig. 1.15 Jupiter's radiation belts, which are driven by the magnetic field, observed at 13 GHz by the *Cassini* orbiter. An optical image of Jupiter is used to show the scale of the belts. Similar emission could be used to obtain a direct detection of exoplanets orbiting distant stars. *Image credit: NASA/JPL.*

be used to directly detect the presence of an exoplanet orbiting a distant star (e.g. Murphy et al. 2015). Figure 1.15 shows an observation of the 13 GHz radiation belts of Jupiter.

Neutron Stars

There are various flavours of neutron stars that show radio transient behaviour. The most famous of these are pulsars: rapidly rotating neutron stars which produce beamed radio emission at the magnetic poles, traversing our line of sight - like a cosmic lighthouse. The rotation period of pulsars is extremely short - often less than 1 s. In addition to this pulsed emission, pulsars are also known to produce giant pulses reaching up to 1000 times the mean pulse intensity. The famous Crab pulsar was first to demonstrate such phenomena, where in one hour of observation the largest peak flux of the source was measured to be $\sim 10^5$ Jy at 430 MHz with a duration of ~ 100 μ s (Hankins & Rickett, 1975). A diagram of a pulsar can be found in Figure 1.16a. The spectral indices of pulsars are steep, ranging from -4 to 0 (Maron et al., 2000).




Fig. 1.16 (a): artists impression of a radio pulsar. It shows the magnetic field lines along with the radiation beams from the surface of the neutron star that are not necessarily aligned with the rotation axis. *Image credit: ESA/ATG medialab.* (b): the dynamic spectrum of FRB 110220 as discovered by Thornton et al. (2013) (Fig. 2) showing the frequency dependent delay. The inset shows how the pulse appears between 1.2–1.5 GHz, with a best fit profile denoted by the black line.

A subclass of pulsars are the Rapidly Rotating Transients (RRATs; McLaughlin et al. 2006). These are pulsars showing single bright flares separated by hours or days. Other sources, defined as intermittent pulsars (Kramer et al., 2006), have periods in the off-state lasting for more than a year. Finally, the most magnetic objects in the Universe, magnetars, are neutron stars with a magnetic strength exceeding 10^{12} G. The McGill Online Magnetar Catalogue (Olausen & Kaspi, 2014)⁵ currently shows that 29 magnetars are known, however, radio pulsed emission has only been detected from four of the 29 (Camilo et al., 2006, 2007; Levin et al., 2010; Shannon & Johnston, 2013; Eatough et al., 2013).

Fast Radio Bursts (FRB)

A newly discovered class of radio transient, FRB events, are single, bright ($\sim$ Jy) bursts, which typically last for a few milliseconds. The first such event was seen by Lorimer et al. (2007) with the Parkes Observatory in Australia. They concluded that because of the high

⁵<http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>

dispersion measure (DM) of $375 \text{ cm}^{-3} \text{ pc}$, the signal was of an extragalactic origin. Since the first discovery by Lorimer et al. (2007), as of writing a further 27 bursts have been discovered: 15 also using the Parkes Observatory (Keane et al., 2011; Thornton et al., 2013; Burke-Spolaor & Bannister, 2014; Ravi et al., 2015; Petroff et al., 2015; Champion et al., 2016; Keane et al., 2016), 11 using the Arecibo Observatory (Spitler et al., 2014, 2016) and one with the Green Bank radio telescope (Masui et al., 2015). These all shared similar characteristics to the Lorimer et al. (2007) burst, including high DM measurements ranging from $470\text{--}1629 \text{ cm}^{-3} \text{ pc}$. Refer to Figure 1.16b for an example of one of the bursts found by Thornton et al. (2013). The implied rate of FRBs based on these events has been predicted to be up to possibly thousands of FRBs occurring every day over the entire sky (Hassall et al., 2013; Lorimer et al., 2013). Hassall et al. (2013) showed that with the correct scattering (high) and dispersion (low) affecting an FRB signal, it would also be possible to detect such events in the image plane.

Two recent results are of high significance. First, Keane et al. (2016) followed up the discovered FRB with the ATCA and discovered a decaying radio transient, which took six days to fade, and were able to identify the host galaxy. This allowed the authors to place a distance to an FRB for the first time which in this case was six billion light years. However, this claim has since been disputed (Williams & Berger, 2016; Williams et al., 2016), with suggestions that the observed radio transient was intrinsic variability of the galaxy. Second, Spitler et al. (2016) discovered 10 bursts with Arecibo which appeared to emanate from the same direction in the sky, and have consistent DM values, as the original Arecibo FRB. This was the first time that an FRB was seen to repeat.

The progenitors of all these events are unknown, leading to a wide range of proposed theories regarding how FRBs are produced, such as: binary neutron star mergers (Totani, 2013); binary white dwarf mergers (Kashiyama et al., 2013); the delayed collapse of ‘supramassive’

neutron stars⁶ (Falcke & Rezzolla, 2014); association with GRBs (Zhang, 2014); and nearby flare stars (Loeb et al., 2014) where the apparent cosmological distance is explained by a blanket of coronal plasma around the host star.

1.6 Radio Transient Surveys

Historically, and still to this day, radio observations have been used to follow-up transient detections made at other wavelengths. Radio facilities generally had a narrow field-of-view (FoV), which made them inadequate to perform rapid transient and variability studies over a large fraction of the sky. However, blind transient surveys have been performed and have produced intriguing results. For example, Bower et al. (2007) (also see Frail et al., 2012) discovered a single epoch millijansky transient at 4.9 GHz while searching 944 epochs of archival VLA data spanning 22 years, with three other possible marginal events. Sky surveys using the Nasu Observatory have also been successful in finding a radio transient source, with Niinuma et al. (2007) having observed a two epoch event, peaking at 3 Jy at 1.42 GHz. Various counterparts were considered at other wavelengths, but the origin of the transient remains unknown. Lastly, Bannister et al. (2011a) surveyed 2775 deg² of sky at 843 MHz using the Molonglo Observatory Synthesis Telescope (MOST), yielding 15 transients at a 5σ level of 14 mJy beam⁻¹, 12 of which had not been previously identified as transient or variable.

Surveys at low frequencies (≤ 330 MHz) have also been completed. Lazio et al. (2010) carried out an all-sky transient survey using the Long Wavelength Demonstrator Array (LWDA) at 73.8 MHz, which detected no transient events to a flux density limit of 500 Jy. In addition, Hyman et al. (2002, 2005, 2006, 2009) discovered three radio transients during monitoring of the Galactic centre at 235 and 330 MHz. These were identified by using

⁶A ‘supramassive’ neutron star refers to a neutron star that exceeds the $1.4 M_{\odot}$ mass limit by spinning very rapidly (Margalit et al., 2015).

archival VLA observations along with regular monitoring using the VLA and the Giant Metrewave Radio Telescope (GMRT). The transients had flux densities in the range of 0.1–1 Jy and occurred on time-scales ranging from minutes to months. Lastly, Jaeger et al. (2012) searched six archival epochs from the VLA at 325 MHz centred on the Spitzer-Space-Telescope Wide-Area Infrared Extragalactic Survey (SWIRE) Deep Field. In an area of 6.5 deg^2 to a 10σ flux limit of $2.1 \text{ mJy beam}^{-1}$, one day-scale transient event was reported with a peak flux density of $1.7 \text{ mJy beam}^{-1}$. Surveys performed with SKA pathfinders and pre-cursors are discussed in Chapter 3.

Some of the historical searches for radio transients discussed above also show how challenging the practise can be. For example, Bower et al. (2007) in fact originally discovered ten single epoch millijansky transients in the VLA archival search. However, these data were re-analysed by Frail et al. (2012) who found that six of these transients were due to data artefacts. Out of the four remaining candidates, only one was considered a valid transient with a good signal-to-noise detection. In addition, there have been 10 other transient detections with the Nasu Observatory (Kuniyoshi et al., 2007; Matsumura et al., 2007; Kida et al., 2008; Matsumura et al., 2009; Niinuma et al., 2009). These all had flux densities exceeding 1 Jy at 1.42 GHz, with no obvious optical counterparts, and had durations ranging from minutes to days. Much like the transients discovered by Bower et al. (2007), the reliability of these transients was investigated by deriving false detection rates by Aoki et al. (2014), who concluded that only one source was statistically significant. As the technology used to observe the radio sky rapidly advances and becomes more complex, one must be very careful in their search for radio transients and validating the reality of any such detections.

It is common to calculate the upper limits of the surface density of transients for surveys such as those discussed above. To do this Poisson statistics are used, specifically:

$$P(n) = e^{-\rho A} \quad (1.2)$$

where ρ is the surface density of sources per square degree and A represents the equivalent solid angle obtained by multiplying the area of the sky surveyed, Ω , by the number of epochs $N - 1$. This approach assumes an isotropic probability of a transient detection, i.e. events are more likely to be extragalactic in origin. $P(n)$ is defined at the Poisson 2σ confidence level, i.e. $P(n) = 0.05$, and by rearranging Equation 1.2 the respective value of ρ for each time-scale can be obtained. When no candidates are found, $P(0) = 0.05$, however, in situations where a transient candidate is found, the upper limit is calculated by changing the n value accordingly. In the case of one transient found, for example, $P(1) = 0.05$.

The transient surface density values for radio transient surveys found in the literature are presented in Table 1.1, which is split into three sections. The first part of the table presents the surveys conducted at low-frequencies (< 325 MHz), the second section then displays the surveys above this frequency value which are mostly at gigahertz frequencies, and finally the third part which presents estimates for the transient survey potential of the SKA, for both shallow and deep strategies. Figures 1.17 and 1.18 then present the ‘Log N, Log S’ parameter space for the low-frequency and high-frequency transient searches presented in Table 1.1 respectively. This work was presented in Fender et al. (2015).

Table 1.1 Summary of radio transient surface densities, ρ , as found in the literature as of 2015 (presented in Fender et al. 2015). Columns: δt is the timescale of each individual epoch, Δt is the time scale(s) between epochs and N_e is the total number of epochs. The first block of the table is considered the low-frequency portion of surface densities, which incorporates surveys conducted at frequencies < 325 MHz. The second block are the densities defined at mainly GHz frequencies. The final and third section of the table provides estimates of transient surface densities that will be achievable with the SKA. An entry containing ‘-’ signifies that the respective value could not be identified from the literature.

Survey/Reference	Flux Limit (mJy)	ρ (deg ⁻²)	δt	Δt	ν (GHz)	N_e
Lazio et al. (2010)	$> 2.5 \times 10^6$	9.5×10^{-8}	5 min	2 min - 5 min	0.074	~ 1272
Obenberger et al. (2014a)*	$> 1.44 \times 10^6$	2.2×10^{-9}	5 sec	5 sec	0.074	$\sim 43\,056$
Cendes et al. (2015)	> 500	10^{-2}	11 min	min - months	0.149	26
Carbone et al. (2016)	> 300	1.3×10^{-3}	11 min	min - months	0.150	151
Bell et al. (2014)	> 5500	7.5×10^{-5}	5 min	min - 1 year	0.154	51
Hyman et al. (2009)* ^a	> 30	0.034	~ 3 hr	days - months	0.235	-
Jaeger et al. (2012)*	> 2.1	0.12	12 hr	1 day - 1 month	0.325	6
MOST/Bannister et al. (2011a,b)*	> 14	1.5×10^{-3}	12 hr	1 day - 20 yr	0.843	2
Ofek et al. (2010)*	> 6	1.5	days	-	1.4	2
ATATS I/Croft et al. (2010)	> 40	0.004	months	1 month	1.4	2
ATATS II/Croft et al. (2011)	> 350	6.0×10^{-4}	1 day	min - days	1.4	12
Bower & Saul (2011)(A)	> 3000	9.0×10^{-4}	2 min	> 1 min	1.4	1852
Bower & Saul (2011)(B)	> 70	3.0×10^{-3}	2 min	1 day - yrs	1.4	1852
Bell et al. (2011)	> 8	0.032	5 min	4 - 45 days	1.4, 4.9, 8.4	5037
Aoki et al. (2014)*	> 3000	2.0×10^{-6}	4 min	day	1.4	1200
ASGARD/Williams et al. (2013)	> 16	0.14	1 hr	days - years	3	29
PiGSS I/Bower et al. (2010)(A)	> 1	1	months	months	3.1	2
PiGSS I/Bower et al. (2010)(B)	> 5	0.3	months	mins	3.1	2
PiGSS II/Bower et al. (2011)(A)	> 15	0.025	hours	days	3.1	78
PiGSS II/Bower et al. (2011)(B)	> 5	0.18	month	month	3.1	5
PiGSS III/Croft et al. (2013)(A)	> 10	0.009	hours	days	3.1	379
PiGSS III/Croft et al. (2013)(B)	> 10	0.02	months	months	3.1	-
PiGSS III/Croft et al. (2013)(C)	> 10	0.08	months	years	3.1	2
Bower et al. (2007)	> 0.9	6	20 min	years	4.8, 8.4	17
Frail et al. (2012)(A)	> 0.2	3	20 min	month	4.8, 8.4	-
Frail et al. (2012)(B)	> 0.37	0.6	20 min	mins - days	4.8, 8.4	-
Ofek et al. (2011)	> 1.8	0.039	50 seconds	day - month	4.9	16
Alexander et al. (2015)	> 0.5	17	15 mins	days - months	4.9	31
SKA-Shallow	1.0×10^{-3}	3.0×10^{-8}	4 hr	-	1.4	~ 5000
SKA-Deep	2.0×10^{-5}	6.0×10^{-3}	5 hr	-	1.4	~ 500

* Limit derived from a transient detection.

^a Using approximate rate derived in Williams et al. (2013) which is inclusive of Hyman et al. (2005, 2006, 2009).




Fig. 1.17 A compilation of rates and upper limits for blind searches for radio transients at MHz frequencies. The surveys included in this plot are: Bell et al. (2014) (Bel2014); Carbone et al. (2016) (Car2016); Cendes et al. (2015) (Cen2015); Hyman et al. (2005,2006,2009) (Hym2009); Jaeger et al. (2012) (Jae2012); Lazio et al. (2010) (Laz2010) and Obenberger et al. (2014a) (Obe2014). Refer to Table 1.1 for the details of all the surveys.




Fig. 1.18 A compilation of rates and upper limits for blind searches for radio transients at GHz frequencies. Multiple rates derived from the same survey are denoted by (A), (B) and (C). The surveys included in this plot are: Alexander et al. (2015) (Ale2014); Aoki et al. (2014) (Aok2014); Bannister et al. (2011a,b) (Ban2011); Bell et al. (2011) (Bel2011); Bower et al. (2007) (Bow2007); Bower et al. (2010) (PiGSS I); Bower et al. (2011) (PiGSS II); Bower & Saul (2011) (B&S2011); Croft et al. (2010) (ATATS I); Croft et al. (2011) (ATATS II); Croft et al. (2013) (PiGSS III); Frail et al. (2012) (Fra2012); Ofek et al. (2010) (Ofe2010); Ofek et al. (2011) (Ofe2011); and Williams et al. (2013) (ASGARD). Also plotted are the estimated transient surface density limits of future deep and shallow (labelled ‘Shlw’) SKA surveys. Refer to Table 1.1 for the details of all the surveys.

1.7 This Work

Presented in this thesis is work that explores the search for, and detection of, radio transients at both the not well explored low-frequency domain and the higher, more probed, gigahertz regime. In addition to this the area of transient classification is explored as we enter an era of sensitive, wide-field surveys at multiple wavelengths.

Chapter 2 provides a description of the International LOFAR Telescope along with details of various commissioning and transient projects that I was significantly involved with.

Chapter 3 presents the analysis of a wide-field, four month transient search using LOFAR, centred at the North Celestial Pole.

Chapter 4 describes a study of the flux properties of radio transient classes with their optical counterparts, exploring if such information could be used to classify newly discovered transient sources.

Chapter 5 presents a multi-epoch transient search at 1.4 GHz using the VLA as part of the CHILES project, where optical observations also took place concurrently.

Chapter 6 summarises the work in this thesis in addition to discussing the possible future direction of any further work.

Chapter 2

The Low-Frequency Array

Officially opened in 2010 by Queen Beatrix of the Netherlands, the LOw-Frequency ARray (LOFAR) is currently the largest radio telescope in the world that operates within the low frequency regime of 10–240 MHz. It is a member of a new generation of radio telescope hardware that utilises arrays composed of thousands of small dipole antennas, as opposed to the more traditional dish based telescope designs. This means that LOFAR is electronically steered as it contains no moving parts, and also allows for the telescope to observe in multiple directions of the sky simultaneously. Combining this feature with the possible high resolution and deep sensitivity that LOFAR can provide, unprecedented views of the low frequency radio sky can be achieved. The use of this new technology also makes LOFAR an important pathfinder to the future Square Kilometre Array (SKA) telescope. This chapter looks in further detail at the capabilities and design of LOFAR, along with details of projects that I have been involved in.

2.1 LOFAR Design & Capabilities

The thousands of antennas which LOFAR is contrived from are arranged into stations that are located around Northern Europe. The core of the telescope is located in the Netherlands,




Fig. 2.1 A schematic map showing the locations of LOFAR stations across Europe. This includes the yet-to-be built, as of early 2016, stations in Poland and the Republic of Ireland. *Image credit: ASTRON.*

specifically near the village of Exloo, and it contains 24 stations within a 2-km-wide area. These are known as the ‘core’ stations. There are a further 18 ‘remote’ stations distributed around the Netherlands, in addition to nine ‘international’ stations located in Germany (6), France (1), United Kingdom (1) and Sweden (1). As of writing in 2016, four additional international stations are being constructed across Poland (3) and in the Republic of Ireland (1). Figure 2.1 shows a schematic map of LOFAR across Europe. This produces a largest distance between antennas of approximately 1500 km. As LOFAR is an interferometric instrument, these large baselines provide it with the potential for very high resolution imaging, recalling that the angular resolution, θ , is given by $\theta = \alpha \frac{\lambda}{D}$, where λ is the observing wavelength, D is the maximum baseline length and α corresponds to a constant (which for LOFAR $\alpha = 0.8$). Thus, for example, with a maximum baseline of 100 km at an observing frequency of 150 MHz ($\lambda = 2$ m), a resolution of 3.3 arcsec can be expected, or 8.3 arcsec at 60 MHz.

The following sections present an overall brief summary of the LOFAR design and capabilities; for full details the reader is referred to van Haarlem et al. (2013).

2.1.1 Low Band Antennas (LBAs)

Due to the broad range of frequency coverage, LOFAR utilises two different types of antenna design. The first, named the LBAs, are capable of operating in the 10–90 MHz range. However, during operations this is limited to a 30–80 MHz range due to strong RFI at the lower end of the spectrum and the proximity of the FM radio band at the upper end. This spectrum, along with a photograph of a single LBA can be found in Figure 2.2.

It can be seen there are four distinct components to an LBA; wires, the head, a vertical tube and a groundsheet. The wires, which are made of copper, detect the two orthogonal linear polarisations (two wires detect each polarisation) and these are connected to a low-noise amplifier (LNA) which is located in the head. The wires also act to hold the head atop a vertical piece of PVC piping, which houses the connecting coaxial cables. The groundsheet consists of a 3×3 metal mesh and plastic sheet, this is to block the antenna from radiation from the Earth and to stop the growth of vegetation respectively. This design met the goals of the antenna being a cost effective instrument that was easy to mass produce, yet sturdy enough to operate with at least a 15 year lifespan. The antenna design also has the advantage that the LBA have an omnidirectional response, which allows for the entire visible sky to be monitored at one time.

2.1.2 High Band Antennas (HBAs)

The HBAs are designed to operate in the frequency range of 120–240 MHz, and a photo of an antenna can be found in Figure 2.3, along with the frequency response of the HBAs in a station. The different design is caused by the fact that the total system noise is no longer




Fig. 2.2 *Left*: A photograph of a single LBA dipole, which shows the wires, central pole and the groundsheet. The inset images show the head of the antenna, as well as the LNA electronics. *Right*: The median averaged spectrum for all of the LBA dipoles for a specific station. Large RFI can be seen below 30 MHz, partly caused by ionosphere reflection, and above 80 MHz due to the FM radio band. It is also clear that the peak response is located at approximately 60 MHz. Figure credit: van Haarlem et al. (2013).

dominated by the sky noise, as is the case for the LBAs. Hence, they were designed to minimise the system noise caused by the electronics.

An HBA consists of 16 bow-tie shaped dual dipole antennas which is called an HBA ‘tile’. The 16 antennas are arranged into a 4×4 grid and are enclosed by two overlapping flexible polypropylene foil layers; this is simply to protect the antennas from the weather. Each tile has a limited FoV of approximately 30 deg, this is due to the fact that a beam is formed on a per tile basis by combining the signal from the sixteen antennas for the direction of sky of interest.

2.1.3 Stations

As previously mentioned, the antennas are grouped into ‘stations’, which essentially act as the ‘dishes’ of LOFAR. The design of each station can be one of three, depending on whether the station is classed as a core, remote or international station (see the beginning of Section 2.1). The different designs exist primarily because of cost reasons, and hence




Fig. 2.3 *Left*: A photograph of a single HBA tile, where some of the protective covering has been removed to display the dipoles underneath. *Right*: The median averaged spectrum for all HBA tiles in a specific station. Alongside minor spikes of RFI, a large spike is seen at approximately 170 MHz caused by this frequency corresponding to an emergency pager signal. Figure credit: van Haarlem et al. (2013).

incorporate designs to limit impact to sensitivity and performance depending on the location of the station. A schematic layout of each type of station is presented in Figure 2.4.

Core stations consist of 96 LBAs and 48 HBAs along with 48 digital receiver units (RCUs; see following Section 2.1.6). Each RCU has three inputs; two of these are connected to the LBAs with the remaining connected to the HBA tiles. This is important as only one of these inputs can be used at any one time, meaning that only 48 LBA signals can be used at a given time. The layout of HBA tiles in core stations is unique in that they are divided into two sub-stations of 24 HBAs, each with a diameter of 31 m. During observations, the two sub-stations can either be used together as a single station, or split, forming two independent LOFAR stations. When used as two separate stations, the primary advantage is that this generates many more short baselines within the core contributing to a more uniform uv coverage. LBAs are located in a single field of diameter 87 m.

At the heart of the LOFAR core exists a 320 m diameter island that contains six core stations named the ‘Superterp’, which is pictured in Figure 2.5. Not only does this arrange-




Fig. 2.4 Diagrams showing the layout of LOFAR core, remote and international stations. LBAs are those which are contained within large circles, whereas the arrays consisting of small squares denote HBAs. *Image credit: ASTRON.*

ment provide the shortest baselines available to LOFAR, but these six stations can also be combined to form one large station.

Remote stations are very similar in construct to core stations, containing the same amount of LBAs, HBAs and RCUs. The major difference is that the HBA tiles are contained within a single field which is 41 m in diameter.

International stations contain double the number of HBA tiles and RCUs than its core and remote counterparts, bringing the total numbers to 96 LBAs, 96 HBAs and 96 RCUs. This allows for the use of all 96 LBAs during an observation. The design of the international stations is in fact the closest to the original design first proposed for LOFAR in which all the stations would have been uniform.



Fig. 2.5 A birds-eye view of the Superterp which lies directly at the centre of the LOFAR core. *Image credit: ASTRON.*

2.1.4 *uv* coverage & Sensitivity

Examples of the *uv* coverage obtained by LOFAR can be seen in Figure 2.6, which have been produced by simulation and uses 40 core and remote stations and 8 international stations. These figures are originally found in van Haarlem et al. (2013). The simulations represent a 6 hour observation of a source at declination 48° between 30 and 78 MHz, using a single beam and 48 MHz bandwidth. The three different plots represent observations using only the core stations, only the core and remote stations and finally all LOFAR stations. In addition, the synthesised beam for each configuration is shown (using uniform weighting) which displays the side-lobe response pattern. The instantaneous response of LOFAR, using the same observational setup described above for an object at 0 hour angle (transit), is also shown in Figure 2.6.




Fig. 2.6 Sample uv coverage plots and synthesised beams for *left*: 6 hour observation of a source located at declination 48° , using a single beam and 48 MHz bandwidth, and *right*: the same observation but the instantaneous coverage at an hour angle of 0 (transit). The top row represents 24 core stations only, the middle is the 40 core and remote stations only and the bottom row is the full 48 station array. *Figure from van Haarlem et al. (2013).*

Table 2.1 LOFAR sensitivities. In these calculations the Superterp is 6 stations, NL Core is 24 stations, Full NL is 40 stations and the Full EU is 48 stations. Values marked with * are preliminary values that will be improved in the future.

Frequency (MHz)	λ (m)	Sensitivity			
		Superterp (mJy)	NL Core (mJy)	Full NL (mJy)	Full EU (mJy)
30	10.0	36	9.0	5.7	3.8
45	6.67	29	7.4	4.7	3.1
60	5.00	25	6.2	3.9	2.6
75	4.00	44	10.8	6.8	4.5
120	2.50	1.5	0.38	0.30	0.20
150	2.00	1.3	0.31	0.24	0.16
180	1.67	1.5	0.38	0.30	0.20
200	1.50	2.5*	0.62*	0.48*	0.32*
210	1.43	2.5*	0.62*	0.48*	0.32*
240	1.25	5.6*	1.4*	1.1*	0.73*

The expected noise values quoted by the observatory¹ are found in Table 2.1. These values are image noise calculated assuming an 8 hour observation and an effective bandwidth of 3.66 MHz. Values at 200 MHz and above are currently preliminary and should improve with future station calibration. While the HBA values presented in the table have been achieved, the LBA values are yet to be reached.

2.1.5 Pointing LOFAR

As LOFAR is a software telescope with no moving parts, it should be noted how the instrument is able to ‘point’ towards a specific target or region on the sky. We can consider the situation presented in Figure 2.7, where two antennas, separated by a distance D , wish to observe a source that is off-axis by an angle θ . The radiation from this source arrives at each antenna at different points in time and out of phase, with the time difference

$$\Delta T = \frac{D \sin \theta}{c}. \quad (2.1)$$

¹<https://www.astron.nl/radio-observatory/astronomers/lofar-imaging-capabilities-sensitivity/sensitivity-lofar-array/sensiti>




Fig. 2.7 Diagram showing how LOFAR is able to point towards specific sources or regions of the sky. The radiation from the source arrives out of the plane of the antennas by an angle θ and hence is detected at different times at each antenna. The time delay can be calculated as $\Delta T = D \sin \theta / c$ and corrected for by software, making LOFAR effectively ‘point’ in the desired direction.

Therefore, by compensating for this delay the antennas can be effectively ‘steered’ into the desired direction. This is done with each element of the array by LOFAR using software to introduce the required delay.

2.1.6 Signal Path

The following is a brief description of the path the radio signal takes from observation to processing. Once again, for full details the reader is referred to van Haarlem et al. (2013).

The radio signals received by the antennas are firstly amplified and then transported by the connected coaxial cables to a cabinet at the edge of each station. This cabinet, which is heavily shielded from outside radiation, houses all the electronics which are used to process the incoming raw signal. Firstly, the RCU provides further amplification of the signal which

is followed by digitisation. The digitised signal is then split into sub bands by Remote Station Processing (RSP) boards, with the final number of sub bands available depending on the sampling mode used. These numbers are 244 sub bands for 16 bit sampling, or 488 sub bands when using 8 bit sampling, resulting in a total bandwidth of 48 MHz and 96 MHz respectively. Next, the signals from all the antennas are combined to form a beam in the desired direction of the sky (see previous Section 2.1.5). Lastly, these data are then transported to ‘Central Processing’ (CEP; located in the Dutch city of Groningen) through dedicated, high speed fibre optic connections.

At CEP, the data from all stations are combined using the method required for the observation (i.e. interferometry or beam-formed), which requires a large amount of computing power. Previously performed by an IBM Blue Gene/L supercomputer (until 2014), the correlating task is now undertaken by an even more powerful GPU based system named COBALT. The data is then moved on for further post processing on a separate computing cluster at CEP.

2.1.7 The Long Term Archive (LTA)

A LOFAR observation can produce data rates up to 80 Gbit/s, however, once the data has passed through the various processing steps described above, their volume can be substantially reduced. Despite this, the data volumes (e.g. on the order of terabytes for a typical multi-hour observation) are still large enough to not allow the indefinite storage of processed data on the CEP computing cluster. Instead the data is transferred to a dedicated storage solution, the LTA, which is based in two Dutch locations, Amsterdam and Groningen, as well as in Jülich, Germany. This system allows for the storage, and retrieval, of observational data by astronomers and is expected to grow by up to 5 Pbyte per year. As of December 2015, the system was storing 20.5 PB of data.

2.2 The Transients Key Science Project (TKP)

The TKP (Fender et al., 2006), along with five other projects: the epoch of reionisation, deep extragalactic surveys, ultra high energy cosmic rays, solar science and space weather, and cosmic magnetism, form the six key science drivers of LOFAR. A wide range of transient and variable astrophysical objects are detectable with LOFAR, and it is highlighted by the five significant areas of research within the TKP: jet sources, flare stars, planets, pulsars and serendipitous discovery of transient sources. LOFAR's wide FoV, along with high sensitivity, make it an excellent instrument to conduct transient studies, especially in the case of blind transient surveys.

The diverse range of transient objects means that various observation methods are employed, and this is largely due to the time scales at which the transient events occur. For example, pulsar and planetary emissions can occur at a sub-second time scale and are usually observed in a mode commonly referred to as 'beam-formed observations' (refer to Stappers et al. 2011). This is where data are recorded as a single beam pointing in the direction of interest, with the data from different stations being summed either coherently or incoherently. No spatial information can be gained from this type of observation. On the other hand, objects that evolve on slower time scales (seconds–years), such as GRBs or variable AGN, are commonly observed in an interferometric mode, producing images of the sky. Though this is not to say that objects are exclusive to only one form of observation, in fact LOFAR makes it possible to record both beam-formed and interferometric data simultaneously.

The work presented in this thesis has used exclusively interferometric observations to carry out transient studies, be it blind surveys or targeted observations. The proceeding three sections present work that I significantly contributed to as part of the TKP.

2.2.1 The Radio Sky Monitor (RSM)

The RSM is one of the main surveys performed by the TKP, dedicated to characterising the transient low-frequency radio sky. It is an image-plane based survey which takes advantage of LOFAR's wide FoV and multi-beam capabilities to observe an entire 'zenith strip' over a 24 hour period. This is repeated on time-scales of months to years.

In further detail, the survey is performed as follows. Firstly, six beams are formed around the zenith ($+52^\circ$ declination for LOFAR) in a mosaic pattern, as seen in Figure 2.8. This covers approximately 60 deg^2 of sky, and this same region is observed twice for 11 min each. Before each of these observations, one of a pre-selected group of calibrator sources is observed for 2 min. As the Earth rotates, the neighbouring region of the zenith strip (from here on, each region of the zenith strip is referred to as a 'pointing') is observed in the same way, until after 24 hr, the whole zenith strip is monitored, covering an area of approximately 1500 deg^2 . Using the HBAs, 48 of these pointings are required to cover the entire strip and the observations cover a frequency range of 124–185 MHz, whereas for the LBAs, only 24 are needed and these cover a frequency range of 43–74 MHz.

As of 2016, 10 HBA and 4 LBA full zenith strip observations have been completed between 2013–2015 with time scales of 1 month to 2 years between runs. Typically the HBA observations have a sensitivity between 1–10 mJy, whereas the LBA observations are less sensitive at around 20–40 mJy. Taking even further advantage of LOFAR's capabilities, a little spare bandwidth available in each observation was used to observe extra targets of interest (e.g. LSI +61 303 and M31) which in turn spawned their own projects, and in some cases published science results.

RSM Processing Pipeline. My personal involvement on the project was to construct a pipeline which could efficiently and swiftly process an entire zenith observation. As discussed by Fender et al. (2007), originally it was envisioned that fully calibrated radio




Fig. 2.8 Diagram displaying how each observation of the RSM is performed. The main zenith monitoring area is observed by covering a large amount of sky using multiple beams. This is repeated numerous times over 24 hours to cover the entire zenith strip. Spare bandwidth can be used to form an additional beam (more than one if desired) to monitor extra targets.

data and images would be delivered by the radio observatory. However, due to the complex nature of commissioning LOFAR plus the challenging nature of calibrating LOFAR data at the time, this was not possible. A pipeline was also necessary because of the amount of data involved. Taking the case of an HBA RSM zenith scan as an example, each of the six beams of an observation consists of 40 sub bands. Two observations make up a pointing along with two calibrator scans, which consist of 40 sub bands each. The extra target observations also consist of 4 sub bands per observation. Hence, in total, there are 568 sub bands per pointing, meaning a total of 27,264 (568×48) sub bands, i.e. measurement sets, per zenith observation. Thus, while the software was available to process LOFAR observations, performing the task manually would have proved too cumbersome, and more importantly too slow, as the aim was to detect transient sources that, potentially, could be followed up by other telescopes.

A series of automated pipeline scripts were developed to automate the process of LOFAR data reduction, which could also handle the unique processing steps of the RSM observations. The full LOFAR data reduction process will not be explained in this section; the standard

procedure can be found in van Haarlem et al. (2013). Additionally the reduction detailed in Chapter 3 of this thesis is the same process which is used for the RSM. The observational RSM data are received having been compressed, flagged for RFI and demixed by the observatory. Demixing refers to the removal of the bright ‘A-team’ sources such as Cas A and Cyg A. Once received, in brief, the pipeline performs the following steps for a single RSM pointing (see Figure 2.9 for a flowchart representation of the process):

1. Each sub band is passed through a piece of software named the ‘Default Pre-Processing Pipeline’, where further compression or flagging steps can be applied if required, and placed in a working directory.
2. Sub bands are checked for corruption in addition to checking for missing sub bands which could sometimes occur. This is important when combining sub bands in a later step.
3. Calibrator sub bands are fully calibrated using the respective calibrator sky model. This is done using software named Black Board Selfcal (BBS; Pandey et al., 2009).
4. The resulting calibrator gains are clipped for outliers (optional), and then these solutions are transferred to the respective target sub band, at the same frequency, per beam.
5. The 40 sub bands from each beam are combined into 4 bands, each consisting of 10 sub bands.
6. A phase-only calibration is then applied to each of the bands based on an assumed sky model of the target field (see Section 3.2.4 for information on the sky model). This is followed by further automatic flagging steps (optional).
7. The two observations that make up a pointing are then combined to create a single pointing measurement set.
8. Imaging is performed, per band, per beam, followed by mosaicing the field (optional). The imaging is performed using the AWimager (Tasse et al., 2013).

The pipeline was designed to streamline LOFAR data reduction, along with adding various tools and additional steps that had been developed as part of the wider LOFAR commissioning




Fig. 2.9 A flowchart of the steps undertaken by the RSM Processing Pipeline to process an RSM observation. ‘SBs’ refer to sub bands and ‘DPPP’ refers to the Default Pre-Processing Pipeline. All calibration was performed with BBS. Note that when the data is combined into bands, this is done on a per frequency per 11 min observation basis, resulting in 24 band datasets. This reduces to 12 once the datasets are combined in time. All steps are performed with no manual intervention.

effort by various people. It also made sure to utilise parallel processing in an attempt to make the processing as fast as possible. Examples of fully processed individual RSM HBA and LBA pointings can be seen in Figure 2.10.

Though mainly used for RSM observations, the pipeline script provided a useful method to process any other LOFAR observations. This was possible because the settings of the pipeline were completely customisable due to the use of configuration files. This included the configuration files of the software being used within the pipeline which were provided by the user and read-in by the script, as opposed to being fixed settings. This allowed the processing to be altered to suit specific science goals. It should be noted that in 2014, a dedicated LOFAR calibration and imaging working group was formed to greatly improve the software and scripts available to calibrate LOFAR data. This resulted in new software




Fig. 2.10 *Upper panel*: A single HBA ‘pointing’ which is constructed from six separate beams at a central frequency of 150 MHz. This particular observation was recorded on 2015 April 10 and has an average sensitivity of 10 mJy beam^{-1} with a resolution of 2.1×2.1 arcmin (BPA 8.9°). At a 5σ level, ~ 800 sources are detected. *Lower panel*: The same field as the upper panel but now observed with the LBA on 2013 October 30 at a central frequency of 57 MHz, note the larger area of sky covered. This image has an average sensitivity of 90 mJy beam^{-1} with an resolution of 6.2×6.2 arcmin (BPA 58.1°). At 5σ , ~ 700 sources are detected. The colour bar units are Jy beam^{-1} .

along with a dedicated user pipeline which uses a facet calibration technique (van Weeren et al., 2016). Outside of such software being available, the RSM processing pipeline, which can be found on Github², provided a very useful method to swiftly calibrate LOFAR data. The following works have benefited from the RSM pipeline script, acknowledged by my co-authorship:

- Broderick J. W., et al., *Low-radio-frequency eclipses of the redback pulsar J2215+5135 observed in the image plane with LOFAR*, 2016, MNRAS, 459, 2681,
- Broderick et al., *LOFAR high-band observations of SS 433 – W50*, in preparation,
- Harwood J. J., et al., *FR II radio galaxies at low frequencies - I. Morphology, magnetic field strength and energetics*, 2016, MNRAS, 458, 4443,
- Heesen et al., *LOFAR reveals the giant: a low-frequency radio continuum study of the outflow in the nearby FR I radio galaxy 3C 31*, submitted, MNRAS,
- Marcote B., et al., *Orbital and superorbital variability of LS I +61 303 at low radio frequencies with GMRT and LOFAR*, 2016, MNRAS, 456, 1791.

2.2.2 The V404 Cygni 2015 Outburst

Discovered by the Japanese GINGA satellite during an outburst in 1989, V404 Cygni, or GS 2023+338, is a low-mass X-ray binary (LMXB) which resides at a distance of 2.4 kpc from Earth (Miller-Jones et al., 2009). The system is believed to consist of a massive stellar black hole that accretes matter from a companion star which is of a spectral type K0; and it is the most luminous black hole LMXB in quiescence.

After lying dormant for 26 years, renewed activity of the source was detected on 2015 June 15 by the Swift Burst Alert Telescope (Barthelmy et al., 2015), and subsequently confirmed by the Monitor of All-sky X-ray Imager which is housed on the International Space Station (Negoro et al., 2015). These alerts prompted a world-wide, multi-wavelength

²<https://github.com/ajstewart/rsmp>

observational campaign aimed at V404, collecting vast amounts of data on what ultimately became a spectacular black hole outburst. Wild flaring events continued for approximately 10 days after this detection, with the source eventually returning to quiescent levels (in X-rays) by August 2015.

LOFAR was triggered to observe V404 Cygni on 2015 June 23 where a six hour observation was performed, resulting in a 30 mJy (8σ) detection of the source at 150 MHz (Broderick et al., 2015). This was the first LOFAR detection of a transient X-ray binary, and the first successful detection of a source following a ToO observation with LOFAR. A total of six further observations were carried out on the dates of June 26, 28 and July 3, 6, 10, 12. Successful detections of the source were made in the June observations, with V404 appearing at its brightest in LOFAR observations in the June 28 run at 130 mJy. After this date the source was no longer detected. Figure 2.11 shows how the source appeared in the brightest detection image along with a light curve in Figure 2.12, which also shows the very well sampled 16 GHz light curve obtained by the Arcminute Microkelvin Imager (AMI; Zwart et al., 2008) as part of the AMI Large Array Rapid Response Mode (ALARRM; Staley et al., 2013).

It is known that synchrotron events evolve much slower at low frequencies than they do at high frequencies (van der Laan, 1966), and this is tentatively demonstrated by the light curve in Figure 2.12. The detected LOFAR flux was rising after a dramatic period of activity, though unfortunately the LOFAR coverage is lacking during the ~ 8 day period where the source may have continued to brighten and then decay. In addition, these observations will also play a crucial role in understanding the self-absorption and electron energy spectrum associated with the synchrotron emission.




Fig. 2.11 Image of how V404 Cygni appeared on 2015 June 28 with LOFAR at 150 MHz. The source itself can be seen clearly within the inset and is denoted by a circle. The sensitivity of the image is approximately 5 mJy beam^{-1} with a resolution of $2.1 \times 2.1 \text{ arcmin}$ (BPA -15.7°). The colour bar units are Jy beam^{-1} .




Fig. 2.12 Plotted are three reported 150 MHz LOFAR flux densities of V404 Cygni, along with a final upper limit measurement, and the well sampled light curve at 16 GHz observed by AMI as part of the ALARRM project. Low-frequency emission evolves on longer time-scales than at high frequencies and we see evidence of this here, as the LOFAR flux is slowly rising after a very active period at high frequencies. *Figure produced by Rob Fender.*

2.3 The Multifrequency Snapshot Sky Survey (MSSS)

MSSS is the first northern-sky imaging survey to be performed by LOFAR, spanning a large frequency range of 30–158 MHz. Before MSSS, no wide-area radio surveys with a large fractional bandwidth had been performed, hence MSSS has the unique ability to provide information on the spectral properties of detected sources over a broad range. The eventual completion of the SKA will mean that these type of multi-frequency imaging surveys will become much more commonplace, utilising multi-beaming techniques to vastly improve survey speed. These two features make MSSS an important first application of methods which are expected to be typical during the SKA era.

