


Astrophysical Radio Transients: Surveys and Techniques

Keith Bannister

Supervised by:
B. M. Gaensler
T. Murphy
T. J. Cornwell



THE UNIVERSITY OF
SYDNEY

*A thesis submitted for the degree of Doctor of
Philosophy at the University of Sydney*

November 2011

Pour Michelle, Arianne et Soren.

Abstract

I present two pilot surveys and a range of new algorithms to aid in planning and implementing wide-field radio surveys for transient and variable sources. The first pilot survey is a blind survey for transient and variable radio sources with the Molonglo Observatory Synthesis Telescope at 843 MHz. This survey discovered 53 highly variable sources and 15 transient sources and determined surface densities of 0.268 deg^{-2} for variable sources and $1.3 \times 10^{-2} \text{ deg}^{-2}$ for transient sources, with associated timescales of between one day and three years. I also propose two new techniques which were developed for this survey. The first technique provides a post-imaging calibration for image gain. The second technique is a statistical method for verifying whether flux error measurements agree with the scatter in light curves over a population of sources.

I also describe a follow-up survey for prompt radio emission from gamma ray bursts (GRBs) at 1.4 GHz. This survey used a single-dish telescope to automatically slew to a GRB position within 2 min of the gamma ray trigger and performed single-pulse, repeating and low-time resolution searches for variability. This survey discovered single, dispersed pulses following two long GRBs, which are possibly related to the delayed formation of a black hole at the centre of the GRBs. The high-time resolution measurements from this survey are some of the most constraining limits on prompt radio emission from GRBs to date.

I also present two efficient new algorithms for detecting dispersed radio emission in interferometric data: the Chirpolator and the Chimageator. These two techniques excel in the regime of sparse arrays, where they both require substantially lower data rates, and the Chirpolator requires a much lower post-integrator operation rate than the existing algorithms. These techniques are well matched to future supercomputing architectures, where the arithmetic capability is outstripping the bandwidth capability, and are therefore suitable for use by interferometer-based fast transient surveys.

The surveys and techniques described in this thesis will have immediate application to upcoming surveys with the next generation of wide-field radio telescopes, such as the two transients surveys proposed for the Australian Square Kilometre Array Pathfinder.

Acknowledgements

This thesis has had contributions from so many people, it is impossible to list them all.

To my supervisors, I cannot express my thanks enough. I feel truly enriched to have had such good relationships with such impressive people.

Bryan, your enthusiasm, skill and work ethic are awe-inspiring. You have showed me the depth and breadth of what is possible in astronomy, offered outstanding advice with humility and influenced so many parts of my life during this thesis: from the data reduction and statistics, to physics, proof-reading and career advice. You recently sent me an email saying “family comes first”, which beautifully illustrates your support for my family over the last three and a half years. I wish you all the very best with your family and with CAASTRO and I look forward to working alongside you in the next few years.

Tara, your mentoring has been a source of great strength for me. Every time I meet with you I have come away even more impressed with your professionalism, expertise and ability to inspire and lead. You have been so patient with my occasional crises of confidence and put me gently and firmly on the right track. You are a truly great mentor and leader and I hope a great many more PhD students will benefit from your talents. I wish you all the very best with VAST and I look forward to working closely with you in the future.

Tim, your mathematical and physical insights are astounding. So many times over the past three and a half years I have walked into your office and walked out with a whole new view of mathematics and interferometry. You have been so patient with my small steps into this mathematical world and given direction with great humour. Wandering into your office and scrawling on your whiteboard has been one of the joys and delights of this PhD. I wish you the all the very best with ASKAP and look forward to cooking up some new ideas from down the corridor.

To my fellow students, your comradeship has been another highlight of the last three and a half years. Elizabeth & Kitty, it has been a great pleasure sharing an office with you. You have joyfully put up with my occasional losses

of coherence at the keyboard, and always been willing to go for an emergency coffee or hot chocolate. You are both inspiring people and I will watch your careers with great interest. Thanks go to the other students with whom I have shared great conversation, warmth, Thai and computer tricks. You all been a great source of encouragement for me and been so generous with answering my questions. Special thanks go to Anant Tanna, Chris Hales, Juliana Kwan, Holly Trowland, Ben Jelliffe, Kate Randall, Scott Wales, Hugh Garsden, Billy Robbins, Madhura Killedar and Matt Francis and Yu Meng for being so helpful and fun to study with.

To the staff and postdocs at the Sydney Institute for Astronomy (SIfA), you set an excellent standard of professionalism and accomplishment and have set a wonderful collegiate atmosphere. Special thanks go to Paul Hancock for being my number one support for ATCA and WiFeS observations. I have greatly enjoyed our two trips and you have been an observing buddy of the highest calibre. I'm looking forward to working more with you on VAST projects. Special thanks go to Dick Hunstead, Elaine Sadler, Scott Croom and Shami Chatterjee with whom I have enjoyed great discussions. You have all been excellent advisers, always ready to listen, talk and help, and always with good humour. I have also very much enjoyed working with the SIfA postdocs and staff: Greg Madsen, Lisa Harvey-Smith (now at CSIRO Astronomy and Space Sciences; CASS), Jay Banyer and Tim Robishaw to name but a few. Finally, Jo Daniels, you've been very helpful and always kept me on my toes. I don't know anyone with a bigger heart.

To my colleagues at CASS, you are a special group of people in a very special place. John Reynolds deserves special mention. Without you and your judicious use of the Sydney Morning Herald rule, the 12 m project would never have been possible. I also must thank Brett Preisig and Ian McRobert for supporting the 12 m operations at Parkes. Thanks for your time and perseverance. I have enjoyed many great conversations on all sorts of topics with a very stimulating group of people including: Ben Humphries, Tony Maher, Malte Maquarding, Emil Lenc, Ron Ekers, Simon Johnston, Baerbel Koribalksi, George Hobbs, Ray Norris, Dick Manchester, Chris Phillips, Ryan Shannon, Mike Keith, Mike Kesterven, John Tuthill, Paul Roberts, Aaron Chippendale and Tim Bateman. Thanks for answering my often naïve questions with such enthusiasm and forbearance.

To my colleagues from other institutes, you have been invaluable during the last few years. Willem van Straten helped setup the IBOB backend for the 12 m and endured many questions with good humour, often aided by Andrew Jameson. Thanks to Matthew Bailes and the HTRU team for generously lending me an IBOB for the 12 m experiment. Thanks to Steven Tingay, Ramesh Baht, Jayanta Roy and Peter Cox who supplied Giant Metrewave Radio Telescope (GMRT) data and help with understanding the GMRT data format. Sadly, those data didn't make it into the thesis but there will be another chance for that soon.

To the Australian Taxpayer for funding Australian astronomy, and The University of Sydney and CSIRO/CASS for being such a wonderful nursery for students. I look forward to giving back what you invested in me, many times over.

To my friends, who have been with me in person or in spirit over the last few years, I couldn't have done it without you. Marcella Kirby, Kev Smith, Amy Allen, Matthew and Jane Holland, you believed in me during my wilderness. I'm out now, and it's beautiful here. Andre and Steph Kyme, without your generous gift of the house at Forestville, we would've really struggled. Thanks to Pete & Eugenie Ree for being patient with me! Mikey Bourne, thanks for hanging out on Saturdays with bacon and egg rolls. You've been a rock. Ricky, those dinners at Blues Point Road were what life is all about. To the big Wednesday crowd, Ger, Jules, Kim, Cazz, Chris, Logan and Elizabeth, you are an oasis in my week. Eating and drinking with you fills my soul.

To the Upper Roomers who have been so much a part of my joy and freedom in the last few years. I'm loving the ride and thanks for sharing in it. Mick & Marg Martin, your wisdom, humour and humility are inspiring and a joy to be around. Matt & Julie Cumberworth, your insightful words put me on a new path, and have totally come true. Aiya Saito, your gentleness is still a source of peace for me. It's been great to share life with so many beautiful people: Chris & Mel, Matt & Ngaire, Pete & Ann, Simon & Kath, Clive & Emma, Pete & Merideth, Pete & Marg, Pru and Dee, just to name but a few. I'm sorry I don't have the space to list you all!

To my family, your support has been rock-solid and beautiful. Without your help we would have all gone mad. Maman, merci pour l'amour inébranlable que tu me portes et pour avoir gardé si gentiment et à multiples reprises nos enfants. Tu sais qu'ils t'adorent. Leur "Kiki" c'est la meilleure au monde! Dad, your humble words of encouragement are always a comfort and discussing maths with you has always been so enjoyable. Thanks for supporting Mum with all this! I know you'd like to see more of her but I know you get a lot of joy out of *zi children* too. Coco, dammit you're *funny*. I love you to bits. You've helped me keep it real over the last few years. Thanks.

To Mich's family, you have kept us sane over the ups and downs. Without you it would have all crumbled. Nat-Nat, you're the biggest hearted, generous and gutsy person I know. Wonderful. So many thanks to Nanna Dee, Mai-Mai, Brett, Leith, Sam & Elliot for helping us keep it together.

To my wife, Michelle, thanks for having the courage to take this journey with me. You told me to follow my dreams and believed it would be OK. You were right. You're amazing. I love you.

Pour mes beaux enfants, Arianne et Soren. Vous me comblez de joie quotidiennement. Vous m'avez appris à vivre avec plus de sagesse tout en me rapelant ce qui importe le plus dans la vie: vous. J'ai hâte de savoir ce qui va vous motiver à l'avenir.

To my Lord Jesus, wild ride eh? I can't believe how good the last few years have been. Wow. We have had *everything* we needed, in abundance. It's humbling and beautiful. Thank you.

These last few pages have been the easiest to write of this entire PhD. It has been such a privilege to do such a thing and have so many people to thank for it. Thank you. Thank you so much.

This research has made use of several facilities and tools including:

1. The MOST, which is operated by the School of Physics, University of Sydney with support from the Australian Research Council.
2. The NASA/IPAC Extragalactic Database (NED) which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.
3. The SIMBAD database, operated at CDS, Strasbourg, France.
4. SAOImage DS9, developed by the Smithsonian Astrophysical Observatory.
5. The Aladin sky atlas (Bonnarel et al., 2000).
6. Data obtained through the High Energy Astrophysics Science Archive Research Center Online Service, provided by the NASA/Goddard Space Flight Center.
7. Data obtained from the Leicester Database and Archive Service at the Department of Physics and Astronomy, Leicester University, UK.
8. The Q3C spatial indexing extension for Postgresql by Koposov & Bartunov (2006).
9. The Cosmology Calculator (Wright, 2006).
10. The PSRCHIVE software suite (Hotan et al., 2004).

The author acknowledges the support of an Australian Postgraduate Award and a CSIRO top-up scholarship. This research has been supported by the Australian Research Council through grant DP0987072.

Statement of Originality

All work performed for this thesis and resulting publications are entirely my own work except as follows.

Chapter 2 : Dick Hunstead supplied the MOST data and provided technical insight into the operations of the MOST, especially the calibration procedures. Barbara Piestrzynski maintained the MOST archive. Duncan Campbell-Wilson supplied some important technical information and maintains the MOST. Elaine Sadler provided useful comments on optical counterparts and Greg Madsen provided help with putting axes on Figure 2.2. The anonymous referee for the resulting paper (Bannister et al., 2011) also provided some extremely useful comments. The co-authors for the paper, Bryan Gaensler, Tara Murphy, Dick Hunstead and Shami Chatterjee, provided general advice on the analysis of the data and comments on the paper drafts.

Chapter 3 : John Reynolds refurbished the ‘bucket feed’, supplied a networking adaptor, and provided a great deal of advice. Brett Preisig and Ian McRobert provided technical support at Parkes for the 12 m antenna. The bucket feed was installed by the staff at Parkes. The pulsar back end (IBOB) was borrowed from Matthew Bailes and the HTRU team (Keith et al., 2010). Willem van Straten and Andrew Jameson provided support for the IBOB and PSRCHIVE software package. Sarah Burke-Spolaor and Lina Levin supplied advice on the HITRUN scripts and processing pipeline. Chris Phillips supplied computers at Parkes and support for those computers.

Chapter 4 : The original idea of gridding voltages (an idea which is related to, but not exactly incorporated in the Chimageator) was first proposed to me by Tim Cornwell. I discovered the Chimageator imaging optimisations and derived the resulting performance equations. The co-author,

Tim Cornwell, provided general advice in formulating the techniques and comments on the paper drafts.

Publications in this thesis

Chapter 2 along with appendices A and B have been published as Bannister et al. (2011).

Chapter 4 along with appendices C and D have been published as Bannister & Cornwell (2011).

Contents

Contents	xiv
1 Introduction	1
1.1 A brief history	1
1.1.1 Monitoring and the discovery of scintillation	1
1.1.2 Serendipity and the dawn of pulsar astronomy	2
1.1.3 Follow-up of transients detected at other wavelengths	2
1.1.4 Blind surveys	3
1.2 The unexplored parameter space	4
1.3 Key questions in the era of wide-field radio astronomy	7
1.4 Sources of transient and variable radio emission	10
1.4.1 Scintillation of compact radio sources	10
1.4.2 Radio supernovae and gamma-ray bursts	11
1.4.3 Pulsars	11
1.4.4 X-ray binaries	13
1.5 Radio Transient Surveys	14
1.5.1 Relationships between source and survey parameters	14
1.5.2 Sensitivity	15
1.5.3 Coverage	17
1.5.4 Cadence	17
1.5.5 The interplay of survey parameters	18
1.5.6 The role of pilot surveys	19
1.6 Techniques	20
1.6.1 Detecting slow transients with interferometers	20
1.6.2 Detecting dispersed transients	20
1.7 Purpose & rationale of this thesis	24
1.7.1 Pilot surveys	25
1.7.2 New algorithms	26
1.7.3 Summary	26

2	Transients and variables with MOST	27
2.1	Introduction	28
2.2	Observations with MOST	30
2.2.1	Observing modes	31
2.2.2	Archive coverage	31
2.2.3	Image artefacts and systematic errors	37
2.3	Light curve extraction	40
2.3.1	Source detection	41
2.3.2	Source association	41
2.3.3	Measurement	41
2.3.4	Post-facto calibration	42
2.3.5	Light curve characterisation	46
2.3.6	List of variable and transient sources	50
2.4	Results	50
2.4.1	Data Quality	51
2.4.2	Variable sources	55
2.4.3	Transient sources	61
2.4.4	Optical Counterparts	61
2.4.5	Gamma Ray Burst Counterparts	62
2.5	Discussion	63
2.5.1	Overall statistics	63
2.5.2	Variability time-scales	63
2.5.3	Source counts for variable sources	65
2.5.4	Interpretation as interstellar scintillation	66
2.5.5	Source counts for transient sources	68
2.6	Discussion of selected sources	69
2.6.1	Variable sources with resolved optical counterparts	69
2.6.2	A compact steep spectrum source	72
2.6.3	Transient sources with resolved optical counterparts	75
2.6.4	Transient sources with point-like optical counterparts	79
2.7	Conclusions	82
3	Automated GRB follow-up at 1.4 GHz	83
3.1	Introduction	84
3.2	Observations	86
3.2.1	Telescope Configuration	86
3.2.2	GRB response	87
3.2.3	Processing	87
3.3	Results	90
3.3.1	Events	90
3.3.2	Single-pulse search	92
3.3.3	Periodicity search	96
3.3.4	Low-time resolution search	96
3.4	Discussion	99

3.4.1	Single-pulse detections	99
3.4.2	Pulsar detection	117
3.4.3	Long time integrations	118
3.5	Conclusion	118
4	The Chirpolator & The Chimageator	121
4.1	Introduction	122
4.1.1	Scientific Motivation	122
4.1.2	Improving pulsar and single-pulse surveys	122
4.1.3	The problem: processing requirements	123
4.1.4	Two new techniques	124
4.2	Background	125
4.2.1	Dispersion in the Interstellar Medium	125
4.2.2	A taxonomy of methods	127
4.2.3	Dedispersion	129
4.2.4	Array processing	130
4.3	The Chirpolator	131
4.3.1	Intuitive Description	131
4.3.2	Two-antenna case	132
4.3.3	Multiple telescopes in 1D	136
4.3.4	Compensating for the smearing in $D_N(x)$	137
4.4	The Chimageator	138
4.4.1	Intuitive Description	138
4.4.2	Formulation for a uniform linear array	139
4.5	Method of comparison	141
4.5.1	A simple model for evaluating algorithm cost	143
4.5.2	Array, survey, and algorithm parameters	144
4.6	Results	144
4.7	Discussion	146
4.7.1	Further work	146
4.8	Summary	150
5	Conclusion	153
5.1	Pilot surveys	154
5.2	Techniques	154
5.3	The Future	155
5.4	Summary	156
	Bibliography	157
A	Variables from the MOST Survey	169
B	Transients from the MOST survey	175
C	The Chirpolator: analysis and implementation	181

C.1	Multiple telescopes in 3D	182
C.2	Non linear dispersion	182
C.2.1	The Problem: Decoherence in the DFT Bins	182
C.2.2	The solution: Oversampling	184
C.3	Implementation optimisations	186
C.3.1	Compute only DFTs Required on a baseline basis	187
C.3.2	Efficient calculation of $X_{pq}[k]$ with sliding DFTs	188
C.3.3	Factorising the DFTs	190
C.3.4	Efficiently computing negative DFT bins	190
C.4	Performance	192
C.4.1	Resolution	192
C.4.2	Operation rates	193
C.4.3	Data rates	195
D	The Chimageator: analysis and implementation	197
D.1	Non linear dispersion	198
D.2	Implementation optimisations	198
D.2.1	Optimising operations in the imaging step	198
D.2.2	Sampling	199
D.2.3	Spectral smearing	199
D.3	Performance	200
D.3.1	Resolution	200
D.3.2	Operation rates	201
D.3.3	Data rates	203

Introduction

1.1 A BRIEF HISTORY

A fascination with the visible sky is as old as humanity, and changes in the sky have been keenly observed. While many cultures followed the movements of the moon and planets, arguably the most famous early example of human interest in the dynamic sky can be found in the multiple Chinese and Arab recordings of the supernova (SN) 1054. SN1054 was a core-collapse supernova that was, at its brightest, visible by eye during the day, and produced what we know as the Crab nebula and pulsar. Multiple documents describing this event have enabled historians to accurately date the supernova and allowed modern astronomers to test important models relating to the Crab nebula and pulsar (Hester, 2008).

The modern age of astronomy began when Galileo turned his refracting optical telescope to the heavens in 1609. With this new telescope technology, astronomers studied a range of recently discovered variable phenomena such as novae, Mira variables, and eclipsing binaries (M. Hoskin, 1997, p. 198). Radio astronomy arrived much later in history, when Karl Jansky discovered radio emission from the Milky Way at 20.5 MHz in 1931 (Jansky, 1933). Radio astronomy was enthusiastically pursued after World War II, due to the availability of surplus military equipment and the desire to look into this never-before used window onto the universe: the radio sky.

Since Jansky's pioneering efforts, radio astronomers have used monitoring, serendipity, follow-up of events from other wavelengths and blind surveys to explore the transient radio sky.

1.1.1 Monitoring and the discovery of scintillation

In 1946, the first variations in discrete radio sources were detected at a wavelength of 5 m using surplus military equipment (Hey et al., 1946). These variations were subsequently attributed to scintillation in the ionosphere, much like

the twinkling of stars is due to scintillation in the Earth's atmosphere. Scintillation requires that sources be compact (which is why planets don't twinkle) and it was through these observations that the relatively small angular size of these radio sources was first determined (Kellermann & Pauliny-Toth, 1968). By 1968, variability had been reported in the radio brightness of the Sun, Jupiter, the supernova remnant Cassiopeia and some quasars (Kellermann & Pauliny-Toth, 1968). All these variations were explained by scintillation in the ionosphere, solar wind, or physical processes at the source. However, the discovery of relatively rapid radio variations in quasars at 408 MHz by Hunstead (1972) led to a perplexing problem. The speed of the variations implied a small causative region due to light travel time, which, when combined with the flux density and known redshift, implied a brightness temperature above the 10^{12} K inverse-Compton limit (Kellermann & Pauliny-Toth, 1969). This situation went unresolved for a number of years, until Shapirovskaya (1985) and Rickett (1986) proposed an extrinsic cause: scintillation due to inhomogeneities in the ionised gas in the interstellar medium (ISM) of the Milky Way.

1.1.2 Serendipity and the dawn of pulsar astronomy

As with so much in science, one of the most important detections of variability was a result of serendipity. While performing a radio survey at the Mullard Radio Astronomy Observatory, Jocelyn Bell noticed 'scruff' on the pages of a chart recorder connected to the telescope (Bell Burnell, 1977). This 'scruff', dubbed 'Little Green Man – 1', appeared at the same sidereal time each day, implying it was astronomical rather than terrestrial. Bell, her supervisor Antony Hewish, and their collaborators then began a long tradition in pulsar astronomy, of accurately measuring the period of the 'rapidly pulsating radio source' as 1.33733 ± 0.00001 s. They proposed a neutron star origin (Hewish et al., 1968), which was soon confirmed by discovery of single pulses from the Crab pulsar (Hewish, 1975; Staelin & Reifenstein, 1968).

1.1.3 Follow-up of transients detected at other wavelengths

Wide-field capabilities of observatories operating at optical, X-ray and gamma-ray wavelengths have led to detections of a range of transient and variable radio sources. Gregory & Kronberg (1972) discovered a giant radio outburst while observing the X-ray microquasar Cygnus X-3. This outburst is thought to have been the result of jets forming from accretion onto a black hole. Follow-up of core-collapse supernovae detected by optical observatories (Weiler et al., 2002, 2004), and of gamma-ray bursts detected by gamma-ray observatories (Frail et al., 2003), have yielded a number of radio afterglows originating in supernova shocks. Similarly, follow-up of X-ray flares from magnetars have detected a number of radio afterglows (Cameron et al., 2005; Frail et al., 1999;

Gaensler et al., 2005), which are thought to be associated with relativistically ejected debris.

1.1.4 Blind surveys

A blind survey is a survey in which a patch of sky is monitored for transient and variable sources. A number of blind surveys for transient and variable radio sources have been performed, albeit over relatively small fields, or with poor sensitivities.

Perhaps the first moderately wide-field survey results were obtained by Gregory & Taylor (1986), who discovered a number of flaring radio stars in a survey of the Galactic plane. Relatively few new transient sources were discovered by blind surveys until Hyman et al. (2005) discovered a high brightness temperature transient source (GCRT J1745-2009) with a period of 77 min (Hyman et al., 2007) during a moderately wide-field monitoring campaign of the Galactic Centre. This mysterious object had unusual properties, and the authors concluded that these were unlike those of any previously known source class. Infrared observations appeared to rule out a nearby ultra-cool dwarf star (Kaplan et al., 2008), while Roy et al. (2010) detected circular polarisation in a reanalysis of the discovery observation, implying that the radio transient could have been a result of electron cyclotron or plasma emission from a magnetically dominated dwarf. The ultimate origin of this transient is still unknown.

In terms of potentially extragalactic transients, the field was revolutionised by the results of Bower et al. (2007), who with 944 images of a single field, discovered ten faint transient sources, eight of which had no optical counterparts, with the remaining two thought to be radio supernovae. Ofek et al. (2010) proposed a neutron star origin for these transients by applying population arguments. Attempts to identify these transients are ongoing (Croft et al., 2011b).

Other surveys have been useful in setting limits in various regimes, but most transient sources have not been conclusively identified. In a very wide-field survey, Levinson et al. (2002) covered a large fraction of the sky with just two epochs by comparing the results of two large Very Large Array (VLA) surveys (Becker et al., 1995; Condon et al., 1998). They discovered no transients but set constraints on the GRB beaming fraction. Gal-Yam et al. (2006) performed a similar analysis, discovering two radio supernovae and a transient source without an optical host, and derived a more stringent limit on the GRB beaming fraction. Croft et al. (2010) set an upper limit on the surface density of 1.4 GHz transients with flux densities greater than 40 millijansky (mJy; $1 \text{ Jy} = 10^{-26} \text{ Wm}^{-2}\text{Hz}$) of 0.004 deg^{-2} , and Croft et al. (2011a) were able to rule out the existence of 1 Jy transients that had been previously reported by Matsumura et al. (2009). The Croft et al. (2011a) results appear to be at odds with a VLA survey at 5 GHz by Ofek et al. (2011), who measured a much

higher surface density of 0.039 deg^{-2} at a lower flux density (1.4 mJy). Becker et al. (2010) found 39 variable sources in the Galactic plane with 3 epochs of 5 GHz VLA observations; most of these sources had no known counterparts at other wavelengths. A number of recent surveys have either yielded no transient detections (Bell et al., 2011; Bower & Saul, 2011; Bower et al., 2011, 2010; Lazio et al., 2010), or none that have been conclusively identified (Ofek et al., 2011)

1.2 THE UNEXPLORED PARAMETER SPACE

Given the relative success of serendipity, monitoring, follow-up and moderately wide-field observations, is there any value in continuing these efforts in a wide-field regime? Fortunately, there are a great number of new discoveries to be made, because the parameter space for radio transients is still poorly explored. To illustrate this parameter space argument, I follow the arguments of Cordes et al. (2004).

We can formulate a phase space for radio transients starting with the Rayleigh-Jeans approximation for a source’s radio emission. We consider a source with brightness temperature T , at a distance D . If the source varies on a timescale W , then the angular size of the source is limited by the light travel time implied by W , which constrains W to:

$$W^2 = \frac{1}{2\pi k} \frac{SD^2}{T} \frac{1}{\nu^2}, \quad (1.1)$$

where S is the flux density of the source, ν is the emission frequency and k is Boltzmann’s constant. Figure 1.1 is a plot of the parameter space described by Equation 1.1, populated with known transient sources. Notable features of Figure 1.1 are that the axes span many orders of magnitude, and that there are large regions of the parameter space that are empty or sparsely populated. The emptiness in this plot is because either nature has not populated this parameter space very evenly, or because our observations have not been sensitive to large parts of this parameter space.

To determine how much of this parameter space is yet to be explored, we need only to compare the capabilities of the surveys which discovered the known sources, with the capabilities of surveys with future instruments. One of the most capable upcoming instruments is the Australian Square Kilometre Array Pathfinder (ASKAP), a 36-antenna interferometer with a 30-square-degree field of view and a frequency range of 700–1800 MHz (Johnston et al., 2008). In Figure 1.2, we compare surveys for slow transients in the parameter space of snapshot rate¹ vs. flux density for a range of existing and future surveys. The surveys shown have been performed with many hundreds of hours of observing time using existing state-of-the-art telescopes. In just 1 hr

¹A strict definition of snapshot rate is given in Section 1.5.3.




Figure 1.1: Phase space for radio transients, adapted from Cordes et al. (2004). Red lines indicate the sensitivity of ASKAP with a 1 s integration time, to sources at distances of 10 pc, 1 kpc and 1 Mpc. Diagonal lines indicate lines of constant brightness temperature. Incoherent synchrotron emission has a brightness temperature $T < 10^{12}$ K. Brightness temperatures higher than this threshold imply coherent processes. A number of source classes are indicated, including intra-day variables (IDV), gamma-ray bursts (GRBs), active galactic nuclei (AGN) and the Galactic Centre (GC). The emission mechanisms for IDV, GRBs and AGN are not intrinsically coherent, but appear so because of variations due to interstellar scintillation. Figure courtesy of D. Kaplan, S. Chatterjee and T. Murphy.

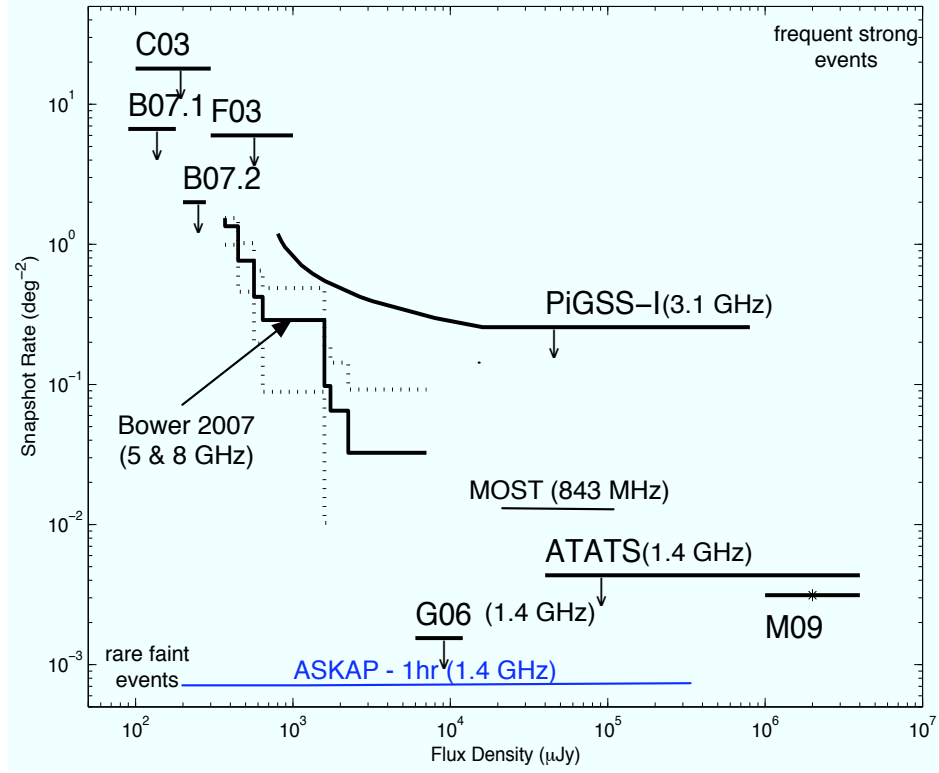


Figure 1.2: Snapshot rate vs. flux density limits and detections for radio transients with timescales longer than about 1 day. Solid lines without arrows are transient detections and solid lines with arrows are upper limits. The blue line indicates the parameter space probed by ASKAP with 1 hr of observing time. A number of surveys are shown, including Carilli et al. (2003) (C03), Frail et al. (2003) (F03), Croft et al. (2010) (ATATS), Gal-Yam et al. (2006) (G06), and Matsumura et al. (2009) (M09). B07.1 and B07.2 refer the results of Bower et al. (2007) in 1 and 2 month averages respectively. The MOST results are reported in Chapter 2 of this thesis. Figure modified from Bower et al. (2010)

of observing, ASKAP will probe deeper into this parameter space in terms of snapshot rate, at almost equivalent sensitivity, than any previous survey. ASKAP will therefore be able to find weaker and rarer transients than is currently possible now.

A similar argument can be made for the capabilities of the telescopes themselves. Figure 1.3 depicts the capabilities of a number of radio telescopes according to their sensitivities and instantaneous fields of view. While traditional single-dish telescopes lie on the diagonal line, a number of existing

multi-beam systems and interferometers have probed further into this parameter space. However, the performance of new wide-field telescopes will easily exceed these existing instruments. For example, ASKAP in a coherent mode will have equivalent sensitivity to the Parkes 64 m but a larger field of view by almost two orders of magnitude, and will therefore be able to search for the much rarer events than previously possible.

As a result, new wide-field telescopes such as ASKAP will have the potential to investigate new parts of transient phase space. If nature has populated those parts of parameter space with as yet undiscovered phenomena, these telescopes will have the potential to uncover them.

1.3 KEY QUESTIONS IN THE ERA OF WIDE-FIELD RADIO ASTRONOMY

Radio astronomy is on the cusp of a new wide-field era. Upcoming wide-field radio telescopes will provide an unprecedented view of transient and variable radio phenomena, providing a regular Cosmic census that has been technologically impossible until recently. Telescopes such as ASKAP, AperTIF (Verheijen et al., 2008), LOFAR (Röttgering, 2003), the Murchison Widefield Array (MWA; Morales, 2005), and the Low Frequency Demonstrator Array (Lazio et al., 2010) are at the forefront of this research.

When planning surveys with these telescopes, a key question is how best to allocate the survey time, given the telescope’s technical capabilities and available observing time. One cannot have unlimited sensitivity, coverage and cadence simultaneously (see Section 1.5.5), meaning survey designers must sacrifice some of these parameters in favour of others. In order to make an informed choice of parameters, it is therefore useful to have real-world measurements of the statistics of transient sources, so that reasonable initial values for these parameters can be chosen.

A common feature of these telescopes is the enormous data volumes that they produce, which in turn require massive supercomputing facilities. The ‘data deluge’ from these instruments poses a substantial problem because traditional techniques do not scale up to such large data volumes. This scaling problem is both because the traditional techniques often require human intervention and because the algorithms themselves are not well matched to implementation on supercomputing facilities.

To address the need to reduce human intervention, survey teams are using two approaches. The first is to implement completely automatic processing pipelines. In order to achieve this, the pipelines need features that allow them to deal with real-world data issues, such as image artefacts, which have traditionally been handled by human intervention. The second approach is to use automatic classification engines based on machine learning, to classify candidates that are detected by the pipelines (e.g. Richards et al., 2011). These

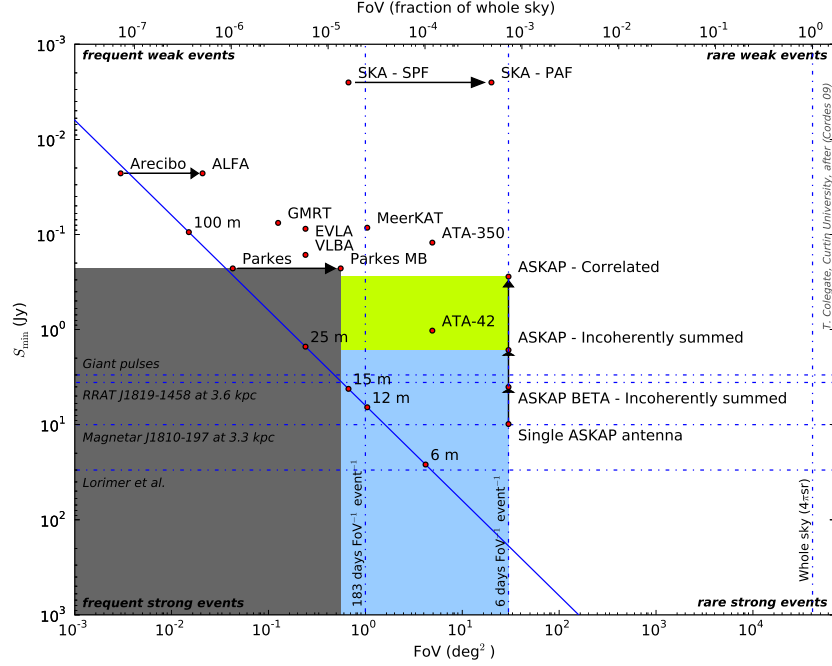


Figure 1.3: Phase space probed by various current or planned radio facilities operating at 1.4 GHz, as depicted by the minimum detectable flux density (S_{\min}) at the 5σ level versus instantaneous accessible field of view. For simplicity in presentation, all telescopes are plotted with the same operating bandwidth (0.3 GHz) and integration time (1 ms). The solid diagonal line refers to single-pixel, single-reflector telescopes all with the same system temperature (25 K) and aperture efficiency (60%). The SKA-SPF (single-pixel feed) line indicates the sensitivity for 2000 single-pixel feed 15-m antennas, extended to phased array feeds (SKA-PAF) with a 20-deg^2 FoV. The horizontal dashed lines show some nominal flux densities of transient sources at 1.4 GHz. The shading represents the parameter space covered by a coherent ASKAP (green), incoherently summed ASKAP (blue) and the Parkes Multibeam (grey). Adapted from Cordes (2009). Figure and caption reproduced from Macquart et al. (2010) (available at <http://www.publish.csiro.au/?paper=AS09082>) with permission from the authors, The Astronomical Society of Australia and CSIRO Publishing.

classification engines aim to ensure that only the most interesting and useful data are presented to humans to assess. Building classification algorithms requires training data such as light curves and brightness distributions that have been classified by humans. It is important that these data are as accurate and realistic as possible in order to maximise the accuracy of the automatic classifications. Currently, there are precious little training data available and more data is required to improve the accuracy of these algorithms. These questions can be summarised as follows:

What is the optimal survey strategy? It is not immediately obvious how one should choose the sensitivity, coverage and cadence to maximise the scientific output of an instrument. For example, if one chooses a shallow large area survey, we might miss important faint events. Conversely, if one chooses a deep, narrow survey, one might miss bright, rare events. Key parameters in this optimisation are the brightness distribution, surface density and timescale of transients, none of which are currently well constrained at any frequency of interest.

What types of sources are most common? With such an unprecedented wide-field, we expect to detect a large number of transient sources. In order to determine which sources will be most interesting, the survey teams aim to deploy automatic classification engines to determine to which class a given source belongs. Training these algorithms requires real-world light curves, brightness distributions and surface densities for known source classes.

How can current the algorithms be improved? Once data are obtained by the telescope hardware, it is up to the processing software to produce scientifically useful data. But these data are subject to a range of real-world problems for which the software must compensate, such as gain variations and image artifacts. These algorithms are yet to be defined. For high time resolution surveys, the sensitivity and coverage will be strongly dependent on the available computing capacity. Increasing the efficiency of algorithms therefore enables larger fields or higher sensitivity while maintaining an affordable computational cost.

In the next sections, I provide a theoretical background to the above questions and describe two ways in which this thesis helps answer them: (1) by conducting pilot surveys and (2) by developing new algorithms.

1.4 SOURCES OF TRANSIENT AND VARIABLE RADIO EMISSION

1.4.1 Scintillation of compact radio sources

Scintillation is a propagation effect due to scattering of incident radiation by turbulent material causing changes in the apparent brightness, shape or position of an astronomical source. It is most often observed as the twinkling of stars due to turbulence in the Earth's atmosphere. At radio wavelengths, scintillation is caused by turbulence in ionised material, either in the Earth's ionosphere, in the Solar wind (interplanetary scintillation) in the interstellar medium (interstellar scintillation, or ISS; Narayan 1992) or in the intergalactic medium (Koay et al., 2011; Lovell et al., 2008).

ISS operates in two regimes, defined according to the physical process that causes the brightness variations (Narayan, 1992). Diffractive scintillation operates over short timescales (minutes to hours) and at high frequencies (> 5 GHz), while refractive scintillation operates over longer timescales (days to years) and at lower frequencies (< 5 GHz). Diffractive scintillation requires a much smaller source size than does refractive scintillation, and has only been observed in pulsars. Both refractive and diffractive scintillation operate in the regime of strong scattering, which can result in very large (up to 100%) fractional changes in apparent brightness.

As well as being an interesting phenomenon in its own right, scintillation can be used to estimate the space velocity of pulsars (Cordes, 1986), measure the distance, velocity and turbulence in otherwise invisible ionised clouds (Bignall et al., 2003), achieve extremely high angular resolution (Briskin et al., 2010; Cornwell & Narayan, 1993; Lazio et al., 2004), and determine the expansion velocity of gamma-ray burst ejecta (Frail et al., 1997). A number of surveys for scintillation have been performed, which have demonstrated the dependence of the strength of the scintillation on the amount of ISM material through which the signal propagates (Gaensler & Hunstead, 2000; Lovell et al., 2008).

In general, ISS operates continuously, although often with a degree that varies with time. But Fiedler et al. (1987) discovered a one-off event, a so-called Extreme Scattering Event (ESE), which dramatically altered the radio brightness of the quasar 0954+658. They attributed these variations to refractive focusing due to small inhomogeneities in the ISM. ESEs have been detected in other quasars (Cimò et al., 2002; Lazio et al., 2001) and pulsars (Maitia et al., 2003), but a comprehensive understanding of ESEs remains a elusive (Walker, 2007).

Monitoring continues to be an important method for discovering and understanding scintillation, with impressively long datasets being collected by a number of groups (Andrew et al., 1978; Lovell et al., 2008; Teräsranta et al., 2004).

1.4.2 Radio supernovae and gamma-ray bursts

The explosive deaths of stars have been observed probably throughout human history, with the first evidence of a historical observation being that of SN 185 by Chinese astronomers (Clark & Stephenson, 1977, p. 87). Core-collapse supernovae, which are usually discovered in the optical band, have been shown in follow-up radio observations to produce radio afterglows (Weiler et al., 2002, 2004) due to the interaction of the blast wave with the progenitor stellar wind and with the interstellar medium. This expanding blast wave brightens and fades with the density of material through which it propagates, which enables the mass loss history of the progenitor to be determined (Ryder et al., 2004; Soderberg et al., 2006).

Gamma ray bursts also leave behind radio afterglows as the shock from the explosion ploughs into the interstellar medium. As opposed to radio supernovae, GRB ejecta are relativistic. The radio afterglows of GRBs are a key argument in establishing the link between gamma-ray bursts and supernovae, as illustrated in the case of GRB980425 and SN1998bw (Kulkarni et al., 1998). Such radio afterglow observations also allow for relatively detailed modelling of blast wave physics in the relativistic regime. The spectral dynamics of radio afterglows have also been used to determine whether an optically detected supernova has relativistic ejecta, and have been used to detect “orphan” gamma-ray burst afterglows (i.e. relativistic supernovae without gamma-ray triggers; Soderberg et al. 2010).

Prompt radio emission accompanying a supernova explosion has yet to be detected. However, in the case of SN1987A a week long burst of radio emission was observed in the days after the optical discovery (Turtle et al., 1987). This burst of radio emission has been explained by synchrotron emission from electrons that had been accelerated in the shock, falling behind the blast wave after 2-3 days and becoming free to emit (Kirk & Wassmann, 1992). The radio afterglow of SN1987A has been studied in detail and is well explained with the usual models of an expanding blast wave (Zanardo et al., 2010).

As yet, no radio emission has been detected from thermonuclear explosions of white dwarfs (type Ia supernovae) in spite of deep observations (Hancock et al., 2011).

1.4.3 Pulsars

Pulsars are rapidly rotating neutron stars with strong magnetic fields anchored to the surface of the star. Beamed, coherent radio emission is thought to originate in the magnetosphere above the star, which appears pulsed when arriving at Earth due to the lighthouse effect from the neutron star rotation.

Pulsars are intriguing astrophysical laboratories. While the pulsar emission mechanism is yet to be fully explained and is the source of a great deal of debate (Hankins et al., 2009), precise measurements of the arrival times

of radio pulses from pulsars have enabled astronomers to test predictions of General Relativity through observations of binary pulsar systems (e.g. Lyne et al., 2004), to search for gravitational waves with widely-spaced pulsars (Yardley et al., 2010, and references therein), and to test theories of matter at the most extreme densities (e.g. Demorest et al., 2010). The short pulses emitted by pulsars undergo propagation effects during their passage through the ISM, which enables measurements of the Galactic magnetic field structure (e.g. Noutsos et al. 2008, Han et al. 1997) and free electron density (e.g. Cordes & Lazio 2002).

Pulsars appear as variable sources due to a number of mechanisms. The primary source of variability is the lighthouse effect due to the rotation of the pulsar, described earlier. Pulsar emission also undergoes diffractive and refractive scintillation, which changes the mean flux density of the pulsar over timescales of minutes to months.

The pulsar emission mechanism itself is also variable. Among the most startling results from single pulse searches was the measurement of the so-called ‘nano-giant’ pulses. While most known pulsars emit regular pulses with each pulse having roughly the same flux density as the last, some pulsars (notably young and millisecond pulsars) occasionally emit a single pulse many thousands of times brighter than the usual pulse brightness, with durations as short as 2 ns (Hankins & Eilek, 2007; Hankins et al., 2003). The implied brightness temperature of these pulses is some 10^{37} K, implying a coherent emission mechanism that is yet to be fully understood.

Pulsar emission varies in other ways. Some pulsars appear to turn off entirely for hours or even weeks at a time, only to reappear again later as more-or-less normal pulsars, in a process known as ‘moding’ (Backer, 1970; Rankin, 1986). An extreme case of moding is observed in the so-called Rotating Radio Transients (RRATs), which were discovered in single pulse searches, and which turn on for a single rotation as rarely as once per hour (McLaughlin et al., 2006). Some pulsars switch between standard pulsars and RRAT type emission (Burke-Spolaor & Bailes, 2010).

Pulsars, therefore, can appear as transient and variable sources over a range of different timescales. At high-time resolution where dispersion is important, the regular lighthouse effect is most evident in periodicity searches, while giant pulses and RRAT type emission are most evident in searches for single pulses. Propagation and mode changing effects can be seen both in high and low time resolution studies.

1.4.3.1 *The Lorimer burst and the Perytons*

While re-processing data from the Parkes Multibeam Pulsar Survey (Manchester et al., 2001), Lorimer et al. (2007) discovered a single extremely bright short-duration (< 5 ms) radio burst in the direction of the Small Magellanic Cloud. This burst had all the hallmarks of having an extragalactic origin.

Firstly, the burst had a large dispersion consistent with the interstellar dispersion law, but with a dispersion measure too large to be explained by material only in the Milky Way. Second, the burst was detected mostly in a single beam of the telescope, which argued against a sidelobe detection from a terrestrial source near the horizon. The most puzzling and surprising aspect of the burst was its flux density, some 30 Jy.

The consequences of this burst, if truly astronomical, are far reaching. Not only is new physics required to explain how such a bright pulse of radio emission could be produced by a natural process, but certain aspects of cosmology can be addressed by measuring the dispersion of extragalactic pulses. A range of experiments could be performed, including directly detecting the majority of the baryons in the Universe (Ginzburg, 1973), and distinguishing between different models of cosmic reionisation (Inoue, 2004). If this radio pulse is related to gamma-ray bursts, which was alluded to by Lorimer et al. (2007) but is by no means certain, then detecting a burst also enables measurements of the physics of the gamma-ray burst, including jet opening angle, and Lorentz factor (Macquart, 2007), density and distance of any scattering material (Lyubarsky, 2008), and structure of the fireball magnetic field (Sagiv et al., 2004).

Some doubt has been cast on the astronomical nature of the Lorimer Burst, however. By re-processing data from the Parkes multi-beam system, Burke-Spolaor et al. (2011a) discovered 16 single bursts with similar spectral characteristics to the Lorimer burst. Unlike the Lorimer Burst, these bursts, which Burke-Spolaor et al. (2011a) call ‘Perytons’, have a number of properties that hint at a terrestrial origin. For example, some Perytons have significant kinks in their dispersion sweeps, rather than the smooth ν^{-2} dispersion law expected from the cold interstellar medium. The Perytons exist within a narrow range of average dispersion times (mostly between 320 and 400 ms), have substantial amplitude modulations, and were detected in the sidelobes of the antenna. All Perytons were detected during the day, in general at mid-morning, and mostly in the mid-winter months of June and July. Thus, the Perytons’ properties are not those expected of astronomical sources, implying a terrestrial origin. The implications for the Lorimer Burst are not entirely clear. While the Lorimer Burst shares many properties of the Perytons, namely the dispersion, time of day, and season, there are other properties that make it unique, including a smooth dispersion sweep, an absence of amplitude modulation, and primary beam detection. At the time of writing, the astronomical nature of the Lorimer Burst is still under debate.

1.4.4 X-ray binaries

X-ray binaries are systems in which a degenerate object (either a black hole or a neutron star) accretes material from a companion star in a relatively short, often elliptical orbit (Fender, 2004). Variable X-rays and jets can be

produced by the accretion process, typically in synchrony with the orbit. The jets produce relativistic particles in a magnetic field, leading to synchrotron emission detectable at radio frequencies. The radio flux during a flare can increase by a factor of a thousand or more (Gregory & Kronberg, 1972).

There is increasing evidence that black hole X-ray binaries are a scaled-down version of the black hole systems powering active galactic nuclei, but with much shorter timescales. Those X-ray binaries that have observable relativistic jets are commonly called ‘microquasars’. To date, the vast majority of X-ray binaries have been detected in the Milky Way by all sky X-ray surveys, and a number of extragalactic X-ray binaries are known (Grimm et al., 2007; Muxlow et al., 2010).

1.4.4.1 *Non-degenerate radio stars*

Radio emission is common in stars. It is observed across a large fraction of the Hertzsprung-Russell diagram (Güdel, 2002), and is due to a range of thermal and non-thermal mechanisms. Flaring is common in cool stars, which show coherent and incoherent radio flares, much like the Sun. The coherent bursts have rise times as short as 5 ms, implying high brightness temperatures ($\sim 10^{16}$ K). Dynamic spectra of the bursts themselves often show narrow ($\Delta\nu/\nu = 0.2\%$) peaks that change with time, implying movement in the emission region.

Incoherent flares typically last from minutes to hours, and have broadband spectra. The flares show evidence of mildly relativistic electrons emitting gyro-synchrotron radiation in magnetic fields. The electrons themselves are postulated to be accelerated by magnetic reconnection.

Wide-field surveys of the Galactic plane are likely to discover a large number of flaring stars.

1.5 RADIO TRANSIENT SURVEYS

In planning radio transient surveys, it is important to understand the relationship between the survey metrics, and the parameters of the transient and variable sources that such a survey aims to detect. This section describes these metrics and parameters, highlights the complex interplay between the survey metrics and argues that pilot surveys can help optimise future survey design.

1.5.1 Relationships between source class and survey parameters

Given a particular class of transient or variable radio source (e.g. gamma-ray burst, pulsar or scintillating source) that is detectable at radio frequencies, there are three general parameters that describe the properties of that class, in terms of its detectability in a blind, wide-field survey. The first property is

the brightness distribution: how many sources of a given brightness would we expect to detect given perfect sampling of the sky. The second is the surface density², which measures the number of examples of such a class that are expected to occur per unit solid angle, per unit of time. The final property is the timescale, i.e. the amount of time over which the brightness changes substantially.

Each of these parameters influences the choice of three corresponding parameters that describe the properties of a wide-field radio survey. The brightness distribution is related to the sensitivity of the survey³, which is the minimum brightness at which a source can be reliably detected. The surface density relates to the survey coverage, which is the amount of sky covered by the survey. Finally, the timescale relates to the survey cadence, which describes how often a given patch of sky is observed.

Clearly, different source classes have different values for their class parameters, and the requirements for designing a survey to detect one class, may be totally contrary to the parameters for another class. For instance, to detect faint common objects we would conduct a high sensitivity survey and sacrifice survey coverage. By contrast, to detect rare bright objects, we would conduct a low sensitivity survey, but would use a wider sky coverage. Similarly, to detect changes on short timescales, we need to observe the same patch of sky often, which is at odds with obtaining high sensitivity or wide coverage.

The key question for the survey designer is how to maximise the scientific output of a survey under the competing demands of sensitivity, coverage and cadence, and the given limitations of a given telescope technology and of the total survey time. While a full treatment of this optimisation is outside the scope of this thesis, I describe aspects of this discussion in the following sections.

1.5.2 Sensitivity

When searching for transient and variable sources with an interferometer, one typically measures visibilities and forms an image through the use of a Fourier transform (Taylor et al., 1999, p. 5). Assuming the observation is limited by thermal noise in the system, the RMS noise in the total intensity image is approximately:

$$\sigma_s = \frac{1}{\eta_s} \frac{\text{SEFD}}{\sqrt{2N(N-1)Bt}}, \quad (1.2)$$

²Unless otherwise stated, the phrase ‘surface density’ applies exclusively to transient or variable sources, and does not apply to background sources.

³Survey sensitivity can be limited by confusion if the angular resolution is insufficient to resolve the surface density of background sources. I assume throughout that any survey has sufficient angular resolution to avoid confusion, and therefore that the sensitivity is limited by other factors such as thermal noise.

where η_s is the system efficiency (which accounts for losses in electronics and the correlator), N is the number of antennas, B is the total observing bandwidth and t is the integration time. SEFD is the system equivalent flux density, defined as

$$\text{SEFD} = \frac{T_{\text{sys}}}{K}, \quad (1.3)$$

where T_{sys} is the system temperature, and K is the measure of antenna performance. Antenna performance is defined as:

$$K = \frac{\eta_a A}{2k_B}, \quad (1.4)$$

where η_a is the antenna illumination efficiency, A is the antenna area and k_B is Boltzmann's constant.

If we want to detect a point source with a signal to noise ratio (S/N) of m , this results in a minimum detectable flux density $S_{\text{min}} = m\sigma_s$, and the time required to reach this flux density can be determined by solving the above equations for t , which gives:

$$t \simeq \left(\frac{kT_{\text{sys}}}{\eta_a \eta_s AN} \right)^2 \frac{2}{S_{\text{min}}^2 B}, \quad (1.5)$$

where we have assumed N is sufficiently large that $N^2 \simeq N(N-1)$. A useful figure of merit here is the so called 'survey speed', which is the amount of sky that can be surveyed per unit of time to a given sensitivity. We can write the survey speed as:

$$\begin{aligned} \text{SS} &= \frac{\Omega_i}{t} \\ &= \Omega_i \left(\frac{\eta_a \eta_s AN}{kT_{\text{sys}}} \right)^2 \frac{S_{\text{min}}^2 B}{2} \\ &= \Omega_i S_{\text{min}}^2 V, \end{aligned} \quad (1.6)$$

where Ω_i is the instantaneous solid angle (or field of view) of the telescope, and

$$V = \left(\frac{\eta_a \eta_s AN}{kT_{\text{sys}}} \right)^2 \frac{B}{2} \quad (1.7)$$

encapsulates the remaining telescope parameters. For parabolic dish-based radio antennas, the full width half maximum of the beam is given by:

$$\theta = 1.22 \frac{\lambda}{D} \quad (\text{radians}), \quad (1.8)$$

where λ is the observing wavelength, and D is the diameter of the dish. If this dish has a multi-beam feed with $N_b > 1$ beams, this implies an instantaneous field of view of:

$$\begin{aligned}\Omega_i &\simeq \pi N_b \left(\frac{\theta}{2}\right)^2 \\ &= 1.16 N_b \left(\frac{\lambda}{D}\right)^2.\end{aligned}\tag{1.9}$$

We can now write the survey speed of a dish-based radio telescope as:

$$\text{SS}_{\text{dish}} = 0.36 N_b \left(\frac{\lambda \eta_a \eta_s D N}{k T_{\text{sys}}}\right)^2 S_{\text{min}}^2 B.\tag{1.10}$$

1.5.3 Coverage

Coverage describes the amount of sky covered by a survey. To measure variability, one requires at least two time-separated measurements. With N_{epoch} such measurements, one has $N_{\text{epoch}} - 1$ opportunities to measure variability. This logic leads to the two-epoch equivalent coverage area as a metric for coverage, defined as:

$$\Omega_t = \sum_{i=1}^{N_{\text{pos}}} \Omega_i (N_{\text{epoch},i} - 1),\tag{1.11}$$

where N_{pos} is the total number of independent directions on the sky observed in the survey, Ω_i is the solid angle of a particular direction of interest, and $N_{\text{epoch},i}$ is the number of epochs for which a particular direction has been observed.

A important related quantity that is often used to describe the surface density of transient sources is the snapshot rate, defined as:

$$R_{\text{snapshot}} = \frac{N_{\text{transients}}}{\Omega_t}.\tag{1.12}$$

Snapshot rate is often quoted in units of deg^{-2} .

1.5.4 Cadence

Cadence is the time interval between two observations. The classical constraint on the cadence is given by the Shannon-Nyquist sampling theorem (Nyquist, 1928; Shannon, 1949), which states that if a signal has no frequencies higher than f Hz, then it can be perfectly reconstructed as long as the interval between measurements of the signal is less than $T_{\text{Nyquist}} < 1/2f$ s.

While this sampling theorem is very important in designing various aspects of radio telescopes, its relevance in survey design is not so certain. In particular, sampling with a period of T_{Nyquist} requires a very high cadence to perfectly reconstruct some known radio light curves. However, it is not clear whether a perfect reconstruction of the radio light curve is required in order to differentiate between different source classes, or perform useful scientific investigations. Therefore, sampling with a longer period $T_{\text{revisit}} > T_{\text{Nyquist}}$ may result in adequate scientific value with lower cadence.

Cadence not only affects the ability to reconstruct all features in the light curve, but also the detectability of transient sources. If a transient source is only detectable for a time $T_{\text{on}} < T_{\text{revisit}}$, then in a given field there is a finite probability that such a transient source will be missed between adjacent observations of that field. Therefore, if a class of transient source has a particularly short timescale, it is desirable to make T_{revisit} as small as possible.

1.5.5 The interplay of survey parameters

Given a telescope with technological parameters V and Ω_i , and a fixed amount of observing time T_{total} , how should one choose values for sensitivity, coverage and cadence? If one chooses to raster scan a solid angle Ω_s of the sky, then the total two-epoch sky coverage is:

$$\begin{aligned}\Omega_t &= \Omega_i(N_{\text{epoch}} - 1) \\ &\simeq \Omega_i N_{\text{epoch}},\end{aligned}\tag{1.13}$$

and the total survey time is:

$$\begin{aligned}T_{\text{total}} &= \frac{\Omega_s}{\Omega_i} N_{\text{epoch}} t \\ &= \frac{\Omega_t}{\Omega_i} t \\ &= \frac{\Omega_t}{SS} \\ &= \frac{1}{V\Omega_i} \frac{\Omega_t}{S_{\text{min}}^2}.\end{aligned}\tag{1.14}$$

From Equation 1.14 we see that the total observation time is fixed by the telescope technology, desired total coverage and minimum flux density. If the observing time is fixed, then we can trade flux density for total coverage and vice versa, but we cannot have both.

There is a similar trade-off with cadence. To illustrate this trade-off, we can calculate the revisit time between raster scans as:

$$\begin{aligned} T_{\text{revisit}} &= \frac{T_{\text{total}}}{N_{\text{epoch}}} \\ &= \frac{1}{V\Omega_i} \frac{\Omega_s}{S_{\text{min}}^2}. \end{aligned} \quad (1.15)$$

For a short revisit time that we might need to catch short timescale transients, we require a small value for Ω_s . On the other hand, to detect rare transients, we need a large value for Ω_s . Equation 1.15 illustrates that these two parameters are in opposition, therefore requiring a trade-off.

A more thorough treatment of survey metrics for transient sources is given by Cordes (2009).

1.5.6 The role of pilot surveys

The key challenge in designing a survey for transient and variable radio sources is the need to maximise the scientific output by choosing a sensitivity, coverage and cadence, under the constraints of telescope technology and survey time. Given the diversity of known sources, the choice of the survey parameters is not obvious. One might choose a wide shallow survey but miss the faint population, or a deep, narrow survey and miss the rare population. If the survey has insufficient cadence, then key information about the light curve is lost, or rare, short duration transients might be missed altogether.

Pilot surveys can provide answers to these questions by providing initial estimates of the source population parameters. Blind, wide-field pilot surveys can furnish overall surface densities and brightness distributions. Follow-up surveys can provide surface densities and brightness distributions for a particular source class.

This thesis contains the results of two pilot surveys that explore new regions of the transient parameter space. Chapter 2 describes a blind, wide-field survey for slow transients, and Chapter 3 describes a directed, follow-up survey of gamma-ray bursts. The role of the pilot surveys in this thesis is to provide initial estimates of brightness distributions, surface densities and timescales to aid in designing surveys for transient and variable radio sources.

1.6 TECHNIQUES

1.6.1 Detecting slow transients with interferometers

Searching for transient and variable sources with interferometers⁴ can be broken down into a number of steps (Figure 1.4). To begin with, we assume a correlating interferometer, which takes voltages from each of the antennas and cross-correlates them to produce ‘visibilities’, which are measurements of the Fourier transform of the sky brightness distribution.

The next step is to find the location and flux density of sources, given the visibility measurements. This is most often performed by Fourier-transforming the visibilities to form an image. Once the images are available, the next step is to search for variability.

When searching for variability in images, one can take one of two approaches. Perhaps the simplest approach is to take the image of interest, and divide or subtract by a reference image that has been obtained previously. The resulting difference or quotient image will have significant deviations at the sky positions where there have been substantial changes in the flux density. The target and reference image must have the same point spread functions, otherwise truly constant emission will not be accurately subtracted between the two images, resulting in false positives. Sophisticated techniques have been developed to match the point spread functions of two images before subtraction (e.g. Alard & Lupton, 1998).

The alternative approach to searching for variability is to perform source finding, either in the image plane (e.g. Hopkins et al. 2002) or in the visibility plane (Trott et al., 2011). Source finding returns a list of the positions and flux densities of all sources in the data. Each measurement is then associated with previous measurements of that source, resulting in a light curve. A list of previously detected sources can be supplied to the source finding step, which enables the source finding step to provide upper limits for sources that are not detected in the data.

With a light curve in hand, the final step is time series analysis. For instance, one can test whether the light curve is consistent with a constant flux density, or is variable. If it is variable, one can attempt to match the light curve against a range of known light curves to obtain a classification.

1.6.2 Detecting dispersed transients

Surveys for dispersed transients are usually performed by single dish telescopes (e.g. Manchester et al. 2001, Nice 2007). Interferometers are occasionally used (e.g. Janssen et al. 2009), and there is great interest in exploiting

⁴We will not discuss searching for slow transients with single dish telescopes, as the advantages of interferometers are so great that single-dish telescopes are rarely used for this purpose.

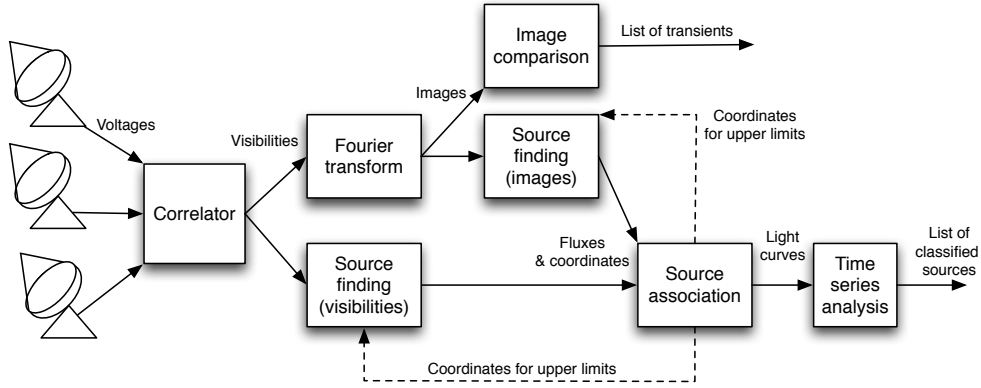


Figure 1.4: A schematic showing methods for performing a search for slow transients with an interferometer. Typically only a single path through the system will be implemented. The dashed lines indicate optional steps.

interferometers more in the future, if a number of technical challenges can be overcome.

The key difficulty in searching for dispersed transients is to handle the dispersion, which I describe below.

1.6.2.1 Dispersion and dedispersion

When a short-duration pulse of radio emission is emitted by a source, it undergoes a frequency-dependent delay during its passage through the ISM, a plasma containing non-relativistic unbound electrons. The delay between two frequencies ν_1 and ν_2 is given by:

$$t_1 - t_2 = \frac{e^2}{2\pi m_e c} \text{DM} (\nu_1^{-2} - \nu_2^{-2}) \quad (1.16)$$

where the physical constant is:

$$\mu = \frac{e^2}{2\pi m_e c} \simeq 4.15 \text{ms}, \quad (1.17)$$

for frequencies in GHz. The dispersion measure (DM) describes the number of electrons between the observer and the emitting source, defined as:

$$\text{DM} = \int_0^d n_e dl, \quad (1.18)$$

where n_e is the electron number and d is the distance to the source. DM is usually quoted in units of pc cm^{-3} . An example of the “dispersion sweep” for the Vela pulsar is shown in Figure 1.5.