In terms of LOFAR, MSSS was conducted with two primary goals. First of all, the imaging survey was designed to provide a ‘global sky model’ which future LOFAR observations would use for calibration purposes. Due to the wide-field imaging nature of LOFAR, it was always envisioned that initial processing and calibration steps would be performed with minimal, if not any, human intervention. A fundamental part of such a process is the need to have a priori sky model of the arbitrary region of sky that is being observed. Hence, given LOFAR’s frequency range, it is important for the sky model to also incorporate spectral properties of the sources in question, and to be of the correct resolution for LOFAR. Second, the survey acted as a focussed commissioning project for LOFAR where the full spectrum of systems and telescope operations were tested. It also helped with focussing development efforts for processing LOFAR data; creating automated processing pipelines and overcoming any unexpected issues experienced with the data.

2.3.1 Survey Design

MSSS is split into two components: MSSS-LBA and MSSS-HBA. These two components make use of the two different types of LOFAR antenna previously discussed in Sections

2.1.1 and 2.1.2. Starting with MSSS-LBA, this portion of the survey covers the range of 30–75 MHz, and consists of 660 fields tiling up the northern sky. Each field is observed nine times over the course of the survey with each observation being 11 min in duration. Calibrators are observed simultaneously. MSSS-HBA covers 119–158 MHz in frequency and has a total of 3616 fields. These fields are visited twice during the course of the survey with each observation lasting 7 min. In this case the calibrators are observed in serial to the target observations. In each portion of the survey, three fields could be observed simultaneously using the 16-bit correlator mode, and this was increased to six once the 8-bit mode was available for use. The full survey is expected to yield the detection of 150,000–200,000 sources.

For full details of MSSS, including results from the first ‘verification fields’, the reader is referred to Heald et al. (2015).

2.3.2 Commissioning Work

The following are various pieces of commissioning MSSS work that I was involved in, as part of a larger MSSS working group. This took place between 2012 and 2013 and my involvement is acknowledged by co-authorship on the Heald et al. (2015) publication.

Testing the calibration schemes - sky model subtraction. The aim for this early MSSS-LBA project was to investigate whether faint sources, that were not present in the calibration sky model, could be uncovered in a particular field by subtracting all the sources from the data after initial calibration. Different sky models were also created from the clean components of the field which had been processed normally, in addition to using a source extractor, PyBDSM (Mohan & Rafferty, 2015). Along with the different sky models, we also varied the subtraction strategy by performing, and not performing, a traditional ‘self-calibration’ step before the subtraction.

We discovered that a full or phase-only self-calibration step, followed by a subtraction using a clean component based sky model produced the most success in subtracting sources. However, the four brightest sources in the field would prove slightly problematic and would commonly leave negative residuals. We were unable to detect any new sources using this technique. We suspected that the ionosphere had a negative influence on this particular dataset as it was observed during dawn, hence we were unable to obtain a perfect subtraction.

Testing the peeling technique. Following on from the previous work, an automated ‘peeling’ process was developed and the task for this investigation was to test how well the process performed on an LBA calibrator field, i.e. a field which contained a dominating bright source. In this case the field was that of the calibrator 3C 196, with the calibrator at the pointing centre of the observations. Peeling refers to the process of solving for directional dependent gains for specific bright sources in the field. These solutions can then be applied to the rest of the field, and the particular source can also be, optionally, well subtracted. The subtraction commonly removes imaging artefacts associated with the bright source and also improves the overall sensitivity of the image.

It was found that peeling 3–5 sources consistently improved the image quality for two different observations which were recorded at different times. An example of the improvement is shown in Figure 2.13. The improvement in sensitivity compared to ‘pre-peeled’ images ranged from 5–20 per cent. In turn, this also meant that more sources could now be detected in the resultant images. However, we also discovered that the peeling process, as it was at the time, seemed to have an adverse effect upon the measured flux of the peeled sources; making them slightly brighter (15 per cent) than what was expected from the sky model. In the following weeks, this was found by other workers to be caused by a portion of the script which attempts to scale the fluxes of the peeled sources, with the gain derived by the peeling process in each direction.




Fig. 2.13 *Left*: The 3C 196 field, with the source itself in the centre of the image, before peeling is applied. *Right*: The same field after 5 sources have been peeled. This is a particularly striking example of the improvement possible with peeling where the image sensitivity improved from approximately $400 \text{ mJy beam}^{-1}$ to $125 \text{ mJy beam}^{-1}$. These observations were recorded in early 2012.

I would later provide results from LBA directional dependent calibration to allow for a ‘phase-screen’ for the imaging process to be developed. A phase screen tracks the ionospheric conditions during the observation in different directions of the sky, allowing for the appropriate corrections to be made. While later MSSS efforts focussed on HBA data, these directional calibration techniques would continue to be researched and eventually evolve into processes such as the facet calibration technique (van Weeren et al., 2016).

Initial HBA data processing. This final project was simply the task of processing some of the first MSSS-HBA observations. This involved checking whether the processing methods could be transferred from LBA to HBA, what tweaks needed to be made to processing settings, and to gauge a general level of data quality. It was very welcome that HBA was slightly less complex to calibrate and process compared to LBA data with data quality at a consistently high level. The peeling technique was also tested and continued to be a useful method in improving data quality, as can be seen in Figure 2.14.




Fig. 2.14 *Upper panel*: An early MSSS-HBA field observed in 2012 after standard processing, with a sensitivity of 15 mJy beam^{-1} . *Lower left*: A zoomed in version of an area from the main field before peeling has taken place. *Lower right*: The same zoomed in area but now peeling has been applied. The sources are much better defined, with the peeling process removing artefacts from around bright sources. Some of these improvements are highlighted by red circles. The colour bar units are Jy beam^{-1} .

2.4 LOFAR Commissioning & TKP Publications

My LOFAR commissioning efforts are acknowledged by co-authorship on the following publications as part of the ‘LOFAR Builders List’ or as a member of the TKP:

- Buitink S., et al., *A large light-mass component of cosmic rays at 10^{17} – $10^{17.5}$ electron-volts from radio observations*, 2016, *Nature*, 531, 70,
- Carbone D., et al., *New methods to constrain the radio transient rate: results from a survey of four fields with LOFAR*, 2016, *MNRAS*, 459, 3161,

- Corstanje A., et al., *The shape of the radio wavefront of extensive air showers as measured with LOFAR*, 2015, APh, 61, 22,
- Garsden H., et al., *LOFAR sparse image reconstruction*, 2015, A&A, 575, A90,
- Heald G. H., et al., *The LOFAR Multifrequency Snapshot Sky Survey (MSSS). I. Survey description and first results*, 2015, A&A, 582, A123,
- Jelić V., et al., *Initial LOFAR observations of epoch of reionization windows. II. Diffuse polarized emission in the ELAIS-N1 field*, 2014, A&A, 568, A101,
- Morosan D. E., et al., *LOFAR tied-array imaging of Type III solar radio bursts*, 2014, A&A, 568, A67,
- Nelles A., et al., *Measuring a Cherenkov ring in the radio emission from air showers at 110-190 MHz with LOFAR*, 2015, APh, 65, 11,
- Noutsos A., et al., *Pulsar polarisation below 200 MHz: Average profiles and propagation effects*, 2015, A&A, 576, A62,
- Orrù E., et al., *Wide-field LOFAR imaging of the field around the double-double radio galaxy B1834+620. A fresh view on a restarted AGN and doubletjes*, 2015, A&A, 584, A112,
- Pilia M., et al., *Wide-band, low-frequency pulse profiles of 100 radio pulsars with LOFAR*, 2016, A&A, 586, A92,
- Sobey C., et al., *LOFAR discovery of a quiet emission mode in PSR B0823+26*, 2015, MNRAS, 451, 2493,
- Swinbank J. D., et al., *The LOFAR Transients Pipeline*, 2015, A&C, 11, 25,
- van Weeren R. J., et al., *LOFAR Low-band Antenna Observations of the 3C 295 and Boötes Fields: Source Counts and Ultra-steep Spectrum Sources*, 2014, ApJ, 793, 82.

2.5 LOFAR Repairs

Gale force winds in early January 2012 caused 1 HBA tile and 15 LBAs to be either destroyed or damaged at the UK LOFAR station which is located in Chilbolton, Hampshire. In some



Fig. 2.15 Myself (left) with two co-workers: Rene Breton (middle) and Małgorzata Piętka, posing after our successful construction of a new LBA at the UK LOFAR station. *Image credit: STFC/LOFAR-UK.*

cases the entire LBA structure had been completely blown away, or damaged by rogue HBA debris. Students were hastily rounded up and sent to repair the damage the following month, which involved replacing the broken (or non-existent) equipment and a lot of hammering of pegs to avoid a repeat. Hence, a small bit of hardware commissioning work was performed, as seen in Figure 2.15 which shows myself and two co-workers looking rather proud after our successful completion of a new low band antenna.

Chapter 3

Searching for Transients at the North Celestial Pole with LOFAR

The work presented in this chapter has been published as

Stewart, A. J., Fender, R. P., Broderick, J. W., et al., 2016, MNRAS, 456, 2321

3.1 Introduction

Radio transient surveys are being revolutionised by the development of the current generation of radio facilities. These include the new low-frequency instruments discussed in Chapters 1 and 2 such as LOFAR, MWA and the Long Wavelength Array (LWA; Ellingson et al., 2013). The telescopes listed offer a large FoV coupled with an enhanced sensitivity, creating the opportunity to survey wide areas of the sky for low frequency (≤ 330 MHz) radio transients and variables.

Previous low frequency surveys using older telescopes were detailed in Chapter 1, however, these new facilities have already produced some interesting results in this largely unexplored parameter space. Bell et al. (2014) searched an area of 1430 deg^2 for transient and variable sources at 154 MHz using the MWA. No transients were found with flux densi-

ties > 5.5 Jy on time-scales of 26 minutes and one year. However, two sources displayed potential intrinsic variability on a one year time-scale. Using the LWA, Obenberger et al. (2014a) detected two kilojansky transient events while using an all-sky monitor to search for prompt low-frequency emission from GRBs. They were found at 37.9 and 29.9 MHz, lasting for 75 and 100 seconds respectively, and were not associated with any known GRBs. This was followed up by Obenberger et al. (2014b) who searched over 11 000 hours of all-sky images for similar events, yielding 49 candidates, all with a duration of tens of seconds. It was discovered that 10 of these events correlated both spatially and temporally with large meteors (or fireballs). This low-frequency emission from fireballs was previously undetected and identifies a new form of naturally occurring radio transient foreground.

Two transient studies have now also been completed using LOFAR. Carbone et al. (2016) searched 2275 deg^2 of sky at 150 MHz, at cadences of 15 minutes and several months, with no transients reported to a flux limit of 0.3 Jy. Cendes et al. (2015) searched through 26, 149-MHz observations centred on the source Swift J1644+57, covering 11.35 deg^2 . No transients were found to a flux limit of 0.5 Jy on a time-scale of 11 minutes.

This chapter presents the results of using the LOFAR telescope to search 400 hours of observations centred at the North Celestial Pole (NCP; $\delta = 90^\circ$), covering 175 deg^2 at a frequency of 60 MHz.

3.1.1 Previous NCP Transient & Variable Studies

A previous study of variable radio sources located near the NCP field ($75^\circ < \delta < 88^\circ$) was carried out by Mingaliev et al. (2009). This study identified 15 objects displaying variability at centimetre wavelengths on time-scales of days or longer. However, the variability amplitude was found to be within seven per cent, which would not be able to be distinguished with LOFAR due to general calibration uncertainties at the time of the project. In addition,

the lower observing frequency used in this work would mean that the expected peak flux densities would be significantly lower, assuming a standard synchrotron event (e.g., van der Laan, 1966), making them challenging to detect. Also, the lower frequency means that the variability would occur over even longer time-scales, again assuming that the emission arises from a synchrotron process.

3.2 LOFAR Observations of the NCP

The monitoring survey of the NCP was performed between 2011 December 23–2012 April 16, resulting in a total of 2609 observations being recorded. The NCP was chosen because it is constantly observable from the Northern Hemisphere, and the centre of the field is located towards constant azimuth and elevation (az/el) coordinates. However, this is not true for sources which lie away from the NCP, where these sources rotate within the LOFAR elliptical beam. Therefore, the transient search was restricted to an area around the NCP where the LOFAR station beam properties are consistent for each epoch observed, avoiding systematic errors in the light curves that might be introduced if this was not the case. It is also an advantage that the line-of-sight ($b=122^{\circ}.93$, $l=+27^{\circ}.13$) is located towards a relatively low column density of Galactic free electrons; the maximum expected dispersion measure (DM) is 55 pc cm^{-3} according to the NE2001 model of the Galactic free electron distribution (Cordes & Lazio, 2003).

The NCP measurements were taken using the LOFAR Low-Band Antennas (LBA) at a single frequency of 60 MHz; the bandwidth was 195 kHz, consisting of 64 channels. The total integration time of each snapshot was 11 min, sampled at 1 s intervals, and data were recorded using the ‘LBA_INNER’ setup, where the beam is formed using the innermost 46 LBA antennas from each station, which gives the largest possible FoV and a full width half maximum (FWHM) of 9.77° .




Fig. 3.1 Histograms giving a general overview of when the 2609 11-minute snapshots of the NCP were observed. The top panel contains a histogram showing how many snapshots were observed on each day over the entire four month period, colour coded by month, which shows the distinct observing blocks in which NCP observations were obtained. The bottom panel displays a ‘zoom-in’ of the date range 21:00 2012/02/10–21:00 2012/02/12 UTC, now showing the number of snapshots per hour. This emphasises further the sometimes fragmented nature of the observing pattern of the NCP, with which careful consideration had to be given on how to combine the observations for the transient search.

3.2.1 Observation Epochs

The programme piggybacked on another commissioning project being performed by LOFAR at the time, the Multifrequency Snapshot Sky Survey (MSSS) – the first major LOFAR observing project surveying the low-frequency sky that was described in Chapter 2 (Heald et al., 2015). With every single MSSS LBA observation that took place, a beam was placed on the NCP using one subband of the full observational setup for MSSS. Figure 3.1 shows a histogram of the number of NCP snapshots observed each day over the duration of the programme, in addition to a similar histogram showing the number of snapshots per hour for a particular set of days. Of the 2609 snapshots, 909 were recorded during the day and 1700 were recorded at night. The MSSS observational set-up also meant that each 11-minute snapshot in the same observation block was separated by a time gap of four minutes.

3.2.2 Calibration & Imaging

3.2.3 Pre-Processing Stages

Prior to any processing, radio-frequency interference (RFI) was removed using AOFLAGGER (Offringa et al., 2010, 2012a,b) with a default strategy, in addition, the two channels at the highest, and lowest, frequency edges of the measurement set were also completely flagged. This reduced the bandwidth to 183 kHz. When using an automatic flagging tool such as AOFLAGGER, it is important to be aware of the fact that transient sources could be mistakenly identified as RFI by the software. This is a complex issue which is beyond the scope of this work. However, an initial investigation for the LOFAR case was carried out by Cendes et al. (2015). In these tests, simulated transient sources, described by a step function, with different flux densities and time durations (from seconds to minutes), were injected into an 11 minute dataset. These datasets were subsequently passed through AOFLAGGER before calibrating and imaging as normal in order to observe how the simulated transient was affected by the automatic flagging, if at all. The authors concluded that transient signals shorter than a duration of two minutes could be partially, or in the case of $\sim$ Jansky level sources, completely flagged. However, there are some caveats to this testing: short time-scale imaging was not tested for short-duration transients, and it remains to be determined how the automatic flagging would treat other types of transients (i.e. a non step function event). Hence, while these results certainly suggest that transients could be affected by AOFLAGGER, further testing is required to completely understand how automatic flagging software can affect the detection of a transient.

At this stage all data from international LOFAR stations were also removed, leaving just the Dutch stations. This was due to the complex challenges in reducing these corresponding data at the time of processing. Following this, the ‘demixing’ technique (described by van der Tol et al., 2007) was used to remove the effects of the bright sources Cassiopeia A and

Cygnus A from the visibilities. Finally, averaging in frequency and time was performed such that each observation consisted of 1 channel and an integration time of 10 seconds per time step. The averaging of the data was necessary to reduce the data volume and computing time required to process the data.

This averaging has the potential to introduce effects caused by bandwidth and time smearing, which are discussed in more detail by Heald et al. (2015) in relation to MSSS data. Following Heald et al. (2015), I used the approximations given by Bridle & Schwab (1999) to calculate the magnitude of the flux loss (S/S_0) in each case, assuming a projected baseline length of 10 km. The bandwidth smearing factor was calculated to equal 7 per cent (using a field radius corresponding to the FWHM) and the time smearing factor was 0.4 per cent. Thus, while the effect of time smearing was negligible, the impact of bandwidth smearing was potentially significant, yet remained within the calibration error margins (10 per cent; see Section 3.4.1).

3.2.4 Main Processing Stages

A selection of flux calibrators, characterised by Scaife & Heald (2012)¹, were used in the main processing of the data and were observed simultaneously utilising LOFAR's multi-beam capability (thus the calibrator scans were also 11 minutes in length). The calibrators and their usage can be found in Table 3.1. The standard LOFAR imaging pipeline was then implemented which consists of the following steps. First, the amplitude and phase gain solutions, using XX and YY correlations, are obtained for each calibrator observation using Black Board Selfcal (BBS; Pandey et al., 2009). These solutions are direction-independent, and are derived for each time step using the full set of visibilities from the Dutch stations, as well as a point source model of the calibrator itself. Beam calibration was also enabled which accounts, and corrects, for elevation and azimuthal effects with the station beam. The

¹Cygnus A is not characterised by Scaife & Heald (2012), but extensive commissioning work (summarised by McKean et al. 2011 and McKean et al. in prep.) has produced a detailed source model.

Table 3.1 The calibrators used for the NCP observations. It was decided early in the MSSS programme that 3C 48 and 3C 147 might not be adequate as calibrators for the LBA portion of the survey, and so these were dropped 8 and 22 days after first use, respectively. Observations using these calibrators displayed no disadvantages over those observed with other calibrators when checked in this project, and hence they were kept as part of the sample.

Calibrator Source	% Use	First Use Date	Last Use Date
3C 48	2%	2011 Dec 24	2012 Jan 01
3C 147	6%	2011 Dec 23	2012 Jan 14
3C 196	43%	2011 Dec 24	2012 Apr 14
3C 295	40%	2011 Dec 24	2012 Apr 01
Cygnus A	9%	2012 Jan 28	2012 Apr 16

amplitudes of these gain solutions were then clipped to a 3σ level to remove significant outliers, which were not uncommon in these early LOFAR data. The gain solutions were then transferred directly from the calibrators to the respective NCP observation.

Second, a phase-only calibration step was performed (also using BBS) to calibrate the phase in the direction of the target field. The solutions were derived using data within a maximum projected uv distance of 4000λ (20 km; 24 core + 10 remote stations). In order to perform this step, a sky model was obtained of the NCP field using data from the global sky model (GSM) developed by Scheers (2011). This model is constructed by firstly gathering sources which are present within a set radius from the target pointing in the 74 MHz VLA Low-Frequency Sky Survey (VLSS; Cohen et al. 2007). In the NCP case, the radius was set to 10 deg. From this basis, sources are then cross-correlated, using a source association radius of 10 arcsec, with the 325 MHz Westerbork Northern Sky Survey (WENSS; Rengelink et al., 1997) and the 1400 MHz NRAO VLA Sky Survey (NVSS; Condon et al., 1998) to obtain spectral index information. In those cases where no match was found, the spectral index, α (using the definition $S_\nu \propto \nu^\alpha$), was set to a canonical value of $\alpha = -0.7$. No self-calibration was performed on the data.

The main MSSS project discovered that observations recorded during this 2011-2012 period potentially contained one or more bad stations, and the data quality would improve if

such stations were removed. LOFAR was still very much in its infancy at the time, and, as a result, was not entirely stable; problems such as network connection issues or bad digital beam forming contributed to the poor performance of some stations. Hence, an automated tool was developed which analysed each station, identifying and flagging those that displayed a significant number of baselines with high measured noise. This tool was utilised in the NCP processing and primarily removed stations with poorly-focussed beam responses (Heald et al., 2015). It should be noted that present LOFAR data no longer require this tool as the issues outlined above have been rectified.

Finally, a FoV of 175 deg^2 was imaged using the AWIMAGER (Tasse et al., 2013), with a robust weighting parameter of 0 (Briggs, 1995), and a primary-beam (PB) correction applied to each image. A maximum projected baseline length of 10 km was used in this study (2000λ ; 24 core + 7 remote stations). This was chosen to obtain good uv coverage and a maximum resolution for which I was confident with the calibration. The typical resolution for the 11-minute snapshots was $5.4 \times 2.3 \text{ arcmin}$.

3.2.5 Quality Control

A number of bad-quality observations were detected and subsequently flagged using two methods: (i) checking the processed visibilities and (ii) inspecting the final images for each 11-minute snapshot. When analysing the visibilities, poor snapshots were flagged when the calibrated visibilities had a mean value greater than the overall mean of the entire four month dataset plus one standard deviation value. A slight, or indeed dramatic rise in the mean of the visibilities does not necessarily imply a completely bad dataset: an extremely bright transient ($> 100 \text{ Jy}$) could have this effect, for example. Such events may have been previously seen from flare stars at low frequencies (Abdul-Aziz et al., 1995), although at shorter time-scales than 11 minutes ($\sim 1 \text{ s}$). However, overall, the survey is less sensitive to extremely bright events because of this quality control step. It was beyond the scope of

this project to fully investigate this possible effect, and so only measurement sets that were deemed to be sufficiently well calibrated were used.

The results from the automated flagging were also checked against a manual, by eye, analysis of the visibility plots and the snapshot images, the latter enabling the detection of more bad observations. In total, 460 (out of 2609) snapshots were marked as bad, and were discarded from the search. The large size of the full dataset meant that there was no single common reason as to why individual snapshots were rejected, but the problems that caused rejection were mostly due to RFI or ionospheric issues. After the quality control was completed, 2149 observations (394 hr) were considered in the analysis.

3.3 Transient & Variability Search Method

3.3.1 Time-scales Searched

As the properties of the target transient population are unknown, the complete dataset was split and combined in various ways to fully explore the transient parameter space available. Along with performing a search on the original snapshots, each with an integration time of 11 minutes, searches were also performed on images with integration times of 30 seconds, 2 minutes, 55 minutes and 297 minutes. For the longer-duration images, only those 11-minute snapshots which were four minutes apart were combined together and imaged. This was to keep the visibilities as continuous as possible in the search for transients. After the quality control step described in Section 3.2.5, 297 minutes was the longest continuous integration time possible. All calibration was performed on each individual 11-minute snapshot; for the longer time-scales the relevant datasets were combined and then imaged.

3.3.2 The Transients Pipeline

The analysis of the data and search for radio transients was performed using software developed by the LOFAR Transients Key Science Project, named the Transients Pipeline (TraP). It is built to search for transients in the image plane, whilst also storing light curves and variability statistics of all detected sources. Moreover, it is designed to cope with large datasets containing thousands of sources such as this NCP project. A full and detailed overview of the TraP can be found in Swinbank et al. (2015)². In brief it performs the following steps:

1. Input images are passed through the TraP quality control which examines two features of the images. Firstly, the rms of the map is compared against the expected theoretical rms of the observation, and if the ratio between the observed and theoretical rms is above a set threshold then the image is flagged as bad. In this case, the threshold was set to the mean ratio value of each time-scale plus one standard deviation. The second test involves checking that the beam is not excessively elliptical by comparing the ratio of the major and minor axes. If this value is over a set threshold then the image is also flagged as bad. All bad images are then rejected and are not analysed by the TraP (see Rowlinson et al. in preparation for methods of setting these thresholds). The number of images accepted by the TraP compared to the total entered can be seen in Table 3.2.
2. Sources are extracted using PYSE - a specially developed source extractor for use in the TraP (Spreeuw, 2010, Carbone et al. in preparation). Importantly, all sources are initially extracted as unresolved point sources, which would be expected from a transient event.
3. For each image, the source extraction data are analysed to associate each source with previous detections of the same source, such that a light curve is constructed. In cases

²The work presented in this chapter primarily used TraP release 1.0. However, the data were re-processed once TraP release 2.0 was available, which is the version described by Swinbank et al. (2015), to confirm results.

where no previous source is associated with an extraction, the source is flagged as a potential ‘new source’ and is continually monitored from the detection epoch onwards.

For the source extraction two variables are defined: an island threshold, which defines the region in which source fitting is performed, and a detection threshold where only islands with peaks above this value are considered. These island and detection thresholds were set to 5σ and 10σ respectively. While the use of a 10σ detection threshold may seem very conservative, I agree with the arguments presented by Metzger et al. (2015) (hereafter MWB15) who advocate this criteria when identifying a transient source. In their paper, the authors’ main motivation for this high threshold is the significant possibility of spurious signals such as those seen in previous radio transient searches (Gal-Yam et al., 2006; Ofek et al., 2010; Croft et al., 2011; Frail et al., 2012; Aoki et al., 2014), arising from calibration artefacts, residual sidelobes and other similar issues. This was in addition to being generally cautious as this survey was one of the first conducted with the new LOFAR telescope. As also stated by MWB15, previous surveys have used 5σ as a detection threshold, which will of course increase the number of potential transient detections; however, this will also yield a high number of false detections, especially with the large number of epochs being used in this survey. Thus, minimising false detections and obtaining a manageable number of transient candidates were further motivations to use a 10σ detection threshold. I refer the reader to MWB15 for further discussion on this topic.

The transient search was also constrained to within a circular area of radius 7.5 deg from the centre of the image. This was to avoid the outer part of the image which was much noisier and did not have reliable flux calibration.

For each light curve, two values are calculated in order to define whether a source is a likely transient or variable: V_v , a coefficient of variation, and η_v , the significance of the

variability (Scheers, 2011). V_v is defined as

$$V_v = \frac{s_v}{\bar{I}_v} = \frac{1}{\bar{I}_v} \sqrt{\frac{N}{N-1} \left(\overline{I_v^2} - \bar{I}_v^2 \right)}, \quad (3.1)$$

where s is the unbiased sample flux standard deviation, $\bar{I}$ is the arithmetic mean flux of the sample, and N is the number of flux measurements obtained for a source. The significance value, η_v , is based on reduced χ^2 statistics and indicates how well a source light curve is modelled by a constant value. It is given by

$$\eta_v = \frac{N}{N-1} \left(\overline{\omega I_v^2} - \frac{\overline{\omega I_v}^2}{\bar{\omega}} \right), \quad (3.2)$$

where ω is a weight which is inversely proportional to the error of a given flux measurement ($\omega = 1/\sigma_{I_v}^2$). Throughout this thesis I define these parameters as the ‘variability parameters’. I also note that these parameters are defined to be able to identify sources that appear variable compared to the rest of the sample, as opposed to defining the variability of an individual source. To describe how the parameters evolve over a light curve, I consider an example where a transient is detected only once and the subsequent measurements are only noise, which is slightly varying. In this situation, as the noise measurements are added to the light curve the V_v and η_v values will generally increase and decrease respectively from their initial values (the initial values are dependent on the specifics of the measurements such as brightness and error estimates). As N epochs of noise are added, the fractional variability becomes larger hence the increase of V_v , while at the same time the light curve becomes flatter hence η_v decreases. For a more detailed discussion on these parameters I refer the reader to Swinbank et al. (2015) and in particular Scheers (2011).

In this chapter, to define a transient or variable source, two methods are used. For clarity a ‘transient’ is defined as a new source that appears within the dataset that was not present before, where as a ‘variable’ is a source detected in all epochs but shows significant

variability. First, with an emphasis on variables, a histogram of each parameter for the sample was created and fitted with a Gaussian in logarithmic space in order to place cuts based upon the Gaussian at $n\sigma$. In particular, a 3σ threshold is used in this project, with any source that exceeds this value in either V_v or η_v being flagged as a potential variable candidate. Rowlinson et al. (in preparation) will offer an in-depth discussion on finding variable and also transient sources using these methods. Second, as previously mentioned the TraP flags all sources which appear during the dataset, for which the transient parameters may not exceed the 3σ cut placed on the data. Thus, all ‘new’ sources are analysed manually without the use of the histogram method. This strategy was refined and improved for the project detailed in Chapter 5.

3.4 Results

3.4.1 Image Quality

Examples of the 30 s, 11 min and 297 min time-scale images can be found in Figure 3.2. Note that imaging the NCP can sometimes cause confusion when displaying the right ascension (RA) and declination (Dec) on the image axis, as the grid lines become circular. The grid lines are shown in all figures to help demonstrate this. The obtained uv coverage of the 11 and 297 min observations can be viewed in Figure 3.3. The average sensitivity reached with each time-scale is summarised in Table 3.2, along with the number of epochs available after the quality control described in Sections 3.2 and 3.3. It is important to note that, as a consequence of the primary beam correction, search areas centred on the NCP do not have a uniform noise level. Larger search areas include noisier regions further from the phase centre, and hence the flux density threshold at which a transient could be detected across the full search area will be higher. Figure 3.4 shows an example of a primary beam map from one of the NCP observations. In order to obtain a noise estimate accounting for the variation




Fig. 3.2 Examples of the NCP field maps at different time-scales. Where present, the area within the black circle indicates the portion of the image searched for transients. This was the same for each time-scale and had a radius of 7.5° . *Upper left panel*: an image on the 30 s time-scale which was observed on 2012 January 9. Using projected baselines of up to 10 km, the map has a resolution of 4.2×2.3 arcmin (synthesized beam position angle [BPA] -39°) with a noise level of 1.9 Jy beam^{-1} . Only the source 3C 61.1 is detected at a 10σ level, and this source is marked on the image. *Upper right panel*: an 11 minute snapshot observed on 2011 December 31. The noise level is $320 \text{ mJy beam}^{-1}$ and the resolution is 5.6×3.6 arcmin (BPA 43°). The number of detected sources at a 10σ level is now ~ 15 . *Lower left panel*: an example of the longest time-scale images available of 297 minutes, constructed by concatenating and imaging 27, 11-minute sequential snapshots. Observed on 2012 February 4, this image has a resolution of 3.5×2.0 arcmin (BPA -6°) and a noise level of $140 \text{ mJy beam}^{-1}$, with ~ 50 sources now detected at a 10σ level. *Lower right panel*: a magnified portion of the lower left panel image. The colour bar units are Jy beam^{-1} .

Table 3.2 The average image sensitivity and number of epochs for each time-scale at which a transient search was performed. The accepted epochs column defines how many of the total number of images passed the TraP image quality control.

Time (min)	Average rms (mJy beam ⁻¹)	Typical Resolution (arcmin)	Total # Epochs	Accepted # Epochs
0.5	3610	4.8×2.2	47 970	41 340
2	2110	4.7×2.1	10 739	9 262
11	790	5.4×2.3	2 149	1 897
55	550	4.9×2.1	371	328
297	250	3.1×1.4	34	32




Fig. 3.3 *Left panel:* The uv coverage obtained with an 11 min snapshot. *Right panel:* The improved uv coverage gained when combining 27 snapshots (297 min). In each case the uv range is limited to $\pm 2 \text{ k}\lambda$ (10 km).




Fig. 3.4 An example of a normalized primary beam map from one of the NCP observations, which has been scaled to 1.0. The bold, outer solid-line circle represents the full extent of the area for which the transient search was performed (radius of 7.5 deg). The inner solid-line circles show how the area was divided in order to gain an estimate of the average rms for each image accounting for the primary beam. The dashed-line circle indicates the position of the primary beam half-power point.

caused by the beam, for each image at each time-scale I split the area into four annuli, equally spaced in radius. These four regions are also marked on Figure 3.4. The rms for each annulus was then measured, using a clipping technique, with the area-weighted average of these four values providing the single value rms estimate for the individual image. The average of each time-scale, is then used as the sensitivity levels stated in Table 3.2. Figure 3.5 shows that these measured rms values of the different time-scales approximately follow a $1/\sqrt{t}$ relation, where t represents the integration time of the observation. It is noted that the longer time-scale rms values appear to lie above the $1/\sqrt{t}$ relation. This is believed to be caused by the clipping technique being less accurate at measuring the rms of the longer time-scale images annuli. This in itself due to the presence of many more sources compared to the relatively source free short time-scale images. In addition to this, it is possible the CLEAN




Fig. 3.5 The average rms obtained from the images produced by combining and splitting the dataset. Also plotted in light grey are the range of noise values for the individual images at their respective time-scales, in addition to the $1/\sqrt{t}$ relation where t is the integration time of the observation. It can be seen that the average rms values approximately follow this relation; the longer time-scale values are likely to be slightly overestimated due to the methods used to estimate the rms. The errors shown on the average points are one standard deviation of the rms measurements from the respective time-scale.

algorithm was not applied to a deep enough level in some cases. Hence, the combination of these two methods means that the longer time-scale rms values are likely to be slightly overestimated, but not at a concerning level in the context of this project.

The transient search could have been limited to a smaller region with the deepest sensitivity; however, when calculating the figure of merit ($\text{FoM} \propto \Omega s^{-\frac{3}{2}}$ where Ω is the FoV and s is the sensitivity) it can be shown that it is more beneficial to extend the area of the search, despite the increase in average rms. This can easily be demonstrated as the full area is 16 times larger but the weighted sensitivity only drops by a factor of about two; hence the FoM is around five times better, illustrating the motivation for searching wide area. I refer the reader to Macquart (2014) for an in-depth discussion of the FoM in the context of transient surveys.

The 55 and 297 min time-scale images offered the best flux calibration stability from image to image due to the better uv coverage achieved on these time-scales. An example

of the general flux calibration quality can be seen in Figure 3.6, which shows the averaged measured flux across all the 297 min snapshots of sources detected at 60 MHz, cross-matched with the VLSS catalogue at 74 MHz. It shows a general agreement with the fluxes that would be expected assuming an average spectral index of $\alpha = -0.7$. If it is assumed that all sources have this spectral index and proceed to calculate the expected VLSS 60 MHz flux for each source, the average ratio of this expected VLSS flux against the measured LOFAR flux is 1.00 ± 0.17 . Also shown in Figure 3.6 is the difference between reported sky locations of the LOFAR and VLSS source positions. It shows there is a good agreement and were consistent within 5.1 arcsec on average; the typical resolution in the LOFAR band is 3.1×1.4 arcmin for the 297 min time-scale.

Overall there was a typical scatter of 10 per cent in each light curve of sources detected, which was measured by the TraP. It was common that fainter sources ($< 10\sigma$) would appear to ‘blink’ in and out of images; this was especially apparent in the 11 minute snapshots. This was likely due to a mixture of varying rms levels and the ionosphere causing phase calibration issues. Such behaviour was a further reason why a 10σ source detection limit was used in the transient search. The sensitivities of the shortest time-scale maps, 30 s and 2 min, were such that only the brightest source in the field, 3C 61.1, was detectable.

It was also important to determine whether the images produced for the transient search are confusion limited. In order to calculate an estimate of the confusion noise for the average resolutions presented in Table 3.2, I followed the same approach as Heald et al. (2015), using VLSS C-configuration estimates (see Cohen, 2004) which I extrapolate to 60 MHz using a typical spectral index of -0.7 . The formula was altered to also account for the non-circular beams:

$$\sigma_{\text{conf,VLSS}} = 29 \left(\frac{\theta_1 \times \theta_2}{1''} \right)^{0.77} \left(\frac{60 \text{ MHz}}{74 \text{ MHz}} \right)^{-0.7} \mu\text{Jy beam}^{-1}, \quad (3.3)$$

where θ_1 is the synthesized beam size major axis and θ_2 is the minor axis. For the five time-scales used in the transient search shown in Table 3.2, beginning with 30 sec, the confusion




Fig. 3.6 *Left*: The mean extracted flux of sources from the 297 minutes NCP survey at 60 MHz against the cross-matched VLSS survey at 74 MHz. The solid line represents the expected LOFAR flux density assuming a spectral index of $\alpha = -0.7$. For illustrative purposes a dashed-line representing $\alpha = 0$ (a 1:1 ratio) is also shown. *Right*: The positional offset for each cross-matched source between the LOFAR NCP and VLSS, with the dashed lines representing the (0,0) offset point.

noise estimates are calculated to be 113, 107, 128, 111 and 57 mJy beam⁻¹ respectively. Thus, due to the simple reduction strategy, the NCP images, at best, are approximately 4× the confusion noise level and hence would not affect the transient search.

Along with these cadences, a deep map was constructed by using all the available 297 min images, reaching a sensitivity of 71 mJy beam⁻¹ (this value was measured using the weighted average method discussed above in this section). The map can be seen in Figure 3.7. This, however, had to be produced by means of image stacking as opposed to direct imaging due to the amount of data involved. A total of 150 sources were detected at a 10σ level within the same 7.5 deg radius circle used for the transient search, with the map primarily being used as a deep reference image for the field. However, I used this deep map to verify the calibration and imaging procedures by comparing the detected source counts to the VLSS. Firstly, using a spectral index of -0.7 , $S_{60} = 710$ mJy corresponds to a flux density at 74 MHz of $S_{74} = 613$ mJy. Using this flux density limit, there are 263 catalogued VLSS




Fig. 3.7 The deepest map produced of the NCP field from the survey. It was constructed by averaging all 31 of the 297-minute-duration images together in the image plane, using inverse-variance weighting. It has a noise level of 71 mJy beam^{-1} and a resolution of $3.1 \times 1.4 \text{ arcmin}$ (BPA 42°). A total of 150 sources are detected at a 10σ level within a radius of 7.5 deg from the centre of the map. While none of these sources are previously undetected, it provided a detailed reference map to check any transient candidates. The colour bar units are Jy beam^{-1} .

sources within 7.5 deg of the phase centre. Cross-correlating the VLSS with the LOFAR 60 MHz detections resulted in 41 per cent of the VLSS sources having a LOFAR match. The factor of ~ 2 discrepancy can be shown to be due to the primary beam attenuation in the deep map. Hence, I was satisfied that the calibration and imaging results were valid and consistent with previous studies, and therefore would not negatively impact any transient searches.

Performed jointly with Jess Broderick (ASTRON), the deep map was further analysed for any previously uncatalogued radio sources, but none were found. However, the direct comparison to VLSS revealed that one source, located at $02^{\text{h}}13^{\text{m}}28^{\text{s}} +84^{\circ}04'18''$, has apparently significantly different 60 and 74 MHz flux densities: the VLSS integrated flux density is 1.49 ± 0.25 Jy, whereas in the LOFAR band it is detected at the 8σ level with a integrated flux density of 236 mJy. There are no detections of the source in WENSS or NVSS. However, this source is located within a stripe feature in the VLSS image, and the source is not present in the VLSS Redux catalogue (VLSSr; Lane et al. 2014); hence we do not pursue this source further. The full MSSS survey will offer further insight into this potential source, confirming its flux density and spectral index, if it is real.

3.4.2 Variability Search Results

The four month dataset provides an opportunity to search for variable sources as well as transient sources. I define variables as sources which are present throughout the entire dataset, taking into consideration varying sensitivity, whose light curve displays significant variability over the period. This is opposed to transient sources, which I define as sources that appear or disappear during the time spanned by the dataset, again taking into account the varying sensitivity. Consulting historical catalogues also helps with the distinction between variables and transients. Due to the higher level of image quality, the variability search was limited to the two longest time-scales of 55 and 297 min. For each detected source in these two sets of images, variability parameters (V_v and η_v) were calculated by the TraP. Figure 3.8

shows the respective distributions of the variability parameters for each time-scale plotted in logarithmic space. In each case, the central panel shows η_v plotted against V_v for each detected source. The top panel displays a histogram representing the distribution of η_v of all the sources along with a fitted Gaussian curve. The right panel contains the distribution and fitted Gaussian curve for V_v . The dashed lines represent a 3σ threshold for each value; any sources with variability parameters exceeding one or both of these values are considered as potentially variable. Candidates also had to show a variability of significantly more than ten per cent, which was the calibrator error of the measurements. This was set at a level of 2σ from this value. An ideal transient would appear in the top-right-hand corner of the central panel scatter plot, exceeding the threshold in each parameter.