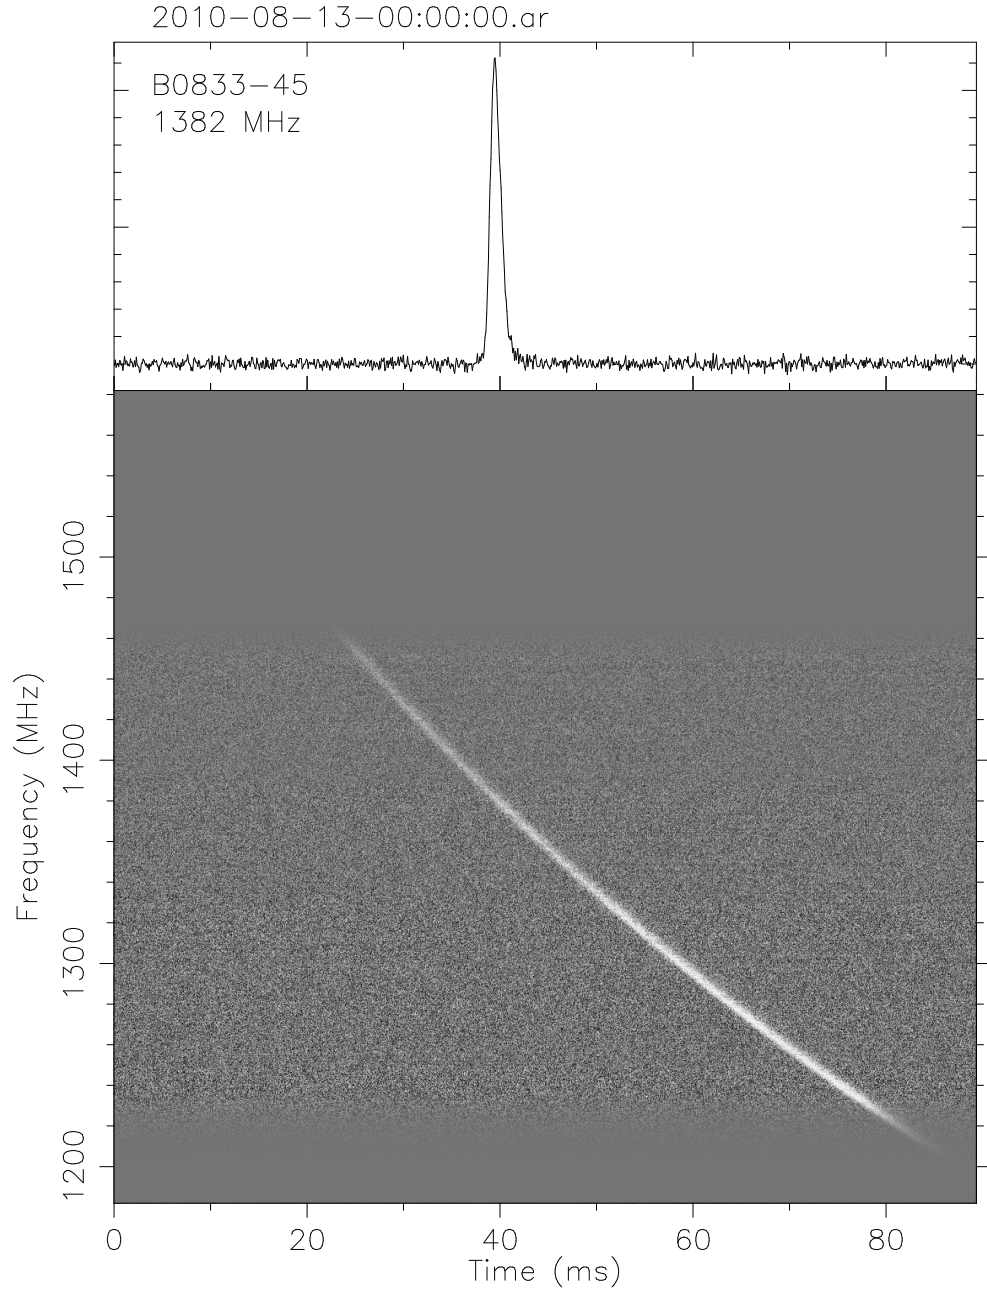


Figure 1.5: Dispersion by the ISM causes frequency a dependent delay. Plotted above is the folded radio pulse from the Vela pulsar (PSR B0833-45) which has a period of 89 ms and a DM of 67.99 pc cm^{-3} . These data were recorded in 60 s with the 12 m telescope used to take the measurements in Chapter 3. Top panel: dedispersed time series showing shape of the pulse emitted by the pulsar, before being dispersed by the ISM. Bottom panel: frequency vs. time image where the gray scale is the intensity of the incoming radiation. The ν^{-2} dispersion sweep is seen between 1220 and 1460 MHz. The telescope feed has no sensitivity outside these frequencies.

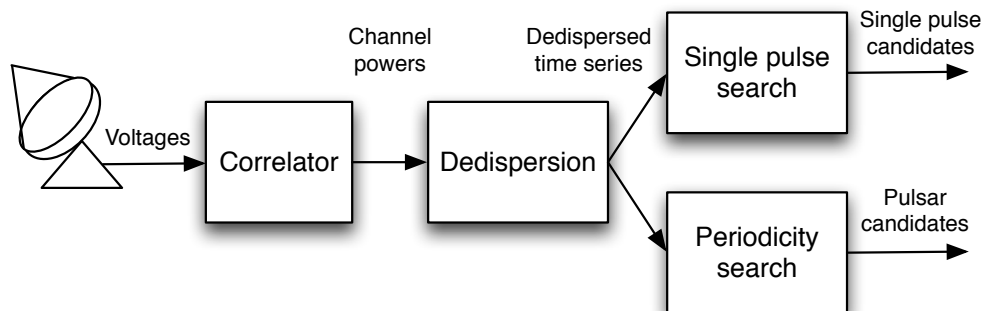


Figure 1.6: Typical steps involved in detecting fast transients with single dish telescopes. Usually, both single pulse and periodicity searches are performed.

This dispersion is a key driver in the design of surveys for dispersed transients. If the dispersion is not accounted for, the sensitivity is dramatically reduced, as the pulse is smeared over a long period of time. Therefore, all modern fast transient surveys perform dedispersion at a range of DMs to maximise the potential for detecting the emission.

A more detailed description of dispersion is given in Section 4.2.1 on page 125.

1.6.2.2 Single dish processing

The voltages from the dish are channelised in a correlator that typically supplies power measurements over a number of channels ($\simeq 1000$) integrated over a few hundred microseconds. This set of channel powers is then dedispersed at a range of trial DMs, to produce a time series at each DM of interest.

Each time series can then be subjected to two types of search. Some types of objects, such as RRATs, “Lorimer Bursts” and giant pulses, are best found using a single pulse search. Regularly pulsing sources such as pulsars are best found using a periodicity search, which is usually implemented with an FFT.

1.6.2.3 Interferometric processing

The standard approach for detecting dispersed transients with interferometers has been to create a set of channel powers using some method, and then to apply the single-dish techniques to this set of channel powers (Figure 1.7).

There are standard methods for creating these channel powers with interferometers (see also Section 4.2.2 on page 127). The first method is to perform the standard correlation and Fourier transform to produce an image. Single-dish processing is then performed on each pixel of the resulting image. The second method is to form ‘tied-array beams’ in a beam-former. Each beam has a small area (the size of a pixel), but it can be placed anywhere in the field

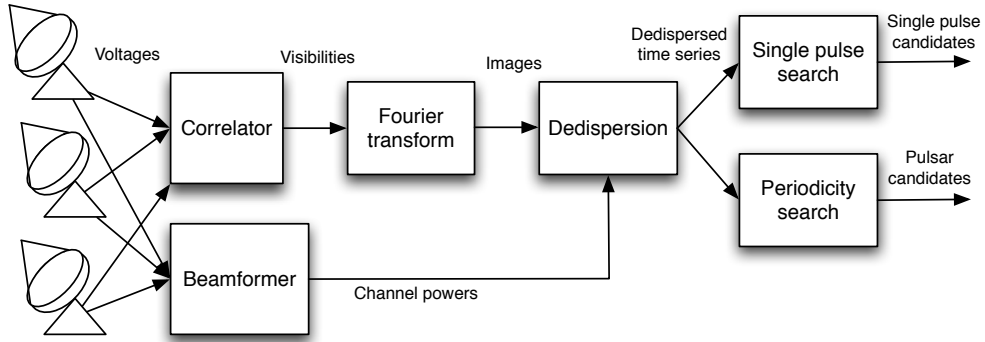


Figure 1.7: Steps involved in detecting fast transients with interferometers using classical approaches. Voltages from the telescopes can be processed either by a beam-former or by a correlator. The beam-former can form either coherent ‘tied-array’ beams or an incoherent ‘power beam’. Beams from the beam-former can be dedispersed directly, while a Fourier transform is required to form images if a correlator is used. The dedispersion and search steps are identical to those used in single dish-surveys.

of view. Single-dish processing is then performed on the tied-array beam. The final method is to form ‘power beams’ in a beam-former. The power beam has an area equal to the field of view, but provides poorer sensitivity than the above methods.

The requirement to either perform dedispersion on an entire image or to form a large number of tied-array beams imposes an impractical computational burden. This computational load has generally limited fast transient surveys to single-dish telescopes, which only have a small number of independent beams that need to be searched. In Chapter 4 of this thesis, I describe two efficient new techniques to detect fast transients with interferometers.

1.7 PURPOSE & RATIONALE OF THIS THESIS

The radio sky contains a veritable zoo of source classes that have been discovered chiefly through serendipity, monitoring, and follow-up of events detected at other wavelengths. Upcoming wide-field radio telescopes, enabled by new technologies, are now ushering in a new age of radio astronomy that promises to revolutionise our understanding of the dynamic radio sky by probing new parts of transient parameter space.

The key telescope in the Australian context is ASKAP. Currently, two survey teams are making preparations for ASKAP first light in 2013, aimed at specifically addressing the theme of transient and variable radio sources. The Variables and Slow Transients survey (VAST; Chatterjee & Murphy 2010) will perform a blind survey for transient and variable radio sources on timescales

of 5 s and longer, and the commensal real-time ASKAP fast transients survey (CRAFT; Macquart et al. 2010), will perform a blind survey for transient sources on timescales from a few hundred microseconds up to 5 s. Both surveys aim to deploy sophisticated and efficient algorithms to help in the processing and classification of the enormous volumes of data expected.

Astronomers are now preparing for this new, wide-field era, but key questions remain: What is the optimal survey strategy? What types of sources will be most commonly detected? How can current algorithms be improved? This thesis addresses these questions in two ways, through conducting pilot surveys and developing new algorithms.

1.7.1 Pilot surveys

Pilot surveys provide answers to all three questions above. While having less sensitivity and sky coverage than the proposed wide-field surveys, pilot surveys will nonetheless furnish estimates of overall transient rates, along with flux distributions of the brightest and most common transient sources. These parameters can inform both survey design, and automatic classification algorithms. In addition, pilot surveys provide experience of real-world data processing problems and follow-up strategies, and provide real-world light curves, surface densities, and brightness distributions. Pilot surveys thus can be used for designing robust transient detection pipelines. Furthermore, pilot surveys have the potential for discovering the brightest and most common sources, providing early scientific gains.

This thesis presents results on two pilot surveys. The first pilot survey, described in Chapter 2, was an archival survey for transient and variable radio sources using the Molonglo Observatory Synthesis Telescope (MOST; Mills 1981, Robertson 1991). This moderately wide-field survey at 843 MHz discovered 53 highly variable sources and 15 transient sources and sets limits on the surface density and brightness distribution of transient sources over a range of timescales. This survey is among the very few radio surveys to date that probes deeply into survey parameter space, with a 5σ sensitivity⁵ of 14 mJy beam^{-1} , a two-epoch coverage of over 2775 deg^2 , and a range of cadences from 1 day to 3 years. It is one of the first blind radio surveys to provide a consistent set of light curves and transient rates at ASKAP operating frequencies, and can be used directly for informing survey designs and classifications algorithms for VAST.

The second pilot survey, described in Chapter 3, involved automatic radio follow-up of gamma-ray bursts. This survey sets limits on prompt radio emission specifically from GRBs at both high- and low-time resolutions. This survey has the fastest prompt high-time resolution radio follow-up of any such

⁵Previous wide-field surveys with MOST such as those reported by Bock et al. (1999), Mauch et al. (2003) and Murphy et al. (2007) had better sensitivity. The higher threshold of this survey was chosen to reduce the rate of false positives.

survey to date, by being on source within 2 minutes of the GRB detection. A single pulse search has detected two single-pulse candidates within 20 min of a GRB. Both candidates have DMs putting them outside the Milky Way, and both are coincident with breaks in the X-ray light curve following the GRB. Bursts of this sort have never previously been detected and await confirmation by other telescopes. This survey also sets limits on the surface density and brightness distribution of prompt radio emission from GRBs, and demonstrates the feasibility of automated follow-up with radio telescopes.

1.7.2 New algorithms

New wide-field telescopes will produce unprecedented volumes of data, and astronomers are being challenged to design algorithms to maximise the scientific output of these new instruments. To that end, minimising or quantifying systematic effects are important in image-based radio surveys, and methods to maximise the sensitivity of the instruments, while working within the constraints of computational cost, are of great interest.

This thesis has two separate contributions of algorithms for use in fast and slow transient surveys. As part of the work on the MOST pilot survey presented in Chapter 2, I propose algorithms to compensate for variations in image gain, to measure image quality, to verify flux error estimates, and to accurately count the number of variable sources in the presence of flux error uncertainties. These algorithms are new to this field and are of direct application to VAST.

In Chapter 4, I present two efficient new algorithms for detecting dispersed radio emission in interferometric data: the Chirpolator and the Chimageator. When compared with classical methods, these techniques excel in the regime of sparse arrays, for which they both require substantially lower data rates, and the Chirpolator requires a much lower post-integrator operation rate. Compared with classical techniques, these approaches are better matched to future supercomputing architectures, for which the arithmetic capability will outstrip the bandwidth capability, and are suitable for use by CRAFT and other interferometer-based fast transient surveys.

1.7.3 Summary

The upcoming age of wide-field astronomy will provide an unprecedented view of the dynamic universe. This thesis is a contribution to the preparations for this new age. The pilot surveys described here provide much needed data in terms of transient rates, brightness distributions, timescales and light curves, all of which will inform survey design and classification algorithms for upcoming surveys. The innovative new algorithms described in this thesis will improve the robustness, sensitivity and computational cost of conducting this new science.

A 22-year southern sky survey for transient and variable radio sources with the Molonglo Observatory Synthesis Telescope

This chapter, along with Appendices A and B has been published as Bannister et al. (2011). It is reproduced here with the following modifications:

- The header of Table 2.2 was incorrect in the original paper. The correct header is ‘Semi-minor Axis ($a_0/2$)’.
- Due to an incorrect matching radius used to associate MOST sources with NVSS, some transient sources were described as having non-detections in NVSS. Table 2.6 and Figure B.1 have been updated with the correct NVSS identifications, and Table 2.7 has been updated with new comments.
- Bias correction had not been applied to the NVSS flux densities which were quoted in the original paper. Tables 2.6 and 2.5 as well as Figures A.1 and B.1 have been updated with the corrected NVSS fluxes.
- Minor formatting changes for the thesis format.

An erratum to Bannister et al. (2011) containing all modifications except the formatting changes was accepted for publication on 31 Aug 2011.

2.1 INTRODUCTION

MANY of the most extreme events in the Universe manifest themselves through variability. For example, explosive events such as Gamma ray bursts (GRBs) are thought to originate from stellar collapses that can momentarily outshine the entire Universe (MacFadyen & Woosley, 1999); the magnetosphere of the Crab pulsar can produce nanosecond-duration pulses with brightness temperatures of 10^{38} K (Hankins et al., 2003); thermonuclear explosions up to a month long have been observed in X-ray binaries (Sokoloski et al. 2006, Bode et al. 2006); and large variations in accretion rate have been observed in accreting black holes, ranging in size from Galactic microquasars (Fender et al., 1999), to supermassive black holes (e.g. Gezari et al. 2006, Błażejowski et al. 2005).

Not only can transient emission signpost extreme events, but many aspects of the interstellar and intergalactic media, and of the Universe as a whole, can be revealed through measurements in the time domain. For example, the presence of massive compact halo objects (Alcock et al., 1997) and planets (Gould, 2005) can be inferred from gravitational microlensing of starlight around massive objects; the expansion history of the Universe can be measured using light curves of Type Ia supernovae (Riess et al. 1998, Perlmutter et al. 1999); and interstellar scintillation (ISS) of compact radio sources can be used to infer the distance, velocity and turbulence of otherwise invisible ionised clouds in the Milky Way (Bignall et al., 2003).

Low-energy radio photons can be easily produced by both thermal and non-thermal processes, and in contrast to optical and X-ray wavelengths, the Universe is essentially transparent to these photons. As a vast range of objects produce radio emission, and these signals are detectable over much of the observable Universe, radio variability surveys are an effective method of discovering and characterising such sources and propagation phenomena. Furthermore, important propagation effects such as interstellar scintillation are only observed in the radio.

While radio observations are often used for follow-up of detections at other wavelengths, blind radio surveys have also been a fruitful method for discovering new objects and phenomena in their own right. Gregory & Taylor (1986) discovered a number of flaring radio stars in a survey of the Galactic plane, while Bower et al. (2007), with 944 images of a single field, discovered ten faint transient sources, eight of which had no optical counterparts. Carilli et al. (2003) also discovered a number of highly variable sources with observations of a single field. Levinson et al. (2002) covered a large fraction of the sky with two epochs by comparing the NVSS (Condon et al., 1998) and FIRST (Becker et al., 1995) images, and was able to set constraints on the GRB beaming fraction. Croft et al. (2010) set an upper limit on the rate of transients with flux densities greater than 40 mJy of 0.004 deg^{-2} by comparing a 690 square degree image obtained with the Allan Telescope Array (ATA) at

1.4 GHz with the NVSS.

A number of transient source detections remain unexplained. Matsumura et al. (2009) have discovered a number of > 1 Jy transient sources with ~ 1 day time-scales at a range of Galactic latitudes. Lorimer et al. (2007) discovered a 30 Jy transient source at 1.4 GHz with a 5 ms time-scale, but there is evidence that this may have been atmospheric (Burke-Spolaor et al., 2011a). Lenc et al. (2008) found a ~ 100 mJy transient at 330 MHz during deep Very Long Baseline Interferometer (VLBI) observations, and Becker et al. (2010) found 39 variable sources in the Galactic plane with 3 epochs of 5 GHz Very Large Array (VLA) observations, most of which had no known counterparts at other wavelengths.

In this Chapter we present a survey for transient and variable sources at 843 MHz, with characteristic time-scales from days to years. We aim to characterise the variable and transient radio sky at this frequency, to enumerate the most extreme variable sources, to find transient sources and to develop techniques suitable for upcoming radio surveys.

For our analysis we are using data taken by the Molonglo Observatory Synthesis Telescope (MOST). The MOST and its predecessor, the Molonglo Cross, have been at the forefront of research in this area over several decades. Using the Molonglo Cross, Hunstead (1972) was the first to observe low frequency variability, at 408 MHz. In an archival survey of calibrator measurements, Gaensler & Hunstead (2000) found that one-third of bright point sources at 843 MHz were variable, and a weak Galactic latitude dependence was demonstrated, indicating that interstellar scintillation was at least partly responsible. The MOST has been used to discover several hundred pulsars in a number of surveys (Manchester, 1985), and a survey for short-duration transients has been performed by Amy et al. (1989) but with a null result. As a follow-up instrument, the MOST was the first telescope to detect prompt radio emission from SN1987A (Turtle et al. 1987), and it has been used to monitor a number of Galactic accreting systems (Hannikainen et al., 1998), a brightening supernova remnant (Murphy et al., 2008) and a magnetar flare (Gaensler et al., 2005).

The MOST archive is a unique resource that is able to address the limitations of sky coverage, sensitivity and cadence that have accompanied other blind surveys. In addition, the experience gained from analysing wide field-of-view images provides an opportunity to develop techniques suitable for upcoming wide-field transient and variability surveys such as the Variables And Slow Transients (VAST) survey for the Australian Square Kilometre Array Pathfinder (ASKAP) (Chatterjee & Murphy, 2010) and the LOFAR Transients Key Project (Fender, 2007).

In §2.2 we describe the MOST and its image archive. In §2.3 we describe our method of extracting light curves from this archive and in §2.4 we present the results of applying our method, including quality checks and selected sources. In §2.5 we discuss our results, and in §2.6 we discuss some

Table 2.1: Technical specifications of the MOST. The Declination range is given for a fully synthesised (12 h) image. The sensitivity and field sizes are mode-dependent and shown in Table 2.2.

Parameter	Value
Centre Frequency	843 MHz
Bandwidth	3 MHz
Polarisation	Right Hand Circular (IEEE)
Declination range for full synthesis	-30° to -90°
Synthesised beam	$43 \times 43 \text{ csc } \delta \text{ arcsec}^2$
Restoring beam	$45 \times 45 \text{ csc } \delta \text{ arcsec}^2$
Area per pixel	$11 \times 11 \text{ csc } \delta \text{ arcsec}^2$
Dynamic Range (typical)	100:1

noteworthy sources. We draw our conclusions in §2.7.

2.2 OBSERVATIONS WITH THE MOLONGLO OBSERVATORY SYNTHESIS TELESCOPE

The MOST is located near Canberra, Australia and was constructed by modification of the East-West arm of the former One-Mile Mills Cross telescope. The MOST is an East-West synthesis array comprising two cylindrical paraboloids each of dimension $778 \text{ m} \times 12 \text{ m}$ separated by 15 m. Radio waves are received by a line feed system of 7744 circular dipoles. The telescope is steered in the North-South axis by mechanical rotation of the paraboloids about their long axis, and in the East-West axis by phasing the feed elements along the arms. By tracking the field over 12 h, a full synthesis image can be formed. The near-continuous $u-v$ coverage from 15 m to 1.6 km results in good response to complex structure and low sidelobe levels. Technical specifications are shown in Table 2.1. The MOST has been described in detail by Mills (1981) and Robertson (1991).

Since 1986, the MOST has observed a single field for a 12 h synthesis almost every weeknight and often during the day on weekends. At the beginning and end of each 12 h synthesis, a set of up to 8 different calibrator sources is observed. Calibrator measurements are discarded if they do not pass a number of checks, and the remaining measurements are averaged to obtain gain and pointing solutions for the beginning and end of the observation. The gain and pointing solutions are linearly interpolated over the synthesis time between two calibrator observations. The full list of calibrators is described by Campbell-Wilson & Hunstead (1994b). Known variable calibrators were removed from the list following the analysis of Gaensler & Hunstead (2000).

Our present analysis is performed on the final images processed according to the procedure described by Green et al. (1999) and Bock et al. (1999).

2.2.1 Observing modes

The MOST is capable of observing in a number of different modes depending on the desired signal-to-noise ratio and field size, as summarised in Table 2.2. The MOST field of view is elliptical with the major axis aligned in the North-South direction and size $a = a_0 \csc |\delta|$, where δ is the declination of the field centre and a_0 is the mode-specific minor axis size. The area of an image is approximately equal to $\pi a_0^2 \csc |\delta|$ and the spatial resolution of each image is the same for each mode at $45 \times 45 \csc |\delta|$ arcsec. Examples of images observed in two different modes are shown in Fig. 2.1.

All MOST images are primary-beam corrected, which results in increased noise towards the edges of images. The increased noise is most pronounced in the wide-field ‘I’ images, which are larger than the half-power point of the primary beam. Therefore, we define an ellipse-shaped available area, parametrised by the minor axis shown in Table 2.2, outside which we ignore any detections and measurements.

We define an image as ‘usable’ if it passes the visual inspection and post-imaging calibration procedure described in Section 2.3.4.

2.2.2 Archive coverage

The primary survey conducted with MOST was the Sydney University Molonglo Sky Survey (SUMSS; Mauch et al. 2003), which covers the sky south of -30° excluding the Galactic plane ($|b| < 10^\circ$), and was carried out between mid-1997 and 2006. The MOST was also used to conduct two Molonglo Galactic Plane Surveys: MGPS-1 (Green et al., 1999), covering the Galactic plane from $245^\circ \leq l \leq 355^\circ$ and $|b| \leq 1.5^\circ$, and MGPS-2 (Murphy et al., 2007) covering $245^\circ < l < 365^\circ$ and $|b| < 10^\circ$. There has also been a wide range of directed observations and monitoring programs. Here we analyse 7227 images, observed between 1986 December 18 and 2008 August 28.

A plot of the total and usable sky coverage of the archive is shown in Fig. 2.2. All of the sky south of -30° has been covered at least once due to the SUMSS survey, and the Galactic plane has been observed twice with the MGPS-1 and MGPS-2 surveys. Some regions have been covered multiple times due to monitoring programs (e.g. the microquasar GX339–4 has been observed 14 times) or deep imaging (e.g. the Hubble Deep Field South, which has 52 usable epochs). A number of images were excluded from further analysis due to poor quality discovered during visual inspection, or difficulty in obtaining a post-imaging calibration, resulting in the usable coverage being significantly less than the total coverage.

For the purposes of comparing survey areas, we define the two-epoch equivalent coverage area as:

Table 2.2: MOST observing mode, image parameters (see Section 2.2.1) and number of images in the archive. The columns are: the mode code, typical Root Mean Square (RMS) of the noise in the image, full and available minor axes, total number of images in the archive and equivalent area (assuming the available minor axis) and number and equivalent area of images for which visual inspection and post-imaging calibration were successful (see Section 2.3.4). The ‘H’ and ‘I’ modes have only been available since the 1993 wide-field upgrade described by Large et al. (1994).

Mode	Typical RMS (mJy beam ⁻¹)	Semi-minor Axis ($a_0/2$)		Total Images		Usable Images	
		Full (deg)	Available (deg)	Count	Area (deg ²)	Count	Area (deg ²)
B	0.8	0.20	0.19	928	220.3	3	0.4
C	1.1	0.39	0.37	85	141.2	8	4.3
D	1.3	0.58	0.52	2292	3045.8	872	1028.4
H	1.8	1.18	1.00	6	29.5	0	0.0
I	2.0	1.36	1.02	3915	18156.4	2128	9266.1




Figure 2.1: Examples of MOST images from the MGPS surveys. Left: An ‘I’ field from MGPS-2, Right: A ‘B’ field from MGPS-1. Both images are centred on $l = 333.9^\circ$, $b = -4.4^\circ$. The grey scale is the same for each image ranging from -5 mJy beam^{-1} to $+15 \text{ mJy beam}^{-1}$. The resolution of each image is the same ($45 \times 61 \text{ arcsec}^2$), and the measured background RMS values are 1.5 and $0.85 \text{ mJy beam}^{-1}$ respectively. The ‘B’ field overlaps the ‘I’ field in the region shown. A number of typical artefacts are visible, including radial spokes from bright sources, regions of negative emission, grating rings and increased noise towards the edge of the ‘I’ field.

$$A_{\text{two-epoch}} = \sum_{i=1}^{N_{\text{pos}}} A_i (N_{\text{epoch},i} - 1), \quad (2.1)$$

where N_{pos} is the total number of independent directions on the sky measured in the survey, A_i is the area of a particular direction of interest, and $N_{\text{epoch},i}$ is the number of epochs for which a particular direction has been measured.

A comparison of the two-epoch equivalent coverage area of the MOST archive with other blind radio variability surveys is shown in Fig. 2.3 and Table 2.3. The MOST archive covers $\sim 10^4 \text{ deg}^2$ once, and $\sim 2 \times 10^3 \text{ deg}^2$ twice. The MOST archive has a maximum of 52 epochs on a small ($\sim 0.1 \text{ deg}^2$) patch of sky. At the large-epoch extreme, Bower et al. (2007) covered a very small (0.024 deg^2) patch of sky 844 times. Plots of the two-epoch equivalent coverage area versus declination and Galactic latitude of the MOST archive are shown in Fig. 2.4.

The interval between epochs covers a wide range. Figure 2.5 shows a

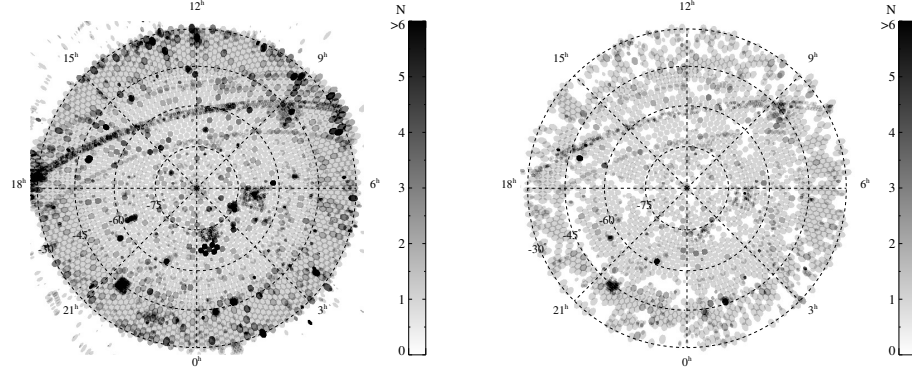


Figure 2.2: Coverage of the MOST archive in a Lambert equal area projection sampled on 0.1° grid and centred on the South Celestial Pole. Left: the entire archive. Right: usable images after visual inspection and post-imaging calibration (see Section 2.3.4). Some 4200 of the 7200 images were not usable and appear only in the left panel. The greyscale is number of epochs with white being no images, grey being one or more images. The shape of each image is the available size shown in Table 2.2, not the full image size. The SUMSS survey covers the sky south of -30° , and the Galactic plane surveys (MGPS-1 and MGPS-2) can be clearly seen as a band from the left-middle to top-right of both images. Repeated fields include the Hubble Deep Field South (52 usable epochs), SN1987A (13 usable epochs), GX339-4 (13 usable epochs) and GRO1655-40 (15 usable epochs).

Table 2.3: Two-epoch coverage area ($A_{\text{two-epoch}}$) for a range of recent and proposed blind radio variability surveys. VAST-Wide and VAST-GP are the wide-field and Galactic plane surveys of the Variables And Slow Transients (VAST) project, planned for the Australian Square Kilometre Array Pathfinder (ASKAP).

Survey	$A_{\text{two-epoch}} (\text{deg}^2)$
Carilli et al. (2003)	0.0008
Bower et al. (2007)	28
This Work	2776
Levinson et al. (2002)	5990
Croft et al. (2010)	6900
PiGSS (Bower et al., 2010)	10000
VAST-GP (proposed)	38250
VAST-Wide (proposed)	7290000

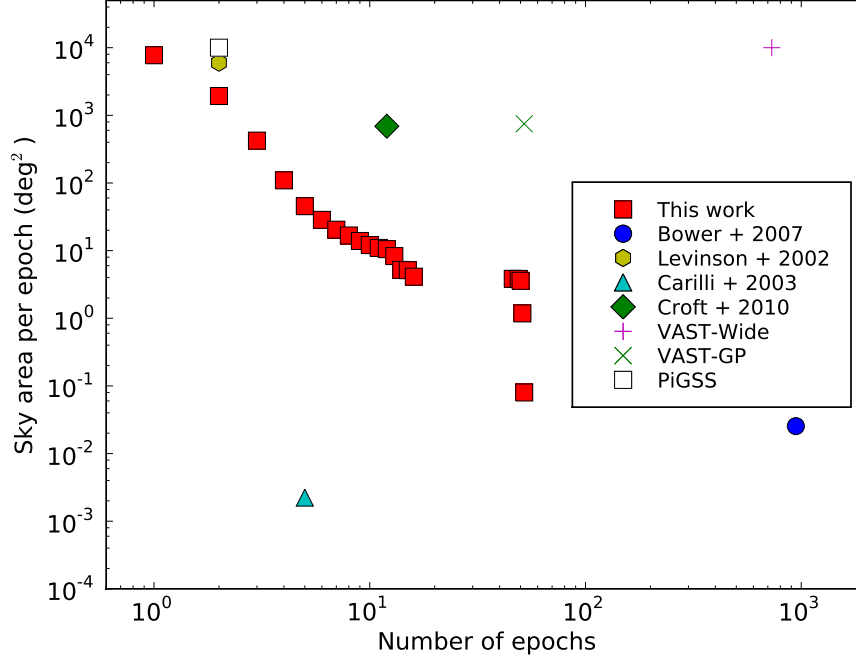


Figure 2.3: Area of epoch vs. number of epochs for blind radio transient surveys. At the extremes, Bower et al. (2007) covered 0.03 deg^2 with 844 epochs, Levinson et al. (2002) covered 5990 deg^2 with 2 epochs, Carilli et al. (2003) covered a 0.002 deg^2 with 5 epochs and Croft et al. (2010) covered 690 deg^2 with 12 epochs. This work, shown in red squares, spans some of the intervening parameter space. VAST-Wide and VAST-GP are the proposed wide-field and Galactic plane surveys of the Variables And Slow Transients (VAST) project, planned for ASKAP, and the Pi GHz Sky Survey (PiGSS) (Bower et al., 2011, 2010) was performed with the Allen Telescope Array.

histogram of the number of intervals between non-redundant pairs of measurements of each source as a function of the interval in days. At the two extremes, Figure 2.5 shows that there are over two thousand pairs of measurements separated by one day, and the same number of measurement pairs separated by over 20 years. The archive is most sensitive to changes in flux density on the time-scale of around one thousand days, where it contains over ten thousand pairs of measurements. The results for transient and variable radio sources are discussed in Section 2.5.2.

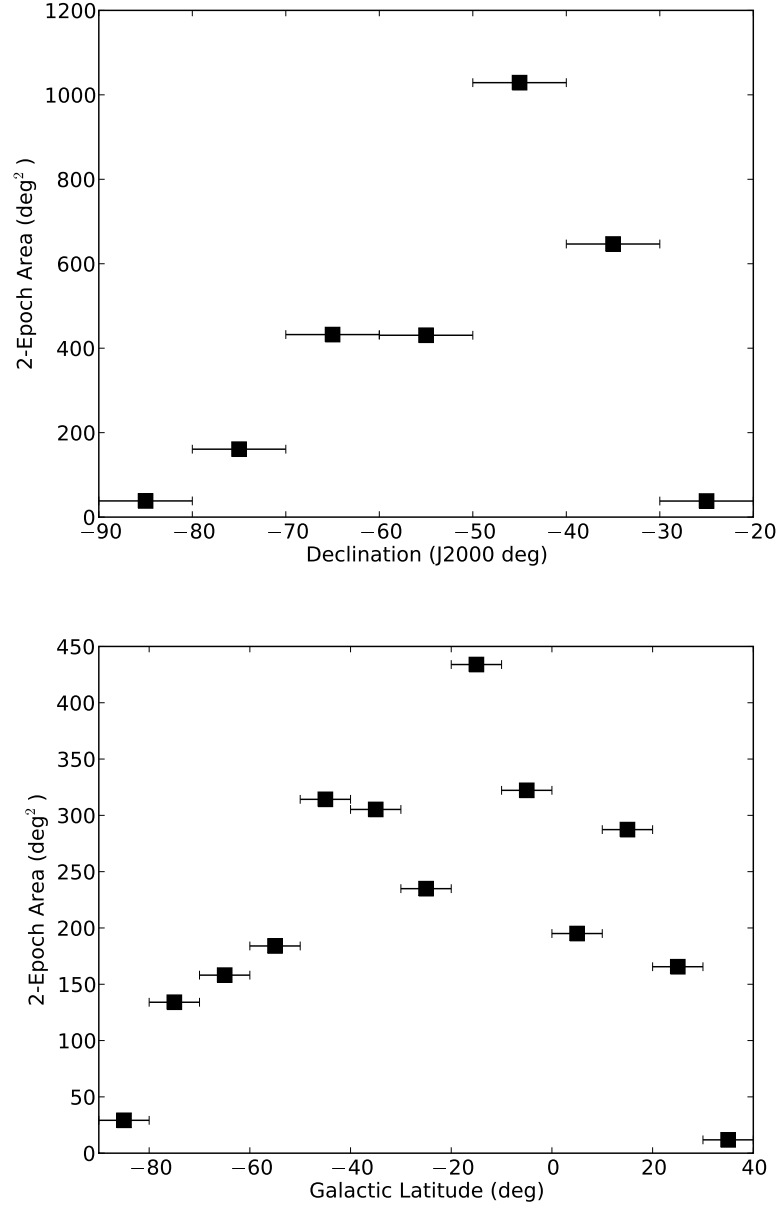


Figure 2.4: Two epoch coverage vs. declination (above) and vs. Galactic latitude (below). The horizontal error bars denote the bin widths of 10° .

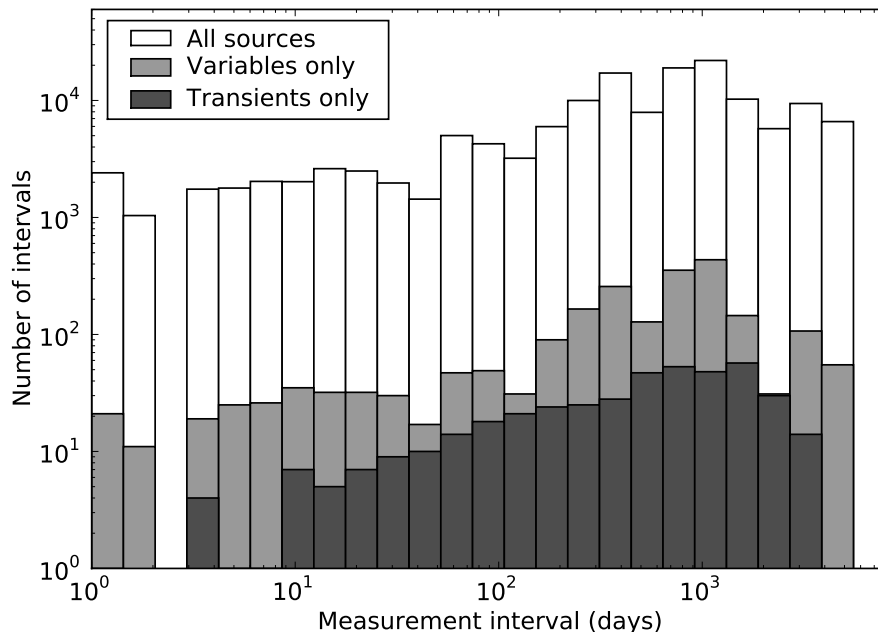


Figure 2.5: Histogram of the number of non-redundant pairs of measurements in the light curves of all sources in the sample, the variable sources only, and the transient sources only, binned by the interval between the measurements in days.

2.2.3 Image artefacts and systematic errors

2.2.3.1 Visually obvious artefacts

MOST images suffer from a number of image artefacts that make accurate automatic point source detection and fitting difficult (Mauch et al., 2003). In particular, some images are contaminated by extended emission (especially in the Galactic plane), with negative bowls or radial spokes around bright sources, grating rings due to the telescope geometry and bands across images due to solar and terrestrial interference (see, for example Fig. 2.1). To mitigate the effects of these artefacts, we discarded clearly affected images by visual inspection, considered only point source fits with small formal errors, and measured the RMS and mean in the region around each detection to correct for localised offsets and noise. Extended emission was not explicitly flagged but point sources with small formal errors of fit were rarely found close to regions of extended emission.

2.2.3.2 CLEAN artefacts

As for any interferometer, poor dirty beam modelling can result in negative flux around sources after CLEANing (see Green et al. 1999 for a discussion). An example of this effect is shown in Fig. 2.6. When performing automatic source fitting, the negative surface brightness around the edge of the source reduces the integrated flux density to a value significantly less than the peak flux density, which should not occur if CLEAN is using the correct beam. In a sample of 6767 otherwise good point source fits in the range $0\text{ h} < \text{RA} < 10\text{ h}$, only 45 per cent had peak and integrated flux density measurements in formal agreement, and more than half of these had a peak value that significantly exceeded the integrated flux density. Assuming that variable and transient sources must be point sources at the MOST resolution, integrated flux densities were not used in the subsequent analysis, with peak values used throughout.

2.2.3.3 Variable calibrators

Gaensler & Hunstead (2000) found that, between 1984 and 1996, 18 of the 55 Molonglo calibrators were variable on characteristic time-scales between 300 and 3000 days. The Molonglo calibration procedure averages the measurements of up to 8 calibrator sources, so variation of a single source is expected to have a relatively small effect on the flux density scale calibration. The characteristic time-scale was much longer than the synthesis time so we expect no variability on that time-scale. In 1996, the variable sources were removed from the calibrator list and the post-1996 flux density calibration (which included SUMSS and MGPS-2) can be considered more reliable. In any case, the flux density scale for every image was calibrated against SUMSS/MGPS-2 using the post-imaging calibration procedure (see §2.3.4), so any residual error in the flux density scale should be removed by this technique.

2.2.3.4 Radio frequency interference

Observations performed after 2007 January at the Molonglo site have been subject to interference from strong local sources of Radio Frequency Interference (RFI). The interference affects almost all wide-field ‘I’ images observed after 2007 January, but ‘B’ fields are much less affected. No other modes were used during this period. The strongest RFI is usually excised during the reduction process, which results in incomplete coverage of the uv plane and higher sidelobe levels. Remaining low-level RFI resulted in significantly increased image noise levels. Occasionally, the interference can manifest itself as gain variations across the image, which is difficult to identify with visual inspection alone. Images affected in this way were detected and removed by the post-imaging calibration procedure (see §2.3.4).




Figure 2.6: Example of poor dirty beam modelling producing negative regions around SUMSS J022129–335825 after CLEANing. Fitting a 2D Gaussian in a $10 \times 10 \text{ pixel}^2$ region centred on the source results in an integrated flux density of 394.0 mJy which is less than peak intensity of $449.3 \pm 16.6 \text{ mJy}$. The size of the fitted Gaussian is consistent with the size of the synthesised beam (shown in the bottom left). The large difference in integrated and peak flux densities is due to poor dirty beam modelling.

2.2.3.5 Missing calibration values

The standard MOST image reduction pipeline assumes default values for calibration parameters if none are supplied. Inspection of image headers revealed that over 600 images had been reduced with default calibration values, rather than the measured values. We were able to re-reduce these images with the same calibration measurements used to produce the results of Gaensler & Hunstead (2000).

In spite of this effort, we found large variations in the flux density scale on some occasions, which we attribute to variable calibrators, wrongly applied calibration and interference during calibrator observations. Rather than discarding images that appeared good apart from their flux density scale, we included them in our analysis and calibrated them with the post-imaging

calibration procedure (see §2.3.4).

2.2.3.6 Resolution effects

Many radio transient and variability surveys, such as that performed by Levinson et al. (2002), have suffered from the potential for mismatched resolutions to cause systematic errors in the analysis. Mismatched resolutions can cause a number of systematic effects. In particular, higher resolutions can resolve a large source into multiple components, as well as resolve out the extended flux, resulting in a lower flux density measurement.

The MOST synthesised beam is fixed in the East-West direction, but elongates in the North-South direction as a function of the declination of the image centre. The difference in resolution is largest when a source sits at the southern tip of an image centred on the upper declination limit ($\delta_{\max} = -30^\circ$), and at the northern edge of an image centred to the south, with declination $\delta_2 \simeq \delta_{\max} - 2a_{0,\max} \csc |\delta_{\max}|$. The largest minor axis occurs for the ‘I’ images, with $a_{0,\max} = 1.36^\circ$ (see Table 2.2). The difference in resolution in the worst case is therefore $45(\csc |\delta_{\max}| - \csc |\delta_2|) \simeq 9.2$ arcsec in the north-south direction, which corresponds to a maximum difference in resolution of overlapping images of approximately 20 per cent.

By way of comparison, Levinson et al. (2002) compared NVSS and FIRST surveys, whose resolutions differed by a factor of 7 (45 to 5 arcsec). We suspect resolution difference affects a very small fraction of our measurements as only those sources in the narrow band of overlap between north-south adjacent images at northern declinations are affected. We also expect that multiple match and extended flux effects will be small, as the change in resolution in our worst case is only 20 per cent.

2.3 LIGHT CURVE EXTRACTION

There are two broad categories of techniques that can be applied to a set of astronomical images to detect variable and transient emission: image comparison and catalogue comparison. Image comparison, as used by Carilli et al. (2003), involves comparing an image of interest with a reference image, either by subtraction or division. Source finding can then be performed on the difference or quotient image. Catalogue comparison makes source catalogues by performing source finding on each image and forming a light curve by measuring the flux density separately on each image. The light curve is then subjected to a statistical analysis. Bower et al. (2007) used a cataloguing approach.

An image comparison approach is not feasible for the MOST archive since MOST image artefacts change as a function of pointing centre and observing mode, which would result in a large number of false detections when a reference

image is compared with a candidate image. Instead we use a catalogue-based approach as discussed in the following sections.

2.3.1 Source detection

Source detection was performed using the SFIND task (Hopkins et al., 2002) in the MIRIAD package (Sault & Killeen, 2009). SFIND uses the False Discovery Rate method for controlling the fraction of false positives when performing multiple hypothesis testing. It accepts two parameters: the size of the square box over which it should measure the background RMS surface brightness, and the percentage of false pixels to accept when applying the False Discovery Method. We chose a box size of 50×50 pixel² and a false pixel rate of 10 per cent. SFIND automatically produces Gaussian source fits and background RMS for each detected source. We used only the fitted positions. We separately measured the flux densities as described in §2.3.3 below, as we found the fitted flux densities and background RMS from SFIND to be occasionally unreliable.

2.3.2 Source association

An initial catalogue was created by merging the SUMSS catalogue⁶ and the MGPS-2 catalogue⁷ and source detection was performed independently on each image as described in § 2.3.1. We associated each detection with a catalogued source if the positions of the detected and catalogued sources were coincident within 10". If there was no previously known source within 10", a new source was added to the catalogue.

2.3.3 Measurement

Once a source had been detected, we measured flux density at the position of the source on all images in the archive for which this position was in the field of view. We did this for a number of reasons:

1. To obtain fitted flux densities at all epochs.
2. To obtain upper limits on flux density for epochs in which a source was not present.
3. To obtain measurements of the surface brightness in a region around a source, which is useful for quantifying the RMS error of the measurement (for error estimates) and correcting for a non-zero mean level.

⁶Version 2.1 dated 2008 Mar 11

<http://www.physics.usyd.edu.au/sifa/Main/SUMSS>

⁷Version dated 2007 Aug 15

<http://www.physics.usyd.edu.au/sifa/Main/MGPS2>

Table 2.4: Threshold values of fitted formal errors used to select point sources for detection and post-imaging calibration.

Parameter	Formal error threshold
Major axis	$< 15''$
Minor axis	$< 12''$
Position in RA	$< 5''$
Position in Dec	$< 6''$
Peak Value	$< 10 \text{ mJy beam}^{-1}$

To obtain a post-imaging calibration solution for an image, we require accurate flux densities of bright point sources. To measure flux density for bright sources we fit a Gaussian source to a $10 \times 10 \text{ pixel}^2$ square box on the detected position. For weak sources ($< 5\sigma$), Gaussian source fitting occasionally produces large errors, partly due to the large number of free parameters. Therefore, to measure flux density of weaker sources, we also performed a 2D parabolic fit to a $3 \times 3 \text{ pixel}^2$ region centred on the source position. To quantify a non-detection, we measured the peak pixel in a $3 \times 3 \text{ pixel}^2$ square region centred on the source position, and to measure the local background we measured the RMS and mean flux density in an annular region of inside radius $66''$ and outside radius $99''$ (6 and 9 pixels in radius, respectively). We did not mask out other bright sources in the annular region, so measurements of the local background can be biased in densely populated areas.

2.3.4 Post-facto calibration

In our initial investigations, we found a large number of sources that appeared to have varying flux densities, which we found was due to variations in the gain of individual images.

In order to reduce the number of false variability detections due to this effect, we implemented a post-imaging calibration scheme by comparing the Gaussian-fitted peak flux density measurements of the bright sources on each image, against the SUMSS flux densities measured by Mauch et al. (2003) or the MGPS-2 flux densities measured by Murphy et al. (2007).

For each image, we performed a 2-stage fit for the gain. First, we selected all point sources that had formal fit errors within the bounds shown in Table 2.4, a peak flux density above 40 mJy, and a counterpart in SUMSS or MGPS-2. We then performed a linear least squares fit to the relationship between peak fitted flux density measurements and catalogued flux densities, weighted by the inverse of the quoted peak flux density measurement error (σ) from the catalogue. The fit was constrained to go through zero. We used the fitted Gaussian measurements because they provided a more accurate flux density measurement than the 2D parabolic fit for the bright sources that we used for the calibration. Gaussian measurements are also more sensitive to

image errors, thus enabling the rejection of images containing subtle image errors that affect a large fraction of sources. All points that deviated from the fit by more than 3σ were removed from the set and the fit was performed again on the remaining points. The gradient of this fit was used as the gain calibration factor, and was applied to all measurements derived from the image in question.

In addition to producing a gain estimate, this procedure also enabled us to remove poor quality images from the remaining analysis by measuring the goodness-of-fit, both before and after the deviant points were removed. As a measure of goodness-of-fit we adopted the reduced χ^2 , defined as

$$\tilde{\chi}_{\text{cal}}^2 = \frac{\chi_{\text{cal}}^2}{N_{\text{dof}}} = \frac{\chi_{\text{cal}}^2}{N_{\text{cal}} - 1}, \quad (2.2)$$

where N_{dof} is the number of degrees of freedom in the fit and N_{cal} is the number of sources being used for the calibration.

We discarded images for which the following conditions were all met:

- The initial value of $\tilde{\chi}_{\text{cal}}^2$ was greater than 30.
- More than 30 per cent of the sources were more than 3σ from the fit.
- The value of $\tilde{\chi}^2$ was greater than 5 after removing the sources which were 3σ from the fit.

The post-imaging calibration procedure assumes that the majority of sources in the sky are static over time-scales of several years, and that variations in flux density are not spatially correlated over the MOST field of view. As such, refractive scintillation caused by a nearby cloud of ionised gas covering a significant fraction of the MOST field of view is removed by the calibration and rendered undetectable. We consider this possibility physically unlikely.

Examples of the post-imaging calibration procedure applied to two images are shown in Fig. 2.7. A plot of $\tilde{\chi}_{\text{cal}}^2$ vs time is shown in Fig. 2.9, from which we were able to detect the interference that began affecting the telescope in 2007 by a group of images with large $\tilde{\chi}_{\text{cal}}^2$. This increased our confidence in the usefulness of the post-imaging calibration, in particular its ability to recognise poor quality images.

A histogram of the post-imaging calibrated gains, shown in Fig. 2.10, shows that the majority of gains are centred around 1 with scatter of the order of 20 per cent. This scatter demonstrates that the image gain would have been a significant source of systematic error, had it not been corrected for by the post-imaging calibration.

An example of the post-imaging calibration procedure applied to the light curve of a single source is shown in Fig. 2.8. The procedure removes the most discrepant points and corrects the remainder, which results in less scatter in the light curve.

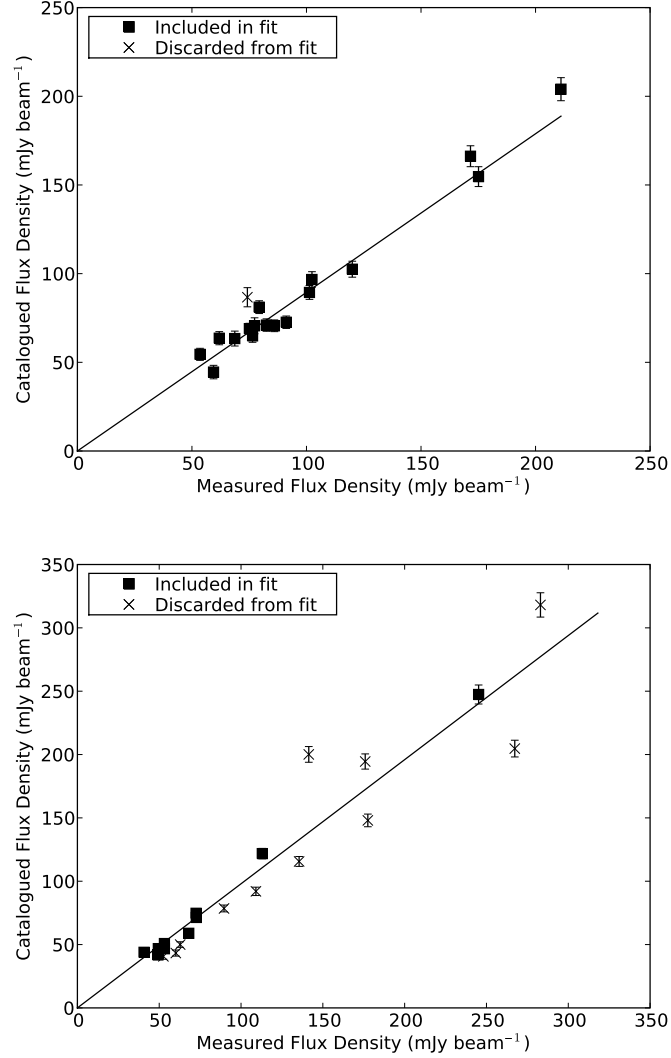


Figure 2.7: Examples of post-imaging calibration of two images. Catalogued flux densities (from SUMSS/MGPS-2) are plotted against the peak flux density measured with a Gaussian fit on a 10×10 pixel² box centred on each source position. The solid line is the line of best fit after removing the sources more than 3σ from the initial best fit line including all points (where σ is the quoted error from SUMSS/MGPS-2). The squares are the points included in the fit and the crosses are the sources removed. Top panel: An ‘I’ image (observed 2000 March 25) with good fit and gain of 0.89. Bottom panel: An ‘I’ image (observed 1999 November 9) with a poor fit due to large discrepancies between measured and catalogued flux densities for many sources, and gain of 0.98.




Figure 2.8: Light curves illustrating the effect of post-imaging calibration on SUMSS J201139–561858. Top panel: Pre-calibrated light curve, centre panel: calibrated light curve, bottom panel: difference between the calibrated and uncalibrated light curves. The uncalibrated light curve has a modulation index (see Equation 2.9) of 0.181 compared to that of the calibrated light curve of 0.0842. Two images were discarded from the light curve by the post-imaging calibration.

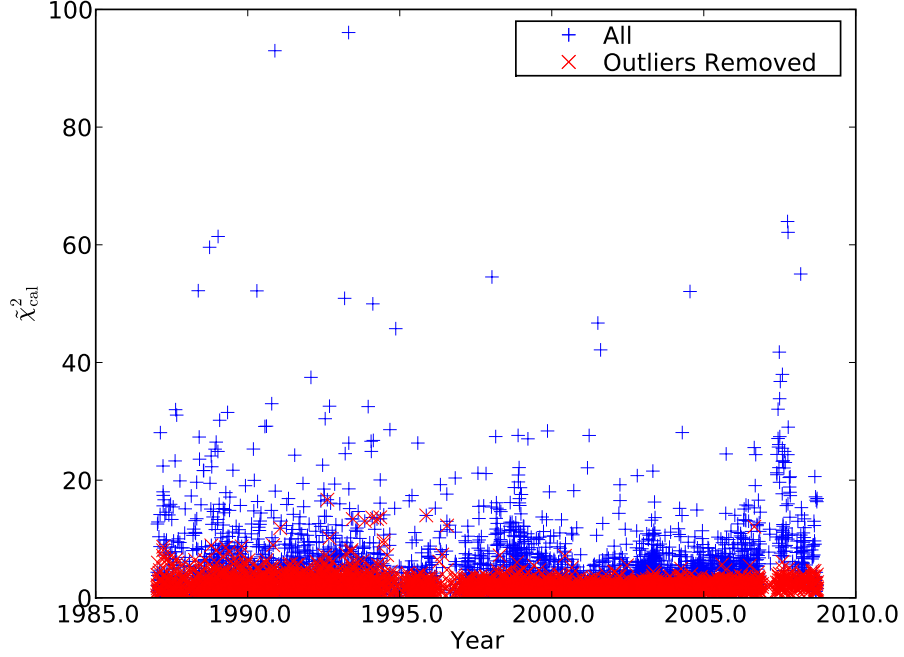


Figure 2.9: $\tilde{\chi}_{\text{cal}}^2$ of the post-imaging calibration fit versus image observation date, before and after 3σ outliers were removed. Note the significant number of images with large value of reduced $\tilde{\chi}_{\text{cal}}^2$ (before outlier removal) in 2007 and 2008, when terrestrial interference was affecting the telescope.

2.3.5 Light curve characterisation

2.3.5.1 Variability Classification

Astronomical sources exhibit a wide variety of light curve shapes. For example, interstellar scintillation produces random variations around a mean, whereas stellar flares are episodic with short duty cycles and flux densities significantly above a quiescent level.

Determining whether a given light curve matches any of a wide variety of light curve shapes, and quantifying the fidelity of the match, are beyond the scope of this work. For simplicity, we adopt the classification measure of Kesteven et al. (1977) to measure random variability in the light curve, namely the χ^2 of the residuals of a weighted fit to a line of constant flux density. This is computed according to the following expressions:

$$\chi_{\text{lc}}^2 = \sum_{i=1}^N \frac{(S_i - \bar{S})^2}{\sigma_i^2}, \quad (2.3)$$




Figure 2.10: Histogram of post-imaging calibrated gains for 2999 usable images. The dashed line is a Gaussian distribution fit to the gains, with a mean gain of 0.985 and standard deviation of 0.056.

where:

$$\bar{S} = \frac{\sum_{i=1}^N S_i / \sigma_i^2}{\sum_{i=1}^N 1 / \sigma_i^2}, \quad (2.4)$$

$$\sigma_i = \sqrt{(g\sigma_{bg,i})^2 + (0.05S_i)^2}, \quad (2.5)$$

$$S_i = g(S_i^{raw} - \rho_i), \quad (2.6)$$

and \bar{S} is the weighted mean flux density, σ_i is the estimated flux density error, S_i^{raw} is the flux density measured on the image with the 2D parabolic fit, g is the gain coefficient calculated by the post-imaging gain calibration and ρ_i and $\sigma_{bg,i}$ are respectively the mean and RMS surface brightness of an annular region around the source. The weighting factor of $1/\sigma_i^2$ captures both the noise level in the surrounding region and the absolute flux density calibration error from the MGPS-2 and SUMSS catalogues, assumed to be 5 per cent. For a source with a peak flux density at the detection limit of 14 mJy, σ_i is dominated by the background RMS term and has a typical value

$\sigma_i \simeq 3 \text{ mJybeam}^{-1}$. For a source with flux density of 100 mJy the flux density calibration term dominates and results in a typical error $\sigma_i \simeq 6 \text{ mJy}$.

To determine whether a source is variable or not, we work from the null hypothesis that all sources are static with added Gaussian measurement error. Applying the measured χ_{lc}^2 to the analytic χ^2 cumulative distribution function (CDF) yields the probability of rejecting the null hypothesis, i.e. the probability that the source is genuinely variable, which we write as $P(\chi_{\text{lc}}^2)$. Working with probabilities rather than χ^2 values also enables us to compare the statistics of light curves with differing number of measurements, which is absorbed in the CDF. Those light curves whose probability exceeds a certain threshold can be classified as variable.

In past investigations such as those of Kesteven et al. (1977), Gregory & Taylor (1986) and Seielstad et al. (1983) a probability threshold of 99 per cent was applied to derive a set of variable sources. In the case of Gregory & Taylor (1986), the 99 per cent threshold left 4 per cent of the sources classified as variable. Of these sources, 1 per cent of the total, or 25 per cent of the classified variables must be false positives, simply from the choice of a low threshold.

As we have almost 30 000 light curves, a 25 per cent false positive rate is too high. To arrive at a more rigorous threshold we define a confidence measure that is the excess sources over expected number given the null hypothesis, expressed as a percentage. Specifically, we use:

$$C(T) = \frac{N_{\text{actual}}(T) - N_{\text{noiseonly}}(T)}{N_{\text{actual}}(T)}, \quad (2.7)$$

where $N_{\text{actual}}(T)$ is the number of sources with light curves whose $P(\chi_{\text{lc}}^2)$ is greater than T , $N_{\text{noiseonly}}(T) = N_{\text{sources}} (1 - P(\chi_{\text{lc}}^2 > T))$ is the number of sources expected if there were only noise in the sample, and N_{sources} is the total number of sources in the sample. The confidence measure $C(T)$ in Equation 2.7 represents the fraction of genuinely variable sources contained in the set of sources with $P(\chi_{\text{lc}}^2) > T$.

2.3.5.2 Variability measures

Previous works have adopted a number of different metrics to quantify variability. To enable comparison with previous work, we adopt two different metrics. The first is the fractional variability defined as:

$$\mathcal{V} = \frac{S_{\text{max}} - S_{\text{min}}}{S_{\text{max}} + S_{\text{min}}}, \quad (2.8)$$

where S_{max} and S_{min} are respectively the maximum and minimum flux densities. This measure is most suited to characterising sources with few measurements or flaring behaviour and was used by Gregory & Taylor (1986).

The second metric is the modulation index, defined as:

$$m = \sigma / \bar{S}, \quad (2.9)$$

where σ is the standard deviation of the light curve and \bar{S} is the weighted mean defined in Equation 2.4. This metric is more suited to measuring random variability in data sets with a large number of measurements and was used by Gaensler & Hunstead (2000).

2.3.5.3 Detections

A measurement must satisfy a number of criteria in order to be considered an acceptable measurement of a point source rather than a result of an image artefact.

The analysis of the SUMSS catalogue by Mauch et al. (2003) showed that the SUMSS survey was complete above 15 mJy and we found a number of convincing detections above 14 mJy. Below 14 mJy the candidates were dominated by image artefacts. Gaussian fits can have major and minor axes significantly different from the synthesised beam, either because a source is partially resolved, or because of a strong image artefact. Restricting the major and minor axes of the fit to less than 20 per cent larger than the synthesised beam ensures the fit is a point source and reduces the number of false detections due to artefacts.

Therefore, we define a ‘detection’ as a Gaussian fit to a source that satisfies all of the following criteria:

- Fit errors are below the thresholds in Table 2.4,
- Fit position is less than $5''$ from the catalogued position (see § 2.3.2),
- Peak flux density ≥ 14 mJy,
- Ratio of fit major axis to synthesised beam major axis is < 1.2 ,
- Ratio of fit minor axis to synthesised beam minor axis is < 1.2 ,
- Image passes the post-imaging calibration.

2.3.5.4 Detection signal-to-noise ratio

To find transient sources, it is particularly important to measure the significance of a given detection or non-detection. We define the detection signal-to-noise ratio as $d = S_{\text{cenmax}} / \sigma_{bg}$ where S_{cenmax} is the maximum surface brightness in a $3 \times 3 \text{ pixel}^2$ box centred on the source position, and σ_{bg} is the RMS surface brightness (in mJy beam^{-1}) in an elliptical region centred on the source position.