It can be seen that at both time-scales, no sources exhibit variable behaviour in V_v above a 3σ level, but one source has a significant η_v value. This source is 3C 61.1, which dominates the field. While the result points towards low-level variability of 3C 61.1, the source is a well resolved radio galaxy (Leahy & Perley, 1991) whose flux is dominated by 100-kpc-scale lobes, making it very unlikely that any intrinsic variability would be detected. It is more likely that this is the result of calibration errors and the source extraction and subsequent calculation of η_v itself. The model for 3C 61.1 used during this investigation is quite basic for such a complex source. This, along with ionospheric effects and the general calibration accuracy of the instrument at the time, can have quite a substantial effect on such a bright source, with such calibration errors not included in this analysis. The source is also spatially extended, but the extraction treats it as a point source (as mentioned in Section 3.3), and this will therefore also have a significant impact on the recorded flux. Removing the point source fitting constraint does indeed move the data point closer back towards the 3σ threshold, but only marginally by 0.1 dex in η_v . As for the η_v value, this parameter is weighted by the flux errors of the source extraction. Bright sources, such as 3C 61.1, are well fitted when they are extracted, which means they have small associated statistical flux errors. This in turn then




Fig. 3.8 The distribution of values obtained for the variability parameters V_v , a coefficient of variation, and η_v , the significance of the variability (see text for full definitions) for each light curve detected. *Left* shows the 55-minute image results and *right* shows the 297-minute time-scale results. In each case, the central panel plots the two values against each other for each source, with the top panel and right side panel displaying the histogram showing the distribution of the η_v and V_v values respectively for all sources. The dotted lines represent a 3σ threshold for each parameter. A very-likely variable or transient source would appear in the top-right of the plot, exceeding a 3σ level in each parameter. At both time-scales, one source (3C 61.1) is found to have a significant value in η_v . However this is likely to arise from fluctuations caused by calibration issues.

causes η_v to rise. If 3C 61.1 is discounted, no sources displayed any significant variability at the 55 and 297 min time-scales.

3.4.3 Transient Search Results

Using the TraP and a manual analysis of its results, searches performed on the time-scales of 0.5, 2, 55 and 297 minutes found no transient candidates. However, nine transient candidates emerged from the analysis of the 11-minute time-scale. At first it appeared strange to achieve nine candidates at one time-scale but none at any other. However, the sensitivity of the shorter time-scales was such that only bright transients (> 25 Jy) would have been confidently detected, and as previously stated no other source, or even artefact, was detected at these flux levels other than 3C 61.1. At the longer time-scales, the improved uv coverage meant that the images improved substantially in quality. This reduced the number of imaging artefacts that could spawn false detections and sources were consistently detected throughout the epochs (as opposed to many sources blinking in and out as discussed in Section 3.4.1). Any sources that were defined as ‘new’ by the TraP (sources which appeared that were not detected in the first image) were in fact association errors and not transient sources.

While the nine candidates could point towards the 11 minute images meeting the required sensitivity and time-scale of a transient population, these images are also the most likely to exhibit misleading artefacts due to the limited uv coverage. Hence, the nine reported candidates were subjected to a series of tests to determine whether they were spurious sources. The following tests were performed:

1. Subtraction of 3C 61.1 from the visibilities using the clean component model from the deconvolution process. The visibilities were then re-imaged.
2. Applying an extra round of RFI removal using AOFLAGGER.

3. Re-running the automated tool to remove perceived bad LOFAR stations from the observations, followed by a manual check.
4. Imaging the data using different weighting schemes and baseline cutoffs.

The tests were applied in the above order, meaning that if one method definitely succeeded in removing the candidate the latter tests were not performed. Two of the nine candidates completely survived all the tests; three were inconclusive but quite doubtful, whereas four were definite artefacts. One of the two passing sources was very marginal; hence efforts were directed towards the candidate of high confidence which is detailed in the next section. The other candidate and the reasoning why it was considered at a lower confidence level is detailed in Section 3.6. Details of all the candidates including information on which tests were passed and failed are presented in Table 3.3.

3.5 Transient Candidate ILT J225347+862146

The first candidate to have passed all the validity checks, was found in a single 11 min snapshot taken on 2011 December 24 at 04:33 UTC. The source was extracted by the TraP with a flux of 7.5 Jy (14σ detection in individual image), at coordinates $22^{\text{h}}53^{\text{m}}47.1^{\text{s}}$ $+86^{\circ}21'46.4''$, with a positional error of $11''$. It was only seen in this one epoch with no detection of the source in the preceding or subsequent snapshots. The observation can be seen in Figure 3.9. Nothing was present at the candidate position in either the relatively deep image constructed from the longer time-scale images (see Section 3.4.1) or the very deep image of the field from the LOFAR Epoch of Reionisation (EoR) group (Yatawatta et al., 2013). Note that the EoR project uses the LOFAR high-band antennas, and hence it is at a higher frequency range of 115–163 MHz.

Table 3.3 The results of the tests performed to determine the reality of the nine transient candidates. The tests were performed in a specific order such that if a particular test proved the transient to be an artefact, then the following tests were not performed. Test (i) was the subtraction of 3C 61.1; test (ii) was extra RFI flagging; (iii) was re-running the automated tool that removes underperforming LOFAR stations (as described in Section 3.2.2); and the last test (iv) was to re-image the observation using a different weighting scheme and baseline cutoff. The symbols defining the results of the tests are given at the foot of the table. Confidence in the transients is rated as very low, low, high or very high, with candidates rated high or very high being considered further. Candidates #1 and #6 are the only objects to pass all the tests. Along with the date, the identification number of each observation is also shown.

Candidate #	Obs. Date & ID	R.A. (J2000.0)	Decl. (J2000.0)	Integrated Flux (Jy)	Test Results				Confidence
					i	ii	iii	iv	
1	2011 Dec 24 L40032	22 ^h 53 ^m 45.1 ^s	+86° 21' 46.4''	7.5 ± 1.0	✓	✓	✓	✓	Very High
2	2012 Jan 01 L41549	08 ^h 01 ^m 29.5 ^s	+86° 36' 12.1''	6.9 ± 1.0	✓	✓	×	-	Very Low
3	2012 Jan 02 L41671	18 ^h 25 ^m 39.1 ^s	+87° 17' 55.8''	6.1 ± 1.0	✓	✓	×	-	Very Low
4	2012 Jan 21 L43496	02 ^h 43 ^m 36.2 ^s	+84° 49' 41.0''	7.0 ± 1.1	✓	○	○	✓	Low
5	2012 Jan 21 L43496	00 ^h 20 ^m 11.1 ^s	+86° 53' 31.6''	5.7 ± 0.9	✓	○	○	✓	Low
6	2012 Feb 11 L46442	13 ^h 06 ^m 41.0 ^s	+86° 59' 46.9''	5.3 ± 0.7	✓	✓	✓	✓	High
7	2012 Feb 11 L46667	02 ^h 32 ^m 22.1 ^s	+86° 59' 41.4''	4.2 ± 0.7	×	-	-	-	Very Low
8	2012 Mar 10 L51787	02 ^h 09 ^m 25.2 ^s	+86° 58' 01.9''	4.2 ± 0.7	×	-	-	-	Very Low
9	2012 Mar 31 L55369	02 ^h 08 ^m 39.0 ^s	+85° 42' 50.3''	7.3 ± 1.2	○	✓	✓	✓	Low

✓ *The transient candidate survived the test.*

○ *The transient candidate was noticeably affected by the test.*

×

3.5.1 A Mirrored Ghost Source

On closer inspection, the transient candidate appeared to have a secondary associated positive ‘ghost’ source mirrored across the brightest source in the field, 3C 61.1 (the transient lies at an angular distance of 3.2° from 3C 61.1), which can also be seen in Figure 3.9. The ghost was not detected by TraP due to the higher rms value in that region, and like the transient candidate it was a ‘new’ source with no previous or subsequent detections. In fact the ghost source was actually nominally brighter than the transient source with a flux density of 13 Jy. However, in the non-primary-beam-corrected map the candidate has a higher peak flux density (9 Jy) than the ghost (6 Jy). This was not the first time we, as a working group, had witnessed this type of effect in LOFAR observations, with previous commissioning data we had obtained in 2010 showing a similar situation. Currently, the exact explanation of why ghosts of this nature, including specifically the ghost presented in this work, are generated in LOFAR data is unknown. It should be noted that none of the other eight transient candidates detailed previously had an associated ghost source. In the following discussions I refer to the original detected transient source ILT J225347+862146, to the west of 3C 61.1, as the ‘transient candidate’ and the source to the east of 3C 61.1 as the ‘ghost’ source (refer to Figure 3.9).

Ghost Artefacts in Radio Interferometry

Calibration artefacts presenting themselves as spurious ‘ghost’ sources is not an entirely new topic to radio interferometry. The topic of ‘spurious symmetrisation’ is discussed in Cornwell & Fomalont (1999); in brief, if a point source model is used for a slightly resolved source, a single iteration of self-calibration can result in features of the image being reflected relative to the point-like object. However, this can be corrected with further iterations of self-calibration which would cause the spurious features to disappear. As will be discussed in Section 3.5.1, the ghost presented in this work can be seen before initiating any kind of self-calibration




Fig. 3.9 *Upper Panel*: Illustrates how the transient source (labelled ‘T’), ILT J225347+862146, was originally detected in the image, along with the associated ghost source (labelled ‘G’) across from 3C 61.1. *Lower Panel*: Now the measurement set has been re-calibrated with the transient included in the sky model; the ghost source has vanished. Upon closer inspection, other faint, source-like features also disappear from the re-calibrated image. These are most likely fainter ghost features which are reduced when the data were calibrated with a more complete sky model. The colour bar units are Jy beam^{-1} .

of the target field, i.e. any calibration using a target field sky model. Therefore, it is highly unlikely that the spurious symmetrisation previously described is the sole cause of the ghost. However, this is not to say that the effect plays no role in its creation.

More recently, Grobler et al. (2014) (hereafter ‘G14’) began a series of investigations dedicated to ghost phenomena. This first study concentrated on ghosts seen in data from the Westerbork Synthesis Radio Telescope (WSRT). In these data, ghost sources appeared as strings of (usually) negative point sources passing through the dominant source(s) in the field. The arrangement of these negative point sources appeared quite regular, along with the fact that the positions were not affected by frequency. In their investigation, G14 were successful in deriving a theoretical framework to predict the appearance of ghosts in WSRT data for a two-source scenario, and were able to confirm what previous work had suggested concerning these ghost sources (see text in G14).

In brief, the main features about ghosts to note are as follows: (i) they are associated with incomplete sky models, for example missing or incorrect flux; (ii) in the WSRT case, the ghosts always formed in a line passing through the poorly modelled or unmodelled source(s) and the dominant source(s) in the field; (iii) the ghosts are mostly negative in flux, while positive ghosts are rare and weaker; and (iv) the general ghost mechanism can also explain the observed flux suppression of unmodelled sources.

G14 also concluded that the simple East-West geometry of the WSRT array is the reason for ghosts appearing in a regular, straight line, pattern. This becomes more complex when a fully 2D/3D array is considered such as LOFAR, where the ghost pattern is expected to become a lot more scattered and noise-like. This subject is the focus of Paper II (Wijnholds et al., 2016) in the series on ghost sources. However, G14 did note that regardless of the array geometry, ghosts are expected to occur at the $n\phi_0$ positions, where ϕ_0 represents the angular separation between the respective bright source and unmodelled source, and n is an integer number. Usually the strongest ghost responses are the $n = 0$ and $n = 1$ positions, i.e. the

suppression ghosts that sit on top of the sources in question. However the case discovered in this work, and also two independent cases (de Bruyn, priv. comm., Clarke, priv. comm.) in LOFAR data suggest that the $n = -1$ position could also generate a strong response. What is significant about the transient presented in this work, however, is that the ghost appears as a positive source.

Investigating the NCP Ghost

Returning to the situation detailed in this section, I was presented with two sources for which either could be the real (transient) source or the ghost. I attempted to simulate the situation within real data, in order to investigate how the different stages of calibration would react to a bright transient, and to see if a positive ghost source could be generated. This was done by taking a different NCP observation and inserting a simulated transient source into the visibilities (the transient was set to be ‘on’ for the entire 11 min) before any calibration had taken place. The snapshot was then calibrated as normal, but importantly the inserted source was *not* included in the NCP sky model used for the phase-only calibration step (refer to Section 3.2.2). This test was repeated using various different sky positions and flux densities for the inserted source.

It was found that a significantly positive ghost source could only be produced if the flux of the simulated transient was relatively bright, ~ 40 Jy. An example can be seen in Figure 3.10. It was common for the total flux to be shared approximately equally between the simulated source and its associated ghost. However, not every position on the sky at which the transient was inserted produced a ghost source, a feature that could not be explained at the time. Yet, when a transient was inserted at the position of ILT J225347+862146, this did produce a ghost source. It was then possible to test what happened when the simulated source was included in the sky model. The result was that when the simulated source was accounted for perfectly in the sky model, the ghost source disappeared. If the sky model




Fig. 3.10 The resultant image after a simulated transient source (labelled ‘ST’) was inserted into the visibilities of a NCP observation and processed without the simulated source in the sky model. A ghost source (labelled ‘G’) appears mirrored across 3C 61.1. The effect is not limited to one specific insertion point of the simulated transient and is more pronounced the brighter the simulated transient. In this example, a source of brightness 80 Jy was inserted, which produces a very significant ghost source. The simulated transient and ghost source each had a measured flux density of ~ 25 Jy, with the remaining ~ 30 Jy being absorbed by 3C 61.1. The colour bar units are Jy beam^{-1} .

component was instead inserted at the location of the ghost source, while the ghost appeared brighter, the simulated transient never fully disappeared.

In light of the results from the simulations, the same sky model test was performed with the transient candidate and ghost in order to determine which source was the ‘real’ source. Recalling that the total flux of the transient candidate and ghost was $\sim 7 \text{ Jy} + \sim 13 \text{ Jy} \approx 20 \text{ Jy}$, I began by inserting a 20 Jy point source into the NCP sky model at the position of the transient candidate and re-calibrated the dataset. In this case it was observed that the flux of the ghost was significantly reduced, by ~ 70 per cent, and the candidate brightened by ~ 100 per cent. Alternatively, if the model component was entered at the ghost location, the candidate source and ghost respective fluxes were only ~ 10 per cent different from their

initial fluxes on discovery, i.e. when they were not in the sky model at all. In fact increasing the sky model component to 25 Jy and placing it back at the position of the transient candidate reduced the ghost such that it was no longer distinguishable from the noise, as seen in the bottom panel of Figure 3.9. Hence, the ‘real’ source was determined to be at the position TraP had originally reported, $22^{\text{h}}53^{\text{m}}47.1^{\text{s}} +86^{\circ}21'46.4''$, to the west of 3C 61.1.

The above tests have concentrated on the target NCP field sky model, but there was also the sky model which was used to calibrate the calibrator observation. The calibrator source for this observation was 3C 295. Considering that ghosts occur because of sky model errors, one could envision a scenario in which the error being transferred from the calibrator to the target field results in the ghost pattern observed. As mentioned in Section 3.2.2, the calibrator sky models only contain the calibrator source itself and not any surrounding field sources. While this generally allows the derivation of sufficiently accurate gain solutions, the missing flux could be attributed to a ghost pattern, which is then transferred to the target field (see also Asad et al. 2015 for a similar discussion regarding the 3C 196 field).

To investigate this, two tests were performed. First, the phase-only calibration step was ignored and the dataset was imaged using the amplitude and phase gain solutions directly from the calibrator. In this case, both the transient and the ghost were present, with no major changes from before (a result which makes ‘spurious symmetrisation’, previously discussed in Section 3.5.1, unlikely to be the sole cause of the ghost). Second, the calibrator observation was not used at all and instead the data were calibrated in both amplitude and phase using the constructed NCP target sky model (described in Section 3.2.2) which importantly did *not* contain the transient source. For this test, the solution interval was increased to one minute (originally 10 s) to gain more signal-to-noise for the calculations. I also had to perform post-processing clipping to the visibilities to eliminate bad amplitude spikes in the calibrated visibilities. In the full 11 min image, while the rms rose to $\sim 800 \text{ mJy beam}^{-1}$, a source was

detected within 1 arcmin (the resolution of the image was 5.6×2.4 arcmin) of the reported transient candidate position with a flux density of 13 Jy. The ghost source was not detected to a 5σ limit of 10 Jy at its expected location, nor was it visible when the map was manually inspected. However, due to the increase of the rms in this case, it cannot be stated with complete confidence that the ghost source is not present at all. Nonetheless, observing the transient source without placing it in the sky model provided additional evidence that I had identified the correct source.

The above result tentatively points to the calibrator having an important role in the ghost creation. However, understanding the exact ghost mechanism is a complex task in the LOFAR case, and each stage of the calibration must be taken into careful consideration. For example, G14 has exclusively investigated situations where full amplitude and phase calibration is used, so the effects of a phase-only calibration is generally unknown at this stage. At the time of writing, it cannot be explained how the ghost is generated; a detailed investigation is under way (Grobler et al. in preparation) in an attempt to resolve the matter.

3.5.2 Transient Flux Density

Obtaining the Correct Flux

The correct flux density of the transient proved difficult to ascertain. I attempted to obtain an estimate by entering flux values of the transient source manually into the calibration sky model, over a range of 7.5–45 Jy, in steps of 2.5 Jy, and proceeded to recalibrate the visibilities (as previously described, this calibration step is phase-only). I then observed the influence this had on the measured flux density of the transient itself, as well as the measured flux densities of some of the surrounding sources, including the ghost source. The reader is reminded that the transient candidate, the source deemed ‘real’, is to the west of 3C 61.1 and the ghost is the source to the east of 3C 61.1. The transient was always placed as a point

source in the sky model. The results of this experiment can be seen in Figure 3.11. It shows that the ghost source became increasingly fainter as the transient flux was increased, right up until the transient was entered as 20 Jy and the ghost could no longer be distinguished from the background. The transient ‘light curve’ itself follows the trend of the increasing sky model flux, but it also exhibits a sudden local maximum when the sky model entry level is changed from 22.5 to 25 Jy. In this instance the extracted flux rises from 16 Jy to 20 Jy. It then proceeds to fall back to an extracted flux level of 18 Jy and continues to rise as before.

As for the other nearby sources, while they are stable prior to the sky model transient component reaching 17.5 Jy, beyond this level they suffer a very noticeable decline that continues as the transient flux is increased. It is also apparent that the other sources in the field are affected by the before mentioned sudden local maximum of the transient light curve around a sky model flux of 25 Jy, with 3C 61.1 also showing a significant flux increase ($\sim 3\sigma$ to the scaled value). However, for VLSS 0110.7+8738 and VLSS 2130.1+8357, which are at a similar flux level to ILT J225347+862146, there is a hint of a decrease, although within the error bars of the flux measurements. In each case, once the sky model flux is increased to the next step, the measured fluxes return to their previous levels. When comparing the fluxes of the field sources with the corresponding averages from the four surrounding snapshots, it is seen that they mostly agree within all the error bars involved. The largest discrepancy comes from 3C 61.1, which appears ~ 10 per cent dimmer in the transient snapshot, which is outside the errors of the average measurement. However, the sudden increase around 25 Jy causes 3C 61.1 to match the surrounding average. This could be seen as a clue that this area represents the real flux of the transient; at this point, with 25 Jy in the sky model, the transient appears as 20 Jy in the image. Hence, with this information, the true flux of the source is associated with the point at which the ghost disappears and the other sources in the field are not heavily affected, which constrains the estimate of the flux density of ILT J225347+862146 to be in the range 15–25 Jy.




Fig. 3.11 The extracted flux of the transient candidate using the PySE source extractor against the manually defined flux entered into the sky model at the transient position when processing. Also shown is a measurement of the ghost source flux obtained by a forced fit at the ghost position. The input flux was defined in steps of 2.5 Jy, from 7.5 Jy to 45 Jy. The plot also shows the extracted fluxes of four other sources in the field in order to monitor any effects to other sources, along with the solid lines which show the average flux of these field sources from the four surrounding snapshots. The error on these averages is shown by the error bar at the beginning and end of the line. Above an input flux value of 20 Jy (extracted transient flux value of 17 Jy) it becomes apparent that the other sources are beginning to be affected. They drop sharply beyond an entered flux of 30 Jy by which point the ghost source is no longer statistically significant. Note that 3C 61.1 has been scaled by subtracting 40 Jy from its flux measurements.

Testing Known Sources

The test detailed above was performed directly on the two field sources that were monitored during the investigation, VLSS 0110.7+8738 and VLSS 2130.1+8357, with 60-MHz flux densities of ~ 9 and ~ 15 Jy respectively. This also included removing the sources from the calibration sky model as well as changing the input flux. Each source was treated as a separate case meaning that both were never subtracted from the sky model, or edited, at the same time. As before, these tests were performed at the transient epoch, but also in the two neighbouring epochs to ensure that any effects were not just local to the transient-containing snapshot. Without the source in the model, the measured flux was reduced by ~ 20 per cent, with the majority of the extra flux in the field being absorbed by 3C 61.1, which appeared slightly brighter. Once the source was reinserted into the sky model, even at a low flux, the source in question returned to the expected level. However, as the sky model input flux was increased, so did the extracted flux, which is consistent with how the transient acted previously. There was also no distinguishing feature that would enable a confident definition of these sources' 'correct flux' without prior knowledge. Thus, it is not a surprise that the transient flux in Section 3.5.2 is hard to identify purely from the behaviour of the source itself during calibration when altering the sky model. Ideally self-calibration would be used, but at the time of processing self-calibration with LOFAR was still a relatively untested technique.

Splitting the Dataset in Time

In the test detailed above, where the transient was inserted into the sky model with various different flux values, it was noticeable that the flux that was inserted was never the flux that was measured. If the transient was not 'on' for the entire 11 min, this could perhaps explain why this was the case. With the transient included in the sky model at a flux of 20 Jy, the observation was firstly split in half and imaged; however, the flux was consistent within the 1σ error bars between each half. To probe deeper, the 2-minute images produced




Fig. 3.12 The extracted flux of the transient candidate in 2-minute intervals, obtained using the TraP, for the cases where the transient is included in the sky model at 10 and 20 Jy. Data points represented by a square signify that the source was extracted from the image with a blind detection. The triangles represent flux values obtained from a forced fit at the source position, where the source was no longer above the source extraction threshold (13 Jy, 5σ). The two light curves follow the same trend, suggesting that the transient is brighter than a 3σ limit of 7.5 Jy in the 4-minute period of 04:34 - 04:38. The 10-Jy input case, returning a measured flux of 16 Jy, also suggests that 16 Jy may be the correct flux during this time frame. These light curves were obtained by extending the phase-calibration time interval to one min. The date of the observation was 2011 December 24.

as part of the transient search, were referred to which did not have the transient included in the sky model. This particular observation, however, was above the average noise level (1.8 Jy beam^{-1}) with an rms of $\sim 2 \text{ Jy beam}^{-1}$. Neither the transient nor ghost source had significant detections (with the significance level now reduced to 5σ in order to try and detect the transient), and even surrounding field sources were hard to distinguish because of the poorer image quality.

In an attempt to improve the situation, using the assumption that the transient should be included in the sky model for the observation, the dataset was phase-calibrated again using a larger solution interval of 1-minute (previously 10 s), to allow more signal to noise for the calculations. Using a source extraction threshold of 5σ , TraP was able to find the transient

source in the second and third of the 5, 2-minute images: the flux densities are 20.9 (8σ) and 18.7 (7σ) Jy, respectively. The light curve can be seen in Figure 3.12.

I was concerned about forcing the flux of the transient to a specific value by simply entering that value into the sky model, especially as in this case the fluxes returned were approximately equal to the flux which was entered (20 Jy). Thus, the test was repeated, but this time entering a 10-Jy transient at the position. In the 10-Jy case the transient was detected in the second image only at a lower flux density. The forced fit performed by the TraP in the third image yields a flux measurement of 13 Jy (just below 5σ), before dropping off, which mimics the characteristics of the 20 Jy sky model case.

The first 2-minute image from each test was of noticeably poorer quality than the other four, 2-minute images of the observation. As seen in Figure 3.12, the forced extraction at the transient location in the first image returns the same flux density value (10 Jy) in each sky model test case. This value hints at the transient being present in this epoch as this flux level is higher than the fourth and fifth epochs, where the transient is no longer detected in both cases. However, due to the uncertainty in this image and the larger error bars associated with this measurement, it cannot be stated for certain that this is the case.

I attempted to split the dataset which had been calibrated directly from the NCP field sky model, as discussed in Section 3.5.1, but the calibration was not of sufficient quality to achieve useful results.

The results therefore suggest that the transient was brightest between the second and sixth minute of the observation, a period of four minutes. However, I was unable to fully characterise the decay, or especially the rise time of the event, and hence cannot rule out the transient being active over a longer, 10-minute time-scale.

3.5.3 Testing if a Source Can be Created by the Sky Model

Because the transient did not correspond to any source contained in the sky model, a major concern was the possibility of ‘creating’ false sources in the field by purely inserting them into the sky model. This could explain the apparent responsiveness of the candidate to an entry in the sky model, and perhaps a source placed anywhere in the field would have the same effect, both in creating a source and causing the ghost source to disappear. This was tested in two ways. First, the snapshot containing the candidate was reprocessed with the candidate component of the sky model moved to an empty, unrelated location on the sky. This resulted in no source being ‘created’ at this location and also left the candidate, and ghost, unaffected from their original detection states.

The second test was to process the two preceding and two subsequent snapshots with the candidate component inserted into the sky model at its correct location. Previously, no detection was made of the candidate in any other snapshot, and as the data were recorded in sequence, the uv coverage of these observations were all very similar. The result was that, once more, no source was present at the candidate location, even when placed in the sky model; this can be seen in Figure 3.13, which shows the detection of the candidate along with the snapshots before and after in time.

These two results meant that simply entering sources into the sky model at an arbitrary position would not ‘create’ an artificial source. In contrast, the responsiveness of the transient candidate to such input at the correct position suggested it was a real source present in the data.

3.5.4 Further Validity Testing

A final set of tests and checks were performed to investigate whether ILT J225347+862146 was an unexpected artefact. With LOFAR being commissioned at the time, an artefact would




Fig. 3.13 A sequence of images in time: the transient detection image along with the snapshots before and after the event, together with a zoom-in of the transient location. Importantly, each observation has been processed with the transient included in the calibration sky model, showing how even with this taken into consideration, there are no significant detections before or after the transient. These images were created using the standard imaging parameters discussed throughout the chapter, including projected baselines of up to 10 km in length. The synthesized beam can be seen in the bottom-left of each image. The colour bar units are Jy beam^{-1} .

not be completely surprising. While the telescope was in a good working state, a lack of optimisation of aspects such as station calibration and beam models could cause issues. A series of tests were devised to rule out certain possible artefact causes, all performed with the source both in and out of the sky model when processing. These tests were:

- **Broadband RFI** - Care was taken to manually reduce the data, removing anything left over that was suspected of being RFI, as well as running another pass of AOFLAGGER on the data after calibration. Neither method affected the transient source.
- **Narrow-band RFI** - To rule out the possibility of narrow-band RFI, the already limited bandwidth was split into two and processed separately. The transient source remained

in each half of the bandwidth, with a flux consistent within the 1σ error bars between the two halves.

- **Calibrator Issues** - The calibrator observation contains nothing peculiar and was of good quality. The calibrator for this observation was 3C 295.
- **Calibrator Gains Only** - Previously discussed in Section 3.5.1, this test meant the phase-only calibration step using the target field sky model was skipped; instead the dataset was imaged with the gain amplitude and phase solutions obtained directly from the calibrator being applied. ILT J225347+862146 was still present in the resulting image along with the ghost.
- **Phase Centre Shift** - The phase centre of the observation was shifted to that of the transient (a shift of ~ 4 deg). The transient was still clearly visible with the shift, especially when the source was included in the sky model.
- **Equal Local Sidereal Time Observations** - Sixteen observations were found to have very similar local sidereal times (LST) to that of the detection measurement, so these were used to check whether the candidate was potentially caused by that particular projection of the baselines on the sky. There was no detection in any of these observations, which covered four months of recording.
- **Bad Station Removal** - Along with the automatic tool that was part of the initial four tests, a manual inspection of the data was also carried out, which was in agreement with the results from the tool: the same two stations were perceived as bad. After these stations were flagged, the image was generally cleaner from artefacts with the transient source unaffected.
- **Random Subset of Stations** - Half of the 33 stations used in the observation were randomly removed, after calibration, with the remaining data being re-imaged. This was repeated three times and the transient source continued to be present in each of three resulting maps.
- **Dirty Map Check** - The source is present in the dirty map.

- **Field Subtraction** - Using a sky model derived from the deeper image, the entire field apart from the transient was subtracted from the dataset. The transient and ghost were clearly visible in this case, with the same flux density.
- **Imaging at a Different Resolution** - Reducing the maximum baseline length used when imaging from 10 km to various lower values had no impact on the transient. An image using a maximum projected baseline length of 15 km was also produced in which the source was still present. However, any images produced with projected baselines longer than 10 km were not considered as scientifically useful, due to concerns regarding the quality of calibration.
- **Different Imaging Weighting Schemes & Imager** - Checking for further side-lobe related issues, the imaging was redone using natural and uniform weighting. The source remained in the resulting maps. The imager itself was also checked by imaging the observation with the ‘Common Astronomy Software Applications’ (CASA; McMullin et al., 2007) software rather than AWIMAGER (this meant that no primary beam correction was made) and the source was still present. In this case, the candidate source was marginally brighter than the ghost: the flux densities were 5.3 and 4.5 Jy respectively.

3.5.5 What is this Transient?

With the candidate successfully passing the numerous exhaustive tests detailed previously, I concluded that the candidate was a real astrophysical source. Hence its possible origin was investigated.

Catalogue Search & Multi-wavelength Follow-up

No source was found within a 2-arcmin radius of the transient position in historical radio catalogues, including VLSS, WENSS and NVSS, and there were no previous observations of

Table 3.4 Details of the follow-up observations performed with the Liverpool Telescope.

Obs. Date	# Epochs	Epoch Exposure Time (s)
2013/11/28	3	100
2013/12/21	3	100
2013/12/30	3	100
2014/01/12	2	300

the transient location in the NRAO (excluding surveys) or GMRT archives within 20 arcmin. Also, no potential counterpart or related object was found in high-energy catalogues, and no published gamma-ray burst or supernova event is known at the position.

Optical follow-up observations of the field using the Liverpool Telescope (LT; Steele et al., 2004) were carried out, though it is worth noting that the transient was discovered two years after the event date. Therefore, while a direct counterpart would not be observed, an object of a certain type may be identified in the vicinity of the transient, and could be potentially associated with the radio emission. These observations with the LT used the r' band filter on four different dates, which are shown in Table 3.4, totalling 11 epochs and 1500 seconds of exposure time. The optical campaign and the following analysis was performed by Teo Muñoz-Darias (IAC).

The combined image, shown in Figure 3.14, was calibrated against USNO-B1 (Monet et al., 2003) and reaches $r' \sim 22$ –22.5 mag. To search for a possible optical counterpart to ILT J225347+862146, we established an error on the positional measurement of the transient from the LOFAR data. This was calculated by accounting for the following uncertainties which contribute to the overall error: the error of the source extraction performed by the TraP which was reported as 11 arcsec and is measured following the error analysis of Condon (1997); the average scatter of the extracted positions of bright sources in the dataset by TraP, measured to be 7 arcsec; and finally the reported 5 arcsec positional error of the VLSS catalogue, which the phase calibration sky model is based upon. Summing these values in quadrature we gain the final positional error of 14 arcsec.




Fig. 3.14 The resultant combined optical image obtained with follow-up observations using the Liverpool Telescope. It reaches a depth of $r' \sim 22\text{--}22.5$ mag and is calibrated against the USNO-B1 catalogue (Monet et al., 2003). The inner circle marked on the image, centred at the reported LOFAR transient position, represents a 1σ positional error of radius 14 arcsec. The outer circle shows the 2σ positional error circle of radius 28 arcsec. The two stars which have proper motions higher than $100 \text{ mas year}^{-1}$, as indicated by the USNO-B1 (Monet et al., 2003) catalogue, are indicated by arrows. *Figure produced by Teo Muñoz-Darias (IAC).*

Using this value, we unambiguously detect four sources within a 2σ error circle centred at the transient position, one of which is fully enclosed by the 1σ uncertainty as shown in Figure 3.14. None of these sources displayed either strong short-term (minutes) or long-term (weeks) variability. We used the USNO-B1 catalogue to look for high-proper-motion stars in the field, with the assumption that these would be nearby objects. Two of the four sources mentioned previously within the 2σ radius, had associated proper motions of higher than $100 \text{ mas year}^{-1}$ and are marked in Figure 3.14. However, neither of these sources exhibited a colour consistent with being an M-dwarf or any other possible transient object (but we note that the error associated to the USNO-B1 catalogue colours is significant). In addition, both sources were detected in our LT images and no associated variability was observed. We also consulted the Wide-Field Infrared Survey Explorer (WISE; Wright et al., 2010) and the Two Micron All Sky Survey (2MASS; Skrutskie et al., 2006); a total of three and two sources

were located within the 2σ error circle in each survey, respectively. In each case, two objects were the previously reported high-proper-motion stars; however neither of these, or the one other source in WISE, were found to be variable.

Given possible further uncertainties of the accuracy of the measured position in the LOFAR band, we extended the error circle to a radius of 1 arcmin from the transient position. A total of 20 sources were within this larger error circle in our LT observations; however, as previously, there was no strong evidence of a possible association with ILT J225347+862146.

I concluded that there is no obvious counterpart candidate. I now consider whether the radio emission could arise from either an incoherent or coherent process.

Incoherent Origin

If an incoherent emission process is considered, a limit can be placed on the maximum distance of the source, by using the known characteristics of the transient along with assuming that its brightness temperature (T_b) is at the maximum $T_b = 10^{12}$ K limit, $T_{b\max}$, for (unbeamed) synchrotron radiation (Kellermann & Pauliny-Toth, 1969). To do this I make use of the Rayleigh-Jeans law, that is

$$T_{b\max} = \frac{\Delta L_\nu}{8\pi k \nu^2 \Delta t^2} , \quad (3.4)$$

where ΔL_ν is the change of the luminosity in time-scale Δt , k is the Boltzmann constant and ν is the observing frequency. The luminosity at frequency ν is defined as $L_\nu = 4\pi d^2 S$, where d is the distance and S is flux density. Using this and re-arranging Equation 3.4 an expression for the distance is obtained as follows:

$$d^2 = \frac{2k \nu^2 \Delta t^2 T_{b\max}}{\Delta S} . \quad (3.5)$$

For the flux density change $\Delta S = 20$ Jy in time $\Delta t = 10$ min, the result is a maximum distance of 13.7 pc. This points to the possibility of the transient being a nearby flare star.

In the case where the source is relativistic, with the synchrotron radiation now beamed, the observed brightness temperature could exceed 10^{12} K. This is seen in populations such as AGN (Horiuchi et al., 2004; Kovalev et al., 2005) and GRBs (Kulkarni et al., 1999; Anderson et al., 2014). From Equation 3.5 it can be seen that $d \propto T^{\frac{1}{2}}$. For example, if a brightness temperature such as 10^{16} K was observed, a value at which these sources can sometimes appear (Pietka et al., 2015), this would place the distance estimate of the transient at 1.4 kpc, an unrealistic distance for such classes of objects.

Various previous studies have investigated the radio transient properties of flare stars at both centimetre and decametre wavelengths (Lovell, 1969; Gudel et al., 1989; Bastian, 1990; Jackson et al., 1990; Abdul-Aziz et al., 1995; Osten et al., 2006). A recent study, Boiko et al. (2012), involved monitoring the flare stars AD Leonis ($d = 4.9$ pc) and EV Lacertae ($d = 5.1$ pc) with the UTR-2 telescope, located in Ukraine, during March 2010 and 2011. These observations, performed in the frequency range of 16.5–33 MHz, yielded a total of 167 and 73 detected radio bursts from the respective stars. In the case of AD Leonis, the average flux of the bursts was in the range of 10–50 Jy, seemingly consistent with the flux measured from ILT J225347+862146 at 60 MHz. However, one discrepancy is that the average duration of these bursts seen from AD Leonis, which is 2–12 seconds, is considerably shorter than the apparent $\sim$ minutes of activity observed for ILT J225347+862146. In addition, the sole detection of the transient is possibly suspicious in this context. One would expect the detection of subsequent transient events from an active flare star over a period of four months (the time-scale of this transient search). Recent results such as Notsu et al. (2013), which show *Kepler* solar type stars exhibiting super-flares, could offer an explanation, with some flare events being active for ~ 10 min time-scales (Schaefer et al., 2000). The flare star hypothesis could be tested further by directing future observations towards the Galactic plane rather than the NCP, as the density of flare stars should be dramatically increased.

If a flare star origin is assumed for ILT J225347+862146, along with the object not being detected in the follow-up optical image to a depth of $r' \sim 22$ mag, this information can be used to make a crude estimate of the distance of such an object such that it is consistent with a non-detection. For this a catalogue of 463 UV Cet-type flare stars compiled by Gershberg et al. (1999) is consulted, selecting objects which have a measured R band magnitude. The stars in the catalogue have a maximum distance of 50 pc, and I also only consider M type stars. These selected stars are then divided into two sub-type groups: early type stars (M0-4) and late type stars (M5-9). In total these groups have 69 and 20 stars respectively. I then calculated the average absolute magnitudes of these two groups, and subsequently at what average distance the population would be if the apparent magnitude became 22 mag. These distances were found to be ~ 2 kpc for the early type stars and ~ 0.3 kpc for the late type stars (consistent with estimated Galactic scale height values defined by Holwerda et al. 2014). Hence, the object responsible for ILT J225347+862146 could be further than these distances, dependent on spectral sub-type, if it is not detected in the deep optical image obtained with the LT. However, at these distances, this would make the radio event extremely luminous, which is unlikely. Therefore, it would become likely that ILT J225347+862146 is a nearby sub-stellar object. These values could be further constrained by a deeper analysis of stellar objects, so this conclusion is made tentatively. It is also worth noting that, although we were unable to pick out a possible responsible object, it is entirely possible that the source is present in the deep optical image.

Coherent Origin

Hassall et al. (2013) showed that when the scattering of a highly-dispersed burst is significant, imaging surveys for FRBs (refer back to Chapter 1 for a description of FRBs) can be competitive with pulsar-like, high-time-resolution surveys, if not more sensitive, in detecting such events. To investigate this possibility, I consider how the scattering time and fluence

of the four FRB events reported in Thornton et al. (2013) at 1.3 GHz, compare to this event at 60 MHz. Firstly, to achieve estimates for the scattering time of the Thornton events at 60 MHz, I use the standard relation of

$$\tau_s(\nu) \propto \nu^\gamma, \quad (3.6)$$

where $\gamma = -4$. For the purposes of this scenario I ignore any dispersion effects and assume that the recorded burst duration is dominated by scattering. This assumption is quite reasonable when considering the value of any dispersion induced smearing of the signal, Δt_{DM} , which can be calculated per MHz of bandwidth by

$$\Delta t_{\text{DM}} = 8.3 \times 10^3 \text{ DM } \nu_{\text{MHz}}^{-3}, \quad (3.7)$$

where DM is the dispersion measure and ν_{MHz} is the observing frequency in MHz. With a bandwidth of 183 kHz and an observing frequency of 60 MHz, $\Delta t_{\text{DM}} = 7 \times 10^{-3} \text{ DM s}$. Thus, even with a DM value of 1000 pc cm^{-3} , Δt_{DM} would only cause 7 s of smearing. Table 3.5 shows that the predicted, scatter-dominated width of the events at 60 MHz range from 242–1234 s (taking the upper limit values), with the highest value belonging to FRB 110220. Hence, the maximum duration of ILT J225347+862146, $< 660 \text{ s}$, is quite consistent with that expected from an FRB at 60 MHz. In reality, only FRB 110220 showed any evidence of scattering.

Next the fluence of the events are compared. Taking a width of 11 min for ILT J225347+862146 ($6.6 \times 10^5 \text{ ms}$), and the flux as 20 Jy, the fluence of ILT J225347+862146 can be stated as $20 \times 6.6 \times 10^5 = 1.3 \times 10^7 \text{ Jy ms}$ at 60 MHz. Taking the shorter time-scale of four minutes (as discussed in Section 3.5.2) at 20 Jy gives a fluence of $4.8 \times 10^6 \text{ Jy ms}$. The event with the highest fluence as reported in Thornton et al. (2013) was FRB 110220 with 8 Jy ms at 1.3 GHz. The fluence values can be compared when assuming different spectral

Table 3.5 The observed width of the four bursts reported in Thornton et al. (2013) along with the estimated width of the event at 60 MHz. To calculate the estimated width the relation $\tau_s(\nu) \propto \nu^\gamma$ is used where $\gamma = -4$. Dispersion effects are ignored and I assume a scenario of the reported widths being dominated by scattering. Of the four reported bursts, FRB 110220 was the sole event to show any evidence of scattering.

Event	Observed width at 1.3 GHz (ms)	Estimated width at 60 MHz (s)
FRB 110220	5.6	1234
FRB 110627	< 1.4	< 309
FRB 110703	< 4.3	< 948
FRB 120127	< 1.1	< 242

indices. In the case of $\alpha = 0$, a direct comparison is possible, showing that the LOFAR event has a vastly greater fluence than the known FRB. Assuming $\alpha = -2$ and extrapolating the peak flux of the LOFAR event to 1.3 GHz, the fluence now becomes 2.81×10^4 Jy ms for the 11 min scenario. This is still much greater than FRB 110220. For the LOFAR event to be consistent with this particular burst, which is by far the highest fluence of the four reported bursts in Thornton et al. (2013), then a spectral index of $\alpha \sim -4.7$ would be required. This implies that the LOFAR event would be an abnormally bright FRB, even more so than the bright Lorimer et al. (2007) burst at 30 Jy. The characteristic spectral index of FRBs is currently not well defined, but a value of -4.7 would be very steep regardless of the population. Although the time-scale of ILT J225347+862146 is consistent with a scattered FRB at low-frequencies, the required steep spectral index along with the inconsistency between the fluence of the known FRB events, casts considerable doubt regarding an FRB origin. With the exact characteristics of FRBs currently unknown, it cannot be stated that ILT J225347+862146 belongs to the same population.

It should be noted, as described by Güdel (2002), that coherent emission from plasma processes can also occur in stellar objects such as flare stars. Such emission has also been seen from the Sun at low-frequencies; for example in type III solar radio bursts (Bastian et al.,

1998). For this hypothesis, the arguments which were presented in Section 3.5.5 concerning a flare star origin also apply here. In particular, if this were the origin of ILT J225347+862146, it appears unusual to not see the event repeat over a four month period, yet this possibility cannot be ruled out.

Other Populations

Other populations such as AGN and X-ray binary systems were considered. However, there is insufficient evidence to confirm or rule out such classes as the origin of ILT J225347+862146.

Over the past decade, a variety of new radio transient sources have been attributed to different kinds of neutron stars. These include populations such as Rotating Radio Transients (RRATs; McLaughlin et al., 2006) and intermittent pulsars (Kramer et al., 2006), with some intermittent pulsars seen to have periods in the off-state of more than a year (Camilo et al., 2012). It is possible that the transient reported in this chapter could be an atypical isolated neutron star such as these described populations; however at this time it is not possible to present any evidence to support this hypothesis.

3.6 Transient Candidate ILT J130641+865946

This was the second candidate to pass all the validity tests presented in Section 3.4.3 (candidate #6). It was observed on 2012 February 11 at 03:18 UTC and was located at $13^{\text{h}}06^{\text{m}}41.0^{\text{s}} +86^{\circ}59'46.9'$. It was also only seen in this single snapshot, with an apparent flux density of 5.3 ± 0.7 Jy. A zoom-in of the source can be seen in Figure 3.15. There were no possible associated sources present in the deep NCP map, VLSS, WENSS or NVSS, no matches to high-energy catalogues, and no association with any GRBs or supernova events.

As the source lies further away from 3C 61.1 than the previous candidate discussed (6.6 deg compared to 3.2 deg for ILT J225347+862146), any associated ghost would lie outside the imaged FoV. A wider image was created but no ghost source was detected.




Fig. 3.15 This map shows transient candidate #6, with an inset showing a zoomed-in section centred on the candidate which is marked by the crosshair. It was observed on 2012 February 11 at 03:18 UTC and has a resolution of 4.7×4.2 arcmin (BPA 61°) with an rms of $950 \text{ mJy beam}^{-1}$. For clarity, this image was created using a maximum projected baseline length of 8 km, due to a very elongated beam when using a maximum of 10 km. The colour bar units are Jy beam^{-1} .

Whether a strong ghost response should be expected or not in this case is currently an open question. The simulations previously discussed in Section 3.5.1 seemed to suggest that some transient locations were more sensitive to ghosts than others, so this candidate could reside within a ‘weak ghost region’.

The source was subjected to further testing similar to that detailed above for the previous transient. It remained unaffected when checked for broadband or narrowband RFI, processed with calibrator amplitudes and phases, imaged using different weighting schemes and resolutions, and was present in the dirty map. The observation was split in half; however, this was inconclusive due to the first half of the measurement set being of poor quality, with real sources along with the candidate disappearing. The candidate was present in the second half. The source was entered into the processing sky model, including a range of different

flux values, to observe how it behaved with this addition. Unlike the previous candidate, the source was unaffected by any sky model input in terms of flux or appearance.

This source was discovered after ILT J225347+862146 and unfortunately due to time restrictions has not been investigated to the same level of detail as ILT J225347+862146. This also is the reason for the lack of multi-wavelength follow-up observations. In spite of this, the tests performed, detailed above, suggest that this source behaves differently to the previous candidate, which currently has no explanation. On the other hand, there is no evidence to suggest that the candidate is a spurious source. Thus, while the detection is noted, I assume that ILT J225347+862146 is the sole confirmed transient discovered in this investigation. If the second transient is indeed real, then it would appear that this object has very similar characteristics to ILT J225347+862146, although significantly fainter at 5 Jy, and could possibly belong to the same population.

3.7 Transient Surface Density & Rates

No transients were found at four of the five time-scales searched, with one detection in the other. This allows us to place upper limits on the rate of low-frequency transient events on the whole sky at these time-scales. To calculate the upper limits of the surface density of transients, Poisson statistics are used, specifically:

$$P(n) = e^{-\rho A} \quad (3.8)$$

where ρ is the surface density of sources per square degree and A represents the equivalent solid angle obtained by multiplying the area of the sky surveyed, Ω , by the number of epochs $N - 1$. In using this approach I am also assuming isotropic probability of a transient detection, i.e. events are assumed to be extragalactic in origin. At the Poisson 2σ confidence level $P(n) = 0.05$ can be defined, and by rearranging Equation 3.8, the respective value of ρ can

Table 3.6 Summary of the transient surface densities (ρ) and general information of the results of this work (*top section*) and other low-frequency (≤ 330 MHz) transient surveys (*bottom section*). A similar approach to Ofek et al. (2011) is followed where δt is the time-scale of each individual epoch searched in the survey, and Δt is the cadence time-scales(s) of the epochs observed. N_e and N_T signifies the number of epochs and transients respectively. A value denoted by ‘-’ signifies that I was unsure of the correct value from the literature. A time value of ‘cont.’ means the observations were continuous. These values are those which are used in Figures 3.16 and 3.17.

Survey/Reference	Flux Limit (Jy)	ρ (deg ⁻²)	δt	Δt	ν (MHz)	N_e	N_T
This work	> 36.1	4.1×10^{-7}	30 s	cont.-4 months	60	41 350	0
This work	> 21.1	1.8×10^{-6}	2 min	cont.-4 months	60	9 262	0
This work	> 7.9	1.4×10^{-5}	11 min	4 min-4 months	60	1 897	1
This work	> 5.5	5.2×10^{-5}	55 min	4 min-4 months	60	328	0
This work	> 2.5	5.3×10^{-4}	297 min	4 min-4 months	60	32	0
Lazio et al. (2010)	> 2500	9.5×10^{-8}	5 min	2 min - 5 min	74	~1 272	0
Obenberger et al. (2014a)*	> 1440	2.2×10^{-9}	5 sec	5 sec	74	~43 056	2
Cendes et al. (2015)	> 0.5	10^{-2}	11 min	min - months	149	26	0
Carbone et al. (2016)	> 0.3	1.3×10^{-3}	11 min	min - months	150	151	0
Bell et al. (2014)	> 5.5	7.5×10^{-5}	5 min	min - 1 year	154	51	0
Hyman et al. (2009)* ^a	> 30×10^{-3}	0.034	~3 hr	days - months	235,330	-	3
Jaeger et al. (2012)*	> 2.1×10^{-3}	0.12	12 hr	1 day - 1 month	325	6	1

^a Reported as $1.4 \times 10^{-2} \text{ yr}^{-1} \text{ deg}^{-2}$. Using the integration time of 5 s, this converts to $2.2 \times 10^{-9} \text{ deg}^{-2}$.

^b Values for this survey are obtained from the calculations performed by Williams et al. (2013) which takes into account results from Hyman et al. (2005, 2006, 2009).

be obtained for each time-scale, recalling that the area of sky searched in each case was 175 deg^2 . For no detections $P(0) = 0.05$ is used, whereas for the 11-minute time-scale this becomes $P(1) = 0.05$ because of the single transient detection. These values can be found in Table 3.6.

ILT J225347+862146 provides us with $1_{-0.95}^{+3.74}$ transient events detected in the 11-minute time-scale search, using upper and lower limits at 95% confidence as defined by Gehrels (1986). As a 10σ limit was used for the source extraction, the flux density limit of this search was 7.9 Jy; moreover, 1897 11-minute epochs are equivalent to 14.5 days of observations. This equates to a transient rate of $3.9_{-3.7}^{+14.7} \times 10^{-4} \text{ day}^{-1} \text{ deg}^{-2}$.

However, it should be noted that the flux density limit of this rate is defined with the assumption that sources are ‘non-ghosted’. As seen in this work, transients with an associated ghost can be reduced in brightness when not accounted for in processing. While the magnitude of the effect is not yet defined, it can be estimated by taking ILT J225347+862146 as an example. In this case the source was originally detected as 7.5 Jy with an accompanying ghost, and when accounted for in processing, the minimum flux estimate was 15 Jy in the 11 minute image (see Section 3.5.2). Therefore a minimum reduction in flux density of 50 per cent is possible, meaning that the flux density limit of the quoted rate would rise to 8 Jy for ghosted sources. For the purposes of comparing rates, in the remainder of the chapter I use the ideal, non-ghosted flux density limit of 7.9 Jy.

3.7.1 Comparison to Other Transient Surveys

I primarily consider how the results relate to other low-frequency surveys (≤ 330 MHz). Also, the results are only compared against transient surveys. Figure 3.16 shows the results in which included are all the surveys that are summarised in Table 3.6. In general, the results are consistent with the previous low-frequency surveys. I was able to improve upon the sensitivity of the Lazio et al. (2010) surface density by at least two orders of magnitude, with the data points from this work providing some of the most extensive searches thus far at low frequencies.

It is also possible to extrapolate our results to gigahertz frequencies, for which extensive reviews of high-frequency surveys have been compiled by Fender & Bell (2011), Ofek et al. (2011) and Fender et al. (2015). It is found that the NCP limits become competitive with previous surveys for transient populations with a spectral index of -1 , and probe deeper than previous surveys if the spectral index is steepened to -2 .




Fig. 3.16 The logarithm of the surface density (deg^{-2}) against the logarithm of the flux density (Jy) of low-frequency transient surveys. Variable source limits are not included. The surface densities for which transients have been detected are marked with a T . The surveys included are as follows: Hyman et al. (2005, 2006, 2009) (Hym2009); Lazio et al. (2010) (Laz2010); Jaeger et al. (2012) (Jae2012); Bell et al. (2014) (Bel2014); Obenberger et al. (2014a) (Obe2014); Carbone et al. (2016) (Car2016) and Cendes et al. (2015) (Cen2015).

3.7.2 Comparison of Time-dependent Surface Densities

It is important to realise that when comparing the transient surface density and flux densities of different surveys, the time-scale at which the survey was performed is just as important. A survey only looking at month-scale epochs will not be sensitive to minute or sub-minute-scale transients. The converse is also true depending on the survey length and sensitivity. Defining the sensitive time-scale of a transient survey is a complex task, with the epoch time (i.e. the time-scale of an individual epoch) and cadence of the epochs usually incorporating a range of values. Generally, it is assumed that the integration time of each observation is the time-scale of a transient on which the survey is most sensitive to. If the integration time matches the duration of the transient event, then the signal-to-noise will, in the majority of cases, be maximised. However as surveys are designed differently, the integration time is not always the equivalent ‘epoch time’. Some epochs are created by averaging many different




Fig. 3.17 *Left panel*: A 3D plot presenting the 3D low-frequency transient search phase space when considering the time-scale of the search. Here, the logarithm of the surface density (deg^{-2}) is plotted against the logarithms of the flux density (Jy) and survey time-scale, with the latter on the z-axis. The viewing angle has been selected to primarily show the flux density and time-scale comparison. Variable source limits are not considered in this plot. The 11-minute limit from this work, as well as the limits of Hyman et al. (2009), Jaeger et al. (2012) and Obenberger et al. (2014a), have been derived from the detection of one or more transients. The others are limits placed with no detections. *Right panel*: The same plot as the left panel, but now the flux density axis has been collapsed to clearly show the surface density and time-scale comparison. The surveys included are as follows: Hyman et al. (2005, 2006, 2009) (Hym2009); Lazio et al. (2010) (Laz2010); Jaeger et al. (2012) (Jae2012); Bell et al. (2014) (Bel2014); Obenberger et al. (2014a) (Obe2014); Carbone et al. (2016) (Car2015) and Cendes et al. (2015) (Cen2015).

observations together for example, or other methods such as creating mosaic fields. This can be especially true in gigahertz surveys. However, the low-frequency searches presented here as a comparison mostly do have an equal integration and epoch time. This is likely due to the large FoV of some of the facilities, minimising the need to use multiple observations to cover a large fraction of the sky.

The cadence of the observations can also be just as valid as a defining characteristic. Taking this work as an example, the NCP search is also sensitive to slow transients that could evolve on time-scales of days, up to the maximum time between two observations of four months. Thus, to avoid the complex visual that would be required to represent all this information, I compare the surveys based upon their respective epoch times (see Carbone

et al. 2016 for the alternate cadence comparison). The left panel of Figure 3.17 shows the same plot as in Figure 3.16; however, a z-axis of time-scales has now been included to display which areas of the time-scale parameter space have been surveyed. The surveys included are the same which have previously been used and are summarised in Table 3.6. The right panel of Figure 3.17 presents the same information but with the flux density axis collapsed, to clearly show the surface density against time-scale comparison.

What is seen from this comparison is a clear definition of surface densities at the time-scales of minutes to hours within a range of sensitivities from the millijansky level to tens of jansky, which does improve upon the Lazio et al. (2010) surface density. It also becomes apparent that there is a region of sensitivities – the jansky–millijansky regime – that is yet to be explored at all time-scales. Jaeger et al. (2012) is currently the only survey to have probed to a ~ 1 mJy depth at low frequencies, and with a detected transient, this perhaps hints at the potential of further discoveries at these depths. Improving the sensitivity at shorter time-scales is also an area that could prove fruitful in transient searches.

3.8 Conclusions

In this chapter, I have presented the results of a search for transient or variable sources at 60 MHz using LOFAR. The search was centred at the NCP, covering 175 deg^2 of sky with a bandwidth of 195 kHz and conducted over the period 2011 December–2012 April. The search for transients and variables was performed using the automated Transients Pipeline (TraP). No transient or variable sources were discovered at time-scales of 30 seconds, 2 minutes, 55 minutes and 297 minutes. However, several candidates were discovered at the 11-minute time-scale. After extensive testing to check if these objects were due to calibration or imaging errors, one of these candidates is considered to be a real astrophysical event, based on the available data. The transient, ILT J225347+862146, was seen only in one 11 min epoch out of 1897, implying a transient rate of $3.9^{+14.7}_{-3.7} \times 10^{-4} \text{ day}^{-1} \text{ deg}^{-2}$. While complicated by the

processing strategy, the flux density of the event is believed to be in the range of 15–25 Jy and was most active for an estimated time of four minutes. However, the rise or decay time-scale of the event is not sufficiently well defined such that the full duration of activity could extend to a 10-minute time-scale.

At present, the astrophysical origin of ILT J225347+862146 is yet to be determined. There are no recorded objects at the transient position in previous radio or high-energy catalogues. Optical follow-up observations were performed at the transient position, with 20 objects detected within the 1-arcminute-radius error circle. None of these optical sources showed any short or long-term variability and no immediately obvious counterpart was identified in the field. However, the discovery of this transient two years after it was active diminishes the effectiveness of follow-up observations, highlighting the need for real-time transient searches and follow-up.

The possibility of the transient being a flare star event was considered, due to the likely close proximity of the object. However, the time-scale of the burst is an order of magnitude longer than what would be expected from previous observations of flare stars at low frequencies. I also considered whether the event could be an FRB. While the duration of the event is consistent with a scattered burst at 1.3 GHz extrapolated to 60 MHz, it was considered unlikely due to the transient exhibiting a much larger fluence than would be expected from previously seen FRBs, which would require a very steep spectral index ($\alpha < -4.7$) to be plausible.

With LOFAR and other instruments now fully operational, the low-frequency transient sky is being probed to unprecedented depths. If the discovered transient presented in this chapter is a member of a real population, then there is no question that more will be found in future transient surveys. This is especially true with ever improving calibration techniques and more accurate sky models at these low frequencies, combined with these surveys taking advantage of the full capabilities of these new, current generation telescopes.

Chapter 4

The Optical Counterparts of Radio Transients and Variables

*The work presented in this chapter has been submitted to the MNRAS journal as
Stewart, A. J., Muñoz-Darias, T. & Fender, R. P.*

4.1 Introduction

The next generation of radio telescopes will see the radio sky surveyed to an unprecedented level, thanks to the expected deeper sensitivities coupled with the ability to swiftly cover large swathes of sky. This has the potential to uncover a wide variety of radio transient and variable phenomena, from both expected, and unexpected, origins. Presently, large area surveys dedicated to searching for such radio sources are becoming ever more commonplace, such as those discussed in Sections 1.6 and 3.1, with some utilising the first wave of new instruments such as the upgraded VLA and LOFAR. These will shortly be joined by the group of new centimetre wavelength instruments previously discussed such as MeerKAT and ASKAP, before the SKA in the 2020s that has the potential to uncover thousands of transients (Metzger et al., 2015; Fender et al., 2015). At optical wavelengths, surveys such

as the Palomar Transient Factory (PTF; Rahmer et al., 2008; Rau et al., 2009; Law et al., 2009) and those using the Pan-STARRS1 telescope (Tonry et al., 2012) have produced many transient sources (e.g. Drout et al., 2014; Cenko et al., 2015). The future Large Synoptic Survey Telescope (LSST; Ivezić et al., 2008) will view the optical sky to an unprecedented level of detail. The quality combined with a fast survey speed and a wide field of view could result in an estimated 10 million events per night (Bloom et al., 2012; Kantor, 2014).

In this chapter we investigate the relationship between the photometric properties of different classes of radio transient and variable objects and their optical counterparts. This is with the eventual aim of deducing whether such data could be used as a rapid classification diagnostic. We primarily focus on both galactic and extragalactic slow radio transient and variable objects that are found in the image plane, such as flare stars and GRBs. However, we do also consider radio pulsars as an example of a fast radio transient which could also be detected in the image plane. As shown in Chapter 1 a handful of radio transient events have been discovered through blind surveys, however, such discoveries in the radio make identifying the class of transient object a challenging task. Although information can be gained from the nature of the radio emission itself, for example the time-scale of the event (Pietka et al., 2015), the use of simultaneous or rapid multi-wavelength follow up data becomes crucial, if not essential, in order to establish the type of the transient source in question. This was demonstrated by Mooley et al. (2016) who used various multi-wavelength follow ups, primarily optical in cohesion with PTF, to identify transient and variable sources detected in a 50 deg^2 pilot survey at 3 GHz with the VLA.

In the SKA era, it is conceivable that the number of radio transient events discovered during surveys will be high enough that detailed individual follow up of every event will not be possible. Rapid classification techniques are widely utilised in optical surveys, where there has been a high yield of transient objects over the past decade (Saglia et al., 2012; Djorgovski et al., 2012). These commonly involve machine learning algorithms based upon

light curve features (Richards et al., 2011) and other measurements such as optical colour (Mahabal et al., 2011). In the radio, attempts have been made to classify objects based upon light curve features and the previously mentioned time-scale of the event (Rebbapragada et al., 2012, Pietka et al., in preparation). However, the need to rapidly classify objects has not been a requirement for radio transient astronomy up to the present day, with efforts instead focussed on the rapid detection of events. Automatic classification has also been investigated at X-ray wavelengths, e.g. Lo et al. (2014) used random forest algorithms to define a classification of X-ray variable sources based upon features derived from time series, spectra and multi-wavelength catalogue data (also see Farrell et al., 2015).

Here, we assume a scenario where rapid, or simultaneous, optical photometry to a discovered radio transient is easily or readily available. Of course the more multi-wavelength data available, such as infra-red or X-ray, the better any classification attempts will be. However, we decided to concentrate on the radio–optical relationship due to the high probability that optical data will be the most abundant. This is due to the growing number of robotic optical telescopes such as the Liverpool Telescope (Steele et al., 2004) in the northern hemisphere, SMARTS¹ in the south, and projects such as the Las Cumbres Observatory Global Telescope Network (LCOGT; Brown et al., 2013). This is in addition to present (e.g. the Sloan Digital Sky Survey (SDSS; Eisenstein et al., 2011)) and future (e.g. LSST) optical catalogues, which also deal with the variable and transient population. Thus, we only consider the two pieces of information that is likely to be first received - the raw radio and optical flux measurements of the object in question. This is where we initiate the investigation, but the optical data also has a wealth of other useful information that can be extracted, as is discussed in the later sections. We also note that only object populations which appear as optical point-like sources are considered, since these are more difficult to classify, and account for the vast majority of the variable and transient populations.

¹SMARTS is the Small and Medium Aperture Research Telescope System; see <http://www.astro.yale.edu/smarts/>.

A project that is very applicable to the above scenario is the MeerLICHT telescope². It is a robotically operated 65 cm telescope with a 2 deg² field of view that will follow MeerKAT observations and provide simultaneous optical coverage. Hence, this presents the ideal situation where any radio transient newly discovered with MeerKAT will potentially have instant optical photometry.

4.2 Data Collection & Methods

The first step of the investigation was to collect radio and optical flux measurements of various types of transient objects from the literature. I define these as F_r and F_o respectively throughout from here on. The classes of object for which these data were gathered were as follows: quasar, stellar, radio pulsar, CV, XRB, GRB and SN.

The large range of objects we wished to collect meant that data would be collected from a range of projects in the literature, all with different methods, instruments and data quality. Hence, in order to obtain a sample of sources which was adequate, we would obtain any F_o or F_r measurements which were available. This meant for example, not being limited to one optical band or radio frequency - it would have been implausible to do so. In the optical case we concentrated on gathering B, V, R, I band measurements, as these were likely to be the most common. Of these bands a preference order of R, V, I and B was used. In cases where only an SDSS *ugriz* measurement was available, we only recorded the *r* value and assigned it as R . While of course there is a difference between the R and r bands, we assume that it is negligible compared to the scatter created by comparing fluxes across the $BVRI$ spectrum. The exception for this was the quasars, where enough information was obtained to perform a conversion - this is detailed in the relative quasar section below. We also do not account for possible extinction on optical magnitudes. In the radio we only used fluxes in the frequency

²MeerLICHT is a prototype telescope for the BlackGEM telescope array; see <http://www.ast.uct.ac.za/meerlicht/MeerLICHT.html>

range of 1–10 GHz, as these were commonly the most abundant measurements. We also only consider detections of objects at both wavelengths and include no upper limit values. In cases where no tabled observational data was available in the literature, measurements were obtained from light curve figures using WEBPLOTDIGITIZER³.

As all the populations can experience rapid flux variations, the timing between the F_r and F_o measurements was also important. Ideally we obtained measurements that were ≤ 1 d apart in observation times, and this was achieved for a significant portion of the sample. However, this was not always possible - especially in classes where the number of sources were low. Hence, while some measurements have a large time difference, we believe that they are still a useful addition to the sample. The time differences between the measurements for each class is detailed in the sections below. There are however, two exceptions to this rule: quasars and stellar sources. While simultaneous measurements would be ideal, it was implausible to obtain such a dataset for these classes that contain a high number of sources. Instead we used catalogue cross-matching to obtain the F_r and F_o measurements, and by doing so, we hope to be independent of the ‘state’ of the quasar or stellar source i.e. we assume that we obtain flux ratios at all different states, providing the entire range of the parameter space occupied by these two classes.

The following sections give specific requirements for the respective class of object, with Table 4.1 showing the final total number of measurements gathered for each class and Figure 4.1 shows the distributions of the frequencies and optical bands used for each class of object. Where applicable, the references for the individual measurements can be found at the end of this chapter.

³Web based tool to extract data from plots, images and maps, authored by Ankit Rohatgi. (<http://arohatgi.info/WebPlotDigitizer>)

Table 4.1 The total number of pairs of F_r and F_o measurements included in the sample for each class of object. Also shown is the total number of optical colour measurements which were obtained where available as an additional reference for our sample. The potential use of colour is discussed in Section 4.4.

Object Class	Total No.	Total No. with colour
Quasar (Optical Sel.)	11,063	11,063
Quasar (Radio Sel.)	788	524
Stellar	685	532
Radio pulsar	7	1
CV	9	2
XRB	6	1
GRB	49	12
SN	26	20
Total	12,633	12,155




Fig. 4.1 *Left*: The distribution of radio frequencies used for the F_r measurements collected for each class of object in the sample. *Right*: The distribution of optical bands used for the F_o measurements collected for each class of object in the sample. The legend in the left plot also applies to the right.

4.2.1 Quasars

In order to gain a firm understanding of the distribution of quasars, we used two different selection methods to populate our sample. First, we selected quasars which had been classified as such through their optical properties, and second, a sample of quasars which had been defined by their radio properties. The two samples are detailed below.

Optically Selected

Two versions of the SDSS quasar catalogue were used: the Seventh Data Release quasar catalogue (DR7QC; Schneider et al., 2010) and the Ninth Data Release quasar catalogue (DR9QC; Pâris et al., 2012). The DR7QC consisted of data from the SDSS-I and SDSS-II surveys (York et al., 2000) and contains 105,783 entries, where as the DR9QC was compiled using SDSS-III (Eisenstein et al., 2011) and contains 87,822 quasars. The two catalogues were then cross-matched with the VLA FIRST 1.4 GHz survey catalogue (White et al., 1997) to obtain radio fluxes along with the SDSS optical magnitudes. If any quasars were present in both the SDSS catalogues, then only the information from the DR9QC was used. In total, 11,063 optically selected quasars are included in our sample. It should be noted that some of the quasar classifications in the SDSS catalogues are due to radio detections - hence they are radio selected. While this will cause some overlap with the radio selected sample below, it otherwise does not affect the use of the population.

Combining the two SDSS quasar catalogues provided us with a diverse range of objects for our sample, particularly in redshift, z , as this was the main difference between the two. The DR9QC was built using data from the Baryon Oscillation Spectroscopic Survey of SDSS-III (BOSS; Dawson et al., 2013), which targeted over 150,000 $z > 2.15$ quasars. This resulted in DR9QC containing 2.6 times more quasars at $z > 2.15$ than the DR7QC, for which the majority of quasars were at $z < 2.15$ (see Figure 22 in Pâris et al. 2012).

The magnitudes obtained from the SDSS surveys were converted from the SDSS *ugriz* filters to $UBVR_cI_c$ using the formulae defined by Jester et al. (2005), specifically for quasars at $z \leq 2.1$ (synthetic). We note that this may not be accurate for the sample found in the DR9QC because of the reasons discussed above, however we deemed the conversion accurate enough for our purpose.

Radio Selected

The basis of our radio selected sample was the Parkes Catalogue (Wright & Otrupcek, 1990), where we selected any objects identified as a quasar or a BL Lacertae. For each object, the radio flux at 5 GHz was obtained (1.4 GHz data were poorly sampled), and then, by using SIMBAD, we gathered *BVRI* photometry measurements where available. In total, 788 radio selected quasars were included in the sample.

4.2.2 Stellar Objects

To characterise the galactic stellar sources, we have used F_o and F_r properties as used by Güdel (2002) (in particular the information presented in their Figure 1). The radio measurements of this catalogue are in turn based upon stellar radio detections between 1–10 GHz as reported by Wendker (1995). F_o measurements in the *B* and *V* bands are also obtained. Contained in this catalogue are the stellar types: RS CVn and Algol-type stars; double stars; magnetic stars; BY Dra-type stars; Symbiotic stars; Herbig stars; WR stars; T Tauri stars; PMS stars; and isolated stars. In total, 706 stellar sources are included. Much like the quasar sample, using a catalogue approach meant that we can sample all stellar objects in a variety of states.

4.2.3 Isolated Neutron Stars

Radio Pulsars

We have included seven radio pulsars which are listed in Table 4.4. As simultaneous observations are rare, we have used optical observations of the pulsars performed by Mignani (2011), paired with radio observations that were compiled in the ATNF Pulsar Catalogue (Manchester et al., 2005). The optical magnitudes are bands *V* or *R*, with the radio measurements all at 1.4 GHz.

Other Isolated Neutron Stars

We also attempted to gather measurements for other types of isolated neutron stars such as magnetars and RRATs, e.g., as of writing, the McGill Online Magnetar Catalogue (Olausen & Kaspi, 2014)⁴ lists four magnetars with radio observations. However, these could not be matched with optical observations as either the optical or radio flux measurements were upper limits, hence we decided not to include them at this stage. The same is true with RRATs where no optical detections have been made (Dhillon et al., 2006, 2011).

4.2.4 Cataclysmic Variables

Five CVs have been included in this work. Two of these are SS Cyg and V3885 Sgr, where the former has two data points representing quiescence and outburst. These measurements are in the *V* and *R* bands and at 5.5–8.6 GHz. The SS Cyg radio and optical data points were within 1 d of each other in time, however the V3885 Sgr measurements are not simultaneous. Reference details of can be found in Table 4.5. Also included are three CVs observed by Coppejans et al. (2015), who detected RW Sex, V603 Aql and TT Ari for the first time in the radio at 6 GHz. The authors recorded two separate measurements for each of these three

⁴<http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>

CVs, and all of them are included in our sample. The optical V band magnitudes of each source, closest to the date of the radio observations, were also obtained from AAVSO⁵ light curves. The average separation in time between the radio and optical measurements was 1 yr, 1 d and 22 d for RW Sex, V603 Aql and TT Ar, respectively.

4.2.5 X-ray Binaries

The X-ray binary sample includes five XRB objects: Scorpius X-1, Cygnus X-2 (V1341 Cyg), SS 433, V404 Cygni and GX 339-4. Table 4.6 gives reference details for these measurements. All the radio and optical measurements are within 1 d of each other, except for Cygnus X-2 where matching fluxes could not be found and thus are ~ 15 yr apart. The optical measurements are bands B and R , with the radio spanning 2.7–9 GHz. Two measurements are included in the sample for Scorpius X-1 as these data covered a radio flare, hence data points at the peak and the quiescent level are recorded. The sources V404 Cygni and GX 339-4 are part of the ‘dynamic data’ sample which is discussed in Section 4.2.7. In these cases the average flux measurements are taken as the single measurement values of the sample.

4.2.6 Supernovae & Gamma-ray Bursts

As these objects have a definite beginning, our approach was to use the first simultaneous (≤ 1 d) measurement available in literature. Note that this is not necessarily the first measurement available at either wavelength. This meant that all the F_r and F_o measurements included in the SNe and GRB sample are ≤ 1 d apart in observation time.

Supernovae

We have used a sample of radio detected core-collapse SNe compiled by Romero-Cañizales et al. (2014). The sample includes a variety of SNe types: Ib, Ic, II, IIn, IIb, IIP and a full

⁵American Association of Variable Star Observers; see <https://www.aavso.org>

list of the 26 SNe can be found in Table 4.7 along with references. The measurements range from 3–743 d after the reported explosion dates for the SNe.

Gamma-ray Bursts

Our GRB sample is based upon a catalogue of radio GRB afterglows that was compiled by Chandra & Frail (2012). We have attempted to gather simultaneous optical and radio measurements for each of the listed GRB events. In total we have obtained recorded measurements for 49 GRBs, with full details presented in Table 4.8. The measurements range from < 1 –26 d post burst.

4.2.7 Dynamic Data

We also wanted to investigate how different classes evolve over time in both radio and optical bands. In order to do this we have gathered simultaneous optical and radio flux measurements of three SNe: SN 1993J, SN 1994I and SN 1998bw; two GRBs: GRB 030329 and GRB 100418A; two XRB flare events: GX 339-4 and V404 Cyg; and two AGN: BL Lacertae and 3C 454.3. In the case of the XRBs, each sample covers a period of time where the source is decaying from a flare event. The coverage of these events all span a minimum of one month, with the full details of the dynamic source sample presented in Table 4.2.

For these sources we have attempted to be as consistent as possible in terms of radio frequency and optical bands used, where the coverage allows. Within the coverage of each source, the entire light curve is in one radio or optical band only. We note that such long coverage of these objects usually signifies a unique characteristic of the source in question. For example, SN 1993J was one of the brightest radio SN to appear in recent history; and the same can be said for GRB 030329 which was one of the brightest radio afterglows ever recorded. Thus, while these sources offer a good example of how classes are likely to evolve

Table 4.2 The information regarding the sample of dynamic sources we have used in the investigation. ‘No. of fluxes’ refers to the number of simultaneous pairs of F_r and F_o that were obtained. The references are as follows: 1. Weiler et al. (2007), 2. Richmond et al. (1996b), 3. Weiler et al. (2011), 4. Richmond et al. (1996a), 5. Yokoo et al. (1994), 6. Lee et al. (1995), 7. Kulkarni et al. (1998), 8. Clocchiatti et al. (2011), 9. Berger et al. (2003a), 10. Lipkin et al. (2004), 11. Moin et al. (2013), 12. Marshall et al. (2011), 13. Han & Hjellming (1992), 14. Casares et al. (1991), 15. Corbel et al. (2013), 16. Buxton et al. (2012), 17. WEBT Archive, Villata et al. (2009) and 18. GASP-WEBT archive, Raiteri et al. (2011). WEBT is the Whole Earth Blazar Telescope and GASP the GLAST-AGILE Support Program.

Object name	Object class	Date of first fluxes	Date of last fluxes	No. of fluxes	Radio freq. (GHz)	Optical band	Radio ref.	Optical ref.
SN 1993J	SN, Type II	1993 Apr 08	1994 Feb 08	38	8.3	R	1	2
SN 1994I	SN, Type Ic	1994 Apr 03	1994 Aug 08	26	8.3	R	3	4,5,6
SN 1998bw	SN, Type Ic	1998 Apr 28	1998 July 15	22	8.64	R	7	8
GRB 030329	GRB	2003 Mar 30	2003 June 04	20	8.64	R	9	10
GRB 100418A	GRB	2010 Apr 20	2010 May 21	7	8.46	white filter	11	12
V404 Cygni	XRB	1989 May 30	1989 July 04	7	8.3	R	13	14
GX 339-4	XRB	2011 Feb 06	2011 Apr 27	15	9.0	V	15	16
BL Lacertae	AGN	1994 June 28	2005 Jan 12	265	5.0	R	17	17
3C 454.3	AGN	2008 May 08	2010 Jan 28	79	8.0	R	18	18

on our tool, a few caveats may be associated with the sample.

4.3 Results

Using the sample collected as described in Section 4.2, we attempted various methods to visualise the collective data to find a scheme which offered the best representation of the parameter space. This included using parameters such as the ratio of fluxes, i.e. F_r/F_o , and optical colours. Our conclusion was that plotting F_o against F_r , in log space, offered the most useful diagram in terms of being able to distinguish the different classes. Further reasoning behind this choice is discussed in Section 4.5. The result is presented in Figure 4.2.

4.3.1 Description of the Base Result

The plot shown in Figure 4.2 provides an overview of where each class of object, given our dataset, typically lies in the F_r - F_o parameter space. In terms of producing the diagram, we have converted the optical magnitudes to millijansky units in order to make a direct comparison to the radio flux measurements. This was performed using zero flux points of 4130, 3781, 2941 and 2635 Jy for bands B , V , R and I respectively (Cox, 2000). It is also important to note that while the optical R band magnitude scale is provided as a reference, the magnitudes used consist of a mixture of optical bands. The same is also true for the radio fluxes, where a range of radio frequencies (1–10 GHz) are utilised in order to obtain radio flux measurements.

A series of lines are also displayed on the plot. The horizontal line highlights the flux density limit of the FIRST survey (1 mJy) - used to obtain flux measurements for the optically selected quasar population. The two vertical lines denote the SDSS saturation limit (14 mag, R band) and the SDSS 95 per cent completeness limit (22 mag, R band). These three limits greatly impact the distribution of quasars and are discussed in the Section 4.4. Next are the




Fig. 4.2 F_r versus F_o of our sample of objects from different classes. The F_o values have been converted from mag to mJy, with the optical magnitude scale given as a reference on the top axis (based upon R band). The dataset contains F_o in bands B, V, R and I and the F_r covers a range of frequencies from 1–10 GHz. Labelled on the plot are: a horizontal line denoting the FIRST survey flux density sensitivity limit (1 mJy); a vertical line denoting the SDSS 95 per cent completeness limit (~ 22 mag), and another showing the SDSS saturation limit (14 mag). The diagonal lines in the background of the plot show the lines of constant ratio between the two fluxes corresponding to changing distance. A 5 mag scale arrow is also shown to show the possible effect of extinction on the optical magnitudes.




Fig. 4.3 A simplistic schematic version of the base result where coloured ellipses represent the space where the majority of different classes of objects occupy in the plot.

background diagonal dashed-lines, which represent the slope of a constant F_r and F_0 ratio. This means that if a particular source was moved closer to Earth, or indeed further away, it would move on the plot along this gradient. Finally, the arrow in the lower right-hand section of the plot demonstrates the movement an object could undergo with a reddening equal to 5 optical magnitudes. Figure 4.3 presents a schematic version of the plot in Figure 4.2, created by placing ellipses centred on the median points of the respective populations dataset and then manually sized to enclose the majority of the populations, which can be seen underneath the ellipses.

From Figure 4.2, we see that there is a relatively good separation between galactic and extragalactic populations, given our current sample. On the right half of the plot resides the stellar population and objects such as CVs, while populations such as quasars and GRBs typically reside on the left hand side. While this is not surprising, it demonstrates that useful information about a transient object can be obtained by referring to these simple properties. Looking closer at the plot, we see that SNe seem to mostly be placed between these two

galactic and extragalactic populations, along with XRBs. However, our XRB sample is very small (five), hence we cannot conclude much from their typical positioning. Found next to these objects, but fainter in optical flux are the GRBs which are commonly below the 1 mJy quasar limit set by the FIRST survey sensitivity. This population is also relatively separated from the rest of the classes, however, this is likely to change once quasars below the 1 mJy limit are considered (we revisit this in Section 4.4.1). As for the quasars, we see no surprise between the optically and radio selected samples, with the radio sample occupying the radio-bright section of the optical sample. Hence, from here on we group the two samples under one class of ‘quasars’. Finally we have the majority of the radio pulsar population, which are faint in the optical and are positioned to the far left-hand side of the plot. Radio pulsars appear to be quite uncontaminated by other classes in our current sample, given their faint optical fluxes.

4.3.2 Dynamic Characteristics

As previously discussed in Section 4.2.7, sources will not stay at a fixed location on the plot and will instead move around over time, which is an important aspect to understand. The paths which are traced out by the dynamic data sample (see Table 4.2), for which the light curves can be seen Figure 4.4, are shown in Figure 4.5 overlaid on a ‘greyed-out’ version of Figure 4.2. Most of these objects cover two to three months of observations, the exceptions being SN 1993J which spans 10 months, and the two AGN sources BL Lacertae and 3C 454.3 10 and 2 yr, respectively. The optical and radio light curves of the sources can be seen in Figure 4.4, where we note that ‘day 0’ in respect to the light curves represents the first measurement we consider in our dynamic sample.

We see that two of the SNe, 1993J (type II) and 1994I (Ic) follow a very similar trend. They both rise steadily in the radio band, whereas this is accompanied by only a slight increase, and then decrease, in the optical band. The reverse scenario then becomes true

for the later part of the tracks where the F_r changes slower than the now decreasing F_o . An exception to this trend is SN 1998bw (Ic), as it exhibits an initial increase in both F_r and F_o , followed by a decrease in radio but continuing to slightly brighten in optical, after which the source starts to decrease in both parameters. SN 1998bw does indeed have a rapid rise to its peak radio F_r (~ 10 d), compared to the other type Ic SN we have sampled, SN 1994I (~ 20 d). This gives rise to this mirrored trend of SN 1998bw compared to the other two SNe because of the radio peak being reached much before the peak in F_o .

Much like the two SNe in the early 90s, GRBs 030329 and 100418A also follow similar trends, as expected for the GRB model of a decrease in F_r followed by a re-brightening, while steadily decreasing in F_o . It should be noted that GRB 100418A only has one month worth of coverage compared to the two months coverage of GRB 030329, hence we do not see the full evolution of the former. Also, GRB 030329 is a bright nearby afterglow ($z = 0.1685$, 587 Mpc; Greiner et al., 2003), hence its position on the plot is boosted in the radio such that it transits the quasar area of the plot. If the GRB was moved further away, e.g. the same as GRB 100418A ($z = 0.624$, 1.4 Gpc; Antonelli et al., 2010), the 10 mJy flux level becomes ~ 1 mJy - consistent with the other GRBs in the sample. In general, the SNe and GRBs follow a simple model in which the optical flare evolves more rapidly than the radio. Hence these, and any other objects that follow a similar model, will display anti-clockwise loops on the plot.

The two XRB events however, do not completely follow the same trend. While the F_o is generally characterised by a decay followed by a small re-brightening before returning to a decay, V404 Cyg exhibits a much sharper decline in magnitude. The radio on the other hand is quite different between the sources at the early stage, with V404 Cyg experiencing a steady decline while GX 339-4 has a small re-brightening phase, before both settling into a steady decay. Thus, given these characteristics of the events it is not completely surprising that common features are hard to extract.