2.3.6 List of variable and transient sources

We computed all metrics described in §2.3.5 for all sources and produced two separate lists of variable and transient sources. The list of variable sources was produced by finding all sources with:

- Light curve probability exceeding a threshold that yields a confidence metric greater than 90 per cent, i.e. $P(\chi_{lc}^2) > T_0$ where $C(T_0) > 90$ percent.
- A ‘detection’ on every image.

A variable source, therefore, has at least two detections and no non-detections. The list of transient sources was produced by finding all sources with:

- Minimum detection signal-to-noise ratio $d_{min} < 5.0$
- Maximum detection signal-to-noise ratio $d_{max} > 6.0$
- Flux density at the maximum detection signal-to-noise ratio epoch (measured with a parabolic fit) > 14 mJy
- $P(\chi_{lc}^2) > 99\%$ computed with S_{cenmax} as the flux density value.

A transient source, therefore, has at least one detection and at least one non-detection. We use the parabolic fit, as this was less prone to error at low signal-to-noise ratio than the Gaussian fit. We compute the χ_{lc}^2 for the transient list using the value of S_{cenmax} rather than the 2D parabolic fitted flux density as the former is the most robust measure of the upper limit to the flux density, and we apply the error as per Equation 2.5. We reduce the threshold on the value of $P(\chi_{lc}^2)$ with respect to the variable sources, because the flux density of the non-detections is not known.

2.4 RESULTS

Throughout this section we assume a Hubble constant $H_0 = 72 \text{ km s}^{-1}\text{Mpc}^{-1}$ and define the spectral index, α , according to $S \propto \nu^\alpha$ where S is the flux density and ν is the observed frequency.

When referring to a radio source, we use the popular name (e.g. SN1987A) if one exists. Otherwise, if the source appears in either the MGPS2 or SUMSS catalogues we use the MGPS2 or SUMSS naming convention, or if not, we use the source coordinates in the form ‘JHHMMSS-DDMMSS’ without prefix.




Figure 2.11: Histogram of the number of epochs measured per source in the variability and transient analysis. At the two extremes, there are 26000 sources with two measurements and 11 sources with 51 measurements. Gaps in the histogram indicate no sources with that number of epochs.

2.4.1 Data Quality

2.4.1.1 Number of usable images

Of the 7227 images in the archive, 2999 were retained after visual inspection and post-imaging calibration. The area of an equivalent two-epoch survey (see Equation 2.1) is 2775.7 deg^2 .

2.4.1.2 Number of measurements

There are 29730 sources that have light curves with measurements at two or more epochs. A histogram of the number of epochs measured on each source is shown in Fig. 2.11.

2.4.1.3 Analysis of the population of light curves

To check that the flux density error estimate is consistent over the population of light curves, we compare the measured and theoretical distributions of the

residuals from a line of constant flux density. As described in Equation 2.3, the residuals of the best-fitting light curve of constant flux density follow a χ^2 distribution with $N_{\text{meas}} - 1$ degrees of freedom. Fig. 2.12 shows the histograms of the χ_{lc}^2 values of the residuals for light curves in our variability sample. The theoretical and actual curves agree reasonably well, with the slightly larger number of sources at small values of χ_{lc}^2 indicating that the flux density errors have likely been marginally overestimated.

In order to compare distributions of sources with differing numbers of measurements, we can compute the probability of a given χ_{lc}^2 value occurring by chance by using the analytic χ^2 cumulative distribution function with the appropriate number of degrees of freedom. If the variation in the light curves is attributed entirely to Gaussian noise, the histogram of the computed probabilities should be uniform and independent of the degrees of freedom. Fig. 2.13 depicts the expected and actual probability histograms. There is a general trend in the actual histogram showing an excess of sources with high $1 - P(\chi_{\text{lc}}^2)$, which we attribute to overestimated flux density errors leading to smaller measured χ_{lc}^2 values. Because a large fraction of sources have a small number of measurements (see Fig. 2.11), at this degree of freedom a small value of χ_{lc}^2 corresponds to a higher value of $1 - P(\chi_{\text{lc}}^2)$ which in turn leads to the observed excess. In addition to the general trend, there is a single bin at the lowest value of $1 - P(\chi_{\text{lc}}^2)$, which we attribute to a population of genuinely variable sources. The genuinely variable sources are clearly seen as an excess of sources, departing from the theoretical curve as the probability is reduced.

Fig. 2.14 shows the confidence measure $C(T)$ vs $1 - P(\chi_{\text{lc}}^2)$ on a log scale. From Fig. 2.14 we deduce that applying a threshold of $T_0 = 1 - P(\chi_{\text{lc}}^2) = 4 \times 10^{-5}$ ensures that the fraction of spurious variable sources in the list of candidate variables is less than 10 per cent.

2.4.1.4 Primary beam effects

In the MOST archive, a significant fraction of the overlapping coverage is at the edges of adjacent images. Errors in primary beam correction, which increase with radial distance from the pointing centre, could therefore introduce false variability. To investigate this effect, we took the light curves of all sources appearing in images centred on $\alpha = 22\text{h}30\text{m}00\text{s}$, $\delta = -61\text{d}06\text{m}36\text{s}$ (J2000), in the Hubble Deep Field South (HDFS) region. We calculated the fraction of measurements that deviated by more than 20 per cent from \bar{S} (see Equation 2.4), binned by radial distance from the pointing centre, as shown in Fig. 2.15. We found no strong dependence with distance; we therefore conclude that the primary beam correction is accurate, and that the adopted available minor axis of the primary beam (see Table 2.2 and § 2.2.1), outside which we ignore measurements, is acceptable.

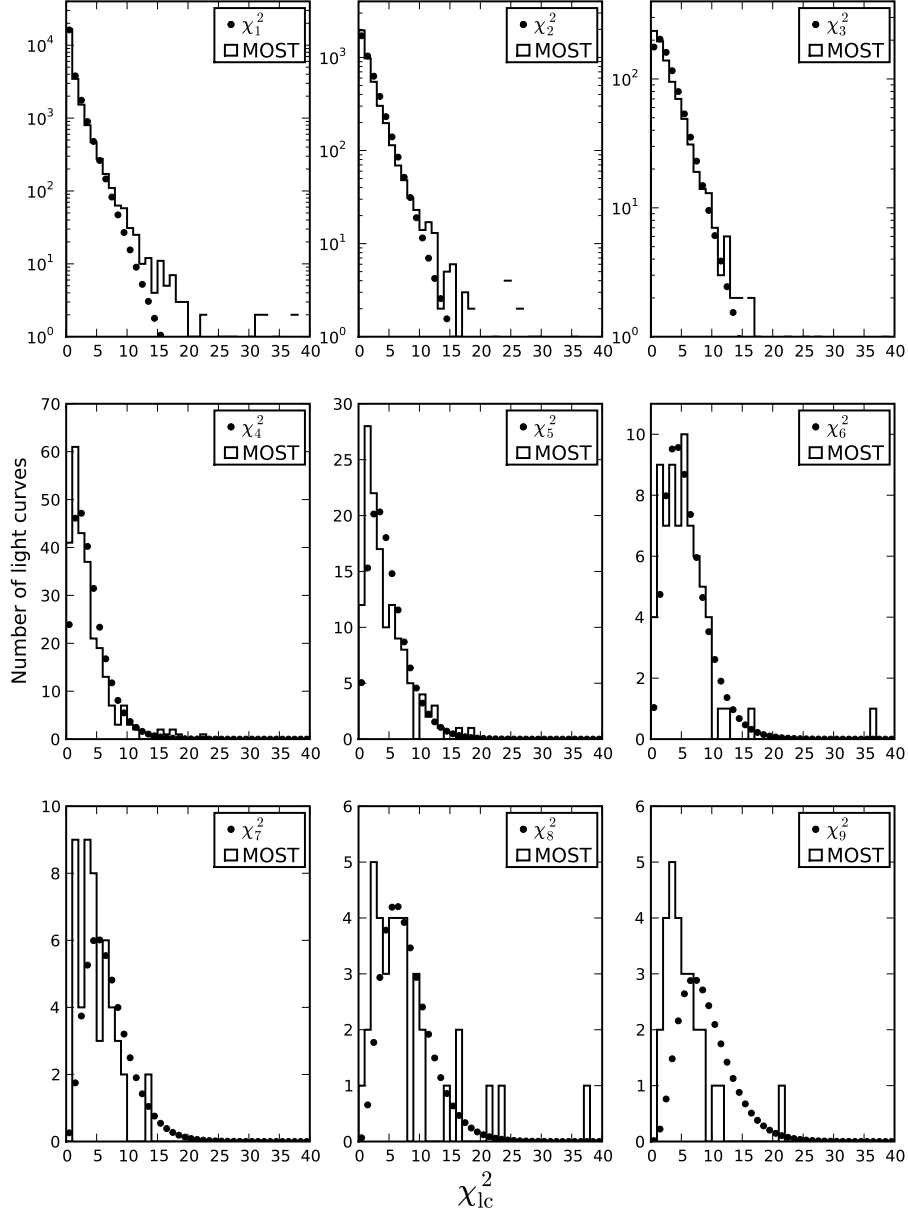


Figure 2.12: Histograms of χ^2_{lc} values for light curves with between 2 and 10 measurements (N_{meas}). Also plotted as filled circles is the theoretical $\chi^2_{N_{meas}-1}$ probability distribution with $N_{dof} = N_{meas} - 1$ degrees of freedom. The theoretical distribution is binned on the same intervals as the measurements.

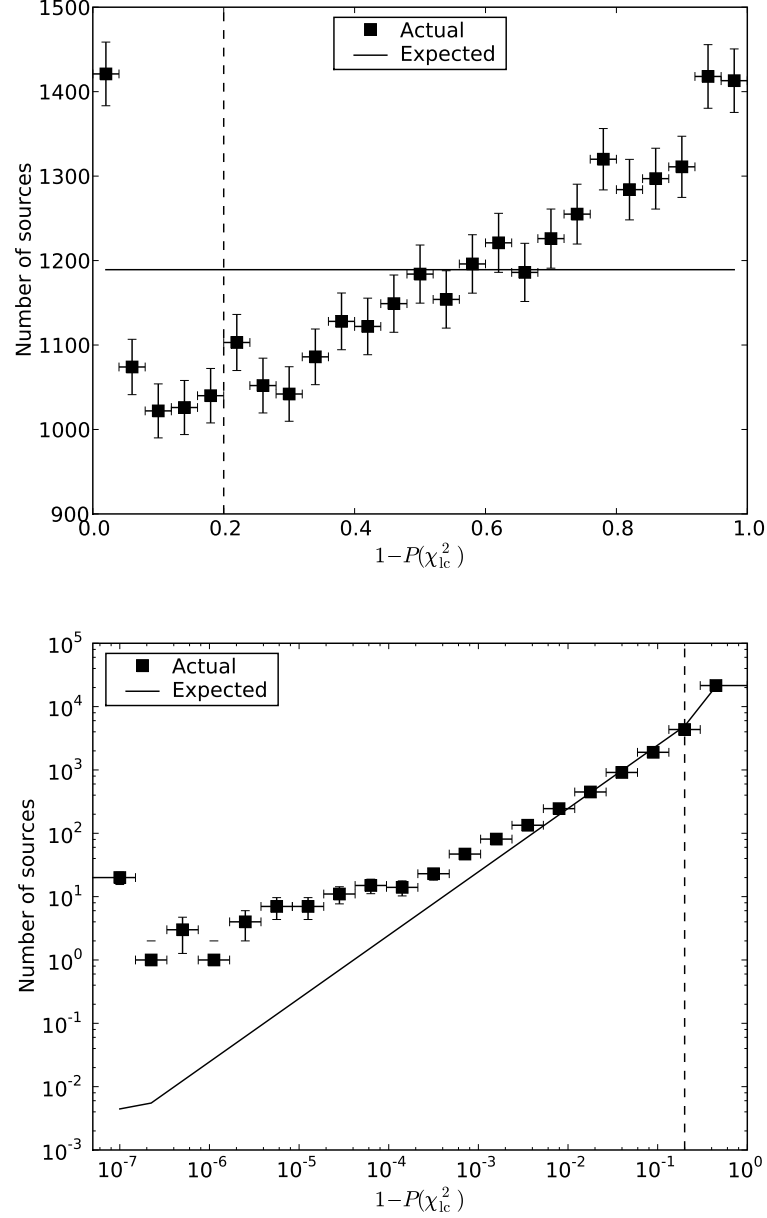


Figure 2.13: Top panel: Number of sources vs $1 - P(\chi_{ic}^2)$ on a linear scale. The y error bars are the Poisson errors and the x error bars indicate the bin widths. The solid line depicts the expected number of sources in each probability bin if the scatter in the light curves was due to Gaussian noise. The point at the far left indicates a population of variable sources and the excess over the solid line indicates that there are approximately 350 genuinely variable sources in our sample. The increasing general trend from left to right indicates that there is a deficiency of moderately variable sources and an excess of static sources, which agrees with Fig. 2.12. The region to the left of the dashed vertical line at $1 - P(\chi_{ic}^2) = 0.2$ indicates the baseline number of sources used to infer the actual number of variable sources (see §2.5.5). Bottom panel: The same data but re-binned on a log scale to illustrate how the number of sources departs from the expected number for small values of $1 - P(\chi_{ic}^2)$.




Figure 2.14: Confidence measure, $C(T)$ vs $1 - P(\chi^2_{lc})$ on a log scale. To achieve a fraction of genuine variable sources of 90 per cent requires $1 - P(\chi^2_{lc}) = 4 \times 10^{-5}$. The y errors are the Poisson errors propagated through the definition of the confidence metric, and the x errors indicate the bin width. The confidence is negative at large values of $1 - P(\chi^2_{lc})$ due to a deficiency of moderately variable sources at large values of $1 - P(\chi^2_{lc})$ (see Fig. 2.13.). The dashed vertical line indicates the adopted threshold of $1 - P(\chi^2_{lc}) = 4 \times 10^{-5} = T_0$, to the left of which the confidence of finding a genuinely variable source is greater than 90 per cent.

2.4.1.5 Declination effects

The sensitivity and resolution of MOST images is strongly declination dependent, with noise and beam size increasing towards Northerly declinations. Therefore, to ensure our sample is not contaminated by declination dependent systematic effects, we plot in Fig. 2.16 the fraction of transient and variable sources as a function of declination. We find no declination dependence.

2.4.2 Variable sources

Table 2.5 lists the 53 sources that we classify as variable according to the criteria described in §2.3.6. The Table also lists radio and optical counterparts




Figure 2.15: Plot of the fraction of flux density measurements that deviate from \bar{S} by more than 20 per cent, binned by distance from the image centre. Only images centred on the HDFs were considered, and 2 known variable sources were removed. The error bars are Poisson errors

where available. Light curves for each of these sources are shown in Appendix A.

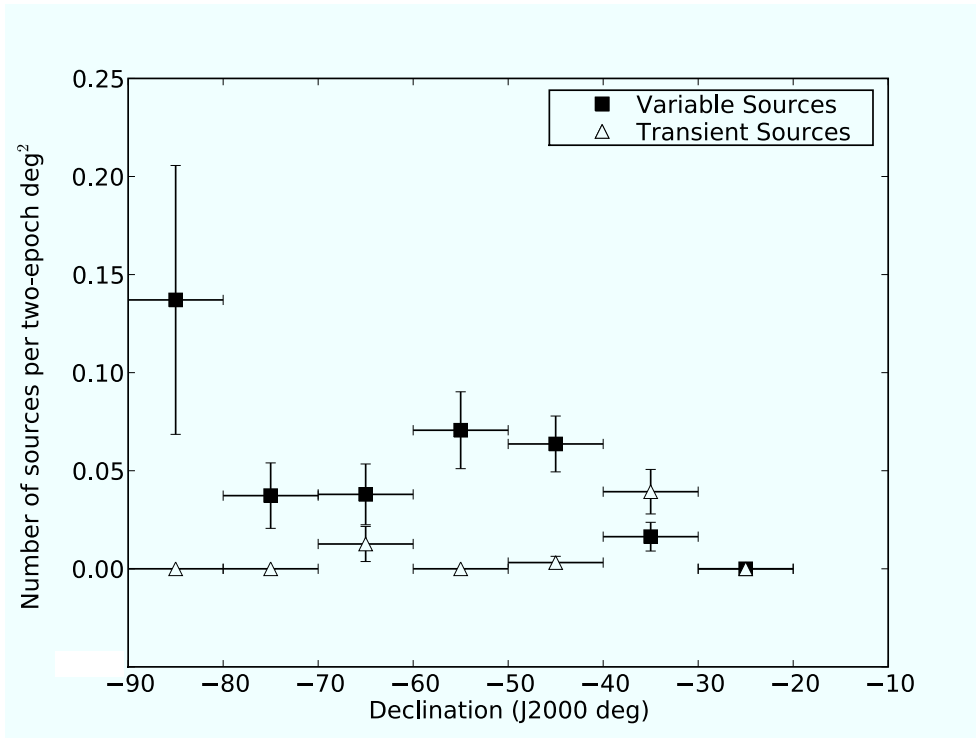


Figure 2.16: Fraction of variable and transient sources vs declination.

Table 2.5: 53 radio sources classified as variable with a confidence greater than 90%. The columns are: the source name, right ascension and declination (J2000), Galactic longitude and latitude, number of accepted measurements, probability of the χ^2_{lc} value, maximum, weighted mean and minimum flux density at 843 MHz, modulation index, ν value, NVSS flux density at 1.4 GHz (if covered by NVSS area), SuperCOSMOS classification: D = point source detected, N = no source detected, C = more than 1 point source in the error box, E = extended according to SuperCOSMOS catalogue, G = clearly resolved galaxy, the SuperCOSMOS optical magnitudes in blue (B), red (R) and infra-red (I), and the B–R colour.

Source name	α (h m s)	δ (d m s)	l (deg)	b (deg)	N_{meas}	$1 - P(\chi^2_{\text{lc}})$	S_{max} (mJy)	\bar{S} (mJy)	S_{min}	m	ν	NVSS mJy	Supercos	B (mag)	R (mag)	I (mag)	B–R (mag)
SUMSS J000716–793206	00 07 16	-79 32 00	305.4	-37.4	3	5.0×10^{-6}	29.4	22.9	18.8	0.2	0.22	-	N	—	—	—	—
SUMSS J001859–811133	00 18 58	-81 11 32	304.5	-35.8	2	2.4×10^{-5}	114.9	94.9	83.8	0.2	0.16	-	D	20.93	20.46	—	.48
SUMSS J003722–421508	00 37 22	-42 15 12	312.8	-74.6	2	6.5×10^{-6}	25.3	19.2	15.2	0.4	0.25	-	D	20.54	20.15	—	.38
SUMSS J004417–375258	00 44 17	-37 52 57	310.4	-79.1	2	2.4×10^{-14}	43.5	23.8	18.7	0.6	0.40	18.4	E	20.30	19.14	18.91	1.16
SUMSS J010426–400412	01 04 26	-40 04 12	292.0	-76.8	2	2.3×10^{-8}	62.6	44.6	37.7	0.4	0.25	-	D	20.00	19.15	18.92	.85
SUMSS J011019–455112	01 10 20	-45 51 16	292.8	-70.9	7	2.3×10^{-6}	20.9	16.0	10.7	0.2	0.32	-	G	—	—	—	—
SUMSS J011136–394340	01 11 36	-39 43 40	285.8	-76.7	2	1.0×10^{-9}	44.5	27.9	23.5	0.4	0.31	22.1	E	19.32	18.35	18.77	.97
SUMSS J011839–452657	01 18 40	-45 26 59	288.2	-70.9	9	9.4×10^{-6}	45.7	37.4	32.0	0.1	0.18	-	N	—	—	—	—
SUMSS J012615–561946	01 26 15	-56 19 47	293.2	-60.1	2	9.5×10^{-6}	14.1	9.3	6.4	0.5	0.38	-	C	—	20.50	18.35	—
SUMSS J013013–742024	01 30 15	-74 20 24	299.4	-42.5	4	1.0×10^{-5}	29.9	24.2	20.9	0.2	0.18	-	N	—	—	—	—
SUMSS J021036–405154	02 10 36	-40 51 55	258.1	-68.7	2	8.8×10^{-6}	40.4	32.6	26.1	0.3	0.22	-	D	21.28	19.96	—	1.32
SUMSS J022223–802356	02 22 23	-80 23 59	298.4	-35.9	2	2.4×10^{-7}	56.7	41.1	35.0	0.3	0.24	-	D	19.95	19.70	—	.24
SUMSS J031201–523432	03 12 00	-52 34 32	266.9	-53.6	2	1.4×10^{-5}	30.8	24.7	18.9	0.3	0.24	-	N	—	—	—	—
SUMSS J032102–405735	03 21 02	-40 57 32	247.0	-56.4	2	2.9×10^{-8}	39.8	27.8	19.4	0.5	0.35	-	D	20.21	19.64	—	.58
SUMSS J040608–540445	04 06 08	-54 04 54	263.9	-45.6	2	1.6×10^{-8}	19.6	12.5	9.0	0.5	0.37	-	C	21.71	19.33	18.49	2.38
SUMSS J043624–563332	04 36 24	-56 33 31	265.7	-40.8	3	9.8×10^{-10}	81.4	61.5	49.6	0.2	0.24	-	D	20.22	19.63	—	.58
SUMSS J044508–541155	04 45 08	-54 11 53	262.4	-40.0	2	1.3×10^{-8}	45.7	33.1	28.2	0.3	0.24	-	D	20.71	19.99	—	.72
SUMSS J044624–391548	04 46 24	-39 15 45	242.8	-40.3	2	2.5×10^{-5}	47.1	37.8	32.2	0.3	0.19	47.7	D	18.41	17.64	17.21	.77
SUMSS J044829–401714	04 48 28	-40 17 12	244.2	-39.9	2	3.4×10^{-9}	58.9	43.8	35.4	0.4	0.25	-	D	19.44	19.39	19.15	.05
SUMSS J045530–720442	04 55 30	-72 04 41	283.9	-34.5	3	$< 1 \times 10^{-14}$	38.5	17.0	11.9	0.6	0.53	-	C	21.50	21.68	—	-18
SUMSS J050230–651321	05 02 30	-65 13 22	275.6	-35.8	4	2.3×10^{-9}	40.4	31.4	25.5	0.2	0.23	-	C	19.47	18.42	—	1.05
SUMSS J061952–643917	06 19 52	-64 39 17	274.3	-27.7	4	2.3×10^{-8}	49.5	33.5	28.5	0.3	0.27	-	E	21.57	19.41	19.32	2.16
SUMSS J062724–405328	06 27 24	-40 53 28	248.9	-21.6	2	2.6×10^{-5}	24.0	16.9	13.7	0.4	0.27	-	N	—	—	—	—
SUMSS J064951–620221	06 49 51	-62 02 21	272.1	-23.9	2	2.4×10^{-6}	46.2	36.4	30.1	0.3	0.21	-	D	21.52	20.84	—	.68
MGPS2 J082437–463505	08 24 38	-46 35 09	263.6	-5.1	2	3.3×10^{-5}	44.3	33.6	29.6	0.3	0.20	-	N	—	—	—	—
MGPS2 J083322–444141	08 33 22	-44 41 39	263.0	-2.8	3	2.0×10^{-8}	102.4	83.3	64.6	0.2	0.23	-	C	—	21.03	—	—

Table 2.5 - *continued*

Source name	α (h m s)	δ (d m s)	l (deg)	b (deg)	N_{meas}	$1 - P(\chi^2_{\text{lc}})$	S_{max} (mJy)	S (mJy)	S_{min}	m	ν	NVSS mJy	SuperCOS	B (mag)	R (mag)	I (mag)	B-R (mag)
MGPS2 J083343-431958	08 33 43	-43 19 58	261.9	-1.9	2	3.8×10^{-6}	26.2	18.3	12.4	0.5	0.36	-	N	-	-	-	-
MGPS2 J083420-385401	08 34 20	-38 53 57	258.4	+0.8	3	3.8×10^{-6}	66.5	48.8	40.8	0.3	0.24	35.1	N	-	-	-	-
MGPS2 J083524-482342	08 35 24	-48 23 42	266.1	-4.7	3	4.6×10^{-7}	51.8	36.9	31.4	0.3	0.24	-	C	22.12	-	19.50	-
MGPS2 J083555-480638	08 35 54	-48 06 38	266.0	-4.5	3	4.6×10^{-6}	41.6	31.1	26.9	0.2	0.21	-	C	19.94	18.88	18.31	1.06
MGPS2 J090857-523612	09 08 57	-52 36 13	272.8	-3.3	3	5.7×10^{-7}	41.7	31.1	25.6	0.3	0.24	-	C	-	-	-	-
MGPS2 J091052-473743	09 10 52	-47 37 41	269.4	+0.3	2	$< 1 \times 10^{-14}$	34.3	15.6	10.8	0.7	0.52	-	N	-	-	-	-
MGPS2 J095141-550630	09 51 41	-55 06 31	279.1	-0.8	2	3.0×10^{-5}	42.4	33.3	29.4	0.3	0.18	-	N	-	-	-	-
SUMSS J100222-423135	10 02 22	-42 31 33	272.7	+10.2	2	2.1×10^{-5}	25.0	16.7	13.6	0.4	0.30	-	N	-	-	-	-
SUMSS J101048-805633	10 10 47	-80 56 32	296.7	-20.1	2	3.7×10^{-5}	19.5	13.4	11.0	0.4	0.28	-	N	-	-	-	-
MGPS2 J104530-505429	10 45 30	-50 54 30	283.5	+7.2	2	9.6×10^{-10}	94.7	68.2	57.3	0.3	0.25	-	C	-	-	-	-
SUMSS J113907-462506	11 39 06	-46 25 06	290.2	+14.7	2	$< 1 \times 10^{-14}$	35.7	18.3	14.3	0.6	0.43	-	C	21.34	19.97	19.48	1.37
SUMSS J120525-470714	12 05 26	-47 07 10	294.9	+15.0	3	1.6×10^{-5}	29.0	22.1	16.4	0.3	0.28	-	N	-	-	-	-
MGPS2 J161340-640022	16 13 40	-64 00 22	323.0	-9.3	2	9.3×10^{-6}	34.9	26.9	23.0	0.3	0.21	-	C	20.09	18.70	18.36	1.39
MGPS2 J170932-441750	17 09 32	-44 17 49	343.2	-2.5	3	2.1×10^{-6}	36.1	30.0	21.4	0.3	0.26	-	C	20.48	19.07	18.52	1.41
MGPS2 J171300-541538	17 12 60	-54 15 39	335.4	-8.8	2	1.7×10^{-6}	33.4	26.9	20.0	0.4	0.25	-	C	20.76	20.16	-	.60
SUMSS J173100-533721	17 31 00	-53 37 21	337.5	-10.7	2	3.2×10^{-5}	55.8	46.0	39.3	0.2	0.17	-	C	20.85	19.73	19.13	1.12
SUMSS J175605-464450	17 56 06	-46 44 49	345.6	-10.7	2	1.1×10^{-8}	34.1	22.3	16.8	0.5	0.34	-	C	16.71	15.56	15.00	1.14
SUMSS J183056-464744	18 30 56	-46 47 43	348.2	-16.1	3	5.2×10^{-6}	16.4	12.1	8.3	0.3	0.33	-	C	-	-	-	-
SUMSS J190344-505108	19 03 44	-50 51 08	346.0	-22.6	3	1.5×10^{-10}	44.9	33.3	23.1	0.3	0.32	-	C	19.73	19.25	19.48	.48
SUMSS J200936-554236	20 09 35	-55 42 31	342.2	-33.0	4	4.2×10^{-6}	21.2	16.8	12.2	0.3	0.27	-	G	-	-	-	-
SUMSS J201524-395949	20 15 24	-39 59 48	0.9	-32.5	2	2.0×10^{-5}	31.5	23.3	16.4	0.4	0.31	9.1	G	-	-	-	-
SUMSS J202536-813918	20 25 36	-81 39 19	311.8	-30.2	2	1.3×10^{-7}	62.1	47.8	39.7	0.3	0.22	-	D	20.26	19.88	-	.39
SUMSS J211009-772600	21 10 09	-77 25 58	315.4	-33.7	2	1.3×10^{-10}	39.0	25.9	20.6	0.4	0.31	-	N	-	-	-	-
SUMSS J214235-724745	21 42 35	-72 47 40	318.9	-37.9	3	1.8×10^{-6}	14.3	9.8	7.8	0.3	0.30	-	N	-	-	-	-
SUMSS J223017-614701	22 30 16	-61 47 03	327.1	-48.2	50	3.0×10^{-5}	58.4	47.0	37.9	0.1	0.21	-	D	21.14	19.16	18.02	1.98
SUMSS J223225-615308	22 32 24	-61 53 13	326.8	-48.3	33	$< 1 \times 10^{-14}$	257.7	221.1	167.5	0.1	0.21	-	N	-	-	-	-
SUMSS J225444-500528	22 54 44	-50 05 30	339.1	-58.0	2	3.5×10^{-7}	24.5	17.0	12.7	0.4	0.31	-	E	20.92	20.46	-	.46

Table 2.6: 15 radio sources classified as transient. The columns are: the source name, right ascension and declination (J2000), Galactic longitude and latitude, number of detections, number of accepted observations, probability of the χ^2_{lc} value, maximum flux density at 843 MHz in mJy, the maximum and minimum detection significance, NVSS flux density at 1.4 GHz in mJy (if covered by NVSS area, ND if not detected), SuperCOSMOS classification: D = point source detected, N = no source detected, C = more than 1 point source in the error box, E = extended according to SuperCOSMOS catalogue, G = clearly resolved galaxy, the SuperCOSMOS optical magnitudes in blue (B), red (R) and infra-red (I), and the B–R colour. More information on these sources, including identifications, can be found in Table 2.7.

Source name	α (h m s)	δ (d m s)	l (deg)	b (deg)	N_{det}	N_{obs}	$1 - P(\chi^2_{lc})$	S_{max} (mJy)	d_{max}	d_{min}	NVSS (mJy)	Supercos	B (mag)	R (mag)	I (mag)	B–R (mag)
SN1987A	05 35 28	-69 16 07	279.7	-31.9	2	13	6.5×10^{-13}	151.8	13.5	-0.4	-	C	18.16	13.25	12.65	4.91
SUMSS J055712-381106	05 57 12	-38 11 05	244.3	-26.5	1	2	2.5×10^{-3}	16.7	7.1	2.9	5.9	E	17.89	16.80	16.18	1.10
J060938-333508	06 09 38	-33 35 08	240.3	-22.7	1	3	1.4×10^{-6}	21.3	6.6	2.3	ND	G	—	—	—	—
J061051-342404	06 10 50	-34 24 04	241.2	-22.8	2	4	3.4×10^{-3}	14.3	8.3	2.2	5.5	N	—	—	—	—
SUMSS J062636-425807	06 26 36	-42 58 07	251.0	-22.4	1	2	7.0×10^{-4}	19.3	15.0	2.9	-	D	19.39	18.18	16.81	1.20
J062716-371736	06 27 16	-37 17 36	245.3	-20.5	2	3	3.1×10^{-3}	14.9	6.6	2.9	3.2	N	—	—	—	—
J064149-371706	06 41 49	-37 17 06	246.3	-17.8	1	2	4.4×10^{-4}	14.4	6.1	2.4	4.1	N	—	—	—	—
SUMSS J102641-333615	10 26 40	-33 36 14	271.3	+20.2	1	2	9.8×10^{-3}	17.2	7.3	2.7	10.0	N	—	—	—	—
SUMSS J112610-330216	11 26 09	-33 02 13	283.0	+26.5	1	2	4.6×10^{-3}	16.6	7.8	2.8	7.5	N	—	—	—	—
Nova Muscae 1991	11 26 27	-68 40 31	295.3	-7.1	1	2	$< 1 \times 10^{-14}$	140.4	67.8	0.1	-	C	21.57	—	18.72	—
J121032-381439	12 10 32	-38 14 39	294.2	+23.9	1	2	7.4×10^{-3}	15.0	9.5	2.8	9.9	N	—	—	—	—
J135304-363726	13 53 04	-36 37 26	316.5	+24.6	1	3	1.0×10^{-3}	16.3	6.6	2.9	5.3	D	—	20.71	—	—
J153613-332915	15 36 13	-33 29 15	338.2	+17.9	1	2	1.5×10^{-3}	18.4	6.2	2.3	8.2	N	—	—	—	—
GRO 1655-40	16 54 00	-39 50 44	345.0	+2.5	2	9	$< 1 \times 10^{-14}$	6504.2	17.4	-0.8	ND	E	18.58	15.26	13.96	3.32
SUMSS J224152-300823	22 41 52	-30 08 23	19.1	-61.5	1	3	9.6×10^{-3}	15.6	9.4	0.6	5.5	N	—	—	—	—

2.4.3 Transient sources

Table 2.6 lists 15 sources we classify as transient by the criteria described in §2.3.6. A summary of the radio light curves, counterparts at other wavelengths and classifications is shown in Table 2.7. Eight sources that satisfied these criteria are not shown in the table, because visual inspection showed incorrect classification due to image artefacts, such as nulls causing non-detections or false detections from grating rings.