Fig. 4.4 The light curves of the dynamic sources that are used in the sample. These cover four different classes of objects for which each class is represented by two plots, the upper panel showing the optical data and the lower panel displaying the radio data. The plots labelled *a* show the supernova events SN 1993J, SN 1994I and SN 1998bw; those labelled *b* contain the GRB events GRB 030329 and GRB 100418A; labelled *c* are the XRB flare events from the sources V404 Cyg and GX 339-4; and lastly labelled *d* are the WEBT AGN monitoring light curves of BL Lacertae and 3C 454.3.




Fig. 4.5 Plotted are the ‘tracks’ which the dynamic sources follow as the astrophysical events evolve over time. The data points representing the beginning and end of each track are labelled with the words ‘Start’ and ‘End’, respectively. This is excluding the tracks for BL Lacertae and 3C 454.3 due to the start and end being at the centre of the respective data points. These are plotted over a ‘greyed-out’ version of the basis diagram such that the positions of the dynamic sources, in relation to the base plot, can be observed. The SN events sampled seem to follow similar paths, as do the GRBs, though it is difficult to see a common area, or trajectory, in the XRBs given the current sample. The AGN sources are relatively stable at their flux level, with BL Lacertae also stable over a 3 month period, shown in dark blue, which enables a better comparison in time to the other objects plotted.

Lastly, the two radio bright AGN sources change very little over time in F_r . This is partly due to the 1–10 Jy brightness of the objects on the log scale, but it also shows that no unique trajectory is preferred, i.e. the flux proceeds to vary back and forth instead of evolving like the other objects. The same is true for F_o , where while variation is seen, it does not follow a clear track. In both cases it is clear that the AGN objects do not show large relative flux variation events like those seen in other classes. As BL Lacertae covers a time scale of 10 yr, also plotted is a 3 month portion of the dataset to allow a better comparison to the other classes of objects. No differing pattern is seen within this time-scale. Simultaneous data for low-luminosity AGN to compare against could not be found.

4.4 Discussion

In the previous section we have presented a relation between the F_r and F_o measurements for a wide range of transient and variable sources. Here we give more detailed discussion of the obtained results.

First, we note that we have not considered measurement uncertainties when compiling our results. We believe that because of the nature of what we are investigating, the intrinsic spread of each population in the F_r - F_o parameter space has a much greater influence than that of measurement error. We also take a similar approach when considering the bias of the sample. These data used in the investigation mean that it is an extremely biased sample, and at this initial stage of the project, we do not attempt to un-bias the dataset due to the complexity of the task. Lastly, we have not corrected our optical sample for extinction. If this was accounted for, this would cause sources to shift left along the x-axis only, as denoted by the black arrow on Figure 4.2. This is a conscious decision in that we do not wish to ‘de-redden’ our sample and rather want to probe measurements that are directly observed, as this could have potential implications on future uses for these data as discussed in Section 4.5.

As briefly discussed in Section 4.3, various different parameters were explored when plotting the sample dataset in order to evaluate which value would provide the most information, yet was also simple in concept. One such parameter that was considered was the ratio of F_r/F_o and plotting this against F_r . As can be interpreted from our final choice, this ratio value also allowed us to somewhat separate classes, especially in the case of galactic or extragalactic objects. This can be seen in Figure 4.6 which shows the distribution of F_r/F_o for each class of object. However, this representation caused the spread of each class to increase and proved more difficult to distinguish. Hence, we favoured plotting F_r and F_o as a much clearer view of the parameter space.

Another parameter that was tested was optical colour. As shown in Table 4.1, a significant fraction of our sample contained optical colour information and we investigated whether it could provide an even better method of dividing the classes. Somewhat surprisingly, we found that plotting optical colour versus F_r , or indeed versus the ratio of F_r/F_o , greatly hindered our ability to separate different classes of objects. This was primarily caused by the narrow range of optical colour values, which in turn limited the parameter space in which the classes could be separated. This was even to the extent where we could no longer distinguish galactic and extragalactic populations.

4.4.1 Extending the Quasar & Stellar Populations

The location of quasars on the plot is currently not a true representation of what would be expected, where the quasar dataset is strictly governed by the sensitivity limits of the surveys we have used in our sample. These are the SDSS, with a 95 per cent completeness limit of ~ 22 mag, and the FIRST survey with a flux density limit of 1 mJy. This can clearly be seen in Figure 4.2, where these limits are represented by labelled dashed-lines. One would expect this population to exceed these limits when observing beyond such constraints, and therefore extend into areas previously only occupied by another class - causing possible confusion. A




Fig. 4.6 The distribution of the ratio F_r/F_0 for each source of the classes included in the sample.

similar notion could be applied to the stellar sample, as this will be another general class of a large size that will continue to be detected as observations probe deeper.

To demonstrate the effect the likely extension of these classes will have, we have done the following. The optically selected SDSS quasars are shifted such that the flux measurements are approximately 10 and 100 times fainter (moved in distance by a factor of $\sqrt{10}$ and 10 respectively). The reason this is only applied to the optically selected quasars is that the redshifts are known for this sample, allowing for the application of K-correction and correcting for bandwidth stretching. ‘K-correction’ refers to the flux correction made when taking into account the effect of redshift on the observed flux, i.e. the observed band corresponds to different parts of the spectrum at different redshifts, and bandwidth stretching refers to effect of the observed bandwidth corresponding to different intrinsic bandwidths at different redshifts. The luminosity with the corrections applied is

$$L = 4\pi F d_L^2 (1+z)^{-\alpha-1}, \quad (4.1)$$

where L is the luminosity, F is the flux, d_L is the luminosity distance, z is the redshift and α is the power-law spectrum index. For the radio flux we use a canonical value of $\alpha_r = -0.7$ and the for the optical a value of $\alpha_o = -0.44$ (Vanden Berk et al., 2001). For the stellar population, our sample includes parallax distance measurements for 533 of the 686 stellar objects plotted, for which the median value is 150 pc. Thus, we use this information to increase the distance of these stellar objects by 1 kpc, demonstrating what this magnitude of distance change does to the distribution. The results of these shifts can be seen in Figure 4.7. Beginning with the quasars, we see that such a shift completely covers the area of the plot that was previously only occupied by GRBs, along with slightly extending into the pulsar region at the far left of the plot. This suggests that at first glance when observing beyond the limits of our current sample, it will prove difficult to differentiate a GRB from a quasar. The stellar sources on the other hand, with a distance boost to 1 kpc, extend down to an F_r of




Fig. 4.7 The effect of probing the quasar and stellar populations to a deeper sensitivity that our sample currently includes. This is overplotted on a ‘greyed-out’ version of the basis plot with the original quasar and stellar sources highlighted by a darker shade. We see that by reducing the measured flux of the quasar sample by approximately one, and then two, orders of magnitude, the population infringes the formerly isolated GRB region of the diagram. These fluxes have been corrected for the ‘K-correction’ and bandwidth stretching. We also scale the flux of a portion of the stellar sample by increasing the distance of the object by 1 kpc, showing that in this case the objects can reach radio flux measurements of 1 μ Jy.

1 μ Jy, with some sources even going beyond this value. The region below the stellar sources on the basis plot was quite void, hence, unlike the quasars, the extension of the stellar sources does not protrude into other classes. On a bigger scale, it is worth noting that the distinction between quasars and stellar sources, i.e. galactic and extragalactic objects, remains relatively clear.

4.5 Transient Diagnostic Application

One of the primary motivations of this project was the potential in using such a multi-wavelength dataset as an early classifier for newly discovered radio transient events. In this section we discuss how the results and discussions presented so far, could be utilised for such a purpose.

The F_r , along with a follow up F_o could be swiftly obtained for a new radio transient event, especially in an era where robotic optical telescopes are becoming common. The need for a swift diagnostic on a transient event becomes even more imperative when considering the high number of transient events that could be found with the SKA (Metzger et al., 2015; Fender et al., 2015), for which follow-up strategies must be implemented and aided by fast classification. Closer than the SKA is the previously discussed MeerLICHT project, which will provide optical coverage for every MeerKAT observation when possible; an ideal scenario where F_r and F_o are obtained automatically for every observation, simultaneously. Therefore, this reasoning was a major factor in our decision to compare F_r with F_o in our final results. This representation provided the desirable combination of being simple to obtain practically and also produced useful separation amongst the transient objects in our sample. This use for these data is also why we did not want to ‘de-redden’ the optical data points; as ideally the sample should mimic data that are produced directly from the telescope with no manipulations.

The base plot described in Section 4.3.1 and seen in Figure 4.2, can separate galactic and extragalactic well. A diagnostic that should not be underestimated, especially when gained swiftly. Concentrating on the individual classes themselves, we refer the reader back to Figure 4.3, which displays schematically how the populations reside. We can see how sources such as GRBs and pulsars are particularly isolated on the base plot, along with the main areas of stellar and quasar sources. Hence, it is plausible to obtain an early estimate of the nature of a transient object from just F_r and F_o .

However, we must be careful to remember the limitations of our sample, particularly the extension of the quasar and stellar populations when deeper observations are performed, shown in Figure 4.7. In Section 4.4.1 we discussed how the quasar population would partially extend into the pulsar region, in addition to covering the GRBs. This shows how the classification of, previously quite isolated, classes can be considerably confused by

considering how the sample will behave beyond our selection limits. The extension of the stellar sources has less impact in terms of classification, but nonetheless is just as important to understand and expect.

The confusion can be countered by considering the dynamic movement of different classes. This was explored and discussed in Section 4.3.2, with examples of how GRBs, SNe, XRBs and AGN move on the diagram in Figure 4.5. It is possible to retain the ability to classify by understanding the unique ways in which the individual classes evolve in the F_r - F_o parameter space. For example, we would expect a GRB to transit across the plot from right to left given our dynamic data sample, and at the same time be located in the GRB area; differentiating the event from a quasar source that is likely to not follow such an evolution. A similar expectation can be set for SNe where the events studied followed a similar trajectory to each other, providing a trait to monitor for. Prior knowledge of the event could also help in such situations, e.g., knowing whether the event in question suddenly appeared, as opposed to whether a source was present previously.

We also expect any transient diagnostic application of the dataset to be used alongside external knowledge of an event - either to decide how to proceed after a classification has been suggested, or to further define a classification that a different method has defined. There are various examples of other techniques, such as diagnostics based on the radio rise and decline times of the event (Pietka et al. 2015; Pietka et al., in preparation); the optical variability time-scale (Kasliwal, 2011); an optical spectrum; and cross-matching with catalogues at other wavelengths other than radio and optical. Even without further data, the optical observation alone may be able to help with problems like quasar confusion by being able to distinguish between a point source and a galaxy. Out of these, the combination with a third parameter, for example X-ray flux, is likely to be the most powerful external association to improve any classification. Ideally, the sample itself could be extended to include a measurement at a different wavelength which is a possibility in future work.

4.5.1 Diagnostic Test

To demonstrate the ability of the dataset to classify an event, we perform a small scale blind test using a catalogue of variable and transient radio sources discovered in the FIRST survey, as compiled by Thyagarajan et al. (2011). We impose the following criteria on the sources we select from the catalogue:

1. sources must have a defined SDSS *i*-band magnitude, which are mostly gathered by the authors from SDSS DR7 catalogue.
2. only sources which are not flagged as being in the vicinity of another source brighter than 500 mJy.

The authors also attempted to classify the candidates by cross-matching with a variety of catalogues and databases, and are also labelled as either transient or variable. For our purposes we select the candidates that have the classifications of:

- **SDSS-*** - an unclassified SDSS counterpart (transient and variable).
- **SDSS-QPO (S) and (P)** - quasars from the SDSS spectroscopic (S; Schneider et al., 2007) and photometric (P; Richards et al., 2009) quasar catalogues (only contain variables).
- **SDSS-GAL** - objects classified as galaxies in SDSS (select transients only).
- **Tycho-Star** - counterpart found in the Tycho star catalogue (Høg et al., 2000) (only contain variables).
- **HIP-Star** - counterpart found in the Hipparcos star catalogue (Perryman et al., 1997) (only contain variables).

We do not consider SDSS-Gal variables due to these sources effectively being already classified, in addition to our previously stated desire to focus on optically point-like transient or variable objects. For the purposes of our test we define our own categories of:

- **Unclassified (Variable)** (SDSS-* variables).

- **SDSS-QSO (Variable)** (SDSS-QPO S,P variables).
- **Star (Variable)** (Tycho/HIP-Stars, variables).
- **SDSS-Gal (Transient)**.
- **Unclassified (Transient)** (SDSS-* transients).

In total we include 198 objects in our testing sample, the by-category totals can be found in Table 4.3. We then take the F_r and F_o values of the 198 transient and variable objects (i.e. the FIRST flux densities and the SDSS- i magnitudes, respectively) and plot them upon the basis diagram. The result is shown in Figure 4.8.

What we immediately see is an agreement with the ‘Star’ and ‘QSO’ objects, in that they lie within the stellar and quasar regions respectively. We then see that the majority of objects lie within the quasar cloud, however we can use the information presented in order to prioritise the objects which are to be studied further. First looking at variables, we see one source which is set apart from the main cloud and lies within the parameter space occupied commonly by pulsars. Second with transients, the ‘Unclassified’ source is tentatively around the GRB area of the plot, hinting at a possible GRB afterglow. While these two objects may of course not be a pulsar or GRB afterglow, it offers an example as to the kind of follow-up strategy one could apply when faced with numerous candidates. The other SDSS-Gal transients are located within the quasar cloud, and given their optical definition, there is a good chance that quasars are indeed the correct classification. While not a rigorous test, the results from this experiment begin to demonstrate the potential usefulness of using these data to obtain an early estimate of the classification of a newly discovered object. We also expect the method to become more sophisticated over time, both in technique and the sample itself. As further measurements are made, these results can be fed back into the sample, improving the accuracy of future classifications. In addition, while we have presented a very simplistic method, a first step in the future would be to translate our findings to probability distributions, once the complex bias of the sample is fully understood. This




Fig. 4.8 The result of plotting the selected radio transients and variables from the FIRST survey (Thyagarajan et al., 2011), sorted by their highest ranked counterpart at other wavelengths deduced by the authors of the catalogue, upon a ‘greyed-out’ version of the basis plot. We see immediately an agreement between likely galactic and extragalactic sources. We also see examples of perhaps which sources could prompt a higher priority of follow-up strategy, such as the unclassified variable to the far left of the plot, an area commonly occupied by neutron stars.

Table 4.3 The total number of each type of object included in the test sample.

Object Type	Total No.
Unclassified (Variable)	89
SDSS-QSO (Variable)	100
Star (Variable)	3
SDSS-Gal (Transient)	5
Unclassified (Transient)	1
Total	198

could eventually lead to incorporating the technique into an automated system, using machine learning tools to create part of an intelligent follow-up scheduler for discovered transients. The upcoming MeerKAT/MeerLICHT project provides an ideal scenario for the sample and technique presented in this chapter to evolve.

4.6 Conclusions

We have presented a study of the optical counterparts to a wide range of radio transients and variables. We collected a sample consisting of 12,633 pairs of F_r and F_o measurements from seven distinct classes of astrophysical objects, including populations such as quasars, GRBs, CVs and SNe. These measurements were recorded by numerous radio and optical facilities, at frequencies ranging from 1–10 GHz and optical filters *BVRI*.

By comparing just the F_r and F_o measurements of each object, we found that this offered a potential method to differentiate the different classes to which they belong. Galactic sources were found to occupy a different region of the parameter space than extragalactic sources, and generally but not perfectly, separate areas occupied by the individual classes could be identified. However, confusion is caused when considering how the sources would appear beyond the limitations of our sample, such as detection limits. For example, the detection of radio fainter quasars than our sample provided, meant clouding the area where GRBs were present. However, we also investigated how a small number of GRBs, SNe, XRBs and AGN evolved in the parameter space over time, observing that tentatively, each class followed

a unique trajectory. A trait which could be used to counter possible confusion caused by multiple classes being present at the same point in the parameter space.

We thus examined the potential of using these data to classify newly discovered radio transient events, finding that at a basic level, a diagnostic based on the measured F_r and F_o could be a useful tool either alone, or combined with other transient classification methods to hone a proposed transient class. In the future, the sample can continually be updated with new discoveries, improving the accuracy, in addition to further studies of the F_r and F_o properties of transients. The latter point allowing for the development of models based upon this parameter space which can feed into a probabilistic classification in an automated system. A possible necessity as we draw closer to an SKA-LSST era.

Table 4.4 The radio pulsars that are included in the sample. The radio measurements were gathered with the use of the ATNF Pulsar Catalogue (Manchester et al., 2005).

Object Name	F_o ref.	F_r ref.
CRAB (B0531+21)	Mignani (2011)	Lorimer et al. (1995)
VELA (B0833-45)	Mignani (2011)	Backer & Fisher (1974)
B0540-69	Mignani (2011)	Johnston et al. (2004)
B1509-58	Mignani (2011)	Hobbs et al. (2004)
B0656+14	Mignani (2011)	Lorimer et al. (1995)
B1133+16	Mignani (2011)	Lorimer et al. (1995)
B0950+08	Mignani (2011)	Lorimer et al. (1995)

Table 4.5 The CVs that are included in the sample. Any object denoted with * signifies that two measurements were included in the sample - at radio outburst and quiescence.

Object Name	F_o ref.	F_r ref.
SS Cyg*	Körding et al. (2008)	Körding et al. (2008)
V3885 Sgr	Skinner et al. (2014)	Körding et al. (2011)
RW Sex	Coppejans et al. (2015)/AAVSO	Coppejans et al. (2015)
V603 Aql	Coppejans et al. (2015)/AAVSO	Coppejans et al. (2015)
TT Ari	Coppejans et al. (2015)/AAVSO	Coppejans et al. (2015)
TT Ari (radio outburst)	Coppejans et al. (2015)/AAVSO	Coppejans et al. (2015)

Table 4.6 The XRB sources that are included in the sample. Some information was gathered with the use of the catalogue of low-mass X-ray binaries in the Galaxy, LMC, and SMC (Liu et al., 2007). Any object denoted with * signifies that two measurements were included in the sample - at radio outburst and quiescence.

Object Name	F_o ref.	F_r ref.
Scorpius X-1*	Canizares et al. (1975)	Canizares et al. (1975)
Cygnus X-2 (V1341 Cyg)	Skrutskie et al. (2006) ^a	Hjellming et al. (1990)
SS 433	Chakrabarti et al. (2005)	Chakrabarti et al. (2005)
V404 Cygni	Casares et al. (1991)	Han & Hjellming (1992)
GX 339-4	Buxton et al. (2012)	Corbel et al. (2013)

^a The Two Micron All Sky Survey (2MASS)

Table 4.7 The SN events that are included in this work. The age column signifies the approximate time elapsed since the accepted explosion date of the respective SN. The sample is based upon the list of radio SNe compiled by Romero-Cañizales et al. (2014).

SN Name	SN Type	F_o ref.	F_r ref.	Age (d)
SN 1980K	Iib	Buta (1982)	Weiler et al. (1986)	37
SN 1984L	Ib	Tsvetkov (1987)	Panagia et al. (1986)	36
SN 1988Z	IIn	Turatto et al. (1993)	Williams et al. (2002)	743
SN 1990B	Ic	Clocchiatti et al. (2001)	van Dyk et al. (1993)	61
SN 1993J	II	Richmond et al. (1996b)	Weiler et al. (2007)	11
SN 1994I	Ic	Richmond et al. (1996a)	Weiler et al. (2011)	5
SN 1996cb	Iib	Qiu et al. (1999)	van Dyk et al. (1996)	9
SN 1998bw	Ic	Clocchiatti et al. (2011)	Kulkarni et al. (1998)	16
SN 1999em	IIP	Leonard et al. (2002)	Pooley et al. (2002)	34
SN 2001ig	Iib	Bembrick et al. (2001)	Ryder et al. (2004)	12
SN 2002ap	Ic	Gal-Yam et al. (2002)	Berger et al. (2002a)	4
SN 2002hh	II	Pozzo et al. (2006)	Stockdale et al. (2002)	17
SN 2003L	Ic	VSNET ^a	Soderberg et al. (2005)	32
SN 2003bg	Iib	Hamuy et al. (2009)	Soderberg et al. (2006c)	10
SN 2004dj	IIP	Schmeer et al. (2004)	Stockdale et al. (2004a)	23
SN 2004dk	Ib	Drout et al. (2011)	Wellons et al. (2012)	8
SN 2004et	II	Misra et al. (2007)	Stockdale et al. (2004b)	14
SN 2004gq	Ib	Drout et al. (2011)	Wellons et al. (2012)	8
SN 2007bg	Ic	Young et al. (2010)	Salas et al. (2013)	19
SN 2007gr	Ic	Hunter et al. (2009)	Soderberg et al. (2010b)	5
SN 2007uy	Ib	Roy et al. (2013)	Roy et al. (2013)	9
SN 2008D	Ib	Mazzali et al. (2008)	van der Horst et al. (2011)	19
SN 2008ax	Iib	Roming et al. (2009)	Roming et al. (2009)	8
SN 2009bb	Ic	Pignata et al. (2011)	Soderberg et al. (2010a)	20
SN 2011hs	Iib	Bufano et al. (2014)	Bufano et al. (2014)	12
SN 2011dh	Iib	Arcavi et al. (2011)	Horesh et al. (2013)	3

^a Variable Star Network (VSNET; Kato et al., 2004)

Table 4.8 GRB data included in this work. The age column signifies the approximate time elapsed since the detection of the burst. The sample is based upon the list of radio GRB-afterglows compiled by Chandra & Frail (2012).

GRB Name	F_o ref.	F_r ref.	Age (d)
GRB 970508	Sokolov et al. (1998)	Frail et al. (2000b)	5
GRB 980329	Palazzi et al. (1998)	Taylor et al. (1998)	< 1
GRB 980519	Gal et al. (1998)	Frail et al. (2000a)	3
GRB 980703	Bloom et al. (1998)	Frail et al. (2003b)	1
GRB 990123	Castro-Tirado et al. (1999)	Kulkarni et al. (1999)	1
GRB 990510	Harrison et al. (1999)	Harrison et al. (1999)	1
GRB 991208	Castro-Tirado et al. (2001)	Galama et al. (2000)	1
GRB 991216	Jensen et al. (1999)	Frail et al. (2000c)	2
GRB 000131	Andersen et al. (2000)	Frail et al. (2003a)	8
GRB 000301C	Veillet (2000)	Berger et al. (2000)	4
GRB 000418	Berger et al. (2001)	Berger et al. (2001)	11
GRB 000911	Price et al. (2002c)	Price et al. (2002c)	4
GRB 000926	Fynbo et al. (2001)	Harrison et al. (2001)	3
GRB 010222	Sagar et al. (2001)	RADW ^a	1
GRB 010921	Price et al. (2002a)	Price et al. (2002a)	26
GRB 011121	Garnavich et al. (2003)	Price et al. (2002b)	4
GRB 020124	Berger et al. (2002b)	Berger et al. (2002b)	2
GRB 020405	Masetti et al. (2003)	Berger et al. (2003b)	1
GRB 020813	Covino et al. (2003)	Frail & Berger (2002)	1
GRB 030115	Levan et al. (2006)	Frail & Berger (2003)	2
GRB 030329	Lipkin et al. (2004)	Berger et al. (2003a)	< 1
GRB 050315	Cobb & Bailyn (2005)	Berger et al. (2005b)	1
GRB 050401	Watson et al. (2006)	Soderberg (2005)	5
GRB 050416A	Holland et al. (2007)	Soderberg et al. (2007)	6
GRB 050525A	Della Valle et al. (2006)	Cameron & Frail (2005) ^b	3
GRB 050603	Grupe et al. (2006)	Cameron (2005)	< 1
GRB 050724	Malesani et al. (2007)	Berger et al. (2005a)	< 1
GRB 050820A	Cenko et al. (2006)	Cenko et al. (2006)	< 1
GRB 051109A	Misra et al. (2005)	Frail et al. (2005)	1
GRB 051221A	Soderberg et al. (2006b)	Soderberg et al. (2006b)	1
GRB 060218	Ferrero et al. (2006)	Soderberg et al. (2006a)	2
GRB 061121	Halpern et al. (2006)	Chandra & Frail (2006)	1
GRB 070125	Updike et al. (2008)	Updike et al. (2008)	2
GRB 071003	Perley et al. (2008)	Perley et al. (2008)	2
GRB 071010B	Kann et al. (2007)	Chandra & Frail (2007a)	3
GRB 071020	Im et al. (2007)	Chandra & Frail (2007b)	2
GRB 080319B	Bloom et al. (2009)	Soderberg et al. (2008)	2

Continued on next page

Table 4.8 – *Continued from previous page*

GRB Name	F_0 ref.	F_r ref.	Age (d)
GRB 080603A	Guidorzi et al. (2011)	Chandra & Frail (2008a)	2
GRB 080810	Page et al. (2009)	Chandra & Frail (2008b)	4
GRB 081203B	Perley & Bloom (2008)	Chandra & Frail (2008c)	5
GRB 090313	Melandri et al. (2010)	Melandri et al. (2010)	6
GRB 090323	Perley (2009a)	Harrison et al. (2009)	5
GRB 090424	Im et al. (2009)	Chandra & Frail (2009a)	2
GRB 090618	Rumyantsev & Pozanenko (2009)	Chandra & Frail (2009b)	1
GRB 090902B	Pandey et al. (2010)	van der Horst et al. (2009)	2
GRB 091020	Perley (2009b)	Frail et al. (2009)	1
GRB 100418A	Moody et al. (2010)	Moin et al. (2013)	2
GRB 100814A	Nardini et al. (2014)	Chandra et al. (2010)	4
GRB 100901A	Rumyantsev et al. (2010)	van der Horst et al. (2010)	3

^a Measurement taken from the ‘Radio Afterglow Data Website’ (RADW;
http://www.aoc.nrao.edu/~dfail/allgrb_table.shtml).

^b Closer measurement in time taken from RADW.

Chapter 5

The Initial Transient Search of the CHILES 1.4 GHz VLA Survey

5.1 Introduction

The ‘COSMOS HI Large Extragalactic Survey’ (CHILES) is a survey dedicated to producing the first ‘HI deep field’ covering a redshift (z) range of $z = 0$ to $z = 0.45$. The field is centred on the extensively observed COSMOS field (Scoville et al., 2007), and is conducted with the VLA in B configuration. Performed concurrently with CHILES is the CHILES Continuum Polarisation (Con Pol) survey: a full polarisation continuum radio survey covering 0.2 deg^2 of the COSMOS field. The survey will achieve a total of approximately 1000 hours on this single field resulting in a final, unprecedented, sensitivity of 500 nJy per 4 arcsec resolving beam at the field centre. This will improve upon the previous radio VLA-COSMOS survey (Schinnerer et al., 2007), that has defined a catalogue of sources with peak flux densities above $60 \mu\text{Jy beam}^{-1}$ at a resolution of 5 arcsec (Schinnerer et al., 2010). Among the primary science goals of CHILES Con Pol, which include investigating the intergalactic magnetic fields and galaxy evolution, is the study of the transient radio sky - named the CHILES Variable & Explosive Radio Dynamic Evolution Survey (CHILES VERDES). The CHILES

Con Pol dataset provides the opportunity to conduct the deepest radio transient search ever performed, despite the relatively small field of view. The 1000 hours of observations are recorded over many epochs, spanning a number of years (2013–2018), each able to achieve a sensitivity of approximately $10 \mu\text{Jy beam}^{-1}$. This means that the transient survey will be sensitive to a wide range of time-scales, from the 8 s integration time to the multiple years time-scale between the first and last observations. The depth of the survey also provides a significant test for future transient surveys with instruments such as the SKA.

This chapter presents the initial transient search of the individual epochs observed between 2013–2015. Details of the observations which form the dataset are given in Section 5.2.

5.1.1 Predicted Numbers of Transient & Variable Objects

Using areal densities reported in the literature, it is possible to estimate the number of variable and transient sources that could be found in the CHILES search. Variables are defined as objects which are always blindly detected but show significant flux variations, whereas transients are events which newly appear. For the calculations it is assumed that the CHILES epoch 5σ sensitivity is $50 \mu\text{Jy beam}^{-1}$, with the transient search covering 0.55 deg^2 (pushing further than the 0.2 deg^2 that will form CHILES Con Pol). A uniformly distributed set of sources is also assumed, i.e. $N \propto S^{-\frac{3}{2}}$, where N represents the number of sources and S the flux density limit. The areal densities used are obtained from Carilli et al. (2003) (variables), Thyagarajan et al. (2011) (variables and transients) and Mooley et al. (2013) (transients). These densities are summarised in Table 5.1, along with how many of the respective types of objects are expected in the CHILES epoch search.

The results in Table 5.1 reflect the range of reported areal densities, suggesting the number of variable or transient objects to be between ~ 1 –30 objects. These values also depend upon the definition of variable sources used by the authors, for example Carilli et al. (2003) require a fractional flux variation of at least 50 per cent. Nonetheless, the previous rates selected here

Table 5.1 Summary of the areal densities used to estimate the expected number of variables (V) and transients (T) in the CHILES epoch dataset. The estimate calculation is performed assuming a CHILES 5σ sensitivity of $50 \mu\text{Jy beam}^{-1}$ and an area of 0.55 deg^2 . A uniformly distributed set of sources is also assumed, i.e. $N \propto S^{-\frac{3}{2}}$.

Survey/Ref.	Type	Areal Density (deg^{-2})	Flux Limit (mJy)	Est. No. in CHILES Dataset
Carilli et al. (2003)	V ^a	18	0.1	28.0
Thyagarajan et al. (2011)	V ^b & T	0.2	1	9.8
Mooley et al. (2013)	T	0.37	0.21	1.8

^a $\Delta S \geq 50\%$

^b $\Delta S \geq 20\%$

indicate that the CHILES epoch search, even at the lowest prediction, will provide transient and/or variable objects.

5.2 The CHILES 2013–2015 Dataset

As of early 2016, two semesters of VLA B-configuration have been completed resulting in a total of 83 CHILES epochs being recorded. Of these, 44 datasets were observed during the period of 2013 October 25 to 2014 January 21, and 39 observations recorded between 2015 February 25 to 2015 May 04. The epochs ranged in integration times from 1–6 hr, but were commonly 6 hr in duration as seen in Figure 5.1 which shows the distribution of integration times found in the dataset. In total, 395.8 hr of data have been recorded. Full details of the observational set-up can be found in Table 5.2.

5.2.1 Calibration & Imaging

The CHILES data used in this chapter were calibrated and imaged by Chris Hales (NRAO). The calibration consists of standard methods, using an amplitude and phase calibrator (J0943-0819) and is performed using a custom built pipeline, which also includes flagging using the PIEFLAG software, also written by Chris Hales. A standard imaging step is then applied to




Fig. 5.1 The distribution of epoch integration times for the 2013–2015 CHILES dataset.

Table 5.2 Details of the observational setup and other characteristics of the 2013–2015 CHILES dataset.

Total data amount (hr)	395.8
No. of epochs	83
Pointing centre (R.A. Decl.)	$10^{\text{h}}01^{\text{m}}24^{\text{s}} +02^{\circ}21'00''$
No. of sub bands	4
Sub band central frequencies (GHz)	1.06, 1.45, 1.70, 1.83
Sub band bandwidth (MHz)	128
No. of channels	64
Resolution (arcsec)	4
Primary Beam FWHM (arcmin)	30
Transient Search Area (deg^2)	0.55

these data using the CLEAN algorithm using multi-frequency synthesis. The data are passed through the pipeline almost immediately after the observation has taken place, hence the pipeline is designed to provide good quality calibration and imaging on rapid time-scales. This enables an image of the sky to be obtained only a few hours after the observation has taken place, allowing for a swift transient search. All the data used in this chapter were calibrated and imaged using this initial calibration procedure. These data will later undergo an improved calibration procedure in preparation for the first CHILES Con Pol data release, incorporating techniques such as ionosphere correction and self-calibration. The current status and calibration quality for each epoch can be found on the CHILES Con Pol calibration history website¹.

Bandwidth Smearing & Time Averaging Loss

The effect of bandwidth smearing (Bridle & Schwab, 1999) on the CHILES observations can be parametrised by the product of the fractional bandwidth, $\Delta\nu/\nu_0$ where ν_0 is the reference frequency and $\Delta\nu$ is the effective bandwidth, and the source offset in synthesised beam widths $\theta_0/\theta_{\text{FWHM}}$ where θ_0 is the distance of the respective source from the phase centre and θ_{FWHM} is the FWHM width (30 arcmin). Using a source distance of $\theta_0 = 25$ arcmin (the edge of the transient search area), $\nu_0 = 1.4$ GHz and $\Delta\nu = 0.512$ GHz, the product value is 0.3, which relates to a $S_{\text{peak}}/S_{\text{peak}}^0$ (the peak surface brightness attenuation) of 0.98². Therefore the effects of bandwidth smearing are minimal.

CHILES is observed in with the VLA in B-configuration with an 8 s time averaging, hence, using the same NRAO reference² as quoted for bandwidth smearing, a 1 per cent loss is expected at the first null of primary beam and is therefore negligible.

¹<http://www.aoc.nrao.edu/~chales/chilesconpol/calhist/>

²Value obtained from NRAO reference table <https://science.nrao.edu/facilities/vla/docs/manuals/oss/performance/fov>.

5.2.2 Radio Data Quality

The average measured sensitivity of the 2013–2015 CHILES epochs is $12.8 \mu\text{Jy beam}^{-1}$, with the 6 hr time-scale images having an average of $10.4 \mu\text{Jy beam}^{-1}$. Figure 5.2 shows the measured rms for each epoch, in addition to displaying the respective integration time. The rms is measured by using a clipping method to 3σ on the inner $1/8$ portion of the image. The ratio of measured/theoretical rms ($\text{rms}_M/\text{rms}_T$) can also be calculated, for which the average value for the dataset is 1.344. This was slightly higher than the expected value derived by Chris Hales of 1.25 (which is the current expected rms due to the current calibration pipeline), but considered reasonable for the dataset. What can be seen from Figure 5.2 is that the data are very stable in quality, with the highest rms, $25 \mu\text{Jy beam}^{-1}$, attributed to the single 1 hr integration observation, closely followed by two, 2 hr integration time epochs. As the main transient search in this work concentrates on the epoch observations, the highest sensitivity is important to note as it can be interpreted as the limiting sensitivity of the survey. It is noted that there is a slight scattering of rms values within the same integration times, this is due to a variety of reasons such as variable strong RFI between observing sessions, excessive flagging by the automated pipeline in some instances and non-ideal imaging for observations at different LSTs. This is expected due to the preliminary nature of the calibration and will improve with future re-processing.

It should be noted that the measured average rms values previously quoted incorporate the full range of integration times. Figure 5.3 plots the average rms value for each hourly integration time, showing that the values follow the expected $\sqrt{t}$ relation, where t represents integration time.




Fig. 5.2 The measured rms of each epoch, where the colour and size of the data point represents the integration time of the observation. Epochs are represented by their number, which is in chronological order, with the boundary between semesters signified by the vertical dashed line. The average rms value is represented by the horizontal red line, with $+1\sigma$ and $+2\sigma$ values represented by the dashed red and green lines, respectively. It is seen that the spread of rms values is primarily defined by the integration times of the observations, with no obvious issues with the dataset.

5.2.3 Calibrators

As well as the analysis of the target field data quality, the amplitude calibrator source J0943-0819 was also examined to confirm that calibration was stable and consistent. The light curve of J0943-0819 for all the observations, which is derived by spectrum fitting across the bands, is displayed in Figure 5.4. It shows that the calibration was stable with negligible difference between the mean flux density of the two semesters (1.009 vs 1.006 Jy). Semester 2 did produce a more stable light curve than semester 1, however there was nothing in the light curves to suggest that amplitude calibration should be a concern during the analysis.




Fig. 5.3 The measured rms for each set of integration times that follow a $\sqrt{t}$ relation, where t represents time. Note that the 1 hr integration time is one single epoch.




Fig. 5.4 The light curve of the derived flux density via spectrum fitting of the calibrator J0943-0819, for the 2013/14 semester (*left*) and the 2015 semester (*right*). It shows that the amplitude calibration is stable, both within and between semesters.

5.2.4 Deep Image

A deep image of the CHILES dataset was created by Chris Hales using data from the 2013/14 observations. This preliminary deep image of the CHILES field has a sensitivity of $\sim 3 \mu\text{Jy beam}^{-1}$, and is used as a reference map for the transient search, helping to identify faint steady sources that could be seen as transients. The deep map is shown in Figure 5.5, which also marks the transient search radius and the area covered by the optical observations (refer to Section 5.2.5). Subtracting a sky model generated from the deep image from each epoch was also attempted, however, the quality of subtraction was not sufficient to perform a transient search at this stage (see Figure 6.2 in Chapter 6).

5.2.5 Optical Data

Beginning on 2013 November 1, optical data were recorded every night the weather permitted during the two radio semesters using the Liverpool Telescope (LT) in La Palma. The field was covered by performing nine pointings, each of a 60 s exposure which can reach an SDSS-r magnitude of 20–21. In total, 145 optical epochs have been recorded (observations were performed past the first semester cut off date), and 30 of these observations were performed on the same night as a VLA radio observation. The optical programme was managed by Teo Muñoz-Darias (IAC). An example mosaic of the optical field is shown in Figure 5.6.

5.3 Transient & Variable Search - Details & Methods

The LOFAR Transients Pipeline (TraP v3.0; Swinbank et al., 2015), described in Chapter 3, was used to search the epoch images for transient and variable sources. Given the quality control analysis in Section 5.2.2, the transient search was conducted on all epochs excluding the three observations³ that were above $20 \mu\text{Jy beam}^{-1}$ in measured rms. Hence this value

³The three epochs excluded were 2014 Jan 12, 2014 Jan 20 and 2014 Jan 21 (04:16).




Fig. 5.5 The deep image produced from the 2013/14 semester of CHILES observations. It reaches a sensitivity of $\sim 3 \mu\text{Jy beam}^{-1}$ with a resolution of 4.9×3.7 arcsec (BPA -41.1°). The green, dashed inner circle represents the FWHM (diameter 30 arcmin); the area enclosed by the solid red square signifies that covered by the optical monitoring campaign; and lastly the blue, solid outer circle signifies the boundary of the transient search area (diameter 50 arcmin). Note the map is not corrected for the primary beam and the colour bar units are $\mu\text{Jy beam}^{-1}$.




Fig. 5.6 The r band optical mosaic of the CHILES field observed on 2015 May 05, which is one of 145 recorded epochs. It reaches a depth of approximately 20 mag and is made up of nine separate telescope pointings which are signified by the dashed lines on the image.

became the highest sensitivity in the search, and reduced the average rms to $12.4 \mu\text{Jy beam}^{-1}$. While all images would usually be used given the small difference in rms, these three particular images were found to have excessive artefacts which the TraP would identify as candidates. Therefore, this was the primary reason for excluding these three images in this preliminary search. All images used in the transient search were not corrected for the primary beam (PB), and the searches were performed twice - once with the deep image (Section 5.2.4) as the first image and once without the deep image.

5.3.1 Time-scales Probed

As previously discussed in this thesis, the time-scale of the transient search is sensitive to the integration time of the observations, the time between consecutive observations and the overall time range on which the entire dataset was recorded. The integration times of the observations have been previously shown in Figure 5.1, signifying that the transient search is most sensitive on 6 hr time-scales - the most common integration time. In terms of the time difference between consecutive observations, Δt_{epochs} , Figure 5.7 shows a histogram showing the range of Δt_{epoch} values for each semester individually and the dataset as a whole. Note that Figure 5.7 does not show the single Δt_{epochs} value between semesters of ~ 1 year, but is included in the median calculation of the complete dataset. Here it is evident that the dominant value of Δt_{epoch} is 20 hr. As the transient search focusses on the individual epochs of the dataset, the 6 hr and $\sim$ day time-scales represent the values for which the search is most sensitive to. Future transients searches will see the dataset manipulated in order to shorten and lengthen the primary time-scales of highest sensitivity.

5.3.2 TraP Settings

Some of the TraP configuration parameters are set to non-default values for the CHILES search. These values are defined in Table 5.3, where any settings not stated in the table are

Table 5.3 The TraP parameters used for source extraction. ‘Auto’ and ‘Man.’ define the settings for the automatic and manual search methods, respectively. A single value means that the setting was the same value for both methods. The pixel size of the CHILES images is 1 arcsec.

TraP Parameter	Value		Description
	Auto	Man.	
detection_threshold	8.0	5.0	Defines the minimum σ value for sources to be extracted after source fitting.
analysis_threshold	5.0	3.0	Defines the minimum σ value for which source fitting is performed.
back_size_x	335 pix		The size of the background grid parallel to the x axis of the pixel grid. Value automatically obtained with the PyBDSM source extractor.
back_size_y	112 pix		The size of the background grid parallel to the y axis of the pixel grid. Value automatically obtained with the PyBDSM source extractor.
deblend_nthresh	32		The number of thresholds used for source deblending, using the same method as used by the optical source extractor, <i>Sextractor</i> (Bertin & Arnouts, 1996).
extraction_radius_pix	1500 pix		The maximum radius for which sources are extracted. In this case, 1500 pixels equates to 25 arcmin.
force_beam	True		When true, detected sources are assumed to have the size and shape of the restoring beam.
expiration	100		Defines that a source is monitored for at least 100 epochs (i.e. the complete CHILES dataset) after a blind extraction.
new_source_sigma_margin	0.0	5.0	A user defined margin of error for the definition of new sources based upon the lowest rms of the best previous image. See the text for a more thorough description.