2.4.4 Optical Counterparts

The deepest optical survey of the area of sky covered by our survey is the SuperCOSMOS Sky Survey (Hambly et al., 2001), which was produced by digitising photographic plates from the UK Schmidt Telescope. The SuperCOSMOS Sky Survey has image and catalogue data in three colours with limits down to $B \sim 22$, $R \sim 20$ and $I \sim 19$. To determine an optical match for each radio source, we visually inspected overlays of the MOST sources with the corresponding SuperCOSMOS B images, and searched the SuperCOSMOS catalogues.

From visual inspection of the overlay plots, we were immediately able to identify four of the 68 transient and variable radio sources as having radio positions directly on, or within 10 arcsec of, the centre of a resolved optical galaxy. These sources have an optical classification of ‘G’ in Table 2.5 and Table 2.6. To determine the probability of random false association, we visually inspected images of 68 random positions on the sky and found that no positions fell near resolved optical galaxies of similar size. We conclude that the radio sources that are positionally coincident with resolved optical galaxies are very likely to be physically associated with these galaxies.

Many of our radio sources are found in optically crowded regions in the Magellanic Clouds. The positional uncertainty for MOST sources is < 5 arcsec (Mauch et al., 2003); therefore, we consider any region with 2 or more optical matches within 5 arcsec to be crowded. Such sources have an optical classification of ‘C’ in Table 2.5 and Table 2.6, and optical magnitudes shown in the tables are those of the nearest candidate to the radio position.

We consider any radio source with a single optical counterpart within 5 arcsec of the radio position to be a detection. These sources have an optical classification of ‘D’ if the match is classified as point-like in the SuperCOSMOS catalogue (mean class is 1), or ‘E’ if the match is extended in the SuperCOSMOS catalogue (mean class is 2). To measure the probability of a false match we compare the number of optical sources contained within a circle centred on the radio source position with the number of such sources on a random position in the sky (Figure 2.17) over a range of radii. We find that, within 5 arcsec, there are 9 optical matches for the random positions, and 27 optical matches for the radio source positions, corresponding to a false match




Figure 2.17: Histogram of the number of SuperCOSMOS B sources as a function of distance from the radio source position for uncrowded MOST transient and variable sources, and randomly chosen positions.

probability at the 5 arcsec match radius of 33 per cent.

The remaining radio sources have no point-like optical counterparts within 5 arcsec to magnitude limits $B \sim 22$, $R \sim 20$ and $I \sim 19$. These sources have an optical classification of ‘N’ in Table 2.5 and Table 2.6.

2.4.5 Gamma Ray Burst Counterparts

Some of our transient and variable sources could be related to GRB afterglows. In the standard picture, a radio afterglow at 843 MHz would be detected several weeks after a GRB event and would fade over several months.

We searched the master catalogue of the *Swift* GRB satellite⁸ (operational since 2005 Jan), for GRBs within 10 arcmin of all transient and variable radio sources. We found no GRBs within the position uncertainties of any transient or variable radio source.

For our transient sources only, we also searched the Stern et al. (2001) catalogue for GRBs detected by the BATSE instrument aboard the *Compton*

⁸http://www.swift.ac.uk/swift_live/arnie5.php

Gamma-Ray Observatory (operational between 1991 and 2000). For each transient source, between 5 and 10 counterparts were identified within the gamma ray positional uncertainties (which are typically tens of degrees). No GRBs were detected in the weeks before the first detection of those transient radio sources that exhibit a typical GRB afterglow light curve.

The vast majority of GRBs are detected at $z > 0.1$ (Jakobsson et al., 2006), with the remainder at much larger distances. All of our transient sources with resolved optical counterparts are at $z < 0.1$. Hence, if our transient sources had a gamma-ray trigger, we would expect them to be detected if a GRB satellite had been operational at the time.

We note that a number of transient sources with radio light curves consistent with GRB afterglows (i.e. J060938–333508, J062636–425807, J064149–371706 and SUMSS J102641–333615), would not have had a GRB detection in either *Swift* BAT or BATSE catalogues as neither instrument was fully operational when the first MOST detection was made.

2.5 DISCUSSION

2.5.1 Overall statistics

Our sample comprises light curves for 29730 sources, with a two-epoch equivalent sky coverage (see Equation 2.1) of 2775.7 deg^2 . The shortest interval between accepted measurements of a single source is 24 h and the longest interval between accepted measurements is 7301 days ($\simeq 20$ years).

For a source to be in our sample, it must appear in at least two images for which the exact two-epoch coverage applies, rather than the equivalent coverage. The two-epoch (exact) coverage is approximately 2000 deg^2 (see Figure 2.3) resulting in a source density of 15 deg^{-2} . The SUMSS source density is 31.6 deg^{-2} meaning the overall completeness of our sample is 47 per cent.

Log N-Log S plots for the SUMSS and MGPS-2 catalogues, our whole sample, the transient sources and the variable sources are shown in Fig. 2.18. At low flux densities the catalogue is incomplete due to variations in the noise floor, and at high flux densities it is due to artefacts around bright sources. In general the low completeness is due to the strict requirements on goodness of fit that we imposed in order to remove false positives due to artefacts.

Fig. 2.19 and 2.20 show histograms of modulation index and \mathcal{V} , respectively, for all sources in our sample.

2.5.2 Variability time-scales

If there is a sufficient number of measurements in the light curve of a source, a structure function can be formed to make an estimate of the variability time-scale (e.g. Gaensler & Hunstead 2000 ; Lovell et al. 2008). As the majority of

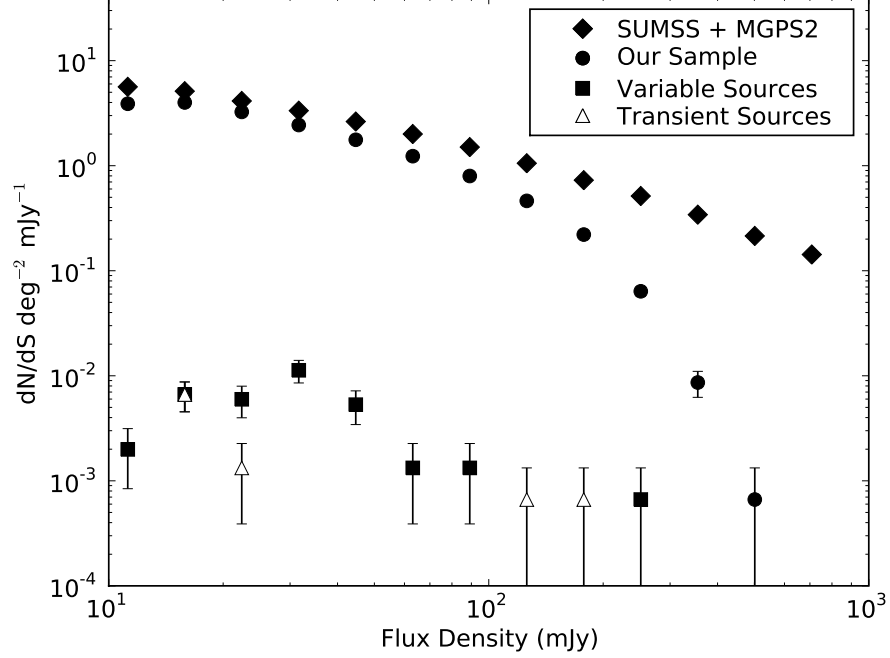


Figure 2.18: Log N-Log S plots for the combined SUMSS and MGPS-2 catalogue, all sources used for our variability analysis, and the sources classified as variable and transient. The error bars are Poisson errors. The sources available for variability analysis are incomplete above 100 mJy, as strong sources have poor fits in MOST images due to artefacts.

our sources have only a small number of measurements (see Fig. 2.11), structure functions cannot be used. To determine if there is a preferred variability time-scale for our sources, we compare the distribution of intervals between measurements of our variable sources, with that of the whole sample as follows: for each source, we take each non-redundant pair of flux density measurements and bin by the time interval between these measurements: i.e. if there are N_{meas} measurements on the light curve, there are $N_{\text{meas}}(N_{\text{meas}} - 1)/2$ intervals. We then form histograms comprised of all sources, and of the subset of variable sources in Table 2.5 only. If the variable sources have a preferred time-scale, this should be evident in the variable sources constituting a larger fraction of their measurement intervals at a particular time-scale. Fig. 2.5 depicts the histograms of intervals described above.

Fig. 2.21 shows the ratio of the number intervals belonging to the variable and transient sources to the number of intervals belonging to the whole sample. The ratio peaks at around 1000 days and then falls rapidly by 2000 days. This suggests that any blind survey for variable radio sources will detect the




Figure 2.19: Histograms of modulation indices ($m = \sigma/\bar{S}$, see Equation 2.9) for all sources and for only variable sources.

majority of variable sources in the first 1000 to 2000 days.

2.5.3 Source counts for variable sources

The number of variable sources in our sample can be determined in several ways. The source list in Table 2.5 contains 55 sources and we are confident that at least 90 per cent of them are truly variable, which implies a source density of 0.020 sources per two-epoch square degree. However, this is not the whole story. Fig. 2.13 shows that there is a clear excess of sources with $P(\chi^2_{lc}) < 0.1$. While we cannot determine whether a particular source's variation is attributable to measurement errors or true variation, we can compare the number of apparently variable sources to the expected number if only Gaussian measurement errors were present. If we conservatively take the 'expected' line as the number of non-varying sources, we arrive at an excess of 200 sources, corresponding to a lower limit on the two-epoch areal density of variable sources at 843 MHz of 0.072 deg^{-2} . This would appear to be an underestimate, as the actual number of sources with $P(\chi^2_{lc}) < 0.2$ is substantially less than the expected number, due to over-estimated errors. If we extrapolate the actual number of sources from $P(\chi^2_{lc}) = 0.2$, we arrive at an excess of

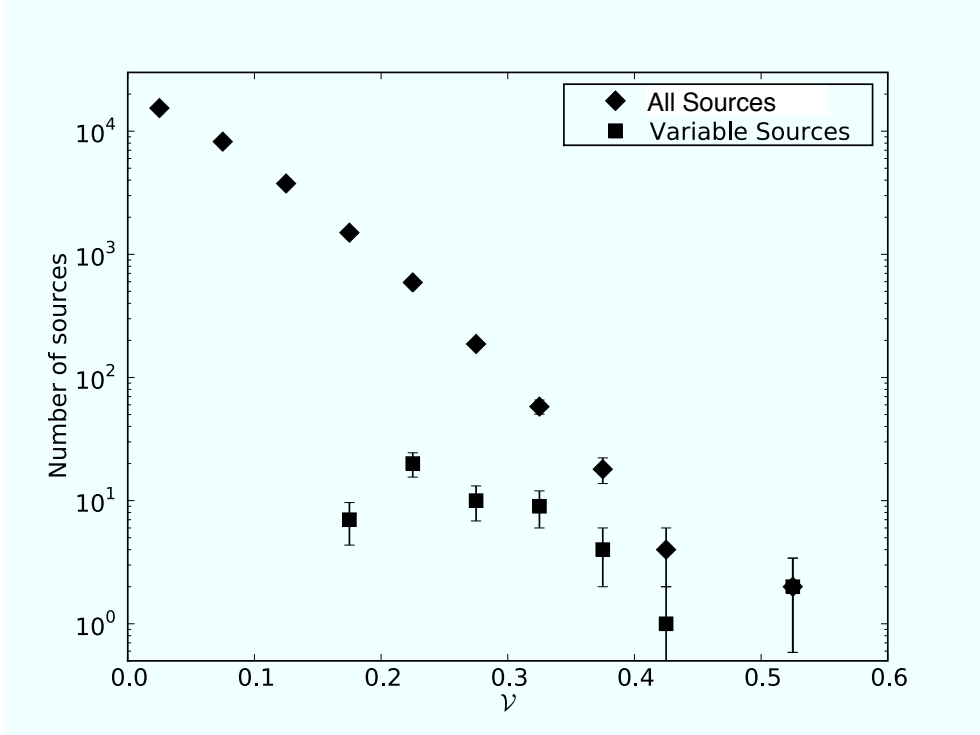


Figure 2.20: Histograms of fractional variability (see Equation 2.8) for all sources and for only variable sources.

approximately 350 sources with $P(\chi^2_{lc}) < 0.2$ at 843 MHz with flux densities between 15 and 100 mJy, which corresponds to a two-epoch areal density of variable sources at 843 MHz of 0.126 deg^{-2} . After accounting for our completeness fraction of 47 per cent, the source density rises to 0.268 deg^{-2} in 2 epochs.

A definitive statement about the number of variable sources requires a caveat on the time-scale of the variability. Our data are not regularly sampled and so we can only make broad statements about variability on time-scales greater than 1 day and less than about 10 years.

2.5.4 Interpretation as interstellar scintillation

At 843 MHz, our observations are firmly in the strong or refractive regime of interstellar scintillation. In these conditions, Rickett (1986) derives the maximum rate of change of flux of sources with a brightness of $T_{12} \times 10^{12} \text{ K}$ at latitudes $|b| \geq 10^\circ$ as

$$(dS/dt)_{max} < S_{rms}/\tau \sim 30T_{12}(\sin b)^{0.5} \text{ mJy day}^{-1}, \quad (2.10)$$

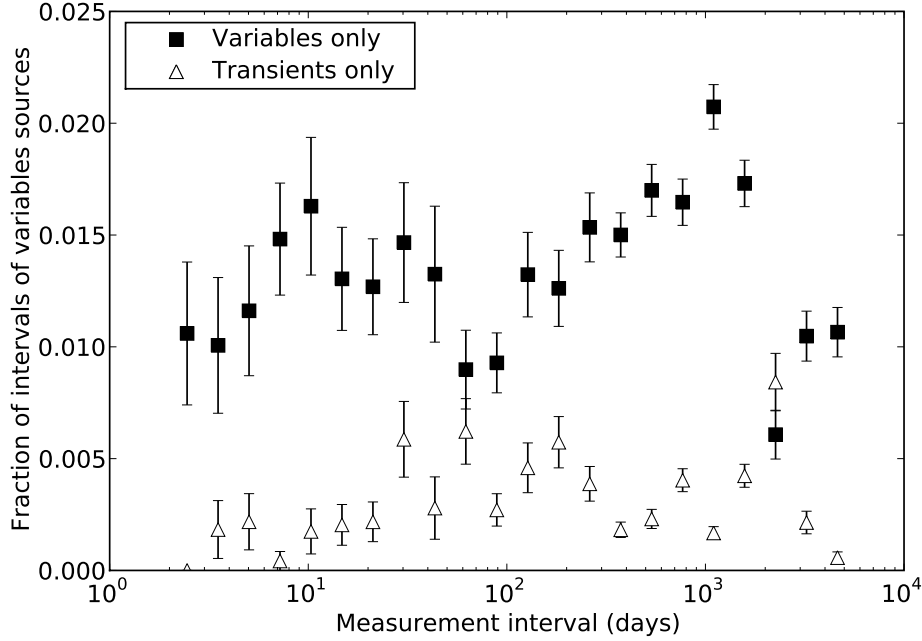


Figure 2.21: The ratio of the number of non-redundant pairs of intervals in the light curves for variable and transient sources, to the total number of intervals for all sources, as a function of the interval in days. The y errors are the Poisson errors.

giving a range of $(dS/dt)_{max}$ for a source with a brightness temperature at the inverse Compton limit of between 12.5 and 30 mJy day $^{-1}$. As our data are sparsely sampled in the time domain, this limit is easily satisfied for all our sources.

The number of sources that scintillate is expected to increase as the line of sight to the source becomes closer to the Galactic plane, and thus becomes more likely to pass through a scintillating screen of material. In Figure 2.22 we plot the fraction of sources that we defined as variable and transient, binned by Galactic latitude. The plot suggests, at best, a mild dependence on Galactic latitude. Nonetheless, we expect that refractive interstellar scintillation of compact radio sources is a likely explanation for a large fraction of our variable sources, as also seen by Gaensler & Hunstead (2000).

Some 30 per cent of the variable sources have no optical counterpart. In the absence of more data, we cannot positively identify these sources.

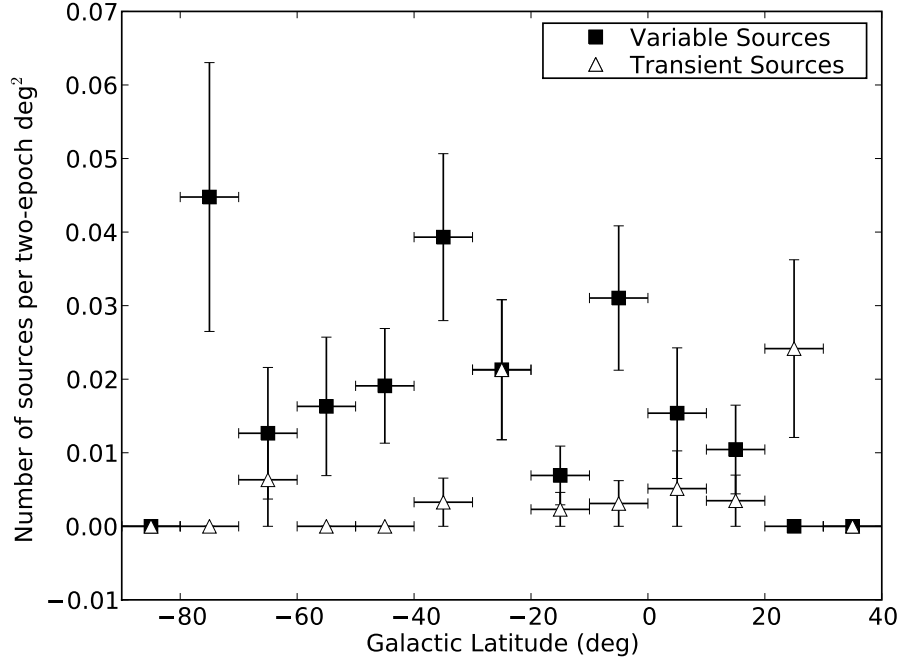


Figure 2.22: Fraction of variable and transient sources vs Galactic latitude.

2.5.5 Source counts for transient sources

Transient emission on characteristic time-scales of less than the full synthesis time is unlikely to be detected by our technique, because a significant change in flux density would result in a distorted synthesised beam for that source, which would result in large errors in a Gaussian fit and ultimately rejection from the analysis. Therefore, we will omit from our discussion source classes such as giant pulses from neutron stars (Hankins et al., 2003), flares from extrasolar planets (Bastian et al., 2000) and fast extragalactic transients (Lorimer et al., 2007), all of which have time-scales much shorter than 12 h. Also, a source that is constant over 12 h but bright (> 100 mJy), or in a region containing bright sources (such as the Galactic plane), is unlikely to be detected due to image artefacts or large formal errors in source fitting.

Table 2.6 shows 15 sources that satisfied our definition of a transient with a peak flux density above 14 mJy, corresponding to an areal density of transient sources at 843 MHz of $6 \times 10^{-3} \text{ deg}^{-2}$ in 2 epochs. After accounting for completeness of 47 per cent, the areal density rises to $1.3 \times 10^{-2} \text{ deg}^{-2}$ in 2 epochs.

Bower et al. (2007) compute a two-epoch areal density of transient sources of 1.5 deg^{-2} above $370 \mu\text{Jy}$ at 5 GHz and a flux density distribution following

$S^{-\gamma}$ with $\gamma = 1.5$, which corresponds to an areal density of $2.0 \times 10^{-2} \text{ deg}^{-2}$, assuming a spectral index of -0.7 and flux density limit of 15 mJy, in broad agreement with our findings. We note that there are aspects of our transients that are not the same as those described by Bower et al. (2007). In particular, our transients are detected at a frequency an order of magnitude lower, have flux densities an order of magnitude higher, and have upper limits on characteristic time-scales of one or two orders of magnitude longer than those discussed by Bower et al. (2007).

2.5.5.1 Known transients

Our search detected three previously identified radio transients:

1. GRO 1655–40 (J165400–395044) is a low mass X-ray binary (or microquasar), detected in our survey because of the images taken while monitoring a dramatic radio flare in 1994 (Campbell-Wilson & Hunstead, 1994a).
2. Nova Muscae 1991 (J112627–684031) is a low mass X-ray binary, detected in our survey from an image taken during the outburst in 1991 (Kesteven & Turtle, 1991).
3. SN1987A (J053528–691607) was a supernova in the Large Magellanic Cloud. It is detected in our survey due to a regular monitoring campaign (Ball et al., 2001) during which its flux density has been steadily increasing.

2.6 DISCUSSION OF SELECTED SOURCES

In this section we discuss sources with noteworthy light curves or optical counterparts.

2.6.1 Variable sources with resolved optical counterparts

2.6.1.1 SUMSS J011019–455112

The MOST light curve of SUMSS J011019–455112 (Fig. 2.23) shows two decreasing flux measurements over 10 years, followed two years later by four significantly higher flux measurements closely spaced in time. We have also obtained flux density and position measurements from archival Australia Telescope Compact Array (ATCA) observations at 1.4 GHz used for the Phoenix Deep Field survey, in which the average flux density for this source over all epochs is 10.6 mJy (Hopkins et al., 2003). An optical overlay of the higher resolution ATCA data shows that the radio source is centred 9.6 arcsec from the centre of the extended 2MASS source 2MASX J01101993–455118, an S0

galaxy in the galaxy cluster Abell 2877 at $z = 0.023$ (Caldwell & Rose, 1997). Using the highest MOST flux density, the implied isotropic radio luminosity is $L_\nu = 2 \times 10^{29} \text{ erg s}^{-1} \text{ Hz}^{-1}$. The spectral index between the mean 1.4 GHz flux density and the minimum and maximum flux densities from MOST is in the range 0 to -1.3 but we note that no measurements are contemporaneous. The radio source detected by ATCA at 1.4 GHz was a point source at the 12 arcsec resolution of the observations, and had a position consistent with $\alpha = 1\text{h}1\text{m}19.442\text{s}$ $\delta = -45^\circ 51\text{m}13.958\text{s}$ (J2000) in all ATCA epochs.

As the MOST contours are offset from the optical centre of 2MASX J01101993–455118, the radio sources is unlikely to be associated with AGN activity, so either a radio supernova (RSN) or GRB afterglow is a possible candidate.

RSNe have typical radio spectral luminosity of $L_\nu \simeq 10^{25} \text{ erg s}^{-1} \text{ Hz}^{-1}$, but can reach as high as $L_\nu \simeq 2 \times 10^{27} \text{ erg s}^{-1} \text{ Hz}^{-1}$ (Eck et al., 2002) and have fading times of 20 years or more. SUMSS J011019–455112 has a much higher radio luminosity than even the most luminous known RSN (1979C), so if it is a radio supernova, it is unusually bright.

The radio luminosity of SUMSS J011019–455112 is closer to those of GRB afterglows, which typically have a range of $L_\nu \simeq 6 \times 10^{28} - 4 \times 10^{31} \text{ erg s}^{-1} \text{ Hz}^{-1}$ (Weiler et al., 2002). On the other hand, the 12 year time-scale is typical for a Type II RSN but too long for a GRB afterglow, which typically fades after several weeks. The approximately constant flux density for the MOST and ATCA observations could be explained by pre-existing disc synchrotron emission from the host galaxy, with the flare in the MOST observations being due a flare or stellar explosion overlaid. In such a case, one would expect the disc synchrotron emission to appear resolved at the ATCA resolution and roughly follow the galaxy’s optical emission, but we find no evidence for resolved emission in the ATCA observations.

We conclude that SUMSS J011019–455112 is associated with an unusual stellar event or explosion in the galaxy 2MASX J01101993–455118.

2.6.1.2 SUMSS J201524–395949

The light curve for SUMSS J201524–395949 (Fig. 2.24) shows an NVSS detection and two MOST detections over 10 years. The MOST flux density doubled over 4 years. The MOST contours are centred on a barred spiral galaxy ESO 340-G06 at redshift $z = 0.02$ (Jones et al., 2004) implying, if the radio source is associated with the spiral galaxy, an isotropic radio luminosity $L_\nu = 3 \times 10^{29} \text{ erg s}^{-1} \text{ Hz}^{-1}$.

The 6dFGS spectrum of ESO 340-G06 (Jones et al., 2004) shows strong Balmer emission lines, indicative of a star-forming galaxy. The *IRAS* 60 μm flux is 1 Jy, but the galaxy is not detected at 12 and 25 μm (Moshir, 1990), suggesting the presence of cold dust. It is rare, but not impossible, for such a star-forming galaxy to contain an AGN.




Figure 2.23: Variable source SUMSS J011019–455112. Top panel: Radio contours from the Phoenix deep field (Hopkins et al., 2003) at 1.4 GHz at 2 to 8 mJy beam⁻¹ in steps of 2 mJy beam⁻¹ overlaid on a SuperCOSMOS B image. Bottom panel: Radio light curves at 1.4 GHz (ATCA archive) and 843 MHz (MOST).

The radio flux in this case is likely due to a stellar explosion, as widespread star formation or disc synchrotron emission cannot explain the radio variability. Once again, the time-scale of 10 years argues strongly against a GRB afterglow. RSN light curves are known to behave differently at different frequencies, with the initial flux density increase being later and slower at lower frequencies. The increasing flux density over 4 years at 843 MHz is consistent with a Type II RSN but the 6 year interval between MOST and NVSS detections is arguably too long to be due entirely to an RSN, in which case it may be that the RSN occurred after 1996, with the NVSS and first MOST measurements being attributed to pre-existing disc synchrotron emission or widespread star formation.

2.6.1.3 SUMSS J200936–554236

The light curve for SUMSS J200936–554236 (Fig. 2.25) shows an increase in flux density over 3 years, and then a steady flux density over 2 years. The MOST contours are centred on the barred spiral galaxy IC4957 at redshift $z = 0.032$ (Strauss et al., 1992). This galaxy is detected across the *IRAS* bands at 0.2 Jy (25 μm), 2.4 Jy (60 μm) and 4.5 Jy (100 μm), implying warm dust. The implied radio luminosity, if the radio source is associated with the spiral galaxy, is $L_\nu = 5 \times 10^{29} \text{ erg s}^{-1} \text{ Hz}^{-1}$.

As the MOST contours are centred on the optical source, AGN activity or ISS of an AGN is a possible explanation except for the fact that AGNs in barred spiral galaxies are rare (Wilson & Colbert, 1995). The warm dust implies star formation, so a stellar event is conceivable. The 3 year time-scale argues strongly against a GRB afterglow interpretation, and the shape and time-scale are consistent with a Type II RSN. The luminosity, once again, is higher than that of any known RSN. We conclude that SUMSS J200936–554236 is associated with an unusual stellar event or explosion in IC4957.

2.6.2 A compact steep spectrum source: SUMSS J223225–615308

The light curve for SUMSS J223225–615308 (Fig. 2.26) shows two pronounced dips in flux density over two years. This source is the most dramatic of some five variable sources in this field. There are other sources at the same flux density level in this field that do not exhibit variability, ruling out a calibration effect. An overlay of SUMSS J223225–615308 and the Parkes-MIT-NRAO (PMN) survey data (Wright et al., 1994) confirms an association with PMNJ2232–6153, whose reported flux density is $40 \pm 7 \text{ mJy}$ at 4.85 GHz in 1990 June. This flux density together with the range of flux densities shown in the 843 MHz light curve, implies a spectral index between -0.8 and -1.0 , although we note that the PMN and MOST observations are not contemporaneous. If the variation is due to scintillation, this implies a compact source

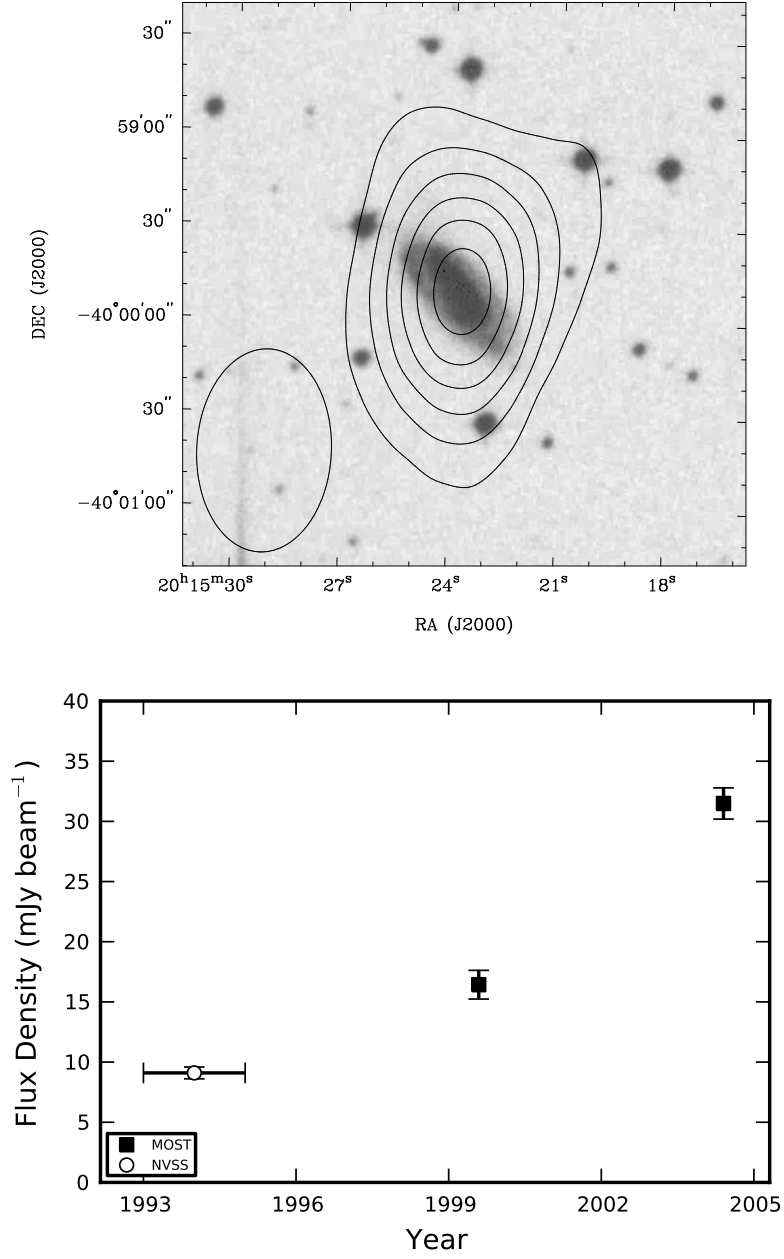


Figure 2.24: Variable radio source J201524–395949. Top panel: MOST contours from 5 to 30 mJy beam⁻¹ in steps of 5 mJy beam⁻¹ for the brightest epoch (2004 May 1) overlaid on a SuperCOSMOS B image. Bottom panel: light curve showing an NVSS detection and a doubling in 843 MHz flux density over 5 years over a 6 year interval.