Fig. 5.7 The distribution of the time difference, Δt_{epoch} , between each consecutive CHILES epoch for the two separate semesters. The time difference between the semesters of ~ 9000 hr is not included on the plot but is included when calculating the overall median.

set to their respective default value. The table also refers to two different transient search methods, *auto* and *manual*, which are discussed in Sections 5.3.4 and 5.3.5, respectively.

The parameter that requires further explanation is *new_source_sigma_margin*, NS_{SM} . This value acts as a margin of error for the signal-to-noise ratio of new sources detected by the TraP and whether they are considered further for analysis. A new source is defined as a source that appears which was not detected in the first image of the dataset. The value is added to the detection threshold value, SNR_{det} (the signal-to-noise ratio of the new source in the detection image) and compared against the signal-to-noise value a source would have in the best rms region of the best previous image, $SNR_{\text{rms,best}}$. For example, if SNR_{det} and NS_{SM} were both set to 5.0, then $SNR_{\text{rms,best}} \geq 10.0$ must be true to be considered as a new source. This is a method to reduce the number of reported new sources, removing those likely to be a steady source that is at the detection threshold. The definition of new sources is discussed further in the following Section 5.3.3.

5.3.3 Interpreting the TraP Results

Candidate transient and variable objects are selected by implementing the same method as was detailed in Chapter 3. The distribution of the two variability parameters calculated by the TraP, η_v and V_v , for all the detected light curves are analysed. These parameters are calculated by also including any forced monitoring extractions when a source is not subsequently detected. Sources that exceed a 1.5σ threshold in either (or both) parameters are selected for further investigation. Any sources that are defined as ‘new’ are also flagged by default regardless of their η_v and V_v values. Thus, sources are selected from the following categories:

- **η_v candidates** - sources which exceed the 1.5σ value of the η_v distribution. These are typically variable sources which were detected blindly in most, if not all, images. One can imagine these sources as those with light curves that are not well described by a constant value.
- **V_v candidates** - sources which exceed the 1.5σ value of the V_v distribution. These are typically sources which were blindly detected only once or twice, causing a large fractional variability. Though could also be a steady source which undergoes a large flaring event. These candidates are more likely to be a transient event.
- **New source candidates** - these are sources which are blindly detected that were not present in the first image of the dataset searched. They may or may not have η_v and V_v values that exceed the 1.5σ limit in each case, but are always considered as interesting objects by default regardless. If a new source also appears in the η_v or V_v candidate list, then it is removed from these respective lists and only considered as a new source. A reminder that these are partially controlled by the *new_source_sigma_margin* parameter detailed in the previous Section 5.3.2.

Defining the New Source Candidates

As the number of discovered new sources can be quite large, it is useful to attempt to further categorise the new sources into those which are likely to be transient against those which could be a steady source at the detection threshold. This sub-classification method was devised by Rowlinson et al., (in preparation) using the TraP results and consists of comparing the detection threshold, to two parameters:

- $SNR_{rms,best}$ - the signal-to-noise value the new source would have if placed in the best rms region of the best previous image. This parameter was previously used in conjunction with *new_source_sigma_margin* to define the new sources which populate the list.
- $SNR_{rms,worst}$ - the signal-to-noise value the new source would have if placed in the worst rms region of the best previous image.

Specifically, the sub-classifications are defined as:

- **New source (transient):** if $SNR_{rms,worst} \geq SNR_{det}$ - the new source is likely to be a transient as the new source definitely should have been seen previously even in the worst rms region of the previous best image.
- **New source (steady):** if $SNR_{rms,best} \geq SNR_{det}$, and, $SNR_{rms,worst} < SNR_{det}$ - the new source could be the detection of a steady source at the threshold level as the source wouldn't have been seen at the worst rms region of the previous best image.

Hence, with these sub-classifications the new sources which are defined as 'transient' can be prioritised over those which could possibly be steady sources.

Effect of the Deep Image on New Source Classifications

It is worth noting however that when the deep image is included in the transient search, this image by default becomes the best previous image. Hence, as the deep map is very sensitive,

$SNR_{\text{rms,best}}$ is always much greater than the detection threshold for any sources detected in the single epoch images. Therefore, all new sources are very likely to be classed as transients when using this classification scheme with the deep image.

Two methods were used to search the CHILES dataset, one performed automatically on the datasets, and a deeper manual search. These two methods are described in the following sections.

5.3.4 Method I - Experimental Automated Rapid Search

As the observations are calibrated and imaged hours after they are recorded, it was possible to perform an initial transient search as soon as the data were available. This process was done with minimal user interaction and was the first project to attempt to use the TraP in an automated way. A flowchart of the automatic pipeline is shown in Figure 5.8. The steps are as follows:

- (i) Once the processed data were available, an automated notification email would be sent. Once received, the download of these data would be triggered, followed by successful transfer checks and housekeeping tasks.
- (ii) The TraP was then automatically launched when new data were received, and when completed, an analysis script was run on the results.
- (iii) Results from the analysis were then emailed to the respective user, including information of any sources which had not been detected in previous images.

The pipeline consisted of a series of python scripts, tying the various steps together⁴. This method was used during the early stages of the 2015 semester, using the previous semester as a reference to define new sources along with the deep image. As this was an initial transient

⁴Making use of *Spotter* python module (authored by Sam Clements; <https://github.com/bornotyping/spotter>) and customised TraP analysis scripts based upon those written by Antonia Rowlinson (University of Amsterdam) (https://github.com/ajstewart/scripts/tree/expanded-trans-tools/TraP_trans_tools).




Fig. 5.8 The processes and data flow of the automatic transient search, for which the only human interaction was at the beginning of the pipeline to launch the calibration once the data have been recorded by the VLA.

search, the detection threshold for source extraction was set to 8σ (5σ analysis threshold; $NS_{SM}=0.0$). This value was to primarily reduce the number of false positives.

The automated method during this time was primarily implemented as a test to demonstrate that such a pipeline could be applied to the incoming datasets. Nonetheless, a number of candidates were flagged by the process that enabled them to be manually checked very swiftly. However, with manual inspection none of these candidates proved to be a real transient detection, with the majority of alerts caused by erroneous fitting on extended sources which were seen as potential new sources. It is envisioned that the method will be iterated upon and improved with future observation runs, including how to combat such false positives.

5.3.5 Method II - Offline Deep Search

The main transient search was performed using a 5σ detection threshold along with a 3σ analysis threshold for source extraction and $NS_{SM} = 5.0$. The dataset was searched twice -

once without using the deep image and once where the deep image was included as the first image to be analysed. Thus, any new sources discovered would be notified to the user if the source appeared with $SNR_{\text{rms,best}} > 10$ (this equates to a peak flux $> \sim 60 \mu\text{Jy beam}^{-1}$ without and $> \sim 30 \mu\text{Jy beam}^{-1}$ with the deep image). The 10σ criteria was to primarily manage the number of candidates reported back to the user.

Candidate sources were identified using the methods described in Section 5.3.3 and split into the respective categories of η_v , V_v , New Source (transient) and New Source (steady). The results of the deep manual search are presented in the next section.

Note that no transient search was performed on the optical data, which were primarily used to help validate the reported radio candidates.

5.4 Offline Transient Search Results

The analysis performed by the TraP provided a number of candidates, as shown by the distribution of variability parameters in Figure 5.9, which shows the results without and with the deep image. The total number of candidates is summarised in Table 5.4. Note that any new sources, either ‘transient’ or ‘steady’ are not included in the total number of η_v or V_v candidates (though they do appear on Figure 5.9). The 1.5σ threshold values were $\eta_v = 1.907$ and $V_v = 0.814$ without the deep image, and $\eta_v = 1.304$ and $V_v = 1.604$ when the deep image was included.

5.4.1 Advantages of the Deep Image

At a 5σ detection level, 2633 sources are detected in the deep image compared to the average of 450 sources in the epoch images. The values in Table 5.4 show that including the deep image had the desired effect of reducing the total number of new sources from 395 to 218, a reduction of 45 per cent. Although, as the deep image has a much improved rms, all




Fig. 5.9 The distribution of η_ν and V_ν for the offline deep transient search that *left* did not include the deep image and *right* included the deep image. Each data point on the scatter plot represents a separate light curve, including the new sources. Marked by the dashed lines are the 1.5σ thresholds for each parameter (*left*: $\eta_\nu = 1.907$, $V_\nu = 0.814$ and *right*: $\eta_\nu = 1.304$, $V_\nu = 1.604$), with the histograms showing the Gaussian fit to the data from which the σ values were calculated. Labelled on the plot are the quadrants where the η_ν and V_ν candidates lie respectively, though new sources do not necessarily exceed either or both thresholds. The effect of including the deep image is to increase the number of V_ν candidates due to the many sources that are only detected in the deep image.

Table 5.4 The total number of CHILES variable and transient candidate objects identified by the analysis of the TraP results. If a source that qualifies as a η_v or V_v is also classified as one of the New source candidates, then it is not included in the total number of the respective η_v or V_v candidates.

Candidate type	Total No. w/o deep image	Total No. w/ deep image
V_v	46	82
η_v	40	67
New Source (steady)	292	0*
New Source (transient)	103	218

* No steady new source candidates were found in this case due to deep image providing a deep previous best rms value (refer to Section 5.3.3).

candidates are now classified as new source (transient) objects, as they definitely should have been previously detected if they were in fact steady sources. As candidates require manual vetting, this is a positive effect that reduces the time required to search the dataset by removing sources that have been associated with deep image sources. The number of η_v and V_v candidates do rise however, especially in the case of V_v which is almost doubled.

Unsurprisingly, the two distributions of η_v and V_v that are displayed in Figure 5.9 are very similar, with the only significant difference between the two being the distribution of the V_v sources around the 1.5σ threshold, which is addressed in Section 5.4.2. Otherwise, including the deep image did not result in any unexpected behaviour when using the TraP, and instead greatly helps in narrowing the transient candidates. Consequently, the following results (excluding the discussion of V_v candidates) are based purely upon the TraP analysis where the deep image is included. Note that the ID numbers assigned to the candidates presented in the following sections are arbitrary numbers designated by the TraP.

5.4.2 V_v Candidates

As mentioned in Section 5.3.3, V_v candidates are those that have high fractional variability, e.g. transient sources that are only detected a few times or strongly flaring objects. Hence,

these candidates overlap significantly with the New Source candidates. In the case of the no deep image search, the high V_V value candidates which were not included in the new source list were new sources that did not meet the 10σ criteria ($\sigma_{\text{det}} + NS_{\text{SM}}$), hence, they were automatically excluded from the analysis.

When the deep image is introduced, the V_V candidates that are not in the new source list are dominated by objects that are *only* marginally detected in the deep image, such as candidate #15741 shown in Figure 5.10. The high V_V value is commonly caused by the variation in the resulting upper limit monitoring measurements that are included in the V_V calculation. This could point towards the source in question undergoing a bright flare, but typically from manual inspection the brightening is caused by non-source-like noise fluctuations, like that shown in Figure 5.10. Therefore, due to the low signal-to-noise of these objects in the deep image, these candidates were not considered further in the analysis. Those that were previously V_V candidates without the deep image have either been associated with a deep image source or placed in the New Source (transient) list - meaning that all sources of interest are now present in the list which is analysed later in this section. This also explains the rise in V_V candidates when including the deep image due to the much higher number of sources detected in the deep image against the individual epochs (~ 2600 vs. 450).

5.4.3 η_V Candidates (Variables)

The primary aim of this analysis of the CHILES dataset is to identify transient events. However, the nature of the TraP means that potential ever-present, but variable candidate sources can also be selected, as these objects account for the η_V candidates found by the analysis. While they are discussed in general here along with two examples, no thorough investigations into any candidates are presented, as this is expected to be the focus of a future CHILES VERDES publication led by another member of the project.




Fig. 5.10 The light curve and radio maps of the V_ν candidate #15741. *Left image*: the source as it appears in the deep image. *Right image*: the source as it appears in the image with the peak measured flux. This particular object represents one of the many candidates which was only detected as a faint source in the deep image. The continued monitoring measures a noise variation which in turn causes the light curve to have a high V_ν value. The light curve is not PB corrected, and the first data point is from the deep image. The grey scale of the radio images ranges from -5 – $30 \mu\text{Jy beam}^{-1}$.

In total, 67 sources lie above the $1.5\sigma \eta_\nu$ limit (see Figure 5.9). Details of all 67 candidates can be found in Table 5.8 located at the end of this chapter. On average, 450 sources are blindly detected at a 5σ level in the individual epochs, thus implying that 15 per cent of the sources in the field are variable at the flux density limit of $\sim 60 \mu\text{Jy}$, with a fractional value (V_ν) ranging from 0.02–0.25. In terms of variable surface density this translates to 122 deg^{-2} . However, this number is $2.4\times$ larger when compared to the highest previous estimate of the density of radio variable sources presented in Section 5.1.1. With further investigation, it is likely that not all 67 sources show real variability for reasons that will be discussed. Also some of the candidates selected here show relatively low levels of variability compared to the criteria set to define the previous surface density ($\Delta S \geq 50$ per cent). Two η_ν candidates, ranging from high to low η_ν values, are described in more detail below.




Fig. 5.11 The PB corrected light curve of the variable η_ν candidate #15701 (blue), along with how the source appears in the CHILES deep image marked by a red circle. The measured flux density of the source from the deep image is represented with the red dashed line. For comparison, the PB corrected light curve of the nearest source with over 70 measurements (#15728, 0.8 arcmin separation) is also plotted (orange). The main quoted η_ν and V_ν values are from the original non-PB corrected light curve, with the values from the PB corrected light curve also quoted (labelled ‘PB’). The grey scale of the radio image ranges from -10 – $70 \mu\text{Jy beam}^{-1}$.

Candidate #15701

Located at $10^{\text{h}}01^{\text{m}}20.1^{\text{s}} +02^{\circ}34'43.6''$, with an η_ν value of 577, source #15701 is the highest rated variable object in the dataset. The light curve of #15701 is shown in Figure 5.11 along with a cut out of how the source appears in the deep image. Also displayed in the figure is a light curve of the nearest source with over 70 measurements (#15728, 0.8 arcmin separation) for comparison (the light curve has been scaled in flux for clarity). The source has a deep image integrated flux density value of 8.9 mJy (PB corrected), and varies between ~ 7.5 – 9.1 mJy in the 2013/14 semester, and begins the 2015 semester below 6.5 mJy before rising to reach 8.0 mJy over two months.

A source is present in NVSS at this location with a integrated flux density of 9.6 mJy (although an $80''$ resolution), and is also detected in the VLA-COSMOS survey with an integrated flux density of 10.6 mJy. Previous analysis have also classified this source as a




Fig. 5.12 The PB corrected light curve of the variable η_v candidate #16037 (blue), along with how the source appears in the CHILES deep image marked by a red circle. The measured flux density of the source from the deep image is represented with the red dashed line. For comparison, the PB corrected light curve of the nearest source with over 70 measurements (#15981, 1.2 arcmin separation) is also plotted (orange). The main quoted η_v and V_v values are from the original non-PB corrected light curve, with the values from the PB corrected light curve also quoted (labelled 'PB'). The grey scale of the radio image ranges from -10 – $70 \mu\text{Jy beam}^{-1}$.

QSO (Brusa et al., 2009) and no optical source is detected in the LT observations. Although this is expected as the COSMOS Optical and Near-IR Catalogue (COSMOS-OIR; Capak et al., 2007) have defined the counterpart to have a V band magnitude of 24.2 mag.

Candidate #16037

This particular source, located at $10^{\text{h}}01^{\text{m}}39.8^{\text{s}} +02^{\circ}25'48.6''$, has an η_v value of 8.3, with a deep image integrated flux density of 0.5 mJy. The light curve of candidate #16037 is displayed in Figure 5.12, along with the light curve of the nearest source with over 70 measurements (#15981, 1.2 arcmin separation) for comparison. It also shows how the source appears in the deep radio image. The light curve shows strong variability, varying in flux between 0.35–0.6 mJy, with the second semester 2015 measurements displaying a clear rise and decline in flux over a period of ~ 1 month.

No source is present in the NVSS catalogue, whereas a source is detected at the same coordinates in the VLA-COSMOS survey with an integrated flux density of 0.6 mJy. The object is also coincident with an extended optical source in the LT observations, with an r band magnitude of 15.7 in the SDSS DR9 catalogue.

Discussion of η_V Candidates

Displayed in Figure 5.13 are cut out images of all 67 η_V candidates as they appear in the CHILES deep map (and a reminder to the reader that details of all 67 sources can be found in Table 5.8 at the end of this chapter). What is noticeable is that some of the variable detections arise from extended sources, for example, #15898. It is extended sources such as these that require careful examination when accurately defining the variability due to the fact that PySE (the source extractor software used by the TraP) is not optimised for extracting extended sources. Using the *force beam* option in the TraP can also lead to false variability being introduced if the morphology of the source in question varies significantly between observations (due to changes in uv coverage or elevation), such that more or less flux is measured. This can be further exacerbated by deblending, where the source varies between a single and double peaked source, hence losing flux if the deblending manages to distinguish between the peaks. The TraP can be run without the force beam option, however, this often leads to various erroneous extractions that do not relate to real sources which can contaminate the variability parameters making the analysis difficult. These will eventually be filtered out with future updates to the TraP.

Alongside a comparison to previous radio catalogues, a clear method to validate the variability of these sources is to extract the light curves using a different source extractor, preferably one which is designed with more emphasis on extended structures, such as PyBDSM (Mohan & Rafferty, 2015) or BLOBCAT (Hales et al., 2012). In fact BLOBCAT employs flood filled algorithms to extract sources providing a useful comparison. For




Fig. 5.13 Postage stamp images of all 67 η_v candidates as they appear in the CHILES deep image, showing how variability is detected from both point-like and extended sources. Each image is 40 arcsec in diameter and the grey scale ranges from -2 – $30 \mu\text{Jy beam}^{-1}$.




Fig. 5.14 The non-PB corrected light curve of the η_v candidate source #15701 as measured by the TraP (blue) and BLOBCAT (orange). It shows how the BLOBCAT light curve, which is extracted by a flood filling algorithm as opposed to Gaussian fitting, is in agreement with that seen from the TraP.

example, Figure 5.14 displays the light curves of the η_v candidate #15701 extracted by the TraP and BLOBCAT. It shows an agreement between the two source extractors, giving no reason to doubt the variability observed. Any false variable sources are commonly very specific cases and any variability detected by the TraP is, and has been proven to be, very reliable (Hobbs et al., 2016; Broderick et al., 2016). Given that a significant percentage of the sources shown in Figure 5.13 are not complex extended sources, there is a high chance that there are 10s of variable sources present in the CHILES field.

5.4.4 New Sources (Transients)

As there were no clear targets to prioritise when beginning the analysis, each of the 218 transient candidates were manually examined by eye in order to identify which objects were of potential interest over those which were very likely to be artefacts or other false positives. After this initial analysis, 75 of the 218 objects were perceived as potential transient

candidates. Before discussing these 75 objects, this section first describes how these sources were selected from the original 218 candidates.

Selecting the Likely False Positives

The manual inspection of the candidates involved examining by eye the candidate in the respective detection image, the position of the candidate in the other epoch images, both before and after, and finally the location of the object in the deep image. The appearance of the source along with the other information allowed the candidate to be categorised into one of six groups (the initial after the group name denotes how the group is referred to):

- **Artefact (A)** - candidates which were clearly most likely to be artefacts, commonly caused by phase errors around bright sources. Also placed in this category were sources which appeared in amongst high noise variation regions (i.e. source like objects also appeared around the selected candidate).
- **Near bright source (B)** - sources which appeared next to or very close to a bright source, but were not obviously an associated artefact.
- **Deep image association (C)** - candidates placed in this category had a source present in the deep image, but it was either not bright enough to be detected in the initial source extraction, or, the source association performed by the TraP had failed to match the candidate to the deep image source.
- **Extended source (E)** - the extracted source was part of an extended structure.
- **Peanut sources (P)** - these are double sources (resembling a peanut) where one of the components brightens in one epoch such that the TraP defines it as a new source. While both components are seen in the deep image, it is likely that the deblending has not been able to differentiate the two peaks.
- **Transient candidates (X)** - candidates which displayed no obvious features that would lead to doubt the reality. Commonly these were very faint sources.

Table 5.5 The total number of candidates placed in each group category. The groups are as follows: artefact (A), near bright source (B), deep image association (C), extended source (E), peanut source (P) and priority transient candidates (X).

Group	# in group
A	51
B	13
C	16
E	46
P	17
X	75

It should be noted that not all the groups mean that the candidate in question is a certain artefact. This is especially true for groups B, C and to some extent E, however these would require a significant amount of testing to ascertain if they were real phenomena. However, the grouping allowed for the priority to be placed on sources with a high chance of being actual transients (group X). The total number of candidates in each group is shown in Table 5.5, and typical examples of each candidate type are shown in Figure 5.15.

The total number, and type, of candidates flagged by the TraP suggests that the extraction threshold could be seen as ambitious for these data. Although, as this analysis uses images that have been produced by an initial calibration pipeline, many of these candidates may be eliminated once the dataset has been put through a more robust version of the pipeline. Also, the number of candidates in this investigation was deemed manageable, but for future transients surveys with better sensitivity and higher source counts, being able to swiftly judge which candidates to prioritise when considering the technical aspects of the detection without manual involvement may become significantly more important.

Analysing the 75 Transient Candidates

The details of the 75 transient candidates (group X) are presented in Table 5.9 at the end of this chapter. As not all the candidates could be studied in detail at this stage, objects were selected for further analysis by utilising the optical dataset to identify those radio candidates




Fig. 5.15 Examples of candidates that have been categorised into each group. The scale is shown in the first panel and is the same for each image. The location of the transient candidate is denoted by a red circle, and blue contours show detections from the deep image at the levels of $15 \mu\text{Jy}$ (5σ), $30 \mu\text{Jy}$ (10σ) and $60 \mu\text{Jy}$ (20σ). *Top row:* #17263 represents a candidate that has been defined as a likely artefact (group A) as it lies upon structure emanating from a bright source. #17279 shows a candidate from group B - a source which lies very close to a bright source yet is not obviously an artefact. #17178 shows a source that has been detected as a new transient source, however, as shown by the contours, a source is detected clearly at its location in the deep image (group C). *Bottom row:* Candidate #17284 is detected as part of an extended structure (group E) which are not considered in this investigation. #17146 is a first detection of the lower component of a peanut shaped source which was not extracted from the deep image (group P). Lastly, #17336 shows a transient candidate which shows no obvious reasons as to why it should not be considered a real event. The colour bar units are $\mu\text{Jy beam}^{-1}$.

that lie on, or near, an optical source. Both the nearest (in time) before and after optical epochs were checked for any differing features between them. As an optical catalogue of the CHILES field had not been produced, this was done manually with attention given to making sure the radio transient appeared relatively clean (no other new point source like structures or noise spikes appearing around the candidate). This was performed assuming that the presence of an optical source provides the best chance that the radio candidate was a real event. It is important to note that this method biases the type of objects that are likely to be highlighted. For example, a radio transient event from an isolated neutron star would not be expected to have an optical counterpart event or host in the LT images. However at this stage of the analysis, the reality of the event was considered the most important aspect.

From the above process, three candidates were selected for initial further analysis - two that were very close to an optical source(s) and one which matched position with an optical object. There were no differences between the before and after optical epochs for all three candidates. The three candidates are detailed below.

Transient Candidate #17230

Detected during a 3 hr observation on 2014 January 02 (05:30 UTC), candidate #17230 was located at $10^{\text{h}}00^{\text{m}}52.1^{\text{s}} +02^{\circ}31'43.7''$ (positional error of 0.5 arcsec) with an integrated flux density of $114.9 \pm 32.6 \mu\text{Jy}$ ($210.1 \pm 57.7 \mu\text{Jy}$ PB corrected). It was only detected in the one epoch with no prior or repeat emission. No source is present in the deep image at the location to a 3σ level, and no object is recorded in the NVSS catalogue or the VLA-COSMOS survey. No GRB or SN events were reported at this location on, or near, this date.

In the LT optical images, a point source lies directly next to the transient candidate, the centre of each source separated by ~ 5 arcsec. The point source is detected in all LT epochs, and is also recorded in the 2MASS All Sky Survey as a star with an R band magnitude of 16.9 mag. The LT data of this object were analysed by Retha Pretorius (University of Oxford),

which included observing a spectrum of the star in question. Both the spectrum and the instrumental light curve of the source provided no indication that the star was variable or a likely flare star.

The radio detection image, along with the nearest subsequent optical epoch and the radio light curve can be seen in Figure 5.16.

Transient Candidate #17289

Candidate #17289 was found during a 6 hr observation on 2015 March 21 (03:04 UTC) located at $10^{\text{h}} 01^{\text{m}} 08.7^{\text{s}} +02^{\circ} 16' 37.3''$ (0.5 arcsec error) with an integrated flux density of $79.4 \pm 19.8 \mu\text{Jy}$ ($92.1 \pm 22.4 \mu\text{Jy}$ PB corrected). As with candidate #17320, the deep image contains no source at the location to a 3σ threshold, there is no object in the NVSS or VLA-COSMOS catalogues and no GRBs or SNe were reported at this location on, or near, this date.

The optical monitoring shows that the radio candidate sits between two sources, at distances of ~ 7 and ~ 9 arcsec. These two objects are also detected in the radio epoch detection image and the deep image, where the nearest source is detected as a double source in the radio. These objects are both classified as galaxies (*Simbad*) with the nearest source having a V band magnitude of 18.7 (COSMOS-OIR) and the other an r band magnitude of 16.9 (SDSS DR7). As both these sources are detected in the radio image, any direct association with #17289 can be ruled out. However, the fact that the transient candidate appears very prominently, and cleanly, alongside these sources suggests that the detection is authentic.

The radio detection image of #17289, along with the nearest subsequent optical epoch and the radio light curve can be seen in Figure 5.17.




Fig. 5.16 *Top left panel:* How the new source transient candidate #17230 appears in the radio detection image, marked by a red dashed circle. The blue contours represent deep image detections at the 5 σ , 10 σ and 15 σ levels. The colour bar units are $\mu\text{Jy beam}^{-1}$. *Top right panel:* The closest optical epoch in time after the radio detection at the location of the transient candidate, which has a depth of ~ 20 mag (r band). The position of the radio candidate is marked with the red dashed circle, with the red contours showing detections from the radio detection image at the 3 σ , 5 σ , 10 σ and 15 σ levels. It shows the radio transient candidate lies ~ 5 arcsec away from an optical source. *Bottom Panel:* The light curve of the candidate, both PB and non-PB corrected. The square data points represent blind detections with triangles signifying forced upper limit extractions. The detection epoch is signified by a red circle, which is the single detection of #17230. A zoom in of the transient detection plus the measurements 20 days before and after is shown in the central sub plot.




Fig. 5.17 *Top left panel*: How the new source transient candidate #17289 appears in the radio detection image, marked by a red dashed circle. The blue contours represent deep image detections at the 5 σ , 10 σ and 15 σ levels. The colour bar units are $\mu\text{Jy beam}^{-1}$. *Top right panel*: The closest optical epoch in time after the radio detection at the location of the transient candidate, which has a depth of ~ 20 mag (r band). The position of the radio candidate is marked with the red dashed circle, with the red contours showing detections from the radio detection image at the 3 σ , 5 σ , 10 σ and 15 σ levels. It shows the radio transient candidate lies between two optical sources at distances of ~ 7 and ~ 9 arcsec. *Bottom Panel*: The light curve of the candidate, both PB and non-PB corrected. The square data points represent blind detections with triangles signifying forced upper limit extractions. The detection epoch is signified by a red circle, which is the single detection of #17289. A zoom in of the transient detection plus the measurements 20 days before and after is shown in the central sub plot.

Transient Candidate #17333

The third candidate was detected on 2015 April 18 (04:33 UTC) during a 3 hr observation. It is located at $10^{\text{h}} 02^{\text{m}} 21.9^{\text{s}} +02^{\circ} 14' 48.7''$ (0.5 arcsec error) and has an integrated flux density of $87.2 \pm 23.7 \mu\text{Jy}$ ($187.3 \pm 53.8 \mu\text{Jy}$ PB corrected). Just as with the other candidates, there is no source in the deep image to a 3σ level, there is no object in the NVSS or VLA-COSMOS catalogues and no GRBs or SNe were reported at this location on, or near, this date. However, unlike the other two candidates, this object is seen to repeat and reappears on 2015 May 03 with a PB corrected integrated flux density of $165.5 \pm 63.0 \mu\text{Jy}$. A manual check by eye also reveals marginal detections in surrounding April and May snapshots and indeed during the first semester, however these are faint and cannot be considered statistically significant.

The LT optical observations showed an extended object coincident with the location of candidate #17333. It has previously been detected in the COSMOS-OIR catalogue with a V band magnitude of 18.8. The galaxy has also been detected as a point source in the *Chandra*-COSMOS X-ray survey (Elvis et al., 2009), with a photometric redshift measurement of $z = 0.124$ derived by Salvato et al. (2011), equating to a distance of ~ 580 Mpc. With this redshift the radio luminosity can be calculated for the object at its brightest epoch ($187.3 \pm 53.8 \mu\text{Jy}$) which is $1.13 \pm 0.32 \times 10^{38} \text{ erg s}^{-1}$. Comparing this value to the luminosity of objects compiled by Pietka et al. (2015), it appears consistent with SNe, GRBs or low-luminosity AGN. Given that #17333 repeats, and lies coincident with an extended optical source, variable AGN activity would appear to be a likely origin for this candidate. However, if such emission is intrinsic, or is instead caused by scintillation (Rickett, 1986), cannot be concluded.

The radio detection image of #17333, along with the nearest subsequent optical epoch and the radio light curve can be seen in Figure 5.18.




Fig. 5.18 *Top left panel*: How the new source transient candidate #17333 appears in the radio detection image, marked by a red dashed circle. The blue contours represent deep image detections at the 5σ , 10σ and 15σ levels. The colour bar units are $\mu\text{Jy beam}^{-1}$. *Top right panel*: The closest optical epoch in time after the radio detection at the location of the transient candidate, which has a depth of ~ 20 mag (r band). The position of the radio candidate is marked with the red dashed circle, with the red contours showing detections from the radio detection image at the 3σ , 5σ , 10σ and 15σ levels. It shows the radio transient candidate lies at the same location as an extended optical object. *Bottom Panel*: The light curve of the candidate, both PB and non-PB corrected. The square data points represent blind detections with triangles signifying forced upper limit extractions. The detection epoch is signified by a red circle, and it can also be seen how #17333 is detected at a later epoch just below the 5σ detection level. A zoom in of the transient detection plus the measurements 20 days before and after is shown in the central sub plot.

Increasing the Time Resolution

In order to try and gain more information about the transient candidates, and to also test authenticity, the datasets were split to increase the time resolution. The datasets were imaged on time-scales of half the total integration time, and then 1 hr and 0.5 hr intervals. Note that not all time intervals are equal in data amount due to calibrator scans during the observation. The results of this test for each candidate are presented in Tables 5.6.

Two candidates, #17289 and #17333, were detected in a portion of their respective observation rather than being present for the entire duration. The same could possibly be true for candidate #17230 however, the first half of the observation was of too poor quality to detect the source. The 0.5 hr images were not sensitive enough to detect any of the candidates significantly, meaning a highest time resolution of 1 hr was probed. Given the results presented in Table 5.6, candidate #17230 was seen for a 1 hr time-scale where as 17289 and 17333 were detected over 2 hr integration times. The brightest candidate was #17289 which was seen in the second hour of the 6 hr observation with a peak flux density of $171 \pm 28 \mu\text{Jy beam}^{-1}$ (6σ). The fact that the sources are clearly detected at higher time resolution provides positive evidence to their reality, however the increase in resolution is not quite enough to judge on further astrophysical origins.

Exploring the frequency behaviour of the candidates was also attempted, however the image quality of the individual frequency bands was not good enough to significantly detect any of the candidates. Producing separate band images will become a product of the re-calibration pipeline, which is expected to provide images of much better quality allowing to definitively perform such a test.

The Remaining 72 Candidates & Future Analysis

As previously mentioned, the selection of the three detailed candidates is biased by the requirement of an incident or nearby optical source, and there are likely to be interesting

Table 5.6 The results of the time-scale analysis on each of the three transient candidates. The column headed ‘Int. Time’ provides the total integration time of the respective observation, with the ‘Half Int. Time’ denoting the column where the dataset is imaged in two halves. All quoted flux densities are peak flux densities.

Candidate #	Int. time (hr)	Half Int. time interval	1 hr intervals	0.5 hr intervals
17230	3.0	First half poor quality, source present in second half (peak flux: $143 \pm 28 \mu\text{Jy beam}^{-1}$).	Poor hours 1 and 2, source present in final hour 3 (peak flux: $150 \pm 29 \mu\text{Jy beam}^{-1}$).	rms too high for analysis.
17289	6.0	Source present in the first half only (peak flux: $115 \pm 16 \mu\text{Jy beam}^{-1}$).	Source unclear in hour 1, but detected very clearly in hour 2. Poor images for 3 and 4, no detection in 5 and 6 (hour 2 peak flux: $171 \pm 28 \mu\text{Jy beam}^{-1}$).	Possibly visible in 1/12 and 2/12 but unclear, likely detected in image 3/12 but too marginal with high rms of $60 \mu\text{Jy beam}^{-1}$. Images generally had too high rms value for detections.
17333	3.0	Source present in the first half only (peak flux: $104 \pm 17 \mu\text{Jy beam}^{-1}$).	Source visible in hours 1 and 2, not in hour 3 (hour 2 peak flux: $105 \pm 23 \mu\text{Jy beam}^{-1}$).	Visible in interval 2/6 by eye but not statistically significant.

events in the remaining 72 candidates. On the other hand the other candidates are just as likely to be spurious sources. While no single method can achieve this definitively, here, two characteristics of the objects are calculated that could be used to effectively rank the candidates. First, the distance from the phase centre, which could highlight regions where candidates are more common hinting at an area of poor quality, and second, the distance to the nearest ‘bright source’, which could point towards the object in question being an associated phase error.

The analysis of the first parameter, the distance of the transient candidate from the phase centre, can be seen in Figure 5.19. The figure shows the position of all 75 candidates in the FoV along with the distribution of distances from the centre. It is shown that 28 of the candidates lie at a distance ≥ 23.7 arcmin, and from the positioning map it can also be seen that many of these sources are grouped together at the edge of the FoV. This signifies that these areas of the radio map are possibly susceptible to noise variations which can be mistaken as transient detections. Such sources can be assigned a very low priority when choosing candidates to follow up, and vice-versa, those towards the phase centre can be regarded as more likely to be true detections.

Before the distance to the nearest ‘bright source’ parameter can be calculated, the sources which are defined as a ‘bright source’ have to be distinguished. This was done by analysing the distribution of all the sources detected in the deep image and placing a mean $(\mu) + 5\sigma$ cut and selecting those sources which exceed this value. This resulted in 10 sources being defined as ‘bright sources’. Their location on the CHILES map is shown in Figure 5.20 along with the flux distribution used to select them. The 10 bright sources account for 20 per cent of the total flux extracted from the image. Figure 5.21 shows the distribution of distances to the nearest bright source for 75 transient candidates. It shows that the distribution is relatively flat, which means not many sources that lie extremely close to a bright source have been missed in the first classification of all 218 candidates. Nine sources have a distance ≤ 2.7 arcmin to




Fig. 5.19 *Left*: The positions of the 75 transient candidates, each marked by a red ‘X’. The transient search area is shown by the outer dashed blue circle (radius 25 arcmin) and the FWHM is denoted by the inner dashed green circle (15 arcmin). *Right*: The distribution of the radial distance for each candidate from the centre of the image, where each bin is of equal area. 28 candidates (37 per cent) are located at a distance of 23.7 arcmin or greater.




Fig. 5.20 *Left*: The positions of the 10 selected bright sources, each marked by a blue ‘X’. The transient search area is shown by the blue circle and the FWHM is denoted by the inner dashed green circle. *Right*: The log flux density distribution of all the sources detected in the deep image, where the mean (μ) and $\mu + 5\sigma$ are marked with green and red dashed lines respectively. The 10 brightest sources are defined as those beyond the $\mu + 5\sigma$ threshold and account for 19 per cent of the total flux density.




Fig. 5.21 The distribution of the distances to the nearest bright source of the 75 transient candidates.

a bright source, which should be treated with caution, otherwise the remaining sources look to be at safe distances from bright sources.

Both parameters have been calculated for all 75 candidates and are included in Table 5.9. It is feasible that tests such as calculating the two parameters detailed above could be performed automatically in large area, deep image-plane transient searches, and are important when analysing these CHILES data which have been initially calibrated, given the number of candidates. It is this type of ‘rough’ calibrated data where tests such as these are required the most. However, once the CHILES dataset has been calibrated to a high level of quality, these new data become the basis of any future analysis and provide the most stringent method to define the reality of the candidates.

5.5 Transient Surface Density

Following the method detailed in Chapter 1, Section 1.6, the upper limits of the transient surface density can be defined for the CHILES dataset using Poisson statistics. As a reminder,

Table 5.7 An extension of Table 1.1 that summarises the calculated radio transient surface densities, ρ , for the CHILES survey as presented in this work. Columns: δt is the timescale of each individual epoch, Δt is the timescale(s) between epochs and N_e is the total number of epochs.

Survey	Flux Limit (mJy)	ρ (deg ⁻²)	δt	Δt	ν (GHz)	N_e
CHILES (this work, 0 transients)	> 0.1	0.0695	2–6 hr	hours - year	1.4	80
CHILES (this work, 75 transients)	> 0.1	2.1	2–6 hr	hours - year	1.4	80

this method uses the equation

$$P(n) = e^{-\rho A}, \quad (5.1)$$

where ρ is the surface density of sources per square degree and A represents the equivalent solid angle obtained by multiplying the area of the sky surveyed, Ω , by the number of epochs $N - 1$. A Poisson 2σ confidence level is used, i.e. $P(n) = 0.05$, and by rearranging Equation 5.1 the value of ρ can be obtained. For the CHILES dataset, $\Omega = 0.55 \text{ deg}^2$ and $N = 80$. Two upper limit surface densities are calculated, first assuming no transients were found in the dataset, i.e. $P(0) = 0.05$ and second, all 75 candidates are considered real, i.e. $P(75) = 0.05$. These two values are $\rho_{0T} < 0.0695 \text{ deg}^{-2}$ and $\rho_{75T} < 2.1 \text{ deg}^{-2}$, respectively, at a flux density limit of 65 μJy which is derived from the 5σ value of the average rms of the dataset (13 μJy). The surface densities calculated for CHILES are summarised in Table 5.7. The CHILES densities compared to the other GHz surveys that were detailed in Chapter 1 can be seen in the Log-N, Log-S plot shown in Figure 5.22.

It can be seen from Figure 5.22 that if all 75 transients were real detections then this limit is consistent with the previous surveys, although no past surveys are yet to reach the sensitivity limit probed by CHILES. If a significant portion of the 75 transients are in fact real, then it could be interpreted as the past surveys having approached, but failed, to exceed the required sensitivity limit to detect numerous transient sources. On the other hand, if there are no real transients detected in these analysed CHILES data then the surface density probes over an order of magnitude deeper than any previous surveys at similar limits. Also shown

on the plot is the surface density when considering the full 1000 hr dataset, assuming no transients are found. This is calculated by assuming these remaining 600 hr of data to be observed results in 100 further 6 hr epochs. It is an aim of the project to probe the dataset at various different time resolutions, both improving/decreasing the sensitivity, at the cost/gain of surface density limits. Hence, many more Log-N, Log-S data points are expected to be obtained from the CHILES dataset. Lastly, the estimated SKA-Deep and SKA-Shallow data points (see Section 1.6) show how the CHILES dataset is the beginning to bridge the gap between these future surveys and those performed previously.

5.6 Conclusions

In this chapter I have presented the initial radio transient search of data which were obtained as part of the CHILES project using the VLA, centred on the COSMOS field, observed during two 3 month periods in 2013/14 and 2015. These data consisted of 83 epochs ranging from 1–6 hr integration times, at a central frequency of 1.4 GHz to an average sensitivity of $13 \mu\text{Jy beam}^{-1}$. The transient search was performed on 80/83 epochs after quality control and covered an area of sky of 0.55 deg^2 . The analysed CHILES dataset consists of ~ 400 hr of the proposed total 1000 hr of observations, and has been calibrated using an initial rapid strategy developed by Chris Hales (NRAO). An optical monitoring campaign was also performed with the Liverpool Telescope, recording r band exposures concurrently with the radio observations. A transient search is yet to be performed on the optical data, but were instead used as a reference for any transient candidates to check for an optical counterpart.