Figure 2.25: Variable radio source SUMSS J200936–554236. Top panel: MOST contours at $-1.5 \text{ mJy beam}^{-1}$ and 6 to 18 mJy beam^{-1} in steps of 3 mJy beam^{-1} for the brightest epoch (1999 March 15) overlaid on a SuperCOSMOS B image. Bottom panel: radio light curve indicating a rising and then constant flux density.




Figure 2.26: Light curve of the candidate compact steep spectrum source: SUMSS J223255–615308

or component, and together with the spectral index implies a classification as a compact steep spectrum source.

2.6.3 Transient sources with resolved optical counterparts

2.6.3.1 J060938–333508

The light curve for J060938–333508 (Fig. 2.27) shows MOST and NVSS non-detections followed by a single MOST detection. Unlike the vast majority of sources in the MOST archive, the MOST contours appear rotated with respect to the MOST beam, possibly indicating a change in flux density over the 12 hr synthesis time, or a marginally resolved jet-like structure. The MOST contours are centred 9 arcsec from the centre of Fairall 1138, a galaxy with spectral type SBab D (Dressler & Shectman, 1988) and with redshift $z = 0.037$ (Dale et al., 1998).

Assuming the radio source and galaxy are associated, the inferred isotropic radio luminosity from the brightest epoch (2004 December 9) is $L_\nu \simeq 6 \times 10^{29} \text{ erg s}^{-1} \text{ Hz}^{-1}$ at 843 MHz.

The offset from the centre of the optical galaxy, and the fact that spiral

galaxies rarely contain a radio-loud AGN argues against an AGN source for the radio variability. The optical morphology of Fairall 1138 appears quite disturbed and the galaxy has a $60\ \mu\text{m}$ flux of 345 mJy (Moshir et al., 1992), which implies the presence of cold dust. The disturbed morphology, spectral type and presence of dust in Fairall 1138 support the interpretation that the galaxy is undergoing star formation, so a stellar explosion, either GRB or RSN are possible explanations.

The spectral luminosity is very high for a RSN and we are unable to discriminate between Type Ib/c or Type II RSNe by the light curve time-scales, as the time interval between the detection and non-detection epochs is too large.

The spectral luminosity of J060938–333508 is within the range of GRB afterglows. It is also within the error circle of GRB 940526B, which occurred after the MOST non-detection in 1993, and 10 years before the MOST detection. If J060938–333508 is the radio afterglow of GRB 940526B, then the radio detection 10 years after the gamma ray event is unlike known GRB afterglows, which typically peak at 843 MHz a few weeks after the explosion and fade over about 3 years. We consider an association of GRB 940526B and J060938–333508 unlikely.

We consider an unusual stellar event in Fairall 1138 as the likely interpretation of this source.

2.6.3.2 SUMSS J055712–381106

The light curve for SUMSS J055712–381106 (Fig. 2.28) shows an NVSS detection followed by a MOST detection approximately 10 years later and then a non-detection 6 days after that, consistent with either a flaring source or a highly variable source occasionally appearing above our sensitivity limit. The MOST contours appear slightly elongated and rotated with respect to the MOST beam, possibly indicating a change in flux density over the 12 hr synthesis time, or a marginally resolved jet-like structure. A flaring or scintillating AGN or a flaring radio star are possible counterparts with these properties.

The SuperCOSMOS B image shows what appears to be a blend of three objects: a star-like object to the south, a faint star-like object immediately to its north, and an extended object to the north-east. The MOST and NVSS radio sources cannot be conclusively associated with any of the three sources in the optical blend, which may in fact be a single interacting system. A more accurate radio position is required for a definitive association.

With no clear optical association we cannot positively classify this source, but the time-scales between the NVSS detection, MOST detection and MOST non-detection rule out RSNe and GRB afterglows.




Figure 2.27: Radio transient J060938–333508. Top panel: MOST contours from 6 to 18 mJy beam⁻¹ in steps of 3 mJy beam⁻¹ for the brightest epoch (2004 December 9) overlaid on a SuperCOSMOS B image. The MOST synthesised beam is drawn at the bottom left. Bottom panel: MOST and NVSS light curves indicating a single detection in 2004 at 21.3 mJy and 3 non-detections.




Figure 2.28: Radio transient SUMSS J055712-381106. Top panel: MOST contours from 3 to 15 mJy beam⁻¹ in steps of 3 mJy beam⁻¹ for the brightest epoch (2004 November 7) overlaid on a SuperCOSMOS B image. Bottom panel: MOST light curve indicating an NVSS and MOST detection, with a non-detection 6 days after the MOST detection.

2.6.4 Transient sources with point-like optical counterparts

2.6.4.1 SUMSS J062636–425807

The light curve for SUMSS J062636–425807 (see Appendix B) has a detection and a non-detection separated by almost a year. The optical counterpart is point-like and the colours are red, with $B-K = 5$. The source is most likely an AGN scintillating above our noise threshold.

2.6.4.2 J135304–363726

The light curve for J135304–363726 (see Appendix B) is consistent with either a flaring source, or a variable source occasionally appearing above our sensitivity limit. It is marginally detected in the SuperCOSMOS R plates and not detected in B or I, or in any 2MASS images. This source is likely an optically faint AGN scintillating above our noise threshold.

2.6.4.3 Unidentified variables

Some 25 per cent of the variable sources have a clear point-like or marginally resolved optical counterpart. Most of these counterparts have red colours, with $0 < B - R < 2$. An optical point source combined with a radio detection would typically imply an AGN, but the colour distribution is unusual for AGN, for which the vast majority have $B - R < 1$ (Croom et al., 2004). The reddest sources are therefore possibly flaring M-dwarfs, but optical spectra are needed to confirm this.

2.6.4.4 Unidentified transients

The spectral luminosities of our transients are $L_\nu \simeq 10^{19} (D_L/1 \text{ kpc})^2 \text{ erg s}^{-1} \text{ Hz}^{-1}$, where D_L is the luminosity distance. This spectral luminosity is an order of magnitude more luminous than those found by Bower et al. (2007) at 5 and 8 GHz, but roughly consistent assuming a spectral index $\alpha = -0.7$. The sampling interval of our light curves is typically much longer than those discussed by Bower et al. (2007) (typically years as opposed to weeks) and we have not yet obtained dedicated optical follow-up of our sources, so our limits on optical counterparts are not very strong. The scarcity of data means we are unable to robustly classify a large fraction of our transient sources.

Of the unidentified transients, none have X-ray counterparts in the *ROSAT* X-ray All-Sky Survey Bright Source Catalogue (Voges et al., 1999). GRB afterglows without a gamma ray trigger (the so-called orphan GRB afterglows) are a possibility and might explain the high luminosities with respect to RSNe. RSNe are also a possibility as long as the hosts are nearby and intrinsically faint, but distant RSNe would appear to be ruled out by the high luminosity. Stellar events are less likely, but flares from nearby late-type stars are not ruled

out. Soft gamma repeaters and X-ray binaries are also possibilities, although the rarity of such objects and our measured transient rate are incompatible.

The most likely source of variability from extragalactic sources at 843 MHz is refractive interstellar scintillation (ISS) of the compact components of AGN. By our definition of a transient source, our transients could well be such scintillating sources occasionally appearing above our flux threshold, although the light curves of the transient sources have $\mathcal{V} > 0.6$, which would make them amongst the most extremely variable sources in our sample (see Figure 2.20). Star-forming galaxies, whose radio flux is from widespread star formation and diffuse synchrotron emission in the disk, are unlikely to be intrinsically variable on our time-scales (due to light travel time constraints), or extrinsically variable (due to larger angular size quenching the interstellar scintillation).

Another explanation for the transient sources is time-integrated emission from pulsars with on-times of greater than ~ 12 h. Ofek et al. (2010) proposed pulsars as the source of the transients discovered by Bower et al. (2007) and Niinuma et al. (2007) based on population arguments, and there is some observational evidence for this. Intermittent pulsars such as PSR B1931+24 (Kramer et al., 2006) have active periods of a few days and then turn completely off. PSR J0941–39, which was discovered as a Rotating Radio Transient (RRAT) with 5 single pulses in one observation, exhibited standard pulsar emission during another observation (Burke-Spolaor & Bailes, 2010). Such objects might appear as transients in our survey if the time between mode-changes has a time-scale of 12 h or longer, and if the time-integrated flux densities in both modes are compatible with our non-detection thresholds. Pulsars are also not likely to be optically detected at the SuperCOSMOS plate limits, which would explain the lack of optical detections of some of our transient sources.

Finally, microlensing is an unlikely explanation for our unknown transients, as the implied number of sources to explain our sample is far too high (Bower et al., 2007). Reflected solar flares are unlikely, as a 1 MJy solar flare would need to reflect off a 1000 km object at a distance of only 7000 km to be detected above our threshold.

Table 2.7: Summary and classifications of the 15 radio sources classified as transient.

Source	Radio Lightcurve	Counterparts at other wavelengths	Classification
SN1987A	Linear increase	Bright point source in SuperCOSMOS R and I bands.	Supernova.
SUMSS J055712–381106	Detection followed by a non-detection 6 days later.	Blend of three objects in SuperCOSMOS and 2MASS.	Unknown
J060938–333507	Single detection. Non-zero PA possibly indicates variation during 12h synthesis.	9 arcsec from the optical centre of Fairall 1138, a barred spiral galaxy at $z = 0.037$.	Luminosity $L_\nu \simeq 6 \times 10^{29} \text{ erg s}^{-1} \text{ Hz}^{-1}$. Possible ultra-bright RSN or GRB afterglow.
J061051–342404	Non-detection between two MOST detections Also detected in NVSS.	No known counterparts.	Possible scintillating AGN appearing above the noise threshold.
SUMSS J062636–425807	Single detection.	SuperCOSMOS and 2MASS colours typical of a quasar.	Probable scintillating AGN.
J062716–371736	Non-detection between two detections.	No known counterparts.	Unknown
J064149–371706	Single detection in MOST and NVSS.	No known counterparts.	Possible scintillating AGN appearing above the noise threshold.
SUMSS J102641–333615	Single detection in MOST and NVSS.	No known counterparts.	Possible scintillating AGN appearing above the noise threshold.
SUMSS J112610–330216	Detection and non-detection separated by 1 day.	No known counterparts.	Less likely to be an AGN or explosion. Possible flare star other short term variable.
Nova Muscae 1991	Single detection.	Bright point source visible only in SuperCOSMOS R image. Crowded field. Point source in RASS.	Low mass X-ray binary.
J121032–381439	Non-detection between two detections.	No known counterparts.	Unknown
J135304–363726	Non-detection between two detections	Faint ($M_{R1} = 20.4$ $M_{R2} = 20.7$) point source in SuperCOSMOS.	Probable scintillating AGN but faint, red detection could be a star.
J153613–332915	Non-detection between two detections. Non-zero PA possibly indicates variation during 12h synthesis.	No known counterparts.	Unknown
GRO 1655–40	Power law decrease over 3 months, then non-detection after 11 years.	Point source in all SuperCOSMOS plates.	GRO J1655–40. A low mass X-ray binary and micro-quasar.
SUMSS J224152–300823	Single detection; also detected in NVSS.	NVSS source is 10 arcsec from a blue point-like optical source, but probably not associated.	Possible scintillating AGN appearing above the noise threshold.

2.7 CONCLUSIONS

We have conducted a 22 year survey for radio variability at 843 MHz and measured 29230 radio light curves. We have discovered 15 candidate transient sources and 53 candidate highly variable sources. Only 3 of the transients were previously known.

We conclude that many variable sources can be explained as scintillating AGN, and some that are associated with nearby galaxies may be over-luminous radio supernovae with atypical light curves. We conclude that at least 3 of the transients are unlike any known source and could belong to the class of radio transients without optical counterpart discovered by Bower et al. (2007), which have a number of possible explanations including giant M-dwarf flares, or flaring Galactic neutron stars.

We have also presented a number of statistical techniques to aid in future radio variability surveys. In particular, we have described techniques to remove systematic gain errors, to confirm that errors in flux density are correctly estimated, and to choose a robust variability threshold.

Upcoming wide-field radio variability surveys will offer a new insight into this parameter space. We have shown that the radio sky is indeed variable so these surveys are likely to uncover many new sources. We have also confirmed that there is at least one class of transient radio source without optical counterparts. The time interval between radio activity and optical measurement in our cases was several years, highlighting the need for prompt multi-wavelength follow-up.

**An automated, high time
resolution search for prompt
radio emission from gamma ray
bursts at 1.4 GHz**

3.1 INTRODUCTION

GAMMA Ray Bursts are flashes of gamma rays that signal the most luminous explosions in the Universe. GRBs fall into two basic types based on the duration of their gamma ray light curves (Kouveliotou et al., 1993): the ‘long-soft’ GRBs which have gamma ray light curves lasting more two or more seconds and relatively soft high energy spectra, and the ‘short-hard’ GRBs with gamma ray light curves lasting less than two seconds, and somewhat harder high energy spectra.

Long GRBs are well explained by the ‘collapsar’ model (MacFadyen & Woosley, 1999), in which the core of a rapidly rotating massive star collapses, producing a compact central engine (which may be a magnetar or black hole) and accompanying supernova explosion. Short GRBs, however, are less well explained. The leading model is a double-degenerate scenario involving the merger of two neutron stars (NS), or a neutron star and black hole (Nakar, 2007, and references therein). Other models for short GRBs have been proposed, for example, the formation of a magnetar by accretion induced collapse of a neutron star (Qin et al., 1998), accretion induced collapse of a white dwarf (Metzger et al., 2008), or merger of two white dwarfs (Levan et al., 2006). Giant magnetar flares have also been proposed as the source of at least some short GRBs (Hurley, 2011), and other models for short GRBs have been proposed, including evaporating primordial black holes (Cline, 1996).

Prompt radio emission is predicted by some of these models. In the long GRB ‘collapsar’ case, Sagiv & Waxman (2002) propose that particle acceleration in collisionless shocks can generate emission in the range 3.5 - 200 MHz, with a flux density of up to 1 Jy (Inoue, 2004) approximately 1 minute after the release of gamma rays. Usov & Katz (2000) propose that scattering of ambient medium electrons by a magnetically dominated wind could produce emission lines at frequencies < 30 MHz for approximately 1 to 100 s after the burst. Finally, Moortgat & Kuijpers (2005) propose that gravitational waves from the system could excite magneto-hydrodynamic waves at kilohertz frequencies, which in turn undergo inverse Compton scattering up to observable (~ 100 MHz) frequencies, in an extremely bright ($> 10^6$ Jy) burst lasting a few hundred seconds.

For the merging neutron star scenario, Lipunov & Panchenko (1996) suggest that a radio pulsar may be revived several orbital periods before the merger. Hansen & Lyutikov (2001) predict that the interaction of the two magnetospheres just before the merger could produce a pair plasma, generating radio emission with a flux density between 1 mJy and 5 Jy. Shibata et al. (2011) predict that the merging NS produce a rapidly and differentially rotating high mass neutron star, whose magnetic field is quickly amplified to produce a 10 to 100 ms burst of emission at quasar-like luminosities, although they do not describe a spectrum for this emission. In the magnetar model for short GRBs, Pshirkov & Postnov (2010) make predictions for low frequency

(100 MHz) emission based on a magnetar-like magnetic field assuming the standard pulsar emission mechanism. Magnetars outside of GRBs are known to produce radio emission (Camilo et al., 2006; Levin et al., 2010). Bright, and very short duration radio emission is also predicted from the evaporation of primordial black holes (Rees, 1977).

The radio emission from these events is not necessarily detectable, however. Macquart (2007) noted that, for small to medium density of circumburst interstellar medium (ISM), induced Compton scattering limits the observable brightness temperature, so that radio emission is only visible if the jet is ultra-relativistic or has a very small opening angle. If the circumburst medium has a high density (e.g. if the progenitor released a dense wind before explosion), induced Raman scattering places an even stronger limit on the brightness temperature of any observed prompt radio emission. On the other hand, Lyubarsky (2008) notes that if the pulse of radio emission is sufficiently short, or if any dense plasma is sufficiently far from the source, then the constraints on brightness temperature are lifted. Therefore, the detectability of any radio emission depends on a complex combination of the emission properties and circumburst medium.

The rewards for detecting such prompt emission are great. Many aspects of the explosion physics can be probed by measurements of the prompt radio emission, such as the jet opening angle, and Lorentz factor (Macquart, 2007), the density and distance of any scattering material (Lyubarsky, 2008), and the structure of the fireball magnetic field (Sagiv et al., 2004). Additionally, a short radio pulse from an extragalactic source is expected to undergo dispersion as it propagates through the intergalactic medium (IGM). Such dispersion would not only provide direct evidence for the existence of the majority of the baryons in the Universe (Ginzburg, 1973), but, for bursts of sufficiently high redshift, would also differentiate between different models of cosmic reionisation history (Inoue, 2004).

These rewards have motivated a number of searches for prompt radio emission. Most searches have been non-directed, in which a large fraction of the sky was monitored, with the hope of a GRB occurring somewhere in this region. Early results at 151 and 408 MHz with an integration time of 300 ms detected some pulses within ± 10 min of the gamma ray trigger, but did not confidently associate any with GRBs (Cortiglioni et al., 1981; Inzani et al., 1982). A survey at 843 MHz, which was sensitive to pulses between 0.001 ms and 800 ms, also made no definitive detections (Amy et al., 1989), although it was not clear if any GRBs were present in the field of view during the observations. More recently, Katz et al. (2003) performed an all-sky survey at 611 MHz, with a time resolution of 125 ms and a flux density detection threshold of 27 kJy. This search detected $\sim 4 \times 10^6$ bursts in 18 months, but rejected 99.9% as RFI and identified the remaining bursts with solar activity. In a similar vein, Lazio et al. (2010) detected no transients above 500 Jy for pulse widths of about 300 s at 73.8 MHz. A number of other surveys have be-

gun but are yet to publish results (Balsano et al., 1998; Morales et al., 2005). The only report of automatic follow-up at radio frequencies is that performed by Koranyi et al. (1995) and Dessenne et al. (1996) at 151 MHz with a time resolution of 1.5 s. Based on observations of two GRBs, Dessenne et al. (1996) report upper limits on any radio emission of 16–73 Jy between 5 hrs before, and 2 hrs after the GRB. A search for the evaporation of primordial black holes at 3 GHz with a time resolution of $2\mu\text{s}$ also failed to detect any radio emission (O’Sullivan et al., 1978).

There are a number of factors that may have contributed to the lack of radio detections by previous surveys. First, most of the surveys described above have been at low frequencies. Low frequency observations have the advantage of a large predicted radio luminosity (e.g. Sagiv & Waxman, 2002), plus a large dispersion delay, that allows for time to slew radio telescopes to observe any radio emission emitted simultaneously with the GRB (Palmer, 1993). However, scatter broadening, which substantially reduces the detectability of radio pulses (Cordes & McLaughlin, 2003), is more severe at low frequencies (Cordes & Lazio, 2002). It is possible that low frequency photons predicted by the models could be inverse-Compton scattered up to GHz frequencies by mildly relativistic outflows. Second, many of the above searches, and the directed searches of Koranyi et al. (1995) and Dessenne et al. (1996) in particular, used relatively long integrations of 1.5 s, which renders any short duration radio pulses, including those required to avoid the brightness temperature constraint from induced scattering (Lyubarsky, 2008), more difficult to detect. Finally, the sensitivities of most of the above surveys are low, and not approaching the flux levels required by some models (e.g. Hansen & Lyutikov 2001).

In this chapter I describe a survey to detect prompt radio emission from GRBs at 1.4 GHz. I used a single 12 m dish that slews automatically to the GRB coordinates, based on the gamma ray position. The aim was twofold: to attempt to detect any GHz radio emission within the first few minutes of the GRB, and to gain experience for a potential future experiment.

3.2 OBSERVATIONS

3.2.1 Telescope Configuration

I used a single 12 m dish antenna at the Parkes radio observatory. The visible declinations ranged between $-90^\circ < \delta < +20^\circ$. The dish had a 1.4 GHz horn at the prime focus, with orthogonal linear polarisations supplying room temperature low noise amplifiers bolted onto the horn, and wrapped in a thermally insulating blanket. The T_{sys} was measured at approximately 100 K and the bandwidth was 220 MHz between 1230 to 1450 MHz. I estimate the system equivalent flux density to be 3.8×10^2 Jy. No radiometer switching was used. Both polarisations were upper-sideband mixed down to a centre fre-

quency of 600 MHz and digitised and channelised by a digital spectrometer based on an IBOB (Internet Break-out Board) (McMahon, 2008). Sampling was performed at 800 MHz, resulting in sampling at the second Nyquist zone. The sampled signal of each polarisation was channelised with a 1024 channel polyphase filterbank spread across 400 MHz, with an approximate frequency resolution of 390 kHz. The cross polarisation product was not formed. The channelised output was detected, integrated and dumped every 64 microseconds with a resolution of 8 bits per channel. The integrated spectrometer data were written to disk in real time and all processing was performed offline.

3.2.2 GRB response

A control computer was connected to the Gamma-ray Coordinate Network (GCN) via the socket distribution method (Barthelmy et al., 1995). The GCN was configured to send only notifications when the time delay between the satellite detection and socket distribution was less than 1 hour, the position of the burst was more than 10 degrees above the visible horizon, and the position error was less than 1 degree (the width of the primary beam of the telescope). The GCN was configured to send notifications from the *Swift*/BAT, *Fermi*, *Integral* and *Agile* missions, although only events from *Swift*/BAT passed the filter constraints.

When a GRB alert arrived from the GCN, the telescope immediately began saving the 8-bit channelised data. It then slewed to the coordinates given in the burst alert, while saving the raw spectra to disk. Typically the slew took two minutes, after which the telescope was “on source”. After 30 minutes on source on the initial GRB position (the tracking position was not updated if additional position refinements were sent by the GCN), data capture was stopped. To check that the receiver system was operating correctly, immediately after the GRB capture stopped, the telescope slewed to a bright pulsar (either the Vela pulsar or PSR B1645-40) and took an additional 10 minutes of raw data once the antenna was on source.

The telescope performed pulsar monitoring when it was not observing a GRB alert, or observing a bright pulsar for the system check. The pulsar monitoring results are not reported here.

3.2.3 Processing

My primary aim was to search for repeating signals, and single pulses at high time resolution (~ 1 ms). My secondary aim was to search for slowly varying signals at low-time resolution (> 1 s). For the slow search, dispersion can be ignored as the dispersion delay across the observing bandwidth is ~ 770 ms for a dispersion measure (DM) of 1000 pc cm^{-3} , which is less than the time resolution. Both timescales were affected by RFI, which can increase the number of false detections.

3.2.3.1 Kurtosis-based RFI excision in the frequency domain

When a channel is affected by impulsive RFI, it typically contains more samples with large values than a channel containing Gaussian noise alone. One method for measuring an excess of large values is to compute the ‘excess kurtosis’ (κ), which compares the number of large values in a time series against the expected number of large values given a Gaussian distribution.

To determine which of the channels contained impulsive RFI, I took the 8-bit raw data and, for each channel computed excess kurtosis in one second intervals. If I define $x[n]$ as the n^{th} sample in a given interval, and N as the number of samples in the interval, then the excess kurtosis is given by the equation:

$$\kappa = \frac{\mu_4}{\sigma^4} - 3 \quad (3.1)$$

where μ_4 is the fourth moment of the distribution, defined as

$$\mu_4 = \frac{1}{N} \sum_{n=0}^{N-1} (x[n] - \mu)^4 \quad (3.2)$$

$$\mu = \frac{1}{N} \sum_{n=0}^{N-1} x[n], \quad (3.3)$$

and σ^4 is the square of the variance of the distribution, given by

$$\sigma^4 = \left(\frac{1}{N-1} \sum_{n=0}^{N-1} (x[n] - \mu)^2 \right)^2. \quad (3.4)$$

The excess kurtosis is calculated for each channel in one-second bins, and displayed in an interactive program. This program plots the maximum and average kurtosis as a function of channel number. The maximum kurtosis for a channel that was not affected by RFI was $\kappa \simeq 0.6$, and the RFI affected channels had larger values of κ (Figure 3.7). An RFI channel mask was created by setting all channels with an average κ less than a threshold to a weight of 1, and the remaining channels to a weight of zero. The threshold could be set interactively in the program, and was set to between 0.6 and 0.8. In a subsequent step, a user-defined number of additional “padding” channels were set to zero on either side of a channel that had zero weight because of the threshold step. The number of additional padding channels was between 1 and 3.

The kurtosis was calculated separately for each polarisation, but the same threshold and number of padding channels were used for each. This resulted in different channel masks for each polarisation. The final channel mask was the logical ‘and’ of both channel masks.

The resulting channel masked was used for both for low-, and high-time resolution processing.

3.2.3.2 Low-time resolution processing

For low-time resolution processing, I took the mean of each channel in one-second bins. To compute a light curve, I weighted the spectra by the channel mask that I calculated according to the method in section 3.2.3.1, and took the mean value across frequency. I then took the mean of the light curves of both polarisations.

The final product was a radio light curve, which I visually inspected for variability or transient emission.

3.2.3.3 High-time resolution processing

Each channel in each polarisation was offset to have zero mean and scaled to a variance of unity. The offset and scale were updated every 10 seconds. Corresponding channels in both polarisations were then added together and the result truncated to 2 bits. A periodicity search was performed in an almost identical manner to Keith et al. (2010) and single-pulse search was performed identically to Burke-Spolaor et al. (2011b). Briefly, the time series was dedispersed with 1196 DM trials between 0 and 1000 pc cm^{-3} . DM steps were chosen so that the smearing in each channel was less than 25%, which results in a roughly quadratically spaced set of DMs, with low spacings at lower DMs and higher spacings at higher DMs. Time-based RFI excision was applied. Unlike the processing of Keith et al. (2010), a frequency mask (calculated in Section 3.2.3.1) was applied, instead of the fast Fourier transform (FFT) based channel masking approach. An FFT based harmonic search was performed to search for periodic or pulsar-like signals, and a single-pulse search was also performed to search for bursty signals. The single-pulse search used the dedispersed time series, and applied 9 different boxcars of width 2^i samples ($0 \leq i \leq 9$) to search for pulse widths ranging from $64 \mu\text{s}$ to 32 ms. The single-pulse candidates were grouped across DM and boxcar trials using a friend-of-friends algorithm (Burke-Spolaor et al., 2011b). Pulses with a significance of $> 6\sigma$ were declared as candidates only if three or more $> 6\sigma$ pulses were detected in a group. Groups with less than three pulses were classified as Gaussian fluctuations, and groups whose peak significance was at a DM of less than 2 pc cm^{-3} were classified as RFI.

I apply an additional filter to reduce the effect of weak, impulsive RFI. Weak impulsive RFI can appear at low DMs and then disappear at intermediate DMs, only to appear again at high DMs. The friend-of-friends algorithm fails to identify this situation because there is not a continuous stream of detections from the high DMs down to low DMs, at which point they are classified as RFI. In addition, I have no interest in low-DM pulses, which are

likely to be Galactic, and some RFI can have peaks at low, but non-zero DM. Therefore, I classified candidates initially produced by the pipeline as RFI if there is a $> 6\sigma$ peak within 3 s of the burst with a DM of $< 25 \text{ pc cm}^{-3}$.

Finally, I visually inspected plots of the candidates which were produced by the pipeline.

3.3 RESULTS

I describe the results of the single pulse, periodic and low-time resolution searches in this section. I note that the dispersion delay between the gamma rays and the 1.4 GHz observing frequency, at the maximum DM of 1000 pc cm^{-3} is $< 3 \text{ s}$, which is small enough to ignore in the following analysis.

3.3.1 Events

The telescope was operating between 2010 June and 2011 February, during which it responded to 16 *Swift*/BAT alerts of which nine were confirmed GRBs (Table 3.1). The remainder were instrumental effects, associated with known recurring objects, or Galactic sources such as X-ray binaries.

Of the nine GRBs, two had a duration $< 2 \text{ s}$, satisfying the definition of short GRBs. The remaining seven GRBs were long GRBs. I detected no radio counterparts in the single pulse, periodic or long integrations for any short GRBs.

Table 3.1: Details of 9 GRBs observed in my experiment. The columns are: the GCN Trigger ID, the source name, the GRB duration (i.e. the duration of activity in the BAT lightcurve), the trigger date, the interval between the GRB detection and the 12 m telescope being ‘on source’ (T_{on}), the GCN circular that describes the detection, the time from the GRB trigger to the arrival of a single pulse where detected (T_{pulse}), the DM of the pulse, the DM distance (D_{DM}) from the NE2001 electron density model of Cordes & Lazio (2002), the S/N of the detected pulse, the pulse width and a comment. The GCN circulars describe a number of attempts to detect host galaxies by a number of groups, but no host galaxies were conclusively identified for any GRBs listed except for GRB 100625A.

Trigger	Source	Duration (s)	Trigger Date	T_{on} (s)	Circular	1.4 GHz Candidate					Comment	
						Type	T_{pulse} (s)	DM (pc cm^{-3})	D_{DM} (kpc)	S/N		Width (ms)
425647	GRB 100625A	0.33	2010-06-25 18:32:28	136	Holland et al. (2010)	-	-	-	-	-	-	Host galaxy detected with $r \sim 23$
426114	GRB 100628A	0.1	2010-06-28 08:16:40	160	Immmler et al. (2010)	-	-	-	-	-	-	
426722	GRB 100704A	280	2010-07-04 03:35:08	187	Grupe et al. (2010a)	-	-	-	-	-	-	
432420	GRB 100823A	16.9	2010-08-23 17:25:35	141	Mangano et al. (2010)	Single	1076	194.57	> 48	6.2	6	
436737	GRB 100928A	3.3	2010-09-28 02:19:52	131	D'Elia et al. (2010)	-	-	-	-	-	-	No <i>Swift</i> follow-up due to sun constraint. Sun angle: 46°
436094	GRB 101011A	50	2010-10-11 16:58:35	105	Cannizzo et al. (2010b)	Single	524	569.98	> 35	6.6	25	
436737	GRB 101020A	177	2010-10-20 23:40:41	317	Saxton et al. (2010)	-	-	-	-	-	-	No <i>Swift</i> follow-up due to sun constraint. Sun angle: 26°
451191	GRB 110412A	30	2011-04-12 07:33:21	214	Gelbord et al. (2011)	Periodic	0-1200	50.73	> 46	6.1	5	No <i>Swift</i> follow-up due to moon constraint. Moon angle: 10° . Candidate period 524.48 ms
451343	GRB 110414A	230	2011-04-14 07:42:14	192	Starling et al. (2011)	-	-	-	-	-	-	

3.3.2 Single-pulse search

My method detected candidates for all GRBs I observed, however visual inspection of the candidate plots indicated that almost all these candidates were probably associated with RFI. I rejected candidates usually because there was strong impulsive RFI at nearby times not covered by my 3 s flagging interval, or because the pulse was detected over a wide range of DMs, but was not classified by RFI by the pipeline because it did not go through the zero DM.

Of the nine GRBs for which I obtained data, two GRBs had one single radio pulse which was not associated with RFI (Figure 3.1). For GRB 100704A, a 6 ms wide single pulse was detected 1076 seconds after the GRB at a DM of 195 pc cm^{-3} , at a significance of 6.2σ . This pulse is fairly indistinct in the time domain, and appears to be stronger at low frequencies (Figure 3.2). For GRB 101011A, a 25 ms single pulse was detected 524 seconds after the trigger at a DM of 570 pc cm^{-3} , at a significance of 6.6σ (Figure 3.3).