The transient search itself was performed in the image-plane using the Transients Pipeline (TraP) to a 5σ detection level, along with utilising a deep image constructed from the 2013/14 dataset that reached a sensitivity of $\sim 3 \mu\text{Jy beam}^{-1}$, to act as a catalogue to cross-match potential transient candidates. At a 5σ level, 2633 sources were extracted from the deep image within the transient search area. As well as transient sources, the analysis was also




Fig. 5.22 A compilation of rates and upper limits for blind searches for radio transients at GHz frequencies, including the CHILES results from this chapter. The three CHILES results are marked with red text and define the upper limits of the analysed dataset assuming all 75 transient candidates are real, the same dataset with no transients found, and finally the full 1000 hr dataset assuming no transients are found. Multiple rates derived from the same survey are denoted by (A), (B) and (C). Along with the estimates of surface densities from the SKA deep and shallow surveys, the other surveys included in this plot are: Alexander et al. (2015) (Ale2014); Aoki et al. (2014) (Aok2014); Bannister et al. (2011a,b) (Ban2011); Bell et al. (2011) (Bel2011); Bower et al. (2007) (Bow2007); Bower et al. (2010) (PiGSS I); Bower et al. (2011) (PiGSS II); Bower & Saul (2011) (B&S2011); Croft et al. (2010) (ATATS I); Croft et al. (2011) (ATATS II); Croft et al. (2013) (PiGSS III); Frail et al. (2012) (Fra2012); Ofek et al. (2010) (Ofe2010); Ofek et al. (2011) (Ofe2011); and Williams et al. (2013) (ASGARD).

able to identify radio objects which potentially showed variability. In total 67, sources were identified as potentially variable, and 218 transient candidates were discovered.

A wide variety of objects account for the 67 variable candidates, from point-like sources to large extended structures. The most variable object in the dataset is a ~ 9 mJy extended source which has been previously detected in other surveys. The source varies from 6.5–9.0 mJy over the course of the observations, a ΔS of 30 per cent. As the TraP is primarily designed to find point-like transient sources, the variable sources should be treated with minor caution as numerous objects, such as the example given, are extended sources for which the TraP source extractor is not optimised. Hence, the main future task for the variable sources is to confirm the variability measured by the TraP. Though these errors are rare, and the work on variable sources from this chapter will form the basis of future work focussed on variable sources in the CHILES field by the CHILES VERDES team.

The 218 transient candidates were analysed in order to eliminate false positives from the results. These were caused by ‘new sources’ being identified as part of large extended structures, phase errors near bright sources and source association errors. This left 75 transients candidates for which no obvious features could be identified to doubt their reality. As a detailed investigation of each candidate could not be performed, the objects were cross-matched with the optical images of the field, selecting those that were incident, or very close to, an optical source. The three candidates had fluxes between 80–115 μ Jy and two were very close to optical sources (< 9 arcsec) whereas one had a matching position with an extended optical object. The latter candidate was also the only transient candidate to be seen in more than one epoch, and is suspected to be the result of AGN activity, either intrinsic or scintillation in origin. The astrophysical origin of the other two candidates studied is unknown. Higher time resolution of 1 hr epochs suggested that the three transient events were 1–2 hr in duration.

If it is assumed that zero of the 75 candidates are real, then the initial CHILES transient search results in an upper limit to the radio transient surface density of $\rho_{0T} < 0.0695 \text{ deg}^{-2}$ at a flux density limit of $65 \text{ } \mu\text{Jy}$, and represents the most extensive search at this sensitivity to date. If all 75 candidates are in fact real then this gives $\rho_{75T} < 2.1 \text{ deg}^{-2}$, which is consistent with previous GHz surveys, though at a deeper sensitivity. If a significant number of the 75 candidates are in fact real, then this suggests that the previous surveys have not yet probed deep enough to detect many transient sources.

The amount, and quality, of the CHILES dataset will undoubtedly improve as the project continues in the future, allowing for even deeper and extensive transient surface density limits, and helping to confirm, or invalidate, the candidates presented in this chapter. The dataset also presents an opportunity to explore the gigahertz transient parameter space at a range of time-scales, from seconds to years. This forms a primary future aim of the project, along with testing different techniques to find transients with minimal imaging or searching the visibilities directly. Such ideas and the near-term future goals of the project are discussed in the *Future Work* section in Chapter 6. The initial transient search of the first 40 per cent of the CHILES dataset described in this chapter gives a tantalising prospect that as the radio sensitivity is pushed ever deeper to μJy levels and beyond, the transient yield could significantly increase as a result.

Table 5.8 A summary of all 67 variable η_V candidates. The columns are as follows: dataset TraP ID number, right ascension, declination, maximum integrated flux, mean integrated flux, detection signal-to-noise (with respect to the deep image), η_V and V_V . The flux densities have been primary beam corrected, and the η_V and V_V values are from the original non-PB light curves used for the TraP analysis.

ID #	R.A.	Decl.	Max Int. Flux (mJy)	Mean Int. Flux (mJy)	SNR_{det}	η_V	V_V
14507	10 ^h 00 ^m 03.4 ^s	+02°07'22.9''	1.74	1.41	69.8	4.58	0.09
14516	10 ^h 00 ^m 05.4 ^s	+02°30'29.0''	4.05	3.34	259.9	15.25	0.08
14546	10 ^h 00 ^m 11.0 ^s	+02°07'24.6''	0.63	0.36	17.9	1.31	0.21
14555	10 ^h 00 ^m 12.1 ^s	+02°08'20.2''	0.71	0.50	33.1	1.32	0.14
14574	10 ^h 00 ^m 13.9 ^s	+02°22'49.4''	0.33	0.22	34.6	1.60	0.25
14578	10 ^h 00 ^m 14.2 ^s	+02°13'12.0''	1.53	1.34	126.8	7.14	0.06
14583	10 ^h 00 ^m 15.7 ^s	+02°07'34.2''	0.71	0.53	28.0	1.64	0.13
14589	10 ^h 00 ^m 16.6 ^s	+02°26'38.3''	4.83	4.08	648.7	51.34	0.07
14682	10 ^h 00 ^m 25.5 ^s	+02°00'51.6''	1.48	1.16	42.3	3.38	0.12
14781	09 ^h 59 ^m 54.7 ^s	+02°30'15.3''	2.57	1.88	100.6	6.67	0.12
14788	09 ^h 59 ^m 58.0 ^s	+02°18'09.2''	5.11	4.01	274.2	25.22	0.09
14795	10 ^h 00 ^m 00.6 ^s	+02°15'31.0''	4.18	3.73	332.3	6.01	0.05
14950	10 ^h 00 ^m 46.9 ^s	+02°07'26.4''	1.97	1.82	284.2	4.34	0.03
14960	10 ^h 00 ^m 47.6 ^s	+02°09'58.5''	0.60	0.41	62.6	2.29	0.15
15039	10 ^h 00 ^m 29.3 ^s	+02°40'45.4''	14.06	10.09	351.8	70.74	0.13
15040	10 ^h 00 ^m 30.0 ^s	+02°41'58.0''	5.48	4.03	123.4	11.09	0.15
15255	10 ^h 01 ^m 04.4 ^s	+02°04'37.1''	0.44	0.28	43.4	1.43	0.18
15260	10 ^h 01 ^m 04.5 ^s	+02°02'03.5''	1.74	1.50	150.9	4.87	0.04
15305	10 ^h 00 ^m 49.8 ^s	+02°16'54.7''	1.12	0.95	258.3	7.29	0.07
15401	10 ^h 01 ^m 09.0 ^s	+02°28'15.8''	0.75	0.62	168.0	3.00	0.07
15404	10 ^h 01 ^m 09.3 ^s	+02°17'21.5''	3.80	3.60	921.8	14.09	0.02
15441	10 ^h 01 ^m 10.6 ^s	+02°24'58.4''	0.71	0.63	164.6	2.61	0.06
15547	10 ^h 01 ^m 01.3 ^s	+02°01'17.9''	1.46	1.29	113.9	5.14	0.05
15552	10 ^h 01 ^m 01.3 ^s	+02°21'59.3''	0.72	0.64	161.7	1.31	0.05
15574	10 ^h 01 ^m 02.4 ^s	+02°05'27.9''	1.81	1.56	220.8	4.12	0.04
15701	10 ^h 01 ^m 20.1 ^s	+02°34'43.6''	9.12	7.84	1728.5	577.41	0.08
15740	10 ^h 01 ^m 22.4 ^s	+02°01'11.9''	10.32	8.59	666.9	77.72	0.05
15757	10 ^h 01 ^m 28.0 ^s	+02°40'29.3''	1.40	1.25	130.9	2.14	0.05
15803	10 ^h 01 ^m 12.1 ^s	+02°41'06.6''	2.44	2.22	198.8	12.74	0.05
15898	10 ^h 01 ^m 31.1 ^s	+02°29'24.7''	4.34	3.51	891.8	92.33	0.07
15924	10 ^h 01 ^m 33.0 ^s	+02°21'09.7''	1.18	1.09	342.3	2.06	0.03
15945	10 ^h 01 ^m 34.2 ^s	+02°09'17.4''	0.84	0.73	162.9	1.77	0.05
16010	10 ^h 01 ^m 36.5 ^s	+02°26'41.6''	0.93	0.71	146.3	3.97	0.09
16037	10 ^h 01 ^m 39.8 ^s	+02°25'48.6''	0.66	0.54	162.8	8.25	0.13
16095	10 ^h 01 ^m 24.0 ^s	+02°20'04.6''	0.74	0.59	182.3	7.89	0.10
16099	10 ^h 01 ^m 24.1 ^s	+02°17'06.3''	1.87	1.69	481.1	4.20	0.03

Continued on next page

Table 5.8 – *Continued from previous page*

ID #	R.A.	Decl.	Max Int. Flux (mJy)	Mean Int. Flux (mJy)	SNR_{det}	η_v	V_v
16191	10 ^h 01 ^m 42.2 ^s	+02°43'30.5''	0.82	0.62	33.5	1.31	0.12
16208	10 ^h 01 ^m 42.6 ^s	+02°07'52.9''	0.87	0.73	137.0	4.78	0.10
16210	10 ^h 01 ^m 42.8 ^s	+02°18'41.3''	0.41	0.25	71.5	5.24	0.22
16229	10 ^h 01 ^m 43.5 ^s	+02°21'34.6''	1.63	1.57	456.4	2.12	0.02
16240	10 ^h 01 ^m 44.1 ^s	+02°42'49.0''	1.08	0.83	56.8	1.37	0.10
16262	10 ^h 01 ^m 44.8 ^s	+02°04'09.1''	1.06	0.87	118.8	2.08	0.08
16341	10 ^h 01 ^m 49.8 ^s	+02°28'31.9''	0.28	0.17	43.9	1.49	0.19
16367	10 ^h 01 ^m 51.5 ^s	+02°25'32.2''	0.80	0.67	151.3	1.45	0.05
16452	10 ^h 01 ^m 52.6 ^s	+02°19'54.1''	0.36	0.30	89.1	1.35	0.09
16489	10 ^h 01 ^m 53.5 ^s	+02°11'52.4''	3.52	2.81	611.4	102.59	0.11
16604	10 ^h 01 ^m 59.8 ^s	+02°39'04.8''	3.48	3.14	299.6	9.10	0.04
16625	10 ^h 02 ^m 10.1 ^s	+02°16'38.0''	0.76	0.63	143.0	2.31	0.08
16667	10 ^h 01 ^m 47.4 ^s	+02°03'14.1''	7.51	6.01	706.7	38.11	0.05
16703	10 ^h 01 ^m 49.6 ^s	+02°33'34.8''	2.02	1.90	350.8	4.49	0.03
16723	10 ^h 02 ^m 14.5 ^s	+02°35'10.0''	4.61	4.16	474.3	11.57	0.03
16821	10 ^h 02 ^m 24.1 ^s	+02°16'21.3''	7.69	6.14	1078.3	118.69	0.07
16847	10 ^h 02 ^m 27.1 ^s	+02°21'19.1''	1.20	0.99	158.9	1.74	0.05
16864	10 ^h 02 ^m 28.4 ^s	+02°32'30.1''	1.59	1.32	107.2	3.17	0.08
16865	10 ^h 02 ^m 28.8 ^s	+02°17'21.8''	3.98	3.49	532.0	13.64	0.04
16878	10 ^h 02 ^m 29.8 ^s	+02°09'10.2''	2.74	2.24	165.2	4.73	0.09
16879	10 ^h 02 ^m 30.7 ^s	+02°09'21.5''	1.72	1.39	102.4	2.95	0.10
16885	10 ^h 02 ^m 29.9 ^s	+02°32'25.1''	2.83	2.47	237.4	2.32	0.05
16891	10 ^h 02 ^m 32.2 ^s	+02°06'32.0''	1.09	0.85	47.9	3.04	0.10
16894	10 ^h 02 ^m 32.7 ^s	+02°20'00.0''	0.74	0.55	82.0	1.55	0.11
16943	10 ^h 02 ^m 01.2 ^s	+02°13'26.9''	3.87	3.21	520.9	28.68	0.06
16965	10 ^h 02 ^m 02.9 ^s	+02°00'26.1''	0.73	0.56	29.0	1.33	0.14
17013	10 ^h 02 ^m 05.0 ^s	+02°37'31.5''	1.06	0.88	88.3	1.49	0.06
17088	10 ^h 02 ^m 08.7 ^s	+02°41'03.1''	8.63	7.71	481.7	52.03	0.05
17090	10 ^h 02 ^m 09.1 ^s	+02°16'02.4''	6.80	5.94	1238.9	76.32	0.04
17114	10 ^h 02 ^m 44.7 ^s	+02°07'34.8''	1.54	1.16	56.0	2.84	0.11
17133	10 ^h 02 ^m 59.1 ^s	+02°20'34.9''	8.94	7.46	439.1	45.94	0.06

Table 5.9 A summary of the 75 transient candidates that were perceived to have no obvious reasons to doubt their reality. The columns are as follows: (i) candidate TraP ID number; (ii) date and time of first detection; (iii) right ascension of the candidate; (iv) declination of the candidate; (v) integrated flux density (non-PB corrected); (vi) the signal-to-noise of the detection (SNR_{det}); (vii) the signal-to-noise of the detection if the source was located in the worst rms region of the deep image ($SNR_{\text{worst,deep}}$); (viii) the distance of the candidate from the phase centre (C_{dist}); (ix) the distance to the nearest bright source (B_{dist}); (x) the η_V value of the light curve; and (xi) the V_V value of the light curve.

ID #	Date & Time (UTC)	R.A.	Decl.	Int. Flux (μJy)	SNR_{det}	$SNR_{\text{worst,deep}}$	C_{dist} (arcmin)	B_{dist} (arcmin)	η_V	V_V
17136	2013-10-25 10:02:52	10 ^h 01 ^m 37.3 ^s	+01°57'10.9"	84.05	5.22	21.11	24.03	5.48	0.73	-5.06
17137	2013-10-29 09:47:00	10 ^h 00 ^m 38.1 ^s	+02°03'08.9"	74.37	5.06	17.87	21.21	11.23	0.73	1.66
17167	2013-11-17 08:31:48	10 ^h 02 ^m 08.3 ^s	+02°17'20.5"	51.05	5.05	12.82	11.63	1.31	0.50	5.16
17168	2013-11-17 08:31:48	10 ^h 02 ^m 38.1 ^s	+02°04'36.8"	66.85	5.28	12.16	24.70	12.25	1.98	1.34
17169	2013-11-18 08:28:04	10 ^h 00 ^m 20.5 ^s	+02°10'35.7"	70.81	5.27	17.78	18.97	13.93	0.71	1.95
17171	2013-11-18 08:28:04	10 ^h 02 ^m 42.5 ^s	+02°06'56.6"	120.68	5.17	15.88	24.09	10.47	1.20	0.64
17172	2013-11-19 08:24:04	10 ^h 00 ^m 07.1 ^s	+02°29'34.0"	57.75	5.37	14.50	21.05	18.94	0.77	1.03
17173	2013-11-19 08:24:04	10 ^h 01 ^m 40.6 ^s	+01°59'00.5"	90.38	5.28	14.35	22.36	4.55	0.81	0.92
17174	2013-11-19 08:24:04	10 ^h 02 ^m 39.3 ^s	+02°05'00.7"	49.18	5.12	11.15	24.67	11.96	1.59	1.94
17175	2013-11-21 08:16:12	10 ^h 00 ^m 28.5 ^s	+02°24'53.8"	56.13	5.18	14.10	14.42	12.67	0.85	1.70
17180	2013-11-21 08:16:12	10 ^h 02 ^m 41.2 ^s	+02°05'10.3"	54.83	5.20	10.17	24.93	11.97	1.50	2.18
17181	2013-11-21 08:16:12	10 ^h 02 ^m 41.3 ^s	+02°05'14.0"	70.39	5.26	12.04	24.91	11.92	1.84	1.57
17186	2013-11-22 08:12:12	10 ^h 00 ^m 37.1 ^s	+02°23'28.0"	87.34	5.36	15.63	12.00	10.10	0.76	0.67
17188	2013-11-22 08:12:12	10 ^h 02 ^m 38.4 ^s	+02°04'33.4"	58.46	5.10	10.58	24.80	12.33	1.42	1.16
17189	2013-11-23 08:08:12	10 ^h 00 ^m 57.0 ^s	+02°01'12.4"	66.40	5.00	12.19	20.90	6.36	0.77	4.39
17192	2013-11-26 07:56:28	10 ^h 00 ^m 13.5 ^s	+02°04'01.5"	50.83	5.17	12.77	24.47	17.45	0.89	2.01
17193	2013-11-26 07:56:28	10 ^h 02 ^m 39.3 ^s	+02°04'57.8"	63.09	5.26	15.84	24.70	12.01	1.54	3.88
17197	2013-12-06 07:32:04	10 ^h 01 ^m 34.4 ^s	+02°03'13.8"	83.32	5.14	20.93	17.94	3.24	0.78	26.69
17199	2013-12-06 07:32:04	10 ^h 02 ^m 37.9 ^s	+02°04'40.3"	82.91	5.99	20.82	24.63	12.18	1.87	2.15
17200	2013-12-10 07:56:20	09 ^h 59 ^m 56.4 ^s	+02°13'15.6"	55.15	5.22	13.85	23.21	18.66	0.94	2.65
17203	2013-12-10 07:56:20	10 ^h 01 ^m 57.8 ^s	+02°21'30.7"	66.95	5.27	15.83	8.44	6.15	0.56	0.80

Continued on next page

Table 5.9 – Continued from previous page

ID #	Date & Time (UTC)	R.A.	Decl.	Int. Flux (μ Jy)	SNR _{det}	SNR _{worst,deep}	C _{dist} (arcmin)	B _{dist} (arcmin)	η_v	V _v
17206	2013-12-12 06:54:04	10 ^h 00 ^m 13.7 ^s	+02°04'00.7"	56.57	5.45	13.74	24.44	17.40	1.02	1.90
17210	2013-12-17 08:03:56	10 ^h 01 ^m 46.4 ^s	+01°56'51.4"	52.51	5.00	13.19	24.76	6.38	0.71	10.24
17214	2013-12-21 06:33:23	10 ^h 00 ^m 13.6 ^s	+02°03'23.2"	44.51	5.15	11.18	24.90	17.34	0.84	2.87
17216	2013-12-22 06:28:44	10 ^h 00 ^m 07.1 ^s	+02°36'38.4"	109.44	5.46	27.49	24.79	18.32	1.05	-32.74
17218	2013-12-26 07:27:48	10 ^h 02 ^m 36.8 ^s	+02°04'40.0"	79.95	5.44	20.08	24.43	12.11	2.03	2.41
17220	2013-12-29 10:46:28	10 ^h 00 ^m 22.4 ^s	+02°22'18.1"	65.78	5.19	16.52	15.47	12.72	0.70	2.36
17221	2013-12-29 10:46:28	10 ^h 01 ^m 22.9 ^s	+01°57'22.2"	64.65	5.28	16.24	23.61	3.83	0.75	1.30
17226	2013-12-30 05:58:04	10 ^h 02 ^m 38.4 ^s	+02°04'40.7"	62.08	5.46	14.35	24.71	12.21	1.32	1.24
17227	2013-12-31 05:38:20	10 ^h 01 ^m 45.5 ^s	+01°59'14.4"	78.53	5.14	19.72	22.39	4.02	0.84	2.92
17229	2014-01-02 05:30:28	10 ^h 00 ^m 38.6 ^s	+02°31'47.5"	113.17	5.17	28.42	15.68	10.76	0.97	0.80
17230	2014-01-02 05:30:28	10 ^h 00 ^m 52.1 ^s	+02°31'43.7"	114.92	5.22	28.86	13.39	7.61	1.12	4.83
17233	2014-01-06 06:15:40	10 ^h 00 ^m 34.0 ^s	+02°36'30.1"	81.54	5.04	20.48	19.92	11.63	0.86	1.74
17235	2014-01-08 06:37:08	10 ^h 00 ^m 08.4 ^s	+02°37'09.9"	90.59	5.32	22.75	24.89	18.07	1.13	3.30
17240	2014-01-19 06:54:43	10 ^h 01 ^m 57.2 ^s	+02°35'25.2"	99.81	5.21	25.07	16.64	5.30	0.80	2.25
17249	2015-03-04 02:21:56	10 ^h 01 ^m 25.9 ^s	+02°00'25.3"	56.67	5.01	14.23	20.57	1.16	0.68	0.58
17250	2015-03-04 02:21:56	10 ^h 01 ^m 50.2 ^s	+01°57'24.8"	56.71	5.31	13.94	24.46	5.87	0.57	0.76
17251	2015-03-04 02:21:56	10 ^h 01 ^m 55.6 ^s	+02°02'12.3"	77.47	5.94	19.46	20.37	2.31	0.59	0.34
17257	2015-03-12 02:30:28	10 ^h 02 ^m 21.5 ^s	+02°07'26.3"	59.80	5.11	15.02	19.72	7.86	0.77	1.38
17258	2015-03-13 01:46:44	10 ^h 00 ^m 37.8 ^s	+02°23'01.6"	65.18	5.17	16.37	11.74	9.69	0.72	0.79
17266	2015-03-15 01:23:48	10 ^h 00 ^m 00.8 ^s	+02°26'31.8"	53.00	5.18	13.31	21.54	19.42	0.68	0.79
17270	2015-03-15 01:23:48	10 ^h 02 ^m 17.3 ^s	+02°18'24.0"	64.28	5.35	16.14	13.54	2.67	1.02	-13.16
17271	2015-03-15 01:23:48	10 ^h 02 ^m 18.0 ^s	+02°00'45.0"	64.20	5.67	16.12	24.31	8.06	0.98	1.46
17273	2015-03-18 01:14:04	09 ^h 59 ^m 59.3 ^s	+02°32'41.6"	45.61	5.01	11.46	24.19	20.27	0.73	1.81
17275	2015-03-18 01:14:04	10 ^h 01 ^m 55.9 ^s	+02°39'05.7"	59.89	5.62	14.99	19.78	8.51	1.15	1.05
17280	2015-03-19 02:07:48	10 ^h 00 ^m 55.6 ^s	+01°57'51.4"	48.16	5.09	12.10	24.20	7.48	0.72	0.98
17282	2015-03-19 02:07:48	10 ^h 02 ^m 36.2 ^s	+02°37'45.0"	45.30	5.20	11.38	24.62	8.67	1.69	-3.00
17283	2015-03-20 01:06:52	10 ^h 00 ^m 10.4 ^s	+02°37'19.5"	53.07	5.20	13.33	24.61	17.59	0.77	0.88
17285	2015-03-20 01:06:52	10 ^h 01 ^m 21.3 ^s	+02°18'43.4"	62.86	5.14	15.79	2.36	3.30	0.86	1.09

Continued on next page

Table 5.9 – Continued from previous page

ID #	Date & Time (UTC)	R.A.	Decl.	Int. Flux (μ Jy)	SNR _{det}	SNR _{worst,deep}	C _{dist} (arcmin)	B _{dist} (arcmin)	η_v	V _v
17286	2015-03-20 01:06:52	10 ^h 02 ^m 15.4 ^s	+01°59'37.8"	51.22	5.21	12.86	24.91	7.89	1.61	1.34
17288	2015-03-21 03:04:28	10 ^h 00 ^m 02.3 ^s	+02°31'33.6"	63.61	5.36	15.98	23.00	19.68	0.65	0.64
17289	2015-03-21 03:04:28	10 ^h 01 ^m 08.7 ^s	+02°16'37.3"	79.42	6.00	19.95	5.82	0.75	0.64	0.66
17290	2015-03-23 00:55:08	10 ^h 00 ^m 12.4 ^s	+02°36'53.8"	40.17	5.29	10.09	23.95	17.04	0.79	4.26
17293	2015-03-23 00:55:08	10 ^h 02 ^m 45.6 ^s	+02°09'07.7"	54.27	5.26	13.63	23.56	9.00	1.01	5.52
17297	2015-03-25 02:44:04	10 ^h 00 ^m 37.1 ^s	+02°26'35.0"	61.64	5.01	15.48	13.01	12.24	0.93	0.74
17300	2015-03-26 00:45:32	10 ^h 02 ^m 32.7 ^s	+02°27'44.8"	58.29	5.11	14.64	18.42	6.46	1.01	2.09
17302	2015-03-27 01:21:23	10 ^h 01 ^m 58.4 ^s	+02°21'12.0"	57.81	5.06	14.52	8.58	5.80	0.83	0.74
17303	2015-03-27 01:21:23	10 ^h 00 ^m 50.0 ^s	+02°04'06.1"	64.22	5.44	16.13	18.91	8.61	1.66	2.12
17309	2015-03-29 00:23:40	10 ^h 02 ^m 18.0 ^s	+02°12'10.9"	54.82	5.07	13.77	16.09	4.39	0.62	0.69
17311	2015-03-30 00:50:12	10 ^h 02 ^m 11.8 ^s	+02°33'55.2"	61.28	5.23	13.00	17.59	2.35	0.98	0.77
17313	2015-03-31 00:15:48	10 ^h 01 ^m 22.0 ^s	+02°09'39.7"	54.66	5.01	13.73	11.33	8.16	0.58	0.57
17315	2015-03-31 00:15:48	10 ^h 00 ^m 47.9 ^s	+02°44'14.9"	45.30	5.38	11.38	24.96	12.46	1.11	1.33
17316	2015-04-02 00:37:56	10 ^h 01 ^m 47.6 ^s	+02°04'51.8"	75.89	5.35	19.06	17.15	1.63	1.42	1.80
17318	2015-04-02 00:37:56	10 ^h 02 ^m 37.2 ^s	+02°37'55.7"	50.08	5.50	12.58	24.92	8.97	2.26	4.03
17321	2015-04-04 00:42:20	10 ^h 01 ^m 41.3 ^s	+02°41'23.0"	51.55	5.12	12.95	20.85	8.50	1.62	1.28
17323	2015-04-04 00:42:20	10 ^h 01 ^m 58.0 ^s	+02°04'11.4"	60.98	5.10	15.31	18.81	2.82	0.94	0.68
17329	2015-04-17 22:35:08	10 ^h 00 ^m 55.6 ^s	+02°17'24.4"	56.68	5.20	14.24	7.97	3.43	1.08	1.64
17330	2015-04-18 04:33:40	10 ^h 01 ^m 46.2 ^s	+02°01'14.6"	74.73	5.23	18.77	20.50	2.02	1.63	2.30
17333	2015-04-18 04:33:40	10 ^h 02 ^m 21.9 ^s	+02°14'48.7"	87.18	5.47	21.89	15.71	1.64	1.53	0.85
17336	2015-04-18 04:33:40	10 ^h 02 ^m 49.9 ^s	+02°27'58.1"	84.95	5.04	21.33	22.54	10.16	1.43	2.14
17338	2015-04-18 22:31:16	10 ^h 00 ^m 46.3 ^s	+02°33'21.6"	60.60	5.70	15.22	15.56	8.54	1.75	1.91
17341	2015-04-25 23:18:36	10 ^h 00 ^m 42.0 ^s	+02°13'00.6"	54.51	5.18	13.69	13.18	8.08	1.21	0.63
17348	2015-05-03 02:05:00	10 ^h 01 ^m 47.4 ^s	+01°57'38.3"	62.40	5.10	15.67	24.06	5.60	3.41	1.77
17352	2015-05-04 00:31:16	10 ^h 00 ^m 45.9 ^s	+02°30'36.4"	86.78	5.10	21.79	13.54	9.47	0.00	0.00
17353	2015-05-04 00:31:16	10 ^h 01 ^m 29.1 ^s	+02°25'21.8"	101.40	5.38	25.47	4.56	4.08	0.00	0.00

Chapter 6

Conclusions & Future Work

6.1 The First LOFAR Radio Transient

While extensive searches for radio transients at gigahertz frequencies have been performed, the low frequency transient sky has remained a poorly explored parameter space. In Chapter 3 I presented one of the initial transient searches to be performed using LOFAR, resulting in the detection of the first radio transient to be discovered with this new telescope.

The survey was centred at the NCP, covering an area of 175 deg^2 at a frequency of 60 MHz, and consisted of 1897 epochs equalling 394 hr of observations. The transient search was performed with the LOFAR Transients Pipeline (TraP) which discovered nine potential transient candidates. After extensive testing, one of these candidates was believed to be a real astrophysical event. The event, ILT J225347+862146, was only seen in a single 11 min epoch, although analysis revealed that the transient may have occurred over a time period of approximately 5 min. It implies a transient rate of $3.9^{+14.7}_{-3.7} \times 10^{-4} \text{ day}^{-1} \text{ deg}^{-2}$ and an upper limit of the transient surface density limit of $\rho_{\text{NCP}} < 1.4 \times 10^{-5}$, at an 11 min time-scale to a flux density limit of 7.9 Jy (10σ). The astrophysical origin of the transient remains unknown.

The transient candidate also showed an accompanying ‘ghost source’ - a second source appearing mirrored across the brightest source in the field from the transients position. The

ghost phenomenon had previously been seen and investigated (Grobler et al., 2014; Wijnholds et al., 2016), and was found to be linked to incomplete sky models when performing calibration. Tests performed on these data suggested that the ghost appearance could be linked to the calibration of the calibrator source in the NCP dataset, although a definitive cause of the ghost is not yet known.

6.1.1 Developments Since the Completion of the NCP Project

Since the completion of the NCP project, two recent pieces of work have been published that begin to constrain both the possible transient population properties and also the origins of the ghost source.

- (i) The narrow bandwidth of the observations meant that limited information - such as the spectral index - could be obtained directly from these radio data. Rowlinson et al. (2016) performed a search for FRB events at 182 MHz using the MWA to a flux density limit of 0.285 Jy. While no transient objects were found, the authors were able to place a limit on the spectral index of any NCP-like transient objects of $\alpha < -2$.
- (ii) Grobler et al. (2016) investigated the properties of ghost sources produced when only phase-only calibration was applied to data using the NCP result as a basis (I was also involved in this project and am listed as a co-author). It was found that while phase-only calibration produces less ghosts, it can result in brighter ghosts than those generated with full-complex calibration. It was also discovered that any primary beam corrections can greatly amplify existing ghost sources. The primary beam correction explained the simulated ghost source I presented in Section 3.5.1, specifically in why the ghost source only appeared when the inserted transient source was in particular locations. This was due to the resulting ghost appearing in an area of the field where a high primary beam amplification was applied. While the simulated ghost was able to be explained by the project, the NCP transient ghost is still under investigation.

6.2 The Optical Counterparts of Radio Transients

As future radio transient surveys continue to exploit the characteristics of new radio facilities such as wide fields of view and deeper sensitivity, the number of blindly detected transients are also likely to increase. At the same time, numerous robotic follow up instruments are either already in use or being developed, and can provide very swift multiwavelength data for newly discovered transient events. In Chapter 4 I analysed the photometric properties of radio transient and variable objects and their optical counterparts, in order to investigate whether they could be used as a method of classification for newly discovered radio, or optical, transient events.

We compiled a large sample of radio and optical flux measurements (F_r and F_o respectively) of a large range of transient and variable objects, containing a total of 12,633 sources. The measurements ranged from 1–10 GHz and used various optical bands. Additionally an emphasis was placed on obtaining radio and optical measurements that were made as simultaneously as possible (ideally < 1 day). We found that plotting F_r against F_o in log space provided the best method to distinguish classes, especially useful in separating galactic and extragalactic objects. Some particular classes of objects such as XRBs and SNe were slightly more scattered in location, but nevertheless the individual class areas could generally be defined.

In some cases the limitations of the sample may have created artificial separation between classes. Quasars, for example, were included to a depth of 1 mJy, however, it is fully expected that quasars below 1 mJy will be detected and hence, will add confusion to the nearby GRB area of the plot. We explored one method of avoiding this confusion by analysing how certain classes evolve on the plot over time. It showed how different classes could potentially follow unique trajectories, opening a further parameter which could be utilised to differentiate the classes.

The ‘transient diagnostic’ application of the plot was tested by using a pre-existing catalogue of transient and variable objects discovered in the FIRST radio survey, that also included optical measurements. The results obtained from our F_r and F_o parameter space agreed with the previous class definitions from the catalogue, and provided potential classifications for those that were unknown. It showed the potential of using this method for a swift classification using just the photometric information of a candidate.

6.3 The Initial Transient Search of CHILES

In Chapter 5 I presented the initial image plane variable and transient search of the first 400 hr of the 1.4 GHz CHILES survey, observed between 2013 and 2015. CHILES will ultimately provide a total of 1000 hr of observations centred on the COSMOS field resulting in a final sensitivity of $500 \text{ nJy beam}^{-1}$. The dataset consisted of 80 epochs ranging from 2–6 hr in integration times over an area of 0.55 deg^2 ($\sim 4 \text{ arcsec}$ resolution), achieving an average sensitivity of $13 \text{ }\mu\text{Jy}$ - the deepest gigahertz transient survey to date. Using the TraP, the search was performed on the individual epochs to a 10σ level with respect to a deep image that has an rms of $3 \text{ }\mu\text{Jy beam}^{-1}$. Optical data of the field were observed concurrently with the radio observations when possible. While not yet searched for transients, the optical data were used as a reference for any radio candidates found.

In total, 218 transient candidates were identified. These were sources which newly appeared in the dataset and were very often only present for one epoch. Each candidate was manually investigated in order to eliminate false positives, which resulted in 75 remaining candidates. Three of these objects were selected for further study by cross matching the sources with the optical images - only choosing those which had an optical counterpart, or were very close to an optical source. These three sources ranged in flux from 80–115 μJy , and while two were close to optical objects, one was incident with an extended optical source.

The latter source is believed to be the result of AGN activity, either intrinsic or scintillation, whereas the origins for the other two transients are unknown.

Two upper limits of the transient surface density are defined from the initial CHILES search, one assuming that zero of the 75 candidates are real, $\rho_{0T} < 0.0695 \text{ deg}^{-2}$, and a limit assuming that all 75 are actual events $\rho_{75T} < 2.1 \text{ deg}^{-2}$. Although further technical authenticity checks need to be applied to the candidates, the results indicate that the CHILES dataset contains multiple transient objects, and could yet further increase as the dataset continues to improve.

6.4 Future Work

In this section I briefly outline the main priorities and future steps of the three main chapters of this thesis.

6.4.1 LOFAR Transients

The low frequency transient search of the early LOFAR dataset proved a challenging task. Therefore the main priority of the LOFAR image plane transient search is to perform a repeat, dedicated transient search of the NCP at 60 MHz in an attempt to further confirm the results found in this thesis. This was started in 2014 with seven, 5 hr observations of the NCP followed by a further four, 6 hr observations in 2015, giving a total of 59 hr of new observations. The advantages of the second transient search are:

- **Full LBA bandwidth.** The new observations use the full 48 MHz bandwidth, enabling the study of the spectral index of any transient events that are found - a crucial factor that was missing in the dataset analysed in Chapter 3. This also means that equivalent 11 min epochs to the original search will be able to probe much deeper.

Table 6.1 Comparing the properties of the old and new NCP surveys. The sensitivity is calculated assuming an improvement of $\sqrt{N}$ where $N = 243$ due to the addition of 243 sub bands.

Survey Detail	NCP Old	NCP New
FoV (deg ²)	175	175
No. of 11 min epochs	1897	320
Sensitivity (10 σ ; Jy)	7.9	0.5

- **Calibration Methods.** The data quality will be further improved because of the advances in LOFAR LBA calibration made since the commissioning phase of the telescope, both in methods and software.
- **Definitive Low Frequency Search.** Transient searches at frequencies < 100 MHz are still relatively uncommon, the new dataset will provide a definitive answer to the transient density at such low frequencies.

The number of expected NCP-like transients that should be found in the new survey can be calculated by comparing the field of view and the number of epochs between the old and new survey, in addition to correcting for the respective sensitivities. Using the parameters shown in Table 6.1, and

$$N_{\text{NCP,new}} = N_{\text{NCP,old}} \times \frac{\text{FoV}_{\text{NCP,new}} \times \text{Epochs}_{\text{NCP,new}}}{\text{FoV}_{\text{NCP,old}} \times \text{Epochs}_{\text{NCP,old}}} \times \left(\frac{S_{\text{NCP,new}}}{S_{\text{NCP,old}}} \right)^{-\frac{3}{2}}, \quad (6.1)$$

where $N_{\text{NCP,old}} = 1$, the number of transients is calculated to be $N_{\text{NCP,new}} = 11$. While the calibration methods have improved for LOFAR LBA datasets, obtaining good quality images still represents a considerable undertaking. Figure 6.1 shows a sample image of one of the new NCP observations from 2014 June 16. It shows that while the image quality is good, it is greatly hindered by phase errors related to 3C 61.1. The source will require careful calibration to ultimately peel and subtract from these data to gain the best possible data quality.




Fig. 6.1 The map produced from a 5 hr integration of the NCP at 60 MHz, using a portion of the full bandwidth (~ 4 MHz). The observation was performed on 2014 June 16 and it reaches a sensitivity of ~ 70 mJy beam $^{-1}$ near the centre of the image with a resolution of 3×2 arcmin (BPA 177°). It shows how the phase calibration of 3C 61.1 is causing issues to the final image, highlighting the need for careful calibration. The colour bar units are Jy beam $^{-1}$.

Lastly, the ghost source also presents a unique aspect to the LOFAR transient discovery, which should not be forgotten. While efforts have so far failed to explain the phenomenon seen in Chapter 3, understanding the ghost will benefit not only LOFAR transient searches but also those performed with future radio telescopes.

6.4.2 Radio & Optical Flux Transient Diagnostic

The classification method presented in Chapter 4 has potential to be continually refined and improved as more transient sources are discovered through future surveys. As the results are iterated upon, any possible diagnostic use for the dataset will become more accurate and reliable.

The eventual aim would be to apply machine learning to these data, building automated classification software that could be implemented into transient detection pipelines. The software would be able to assign a probability of class of a newly discovered object, and continuously update the likelihood as more data are obtained. Such software, combined with other classifying techniques, would aid the decision making processes on selecting which newly discovered transient source requires continued monitoring. Importantly, this could be achieved without the need for any human interaction. An automatic process that may become very important in the SKA-LSST era of astronomy.

6.4.3 CHILES Transients

The main immediate future aim of the work presented in Chapter 5 is to confirm which of the 75 transient candidates are real astrophysical sources. Methods to do this (described in the chapter) include analysing the distance of each transient candidate from the phase centre and the nearest bright source. In addition to this, I anticipate that once the data are re-calibrated to a higher quality, including self-calibration and ionosphere corrections, a cross check between the current and new datasets will be the definitive test that will further narrow down

the candidates. Once this has been performed, and if required, transients can be individually tested using similar methods that were used to test the LOFAR transient in Chapter 3.

The transient search itself will also be enhanced and the following lists the priority aims of the project.

- (i) **Expanding the time-scales searched.** The transient search will be extended to a range of time-scales, from the 8 s integration sampling time to splitting the dataset in half for increased sensitivity. When the survey is complete this will theoretically enable a two epoch transient search at approximately $700 \text{ nJy beam}^{-1}$. The short time-scale offers challenges not only in calibration but also the computational aspect of imaging and storing a vast amount of short time-scale epochs. Points (ii) and (iii) are devised in part to help the short time-scale imaging solve these issues.
- (ii) **Subtracting the deep model.** Routinely being able to subtract the sky model generated from the deep image will make the transient search more efficient. Epochs with the sky subtracted can be analysed by the TraP which will significantly reduce the time required to process the image, due to the minimal source extraction required. This may become essential in the short time-scale imaging searches. Subtracting was attempted during the project in Chapter 5, however the quality was not at a required standard, as Figure 6.2 demonstrates. The re-calibrated dataset should see an increase in success and the subtraction could become a standard part of the calibration pipeline.
- (iii) **Dirty image pixel analysis.** Instead of performing a full CLEAN routine on the epochs, it may be more efficient to extract average and minimum pixel values of the dirty maps produced from the subtracted epochs. In this case a large maximum pixel value, compared to the average, could indicate the presence of a transient source. This will reduce both the processing and analysis time, and enable the identification of epochs of interest that can be examined in further detail. The images are also not




Fig. 6.2 A portion of the CHILES observation recorded on 2015 April 17, where an attempted subtraction of the deep image sky model has been applied. It shows that strong subtraction residuals still remain, in particular around bright sources, and hence requires further calibration efforts to prepare the images for an image-plane transient search. The colour bar units are $\mu\text{Jy beam}^{-1}$.

required to be permanently held on disk once the parameters are gained, reducing the computational requirements.

6.5 Final Remarks

In this thesis I have presented two transient searches; first at low frequencies using the LOFAR telescope, and second at gigahertz frequencies with the VLA. In addition, I have also described an investigation into the relationship between radio transient sources and their optical counterparts, studying whether such a relation could be used for rapid transient classification.

Both transient searches showed that bright, unquestionable radio transients are relatively rare at the sensitivity limits probed, especially when comparing to transient rates at other wavelengths such as optical. The megahertz sky was noticeably quiet and while one low

frequency transient was discovered in this work, it was made challenging by the appearance of the ghost source and the effort required to prove its reality. As radio telescopes become more advanced and use new technology, understanding the artefacts and other phenomena produced is also becoming ever more important. This applies especially to short time-scale transient surveys, as shallow uv coverage could exaggerate unexpected malicious artefacts. The gigahertz CHILES survey on the other hand did offer a high number of potential transient candidate sources, but all were very low signal to noise detections. Currently, the candidate sources have an equal chance of being real or fake, hence, it is possible that the CHILES search has started to approach the sensitivity threshold where radio transient sources are abundant.

The possibility of a future high radio transient rate, in particular with reference to the SKA, creates a transient management issue. Using the radio and optical relation to generate an initial classification could prove to be a significant component of any future transient detection pipelines, be it an interactive process or ideally, an automatic decision making pipeline. The MeerKAT/MeerLICHT project will provide the perfect opportunity for such a multiwavelength tool to evolve and become more accurate and defined. In any case, it is a certainty that the radio transient sky will be redefined in the decade to come.

References

- Abdul-Aziz H., et al., 1995, A&AS, 114, 509
- Alexander K. D., Soderberg A. M., Chomiuk L. B., 2015, ApJ, 806, 106
- Aller H. D., Aller M. F., Latimer G. E., Hodge P. E., 1985, ApJS, 59, 513
- Andersen M. I., et al., 2000, A&A, 364, L54
- Anderson G. E., et al., 2014, MNRAS, 440, 2059
- Antonelli L. A., et al., 2010, GRB Coordinates Network, 10620
- Aoki T., et al., 2014, ApJ, 781, 10
- Arcavi I., et al., 2011, ApJL, 742, L18
- Asad K. M. B., et al., 2015, MNRAS, 451, 3709
- Baars J. W. M., Hooghoudt B. G., 1974, A&A, 31, 323
- Backer D. C., Fisher J. R., 1974, ApJ, 189, 137
- Bannister K. W., Murphy T., Gaensler B. M., Hunstead R. W., Chatterjee S., 2011a, MNRAS, 412, 634
- Bannister K. W., Murphy T., Gaensler B. M., Hunstead R. W., Chatterjee S., 2011b, MNRAS, 418, 2813
- Barthelmy S. D., D’Ai A., D’Avanzo P., Krimm H. A., Lien A. Y., Marshall F. E., Maselli A., Siegel M. H., 2015, GRB Coordinates Network, 17929
- Bastian T. S., 1990, Solar Physics, 130, 265
- Bastian T. S., Benz A. O., Gary D. E., 1998, ARA&A, 36, 131
- Bell M. E., et al., 2011, MNRAS, 415, 2
- Bell M. E., et al., 2014, MNRAS, 438, 352
- Bembrick C., Pearce A., Evans R. O., 2001, IAU Circ., 7781, 3
- Berger E., et al., 2000, ApJ, 545, 56
- Berger E., et al., 2001, ApJ, 556, 556
- Berger E., Kulkarni S. R., Chevalier R. A., 2002a, ApJL, 577, L5
- Berger E., et al., 2002b, ApJ, 581, 981

- Berger E., et al., 2003a, *Nature*, 426, 154
- Berger E., Soderberg A. M., Frail D. A., Kulkarni S. R., 2003b, *ApJL*, 587, L5
- Berger E., et al., 2005a, *Nature*, 438, 988
- Berger E., et al., 2005b, *ApJ*, 634, 501
- Bertin E., Arnouts S., 1996, *A&AS*, 117, 393
- Blandford R. D., Rees M. J., 1974, *MNRAS*, 169, 395
- Bloom J. S., et al., 1998, *ApJL*, 508, L21
- Bloom J. S., et al., 2009, *ApJ*, 691, 723
- Bloom J. S., et al., 2012, *PASP*, 124, 1175
- Boiko A. I., Konovalenko A. A., Koliadin V. L., Melnik V. N., 2012, *Advances in Astronomy and Space Physics*, 2, 121
- Booth R. S., Jonas J. L., 2012, *African Skies*, 16, 101
- Bower G. C., Saul D., 2011, *ApJL*, 728, L14
- 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., 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., Whysong D., Blair S., Croft S., Keating G., Law C., Williams P. K. G., Wright M. C. H., 2011, *ApJ*, 739, 76
- Bridle A. H., Schwab F. R., 1999, in Taylor G. B., Carilli C. L., Perley R. A., eds, *Astronomical Society of the Pacific Conference Series Vol. 180, Synthesis Imaging in Radio Astronomy II*. p. 371
- Briggs D. S., 1995, in *American Astronomical Society Meeting Abstracts*. p. 112.02
- Broderick J., Stewart A., Fender R., Miller-Jones J., Mooley K., Pietka M., 2015, *The Astronomer's Telegram*, 7720
- Broderick J. W., et al., 2016, *MNRAS*, 459, 2681
- Brown T. M., et al., 2013, *PASP*, 125, 1031
- Brusa M., et al., 2009, *ApJ*, 693, 8
- Bufano F., et al., 2014, *MNRAS*, 439, 1807
- Burke-Spolaor S., Bannister K. W., 2014, *ApJ*, 792, 19
- Buta R. J., 1982, *PASP*, 94, 578
- Buxton M. M., Bailyn C. D., Capelo H. L., Chatterjee R., Dinçer T., Kalemci E., Tomsick J. A., 2012, *AJ*, 143, 130

- Cameron P. B., 2005, GRB Coordinates Network, 3513, 1
- Cameron P. B., Frail D. A., 2005, GRB Coordinates Network, 3495
- Camilo F., Ransom S. M., Halpern J. P., Reynolds J., Helfand D. J., Zimmerman N., Sarkissian J., 2006, *Nature*, 442, 892
- Camilo F., Ransom S. M., Halpern J. P., Reynolds J., 2007, *ApJL*, 666, L93
- Camilo F., Ransom S. M., Chatterjee S., Johnston S., Demorest P., 2012, *ApJ*, 746, 63
- Canizares C. R., et al., 1975, *ApJ*, 197, 457
- Capak P., et al., 2007, *ApJS*, 172, 99
- Carbone D., et al., 2016, *MNRAS*, 459, 3161
- Carilli C. L., Ivison R. J., Frail D. A., 2003, *ApJ*, 590, 192
- Casares J., Charles P. A., Jones D. H. P., Rutten R. G. M., Callanan P. J., 1991, *MNRAS*, 250, 712
- Castro-Tirado A. J., et al., 1999, *Science*, 283, 2069
- Castro-Tirado A. J., et al., 2001, *A&A*, 370, 398
- Cendes Y., et al., 2015, preprint, ([arXiv:1412.3986](https://arxiv.org/abs/1412.3986))
- Cenko S. B., et al., 2006, *ApJ*, 652, 490
- Cenko S. B., et al., 2015, *ApJL*, 803, L24
- Chakrabarti S. K., et al., 2005, *MNRAS*, 362, 957
- Champion D. J., et al., 2016, *MNRAS*, 460, L30
- Chandra P., Frail D. A., 2006, GRB Coordinates Network, 5843, 1
- Chandra P., Frail D. A., 2007a, GRB Coordinates Network, 6915, 1
- Chandra P., Frail D. A., 2007b, GRB Coordinates Network, 6978, 1
- Chandra P., Frail D. A., 2008a, GRB Coordinates Network, 7843, 1
- Chandra P., Frail D. A., 2008b, GRB Coordinates Network, 8103, 1
- Chandra P., Frail D. A., 2008c, GRB Coordinates Network, 8625, 1
- Chandra P., Frail D. A., 2009a, GRB Coordinates Network, 9260, 1
- Chandra P., Frail D. A., 2009b, GRB Coordinates Network, 9533, 1
- Chandra P., Frail D. A., 2012, *ApJ*, 746, 156
- Chandra P., Frail D. A., Cenko S. B., 2010, GRB Coordinates Network, 11131, 1
- Clark B. G., 1980, *A&A*, 89, 377
- Clocchiatti A., et al., 2001, *ApJ*, 553, 886

- Clocchiatti A., Suntzeff N. B., Covarrubias R., Candia P., 2011, *AJ*, 141, 163
- Cobb B. E., Bailyn C. D., 2005, *GRB Coordinates Network*, 3104, 1
- Cohen A., 2004, *Long Wavelength Array Memo*, 17
- 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
- Condon J. J., 1997, *PASP*, 109, 166
- Condon J. J., Ledden J. E., Odell S. L., Dennison B., 1979, *AJ*, 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
- Coppejans D. L., Körding E. G., Miller-Jones J. C. A., Rupen M. P., Knigge C., Sivakoff G. R., Groot P. J., 2015, *MNRAS*, 451, 3801
- Corbel S., et al., 2013, *MNRAS*, 431, L107
- Cordes J. M., Lazio T. J. W., 2003, *ArXiv Astrophysics e-prints*,
- Cornwell T., Fomalont E. B., 1999, in Taylor G. B., Carilli C. L., Perley R. A., eds, *Astronomical Society of the Pacific Conference Series Vol. 180, Synthesis Imaging in Radio Astronomy II*. p. 187
- Covino S., et al., 2003, *A&A*, 404, L5
- Cox A., 2000, *Allen's astrophysical quantities*, *Sky & Telesc.*, 100, 72
- Croft S., et al., 2010, *ApJ*, 719, 45
- Croft S., Bower G. C., Keating G., Law C., Whysong D., Williams P. K. G., Wright M., 2011, *ApJ*, 731, 34
- Croft S., Bower G. C., Whysong D., 2013, *ApJ*, 762, 93
- Dawson K. S., et al., 2013, *AJ*, 145, 10
- Della Valle M., et al., 2006, *ApJL*, 642, L103
- Dewdney P. E., Hall P. J., Schilizzi R. T., Lazio T. J. L. W., 2009, *IEEE Proceedings*, 97, 1482
- Dhillon V. S., Marsh T. R., Littlefair S. P., 2006, *MNRAS*, 372, 209
- Dhillon V. S., et al., 2011, *MNRAS*, 414, 3627
- Djorgovski S. G., Mahabal A. A., Donalek C., Graham M. J., Drake A. J., Moghaddam B., Turmon M., 2012, preprint, ([arXiv:1209.1681](https://arxiv.org/abs/1209.1681))
- Drout M. R., et al., 2011, *ApJ*, 741, 97
- Drout M. R., et al., 2014, *ApJ*, 794, 23
- Eatough R. P., et al., 2013, *Nature*, 501, 391
- Eichler D., Livio M., Piran T., Schramm D. N., 1989, *Nature*, 340, 126

- Eisenstein D. J., et al., 2011, *AJ*, 142, 72
- Ekers R., 2012, preprint, (arXiv:1212.3497)
- Ellingson S. W., et al., 2013, *IEEE Transactions on Antennas and Propagation*, 61, 2540
- Elvis M., et al., 2009, *ApJS*, 184, 158
- Ewen H. I., Purcell E. M., 1951, *Nature*, 168, 356
- Eyres S. P. S., Heywood I., O'Brien T. J., Ivison R. J., Muxlow T. W. B., Elkin V. G., 2005, *MNRAS*, 358, 1019
- Falcke H., Rezzolla L., 2014, *A&A*, 562, A137
- Farrell S. A., Murphy T., Lo K. K., 2015, *ApJ*, 813, 28
- Fender R., 2006, *Jets from X-ray binaries*. Cambridge University Press, pp 381–419
- Fender R. P., Bell M. E., 2011, *Bulletin of the Astronomical Society of India*, 39, 315
- Fender R. P., Bell Burnell S. J., Waltman E. B., Pooley G. G., Ghigo F. D., Foster R. S., 1997, *MNRAS*, 288, 849
- Fender R. P., et al., 2006, in *VI Microquasar Workshop: Microquasars and Beyond*. p. 104.1 (arXiv:astro-ph/0611298)
- Fender R., Wijers R., Stappers B., LOFAR Transients Key Science Project t., 2007, in *Bursts, Pulses and Flickering: Wide-Field Monitoring of the Dynamic Radio Sky*. p. 30
- Fender R., Stewart A., Macquart J. P., Donnarumma I., Murphy T., Deller A., Paragi Z., Chatterjee S., 2015, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, p. 51
- Ferrero P., et al., 2006, *A&A*, 457, 857
- Fox D. B., et al., 2005, *Nature*, 437, 845
- Frail D. A., Berger E., 2002, *GRB Coordinates Network*, 1490, 1
- Frail D. A., Berger E., 2003, *GRB Coordinates Network*, 1827, 1
- Frail D. A., et al., 2000a, *ApJ*, 534, 559
- Frail D. A., Waxman E., Kulkarni S. R., 2000b, *ApJ*, 537, 191
- Frail D. A., et al., 2000c, *ApJL*, 538, L129
- Frail D. A., Kulkarni S. R., Berger E., Wieringa M. H., 2003a, *AJ*, 125, 2299
- Frail D. A., et al., 2003b, *ApJ*, 590, 992
- Frail D. A., Cameron P. B., Soderberg A. M., 2005, *GRB Coordinates Network*, 4244, 1
- Frail D. A., Chandra P., Singh T., 2009, *GRB Coordinates Network*, 10088, 1
- Frail D. A., Kulkarni S. R., Ofek E. O., Bower G. C., Nakar E., 2012, *ApJ*, 747, 70

- Frater R. H., Brooks J. W., Whiteoak J. B., 1992, *Journal of Electrical and Electronics Engineering Australia*, 12, 103
- Fynbo J. U., et al., 2001, *A&A*, 373, 796
- Gal-Yam A., Ofek E. O., Shemmer O., 2002, *MNRAS*, 332, L73
- Gal-Yam A., et al., 2006, *ApJ*, 639, 331
- Gal-Yam A., et al., 2014, *Nature*, 509, 471
- Gal R. R., Bloom J. S., Steidel C., Adelberger K. L., Djorgovski S. G., Kulkarni S. R., 1998, *GRB Coordinates Network*, 88
- Galama T. J., et al., 2000, *ApJL*, 541, L45
- Garnavich P. M., et al., 2003, *ApJ*, 582, 924
- Gehrels N., 1986, *ApJ*, 303, 336
- Gershberg R. E., Katsova M. M., Lovkaya M. N., Terebizh A. V., Shakhovskaya N. I., 1999, *A&AS*, 139, 555
- Greiner J., Peimbert M., Esteban C., Kaufer A., Jaunsen A., Smoke J., Klose S., Reimer O., 2003, *GRB Coordinates Network*, 2020
- Grobler T. L., Nunhokee C. D., Smirnov O. M., van Zyl A. J., de Bruyn A. G., 2014, *MNRAS*, 439, 4030
- Grobler T. L., Stewart A. J., Wijnholds S. J., Kenyon J. S., Smirnov O. M., 2016, *MNRAS*, 461, 2975
- Grupe D., et al., 2006, *ApJ*, 645, 464
- Güdel M., 2002, *ARA&A*, 40, 217
- Gudel M., Benz A. O., Bastian T. S., Furst E., Simnett G. M., Davis R. J., 1989, *A&A*, 220, L5
- Guidorzi C., et al., 2011, *MNRAS*, 417, 2124
- Hales C. A., Murphy T., Curran J. R., Middelberg E., Gaensler B. M., Norris R. P., 2012, *MNRAS*, 425, 979
- Halpern J. P., Mirabal N., Armstrong E., 2006, *GRB Coordinates Network*, 5840, 1
- Hamuy M., et al., 2009, *ApJ*, 703, 1612
- Han X., Hjellming R. M., 1992, *ApJ*, 400, 304
- Hankins T. H., Rickett B. J., 1975, *Methods in Computational Physics*, 14, 55
- Harrison F. A., et al., 1999, *ApJL*, 523, L121
- Harrison F. A., et al., 2001, *ApJ*, 559, 123
- Harrison F., Cenko B., Frail D. A., Chandra P., Kulkarni S., 2009, *GRB Coordinates Network*, 9043, 1

- Hassall T. E., Keane E. F., Fender R. P., 2013, MNRAS, 436, 371
- Heald G. H., et al., 2015, A&A, 582, A123
- Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, Nature, 217, 709
- Hjellming R. M., Han X. H., Cordova F. A., Hasinger G., 1990, A&A, 235, 147
- Hobbs G., et al., 2004, MNRAS, 352, 1439
- Hobbs G., et al., 2016, MNRAS, 456, 3948
- Høg E., et al., 2000, A&A, 355, L27
- Högbom J. A., 1974, A&AS, 15, 417
- Holland S. T., et al., 2007, AJ, 133, 122
- Holwerda B. W., et al., 2014, ApJ, 788, 77
- Horesh A., et al., 2012, ApJ, 746, 21
- Horesh A., et al., 2013, MNRAS, 436, 1258
- Horiuchi S., et al., 2004, ApJ, 616, 110
- Hughes S. A., 2005, ArXiv High Energy Physics - Phenomenology e-prints,
- Hunter D. J., et al., 2009, A&A, 508, 371
- Hyman S. D., Lazio T. J. W., Kassim N. E., Bartleson A. L., 2002, AJ, 123, 1497
- Hyman S. D., Lazio T. J. W., Kassim N. E., Ray P. S., Markwardt C. B., Yusef-Zadeh F., 2005, Nature, 434, 50
- Hyman S. D., Lazio T. J. W., Roy S., Ray P. S., Kassim N. E., Neureuther J. L., 2006, ApJ, 639, 348
- Hyman S. D., Wijnands R., Lazio T. J. W., Pal S., Starling R., Kassim N. E., Ray P. S., 2009, ApJ, 696, 280
- Im M., Lee I., Urata Y., 2007, GRB Coordinates Network, 6970, 1
- Im M., Jeon Y., Park W., Lee I., Jeon Y.-B., Urata Y., 2009, GRB Coordinates Network, 9275, 1
- Ivezic Z., et al., 2008, preprint, (arXiv:0805.2366)
- Jackson P. D., Kundu M. R., Kassim N., 1990, Solar Physics, 130, 391
- Jaeger T. R., Hyman S. D., Kassim N. E., Lazio T. J. W., 2012, AJ, 143, 96
- Jansky K. G., 1933, Nature, 132, 66
- Jensen B. L., et al., 1999, GRB Coordinates Network, 498
- Jester S., et al., 2005, AJ, 130, 873
- Johnston S., Romani R. W., Marshall F. E., Zhang W., 2004, MNRAS, 355, 31

- Johnston S., et al., 2007, PASA, 24, 174
- Johnston S., et al., 2008, Experimental Astronomy, 22, 151
- Kann D. A., Hoegner C., Filgas R., 2007, GRB Coordinates Network, 6918, 1
- Kantor J., 2014, in Wozniak P. R., Graham M. J., Mahabal A. A., Seaman R., eds, The Third Hot-wiring the Transient Universe Workshop. pp 19–26
- Kashiyama K., Ioka K., Mészáros P., 2013, ApJL, 776, L39
- Kasliwal M. M., 2011, Bulletin of the Astronomical Society of India, 39, 375
- Kato T., Uemura M., Ishioka R., Nogami D., Kunjaya C., Baba H., Yamaoka H., 2004, PASJ, 56, S1
- Keane E. F., Kramer M., Lyne A. G., Stappers B. W., McLaughlin M. A., 2011, MNRAS, 415, 3065
- Keane E. F., et al., 2016, Nature, 530, 453
- Kellermann K. I., Pauliny-Toth I. I. K., 1969, ApJL, 155, L71
- Kida S., et al., 2008, New Astron., 13, 519
- Körding E., Rupen M., Knigge C., Fender R., Dhawan V., Templeton M., Muxlow T., 2008, Science, 320, 1318
- Körding E. G., Knigge C., Tzioumis T., Fender R., 2011, MNRAS, 418, L129
- Kovalev Y. Y., et al., 2005, AJ, 130, 2473
- Kramer M., Lyne A. G., O’Brien J. T., Jordan C. A., Lorimer D. R., 2006, Science, 312, 549
- Kulkarni S. R., et al., 1998, Nature, 395, 663
- Kulkarni S. R., et al., 1999, ApJL, 522, L97
- Kuniyoshi M., et al., 2007, PASP, 119, 122
- Lainela M., 1994, A&A, 286, 408
- Lane W. M., Cotton W. D., van Velzen S., Clarke T. E., Kassim N. E., Helmboldt J. F., Lazio T. J. W., Cohen A. S., 2014, MNRAS, 440, 327
- Law N. M., et al., 2009, PASP, 121, 1395
- Lazio T. J. W., et al., 2010, AJ, 140, 1995
- Leahy J. P., Perley R. A., 1991, AJ, 102, 537
- Lee M. G., Kim E., Kim S. C., Kim S. L., Park W.-K., Pyo T. S., 1995, Journal of Korean Astronomical Society, 28, 31
- Leonard D. C., et al., 2002, PASP, 114, 35
- Levan A., et al., 2006, ApJ, 647, 471
- Levin L., et al., 2010, ApJL, 721, L33

- Lipkin Y. M., et al., 2004, *ApJ*, 606, 381
- Liu Q. Z., van Paradijs J., van den Heuvel E. P. J., 2007, *A&A*, 469, 807
- Lo K. K., Farrell S., Murphy T., Gaensler B. M., 2014, *ApJ*, 786, 20
- Loeb A., Shvartzvald Y., Maoz D., 2014, *MNRAS*, 439, L46
- Longair M. S., 2011, *High Energy Astrophysics*
- Lorimer D. R., Kramer M., 2012, *Handbook of Pulsar Astronomy*
- Lorimer D. R., Yates J. A., Lyne A. G., Gould D. M., 1995, *MNRAS*, 273, 411
- Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, *Science*, 318, 777
- Lorimer D. R., Karastergiou A., McLaughlin M. A., Johnston S., 2013, *MNRAS*, 436, L5
- Lovell B., 1969, *Nature*, 222, 1126
- MacFadyen A. I., Woosley S. E., Heger A., 2001, *ApJ*, 550, 410
- Macquart J.-P., 2014, *PASA*, 31, 31
- Mahabal A. A., et al., 2011, *Bulletin of the Astronomical Society of India*, 39, 387
- Malesani D., et al., 2007, *A&A*, 473, 77
- Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, *AJ*, 129, 1993
- Margalit B., Metzger B. D., Beloborodov A. M., 2015, *Physical Review Letters*, 115, 171101
- Maron O., Kijak J., Kramer M., Wielebinski R., 2000, *A&AS*, 147, 195
- Marshall F. E., et al., 2011, *ApJ*, 727, 132
- Masetti N., et al., 2003, *A&A*, 404, 465
- Masui K., et al., 2015, *Nature*, 528, 523
- Matsumura N., et al., 2007, *AJ*, 133, 1441
- Matsumura N., et al., 2009, *AJ*, 138, 787
- Mazzali P. A., et al., 2008, *Science*, 321, 1185
- McKean J., et al., 2011, preprint, (arXiv:1106.1041)
- McLaughlin M. A., et al., 2006, *Nature*, 439, 817
- McMullin J. P., Waters B., Schiebel D., Young W., Golap K., 2007, in Shaw R. A., Hill F., Bell D. J., eds, *Astronomical Society of the Pacific Conference Series Vol. 376, Astronomical Data Analysis Software and Systems XVI*. p. 127
- Melandri A., et al., 2010, *ApJ*, 723, 1331
- Metzger B. D., Williams P. K. G., Berger E., 2015, *ApJ*, 806, 224

- Mignani R. P., 2011, *Advances in Space Research*, 47, 1281
- Miller-Jones J. C. A., Jonker P. G., Dhawan V., Briskeen W., Rupen M. P., Nelemans G., Gallo E., 2009, *ApJL*, 706, L230
- Mingaliev M. G., Sotnikova Y. V., Kardashev N. S., Larionov M. G., 2009, *Astronomy Reports*, 53, 487
- Misra K., Kamble A. P., Sahu D. K., Srividya S., Bama P., Anupama G. C., Vanniarajan M. S., 2005, *GRB Coordinates Network*, 4259, 1
- Misra K., Pooley D., Chandra P., Bhattacharya D., Ray A. K., Sagar R., Lewin W. H. G., 2007, *MNRAS*, 381, 280
- Mohan N., Rafferty D., 2015, *PyBDSM: Python Blob Detection and Source Measurement*, *Astrophysics Source Code Library* (ascl:1502.007)
- Moin A., et al., 2013, *ApJ*, 779, 105
- Monet D. G., et al., 2003, *AJ*, 125, 984
- Moody J. W., Laney D., Pearson R., Pace C., 2010, *GRB Coordinates Network*, 10665, 1
- Mooley K. P., Frail D. A., Ofek E. O., Miller N. A., Kulkarni S. R., Horesh A., 2013, *ApJ*, 768, 165
- Mooley K. P., et al., 2016, *ApJ*, 818, 105
- Muller C. A., Oort J. H., 1951, *Nature*, 168, 357
- Murphy T., et al., 2015, *MNRAS*, 446, 2560
- Mutel R. L., Weisberg J. M., 1978, *AJ*, 83, 1499
- Narayan R., Paczynski B., Piran T., 1992, *ApJL*, 395, L83
- Nardini M., et al., 2014, *A&A*, 562, A29
- Negoro H., et al., 2015, *The Astronomer's Telegram*, 7646
- Nelson G. J., Robinson R. D., Slee O. B., Fielding G., Page A. A., Walker W. S. G., 1979, *MNRAS*, 187, 405
- Niinuma K., et al., 2007, *ApJL*, 657, L37
- Niinuma K., et al., 2009, *ApJ*, 704, 652
- Notsu Y., et al., 2013, *ApJ*, 771, 127
- O'Brien T. J., 2009, in *8th International e-VLBI Workshop*. p. 3
- Obenberger K. S., et al., 2014a, *ApJ*, 785, 27
- Obenberger K. S., et al., 2014b, *ApJL*, 788, L26
- Ofek E. O., Breslauer B., Gal-Yam A., Frail D., Kasliwal M. M., Kulkarni S. R., Waxman E., 2010, *ApJ*, 711, 517

- Ofek E. O., Frail D. A., Breslauer B., Kulkarni S. R., Chandra P., Gal-Yam A., Kasliwal M. M., Gehrels N., 2011, *ApJ*, 740, 65
- Offringa A. R., de Bruyn A. G., Biehl M., Zaroubi S., Bernardi G., Pandey V. N., 2010, *MNRAS*, 405, 155
- Offringa A. R., de Bruyn A. G., Zaroubi S., 2012a, *MNRAS*, 422, 563
- Offringa A. R., van de Gronde J. J., Roerdink J. B. T. M., 2012b, *A&A*, 539, A95
- Olausen S. A., Kaspi V. M., 2014, *ApJS*, 212, 6
- Osten R. A., Hawley S. L., Allred J. C., Johns-Krull C. M., Roark C., 2005, *ApJ*, 621, 398
- Osten R. A., Hawley S. L., Allred J., Johns-Krull C. M., Brown A., Harper G. M., 2006, *ApJ*, 647, 1349
- Page K. L., et al., 2009, *MNRAS*, 400, 134
- Palazzi E., et al., 1998, *A&A*, 336, L95
- Panagia N., Sramek R. A., Weiler K. W., 1986, *ApJL*, 300, L55
- Pandey V. N., van Zwieten J. E., de Bruyn A. G., Nijboer R., 2009, in Saikia D. J., Green D. A., Gupta Y., Venturi T., eds, *Astronomical Society of the Pacific Conference Series Vol. 407, The Low-Frequency Radio Universe*. p. 384
- Pandey S. B., et al., 2010, *ApJ*, 714, 799
- Pâris I., et al., 2012, *A&A*, 548, A66
- Pawsey J. L., 1951, *Nature*, 168, 358
- Penzias A. A., Wilson R. W., 1965, *ApJ*, 142, 419
- Perley D. A., 2009a, *GRB Coordinates Network*, 9042, 1
- Perley D. A., 2009b, *GRB Coordinates Network*, 10058, 1
- Perley D. A., Bloom J. S., 2008, *GRB Coordinates Network*, 8631, 1
- Perley D. A., et al., 2008, *ApJ*, 688, 470
- Perley R. A., Chandler C. J., Butler B. J., Wrobel J. M., 2011, *ApJL*, 739, L1
- Perryman M. A. C., et al., 1997, *A&A*, 323, L49
- Petroff E., et al., 2015, *MNRAS*, 447, 246
- Pietka M., Fender R. P., Keane E. F., 2015, *MNRAS*, 446, 3687
- Pignata G., et al., 2011, *ApJ*, 728, 14
- Pooley D., et al., 2002, *ApJ*, 572, 932
- Pozzo M., et al., 2006, *MNRAS*, 368, 1169
- Price P. A., et al., 2002a, *ApJL*, 571, L121

- Price P. A., et al., 2002b, *ApJL*, 572, L51
- Price P. A., et al., 2002c, *ApJ*, 573, 85
- Qiu Y., Li W., Qiao Q., Hu J., 1999, *AJ*, 117, 736
- Rahmer G., Smith R., Velur V., Hale D., Law N., Bui K., Petrie H., Dekany R., 2008, in *Ground-based and Airborne Instrumentation for Astronomy II*. p. 70144Y, doi:10.1117/12.788086
- Raiteri C. M., et al., 2011, *A&A*, 534, A87
- Rau A., et al., 2009, *PASP*, 121, 1334
- Ravi V., Shannon R. M., Jameson A., 2015, *ApJL*, 799, L5
- Rebbapragada U., Lo K., Wagstaff K. L., Reed C., Murphy T., Thompson D. R., 2012, in Griffin E., Hanisch R., Seaman R., eds, *IAU Symposium Vol. 285, New Horizons in Time Domain Astronomy*. pp 397–399, doi:10.1017/S1743921312001196
- 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&AS*, 124, 259
- Richards G. T., et al., 2009, *ApJS*, 180, 67
- Richards J. W., et al., 2011, *ApJ*, 733, 10
- Richmond M. W., et al., 1996a, *AJ*, 111, 327
- Richmond M. W., Treffers R. R., Filippenko A. V., Paik Y., 1996b, *AJ*, 112, 732
- Rickett B. J., 1986, *ApJ*, 307, 564
- Rivilla V. M., Chandler C. J., Sanz-Forcada J., Jiménez-Serra I., Forbrich J., Martín-Pintado J., 2015, *ApJ*, 808, 146
- Romero-Cañizales C., et al., 2014, *MNRAS*, 440, 1067
- Roming P. W. A., et al., 2009, *ApJL*, 704, L118
- Rowlinson A., et al., 2016, *MNRAS*, 458, 3506
- Roy R., et al., 2013, *MNRAS*, 434, 2032
- Rumyantsev V., Pozanenko A., 2009, *GRB Coordinates Network*, 9539, 1
- Rumyantsev V., Shakhovkoy D., Pozanenko A., 2010, *GRB Coordinates Network*, 11255, 1
- Rybicki G. B., Lightman A. P., 1986, *Radiative Processes in Astrophysics*
- Ryder S. D., Sadler E. M., Subrahmanyam R., Weiler K. W., Panagia N., Stockdale C., 2004, *MNRAS*, 349, 1093
- Ryle M., Hewish A., 1960, *MNRAS*, 120, 220
- Ryle M., Vonberg D. D., 1946, *Nature*, 158, 339
- Sagar R., et al., 2001, *Bulletin of the Astronomical Society of India*, 29, 91

- Saglia R. P., et al., 2012, *ApJ*, 746, 128
- Salas P., Bauer F. E., Stockdale C., Prieto J. L., 2013, *MNRAS*, 428, 1207
- Salvato M., et al., 2011, *ApJ*, 742, 61
- Scaife A. M. M., Heald G. H., 2012, *MNRAS*, 423, L30
- Schaefer B. E., King J. R., Deliyannis C. P., 2000, *ApJ*, 529, 1026
- Scheers L. H. A., 2011, PhD thesis, University of Amsterdam
- Schinnerer E., et al., 2007, *ApJS*, 172, 46
- Schinnerer E., et al., 2010, *ApJS*, 188, 384
- Schmeer P., Hornoch K., Lehky M., 2004, *IAU Circ.*, 8379, 2
- Schmidt M., 1963, *Nature*, 197, 1040
- Schneider D. P., et al., 2007, *AJ*, 134, 102
- Schneider D. P., et al., 2010, *AJ*, 139, 2360
- Scoville N., et al., 2007, *ApJS*, 172, 1
- Shannon R. M., Johnston S., 2013, *MNRAS*, 435, L29
- Skinner J. N., Thorstensen J. R., Lépine S., 2014, *AJ*, 148, 115
- Skrutskie M. F., et al., 2006, *AJ*, 131, 1163
- Slee O. B., Willes A. J., Robinson R. D., 2003, *PASA*, 20, 257
- Soderberg A. M., 2005, *GRB Coordinates Network*, 3187, 1
- Soderberg A. M., Kulkarni S. R., Berger E., Chevalier R. A., Frail D. A., Fox D. B., Walker R. C., 2005, *ApJ*, 621, 908
- Soderberg A. M., et al., 2006a, *Nature*, 442, 1014
- Soderberg A. M., et al., 2006b, *ApJ*, 650, 261
- Soderberg A. M., Chevalier R. A., Kulkarni S. R., Frail D. A., 2006c, *ApJ*, 651, 1005
- Soderberg A. M., et al., 2007, *ApJ*, 661, 982
- Soderberg A., Chandra P., Frail D., 2008, *GRB Coordinates Network*, 7506
- Soderberg A. M., et al., 2010a, *Nature*, 463, 513
- Soderberg A. M., Brunthaler A., Nakar E., Chevalier R. A., Bietenholz M. F., 2010b, *ApJ*, 725, 922
- Sokolov V. V., Kopylov A. I., Zharikov S. V., Feroci M., Nicastro L., Palazzi E., 1998, *A&A*, 334, 117
- Spitler L. G., et al., 2014, *ApJ*, 790, 101

- Spitler L. G., et al., 2016, *Nature*, 531, 202
- Spitzer L., 1998, *Physical Processes in the Interstellar Medium*
- Spreeuw J. N., 2010, PhD thesis, University of Amsterdam
- Staley T. D., et al., 2013, *MNRAS*, 428, 3114
- Stappers B. W., et al., 2011, *A&A*, 530, A80
- Steele I. A., et al., 2004, in Oschmann Jr. J. M., ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 5489, Ground-based Telescopes*. pp 679–692, doi:10.1117/12.551456
- Stockdale C. J., et al., 2002, *IAU Circ.*, 8018, 1
- Stockdale C. J., Sramek R. A., Weiler K. W., van Dyk S. D., Panagia N., Pooley D., Lewin W., Marcaide J. M., 2004a, *IAU Circ.*, 8379, 1
- Stockdale C. J., Weiler K. W., van Dyk S. D., Sramek R. A., Panagia N., Marcaide J. M., 2004b, *IAU Circ.*, 8415, 1
- Swinbank J. D., et al., 2015, *Astronomy and Computing*, 11, 25
- Tasse C., van der Tol S., van Zwieten J., van Diepen G., Bhatnagar S., 2013, *A&A*, 553, A105
- Taylor A. R., Gregory P. C., 1983, *AJ*, 88, 1784
- Taylor G. B., Frail D. A., Kulkarni S. R., Shepherd D. S., Feroci M., Frontera F., 1998, *ApJL*, 502, L115
- Thornton D., et al., 2013, *Science*, 341, 53
- Thyagarajan N., Helfand D. J., White R. L., Becker R. H., 2011, *ApJ*, 742, 49
- Tingay S. J., et al., 2013, *PASA*, 30, e007
- Tonry J. L., et al., 2012, *ApJ*, 745, 42
- Totani T., 2013, *PASJ*, 65, L12
- Tsvetkov D. Y., 1987, *Soviet Astronomy Letters*, 13, 376
- Turatto M., Cappellaro E., Danziger I. J., Benetti S., Gouiffes C., della Valle M., 1993, *MNRAS*, 262, 128
- Udike A. C., et al., 2008, *ApJ*, 685, 361
- Vanden Berk D. E., et al., 2001, *AJ*, 122, 549
- Veillet C., 2000, *GRB Coordinates Network*, 588
- Villata M., et al., 2009, *A&A*, 501, 455
- Waltman E. B., Ghigo F. D., Johnston K. J., Foster R. S., Fiedler R. L., Spencer J. H., 1995, *AJ*, 110, 290
- Watson D., et al., 2006, *ApJ*, 652, 1011

- Weiler K. W., Sramek R. A., Panagia N., van der Hulst J. M., Salvati M., 1986, *ApJ*, 301, 790
- Weiler K. W., Panagia N., Montes M. J., Sramek R. A., 2002, *ARA&A*, 40, 387
- Weiler K. W., Williams C. L., Panagia N., Stockdale C. J., Kelley M. T., Sramek R. A., Van Dyk S. D., Marcaide J. M., 2007, *ApJ*, 671, 1959
- Weiler K. W., Panagia N., Stockdale C., Rupen M., Sramek R. A., Williams C. L., 2011, *ApJ*, 740, 79
- Wellons S., Soderberg A. M., Chevalier R. A., 2012, *ApJ*, 752, 17
- Wendker H. J., 1995, *A&AS*, 109, 177
- White S. M., 2000, in Finley D. G., Goss W. M., eds, *Radio interferometry : the saga and the science*. p. 86
- White R. L., Becker R. H., Helfand D. J., Gregg M. D., 1997, *ApJ*, 475, 479
- Wijnholds S. J., Grobler T. L., Smirnov O. M., 2016, *MNRAS*, 457, 2331
- Wilkinson P. N., 1991, in Cornwell T. J., Perley R. A., eds, *Astronomical Society of the Pacific Conference Series Vol. 19, IAU Colloq. 131: Radio Interferometry. Theory, Techniques, and Applications*. pp 428–432
- Williams P. K. G., Berger E., 2016, *ApJL*, 821, L22
- Williams C. L., Panagia N., Van Dyk S. D., Lacey C. K., Weiler K. W., Sramek R. A., 2002, *ApJ*, 581, 396
- Williams P. K. G., Bower G. C., Croft S., Keating G. K., Law C. J., Wright M. C. H., 2013, *ApJ*, 762, 85
- Williams P. K. G., Berger E., Chornock R., 2016, *The Astronomer's Telegram*, 8752
- Wright A., Otrupcek R., 1990, in *PKS Catalog (1990)*. p. 0
- Wright E. L., et al., 2010, *AJ*, 140, 1868
- Yatawatta S., et al., 2013, *A&A*, 550, A136
- Yokoo T., Arimoto J., Matsumoto K., Takahashi A., Sadakane K., 1994, *PASJ*, 46, L191
- York D. G., et al., 2000, *AJ*, 120, 1579
- Young D. R., et al., 2010, *A&A*, 512, A70
- Zarka P., Treumann R. A., Ryabov B. P., Ryabov V. B., 2001, *Ap&SS*, 277, 293
- Zauderer B. A., et al., 2011, *Nature*, 476, 425
- Zhang B., 2014, *ApJL*, 780, L21
- Zwart J. T. L., et al., 2008, *MNRAS*, 391, 1545
- van Dyk S. D., Sramek R. A., Weiler K. W., Panagia N., 1993, *ApJ*, 409, 162
- van Dyk S. D., Sramek R. A., Montes M. J., Weiler K. W., Panagia N., 1996, *IAU Circ.*, 6528, 1

-
- van Haarlem M. P., et al., 2013, A&A, 556, A2
- van Velzen S., et al., 2016, Science, 351, 62
- van Weeren R. J., et al., 2016, ApJS, 223, 2
- van der Horst A. J., Kamble A. P., Wijers R. A. M. J., Kouveliotou C., 2009, GRB Coordinates Network, 9883, 1
- van der Horst A. J., Wiersema K., Kamble A. P., Wijers R. A. M. J., Rol E., Kouveliotou C., 2010, GRB Coordinates Network, 11221, 1
- van der Horst A. J., et al., 2011, ApJ, 726, 99
- van der Laan H., 1966, Nature, 211, 1131
- van der Tol S. ., Jeffs B. D., van der Veen A.-J. ., 2007, IEEE Transactions on Signal Processing, 55, 4497