3.3.2.1 Event rate

I detected two pulses in 4.5 hrs of observing. The full width, half maximum beam width of the 12 m antenna at 1.4 GHz is approximately 1.3 deg , which implies an area of 1.3 deg^2 . If these events are not related to the GRBs and could have been discovered in a blind search, then the implied event rate is $3.3 \times 10^{-1} \text{ deg}^{-2} \text{ hr}^{-1}$.

3.3.2.2 Correspondence with X-ray light curves

Swift/*XRT* light curves are available for both GRBs for which I detected single pulses. For GRB 101011A, Cannizzo et al. (2010a) fit a three component broken power law to the 0.3-10 keV light curve, and derive a value of the final break in the power law of $t_{\text{break},2} = 602^{+175}_{-88} \text{ s}$. Cannizzo et al. (2010a) do not describe their exact fitting method. I fit the same data with a broken power law, with each piece having a form $S(t) \propto t^{\alpha_i}$ for $0 \leq i < N_{\text{break}}$, and with each data point weighted by the flux error. I set the initial break times at $t = 116$ and $t = 602$, as derived by Cannizzo et al. (2010a). From this fit ($\chi^2/N_{\text{d.o.f}} = 1.6$) I derive a break time of $t_{\text{break},2} = 707 \pm 173$. The arrival time of the radio pulse coincides with the break in the power law, to within the statistical errors for both fitting methods (Figure 3.4). I note that the dispersion delay between the gamma rays and my radio frequencies is less than 2 s and therefore insignificant for this calculation.

For GRB 100704A, no fit is derived in the GCN report (Grupe et al., 2010b). I fit a four component power law to the X-ray light curve for the data after $t = 400 \text{ s}$ (Figure 3.4), corresponding with the beginning of the proportional-counting mode data set. I used initial breaks times at 500, 1000 and $3 \times 10^5 \text{ s}$. The resulting fit ($\chi^2/N_{\text{d.o.f}} = 1.2$), has a break at $t = 1700 \pm 410 \text{ s}$, which is within 1.5σ of the pulse arrival time at $t = 1076 \text{ s}$. I note that

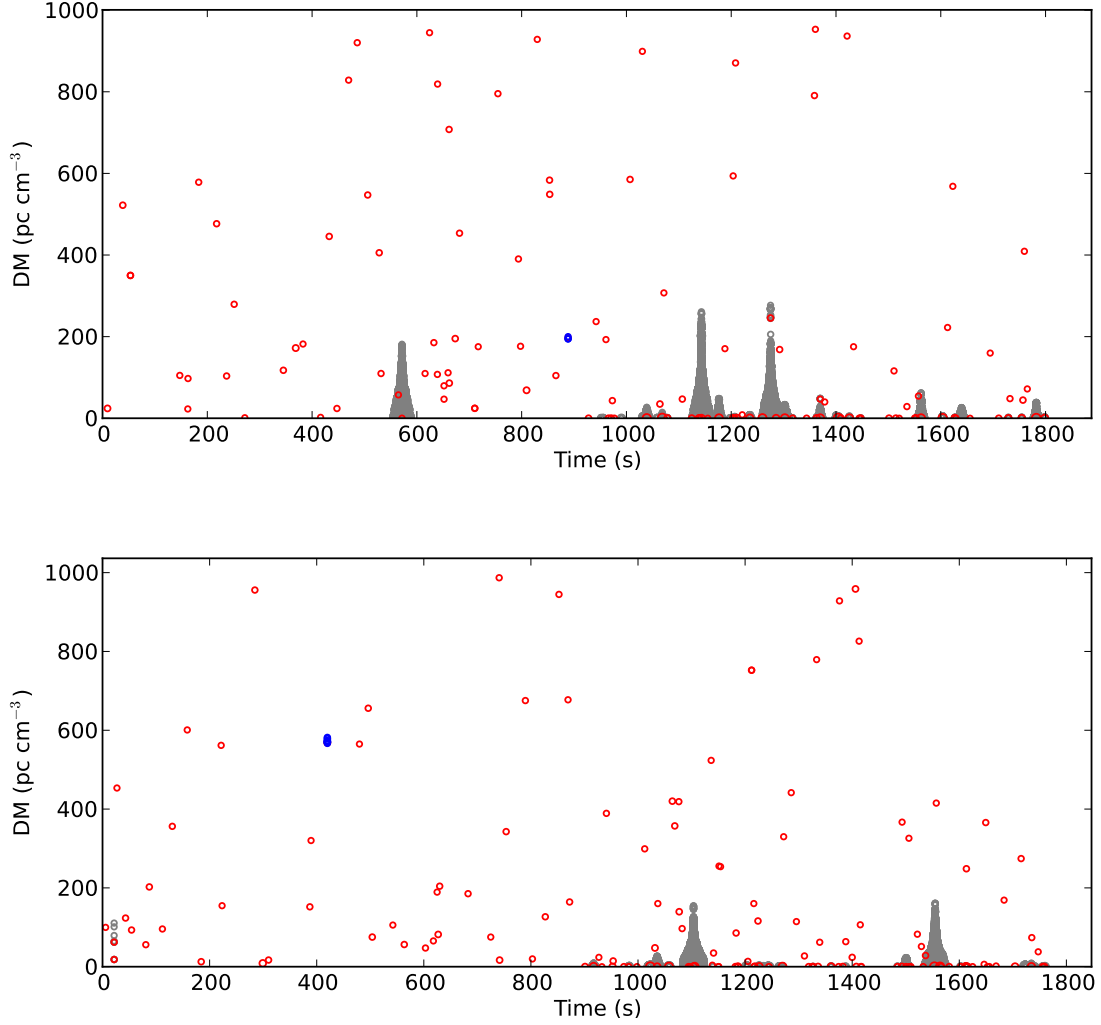


Figure 3.1: DM vs. time for the two GRBs with single-pulse candidates. Single-pulse detections with a significance $\geq 6\sigma$ appear as circles in this plot, with the size proportional to the S/N. The detections are colour coded according to their classification by the friends-of-friends algorithm as candidates (blue), false positives (red) and RFI (grey). Top panel: GRB 100704A with a single pulse candidate 1076 seconds after the GRB at a DM of 195 pc cm^{-3} , with a significance of 6.2σ and width of 6 ms. Bottom panel: GRB 101011A with a single pulse candidate 524 seconds after the GRB at a DM of 570 pc cm^{-3} , with a significance of 6.6σ and width of 25 ms. The time origin of these plots is the time that the telescope first arrived on source (T_{on}).




Figure 3.2: Detail of the 6.2σ pulse detected 1076 seconds after GRB 100704A. The top three panels are the dedispersed time series of 4 frequency channels spread across the band. The pulse appears clearly in both polarisations (the top 2 panels) and in the sum of the two polarisations (third panel). The bottom panel is the time series where all the frequency channels have been summed. The origin of the time axis is the pulse arrival time (T_{pulse}).




Figure 3.3: Detail of the 6.6σ pulse detected 524 seconds after GRB 101011A. The top three panels are the dedispersed time series of 4 frequency channels spread across the band. The pulse is clearly detected across all channels, in both polarisations (the top 2 panels) and in the sum of the two polarisations (third panel). The bottom panel is the time series where all the frequency channels have been summed. The origin of the time axis is the pulse arrival time (T_{pulse}).

the time of the fitted breaks is quite sensitive to the choice of initial break times.

3.3.3 Periodicity search

My method detected candidate repeating signals for all GRBs I observed, however, visual inspection of the candidate plots indicated that almost all of these candidates were associated with RFI.

The pulsar search yielded a single candidate for a repeating signal that was not conclusively associated with RFI. The candidate was detected in the data following GRB 110412A (Figure 3.5) with a S/N of 5.3 in the periodicity search, a topocentric period 524.48 ms and a DM of 51 pc cm⁻³. After optimisation for DM and period, the S/N of the candidate increases to 6.1. The candidate is brightest during the first 20 minutes of the observation. The signal is clearly detected in all sub-bands, as expected for a broad-band pulsar signal, and which also rules out narrow-band RFI. The profile comprises a single component with a width of about 5 ms. The signal is not very well defined in the period-DM plane or in the DM-S/N plot, but this is not unusual as the S/N is relatively low.

3.3.4 Low-time resolution search

We did not detect any significant variability in the low-time resolution search.

The light curves at 1 s integration are shown in Figure 3.6. The light curves with rapid changes in flux density can be attributed to interference. For example, the flux density variations during observations of GRB 1000823A are due to RFI from the *Beidou G1* satellite.⁹ This can be seen from the modulation spectrum centred at 1270 MHz in the bandpass and variance plots. During the observations, the telescope pointing direction was within 10 degrees of the satellite position when the interference was most severe (Figure 3.8). Although my method flagged the worst channels, it is clear that additional power leaked through the unflagged channels, affecting the light curve.

The substantial flux variations during observations of GRB101020A can be attributed to the proximity of the sun to the observing direction, which was 25 degrees from the sun at the beginning of the observation. None of the narrow-band RFI coincides with the large changes in flux density, indicating that the changes in flux are broad-band and are likely to be a result of the sun moving through the far sidelobes.

⁹Refereed, English information about the *Beidou* global satellite system is non-existent. The only available information available in English is on general news and satellite navigation websites. Perhaps the best source of information can be found at http://en.wikipedia.org/wiki/Beidou_navigation_system. The official Chinese website is <http://www.beidou.gov.cn>. *Beidou* satellite ephemerides are provided by the US Space Command and the real time position was computed using data provided at <http://www.n2yo.com/>.

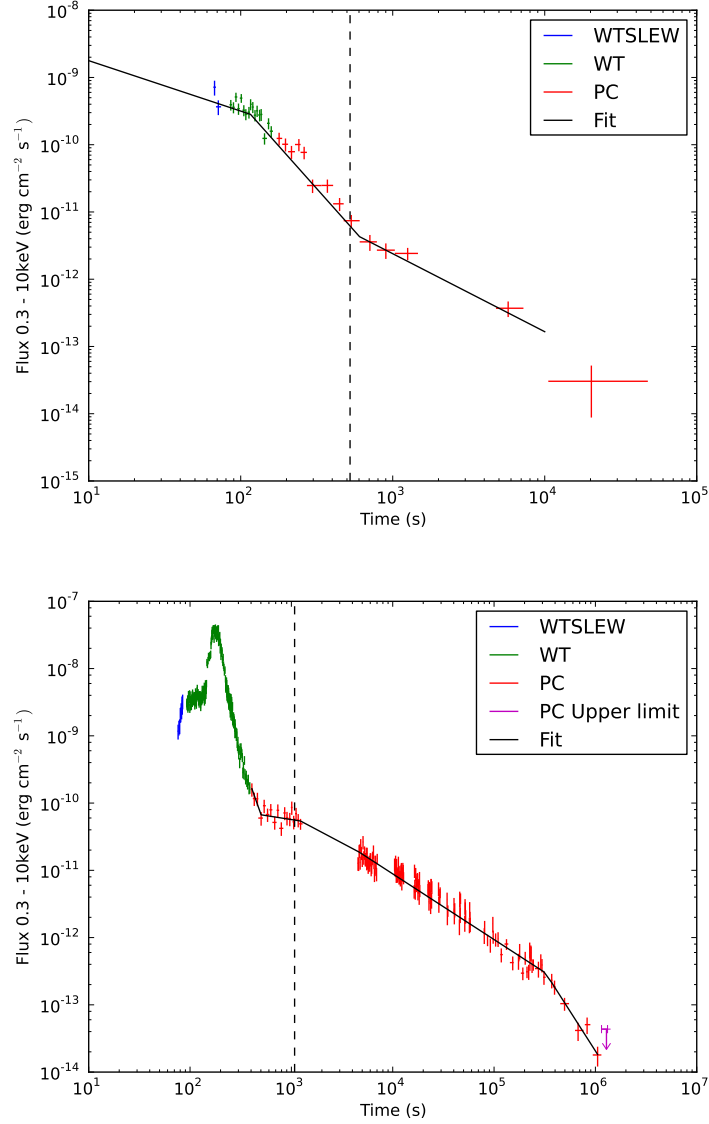


Figure 3.4: The XRT (0.3–10 keV) light curves of the two GRBs after which I detected a single radio pulse. The arrival time of the radio pulse (T_{pulse} ; see Table 3.1) is marked with a vertical line. Top panel: the X-ray light curve of GRB 101011A. The black line is the fit derived by Cannizzo et al. (2010a). Bottom panel: the X-ray light curve of GRB 10704A. The black line is the fit I derived. WTSLEW is the windowed timing mode during the slew of the spacecraft, during which flux measurements can be less reliable. WT is windowed timing mode. PC is photon counting mode.

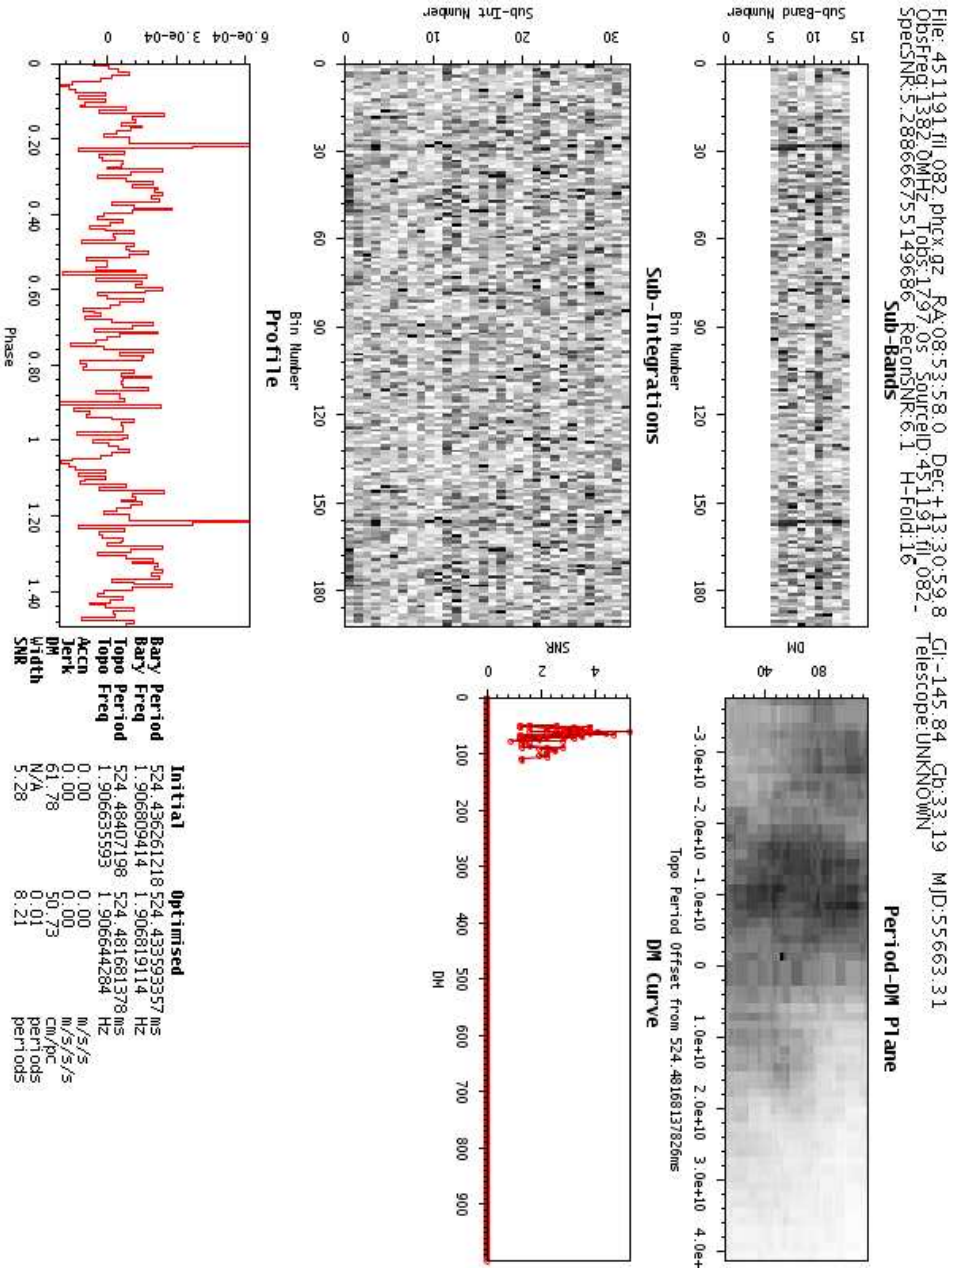


Figure 3.5: Pulsar processing diagnostic plots for the best pulsar candidate detected during observations of GRB 110412A. Top left panel: an image of the dedispersed pulse amplitude vs bin number and frequency. Broad-band, dispersed signals appear as vertical lines in this plot. Centre-left panel: An image of pulse amplitude vs bin number and time. Signals with fixed period appear as vertical lines in this plot. Bottom-left panel: pulse profile integrated over all times and frequencies. Top right panel: an image of S/N vs period and DM before optimisation. High-S/N, dispersed and periodic signals appear as well-localised points in this plot. Centre-right panel: DM vs S/N at the best fitting period. A dispersed signal with high S/N usually has a single peak at the correct DM, although wide pulse widths can substantially broaden the peak in this plot.

The variations during observations of GRB 110412A are likely due to unflagged narrow-band interference at around 1230 MHz (Figure 3.9). These channels had no substantial kurtosis and so were not flagged, but had slightly higher variance than the rest of the band (Figure 3.7). Nonetheless, the fact that the interference is narrow band strongly suggests RFI as the source of the variation.

The slow variations (e.g. GRB 100625A) are due to gain variations in the system, which is stable to about 5% in 30 minutes. I were unable to calibrate out these gain variations, as my feed was not equipped with a switched radiometer.

3.4 DISCUSSION

3.4.1 Single-pulse detections

I have two single-pulse candidates, the brightest of which is a 6.6σ pulse 524 s after GRB 101011A. Below I discuss possible origins for this pulse. My conclusions also broadly apply to the other candidate single-pulse (the 6.2σ pulse 1076 s after GRB 100704A), but I restrict my attention to the brightest pulse.

The DMs of both pulses are much larger than the predicted DM in the GRB direction predicted by the NE2001 model (Cordes & Lazio, 2002). For GRB 101011A, NE2001 predicts a maximum DM of 39 pc cm^{-3} at a DM distance of 3 kpc, which implies an excess DM of 530 pc cm^{-3} . For GRB 100704A, NE2001 predicts a maximum DM of 106 pc cm^{-3} at a DM distance of 9 kpc, which implies an excess DM of 89 pc cm^{-3} . The large excess DMs for both radio pulses are well above the uncertainties in the NE2001 model and therefore exclude Galactic sources such as pulsars or RRATs. The excess DM must therefore be due either to intergalactic material or a possible host galaxy.

3.4.1.1 Random noise fluctuation - theoretical arguments

In principle I can apply statistical arguments to derive the probability of a 6.6σ pulse occurring by chance in my data. In practise, such arguments are not straightforward for a number of reasons. Firstly, my data are not easily described by analytic probability distribution functions, chiefly because of the presence of RFI. Secondly, the processing pipeline applies a number of RFI mitigation strategies and matching algorithms that are difficult to model in a statistical sense. Finally, the search over a range of DMs and pulse widths, which necessarily sample overlapping regions of parameter space, is not independent. That is, a genuine astronomical signal will appear in a number of adjacent DMs and boxcar trials, implying that the simplifying assumption of independent trials is not valid.

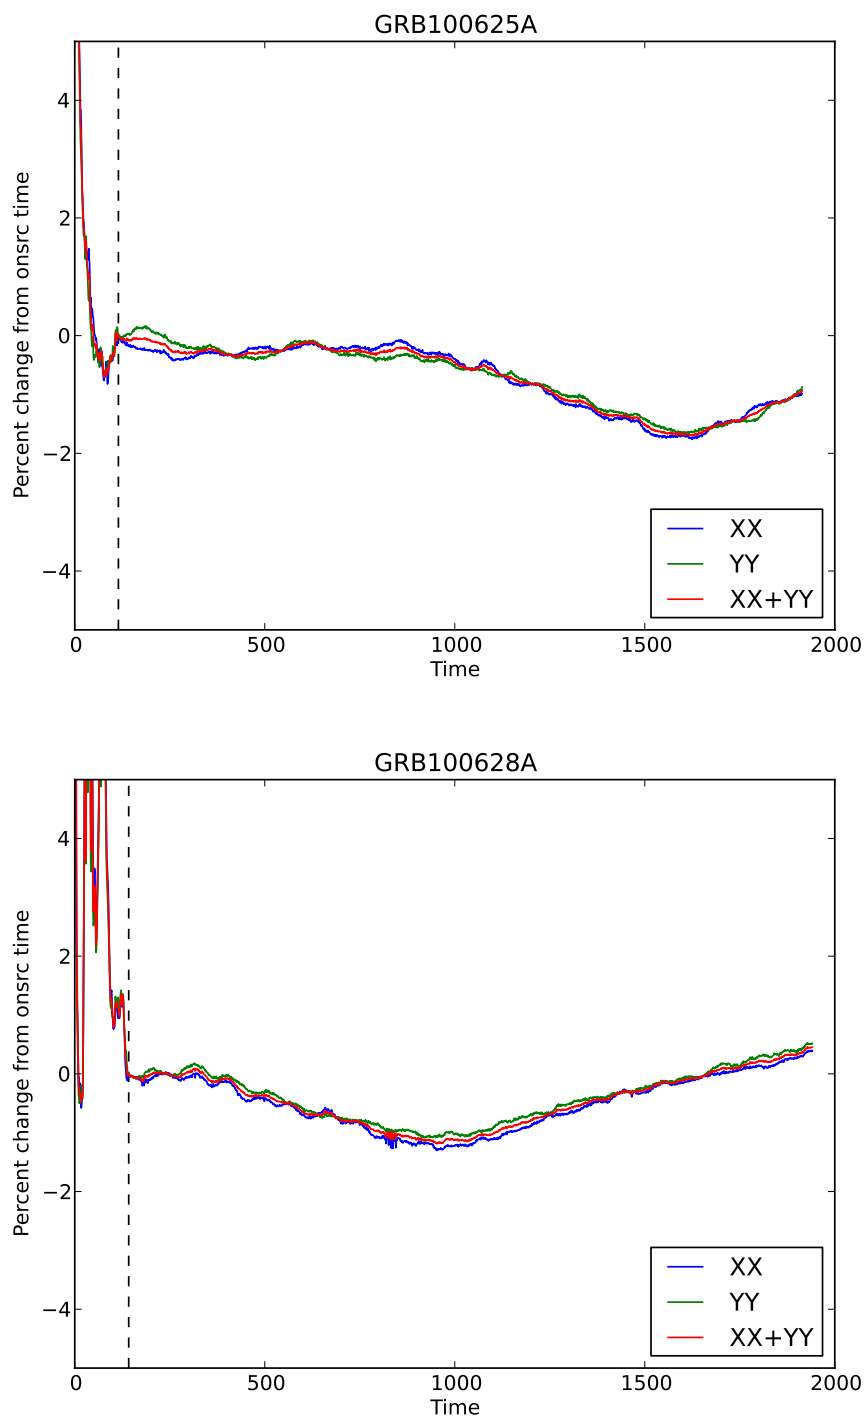


Figure 3.6: Low-time resolution 1.4 GHz light curves for all GRBs. The vertical dashed line indicates when the antenna was on source, with the interval preceding this line being measurements during the slew.

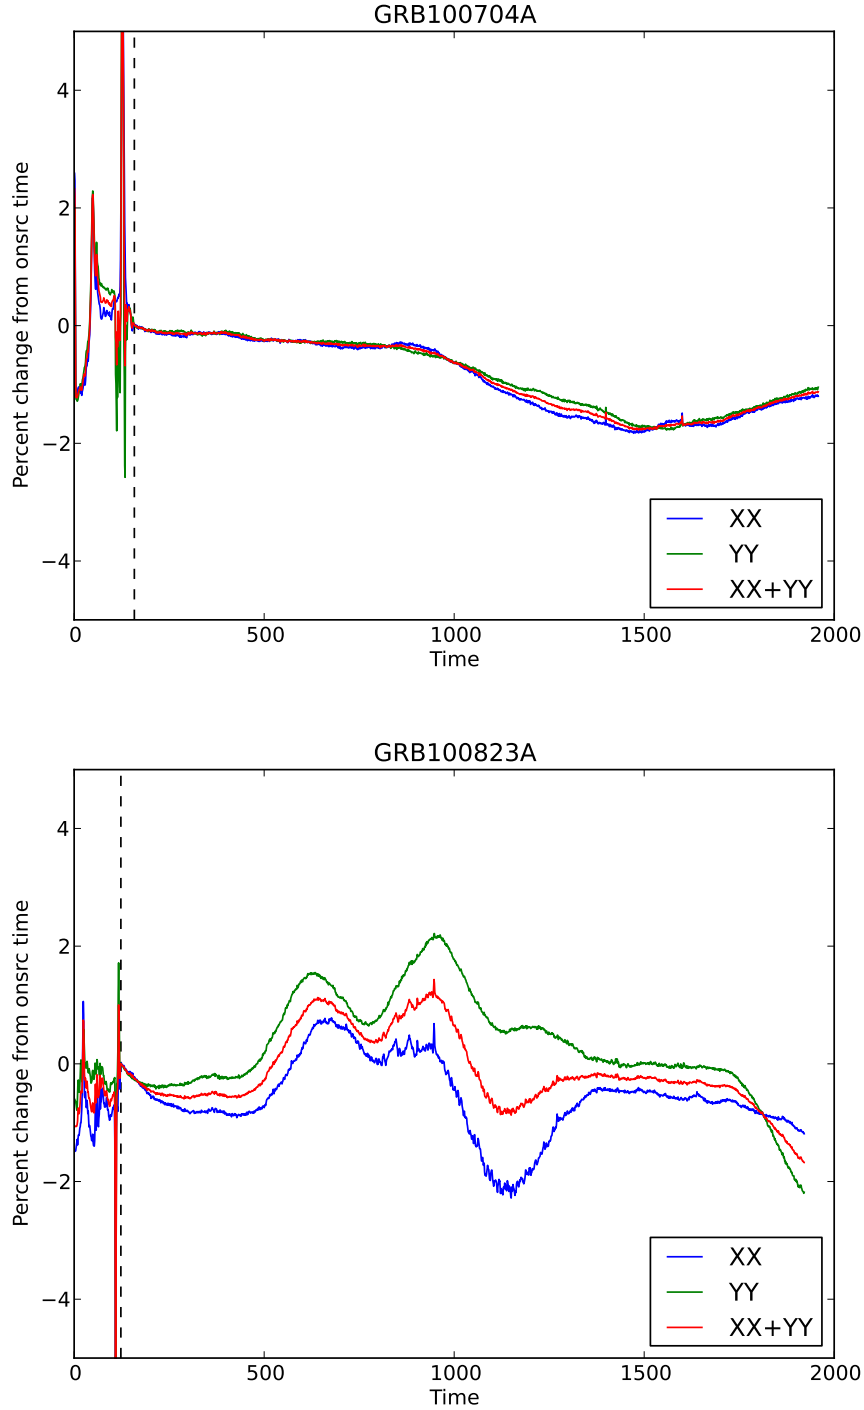


Fig. 3.6 continued.

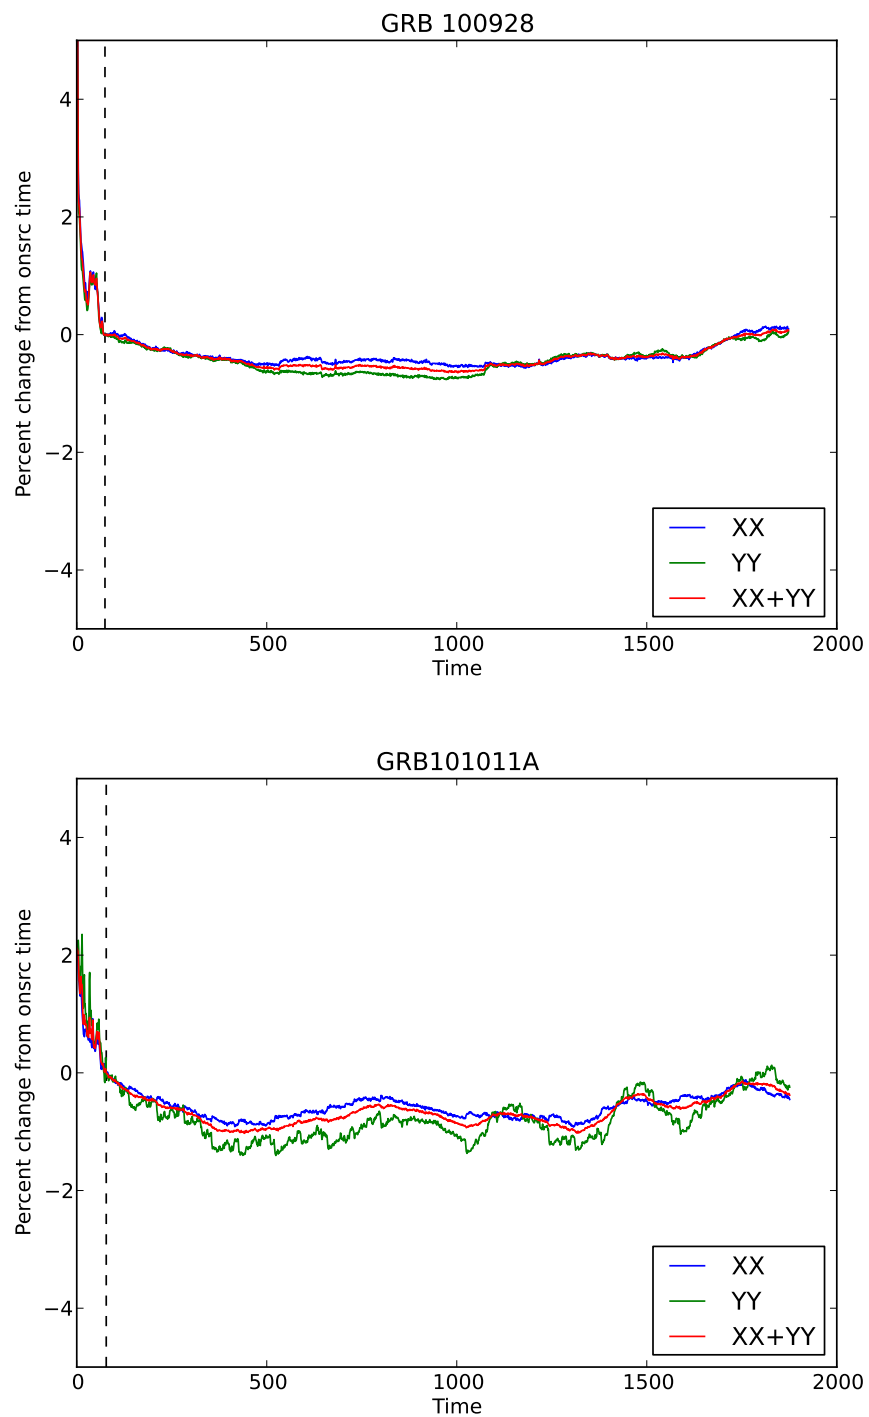


Fig. 3.6 continued.




Fig. 3.6 continued.




Fig. 3.6 continued.

Nonetheless, I can construct a simple model with which to evaluate my real-world results, assuming Gaussian noise and independent trials. Here I follow the logic of Burke-Spolaor et al. (2011b). For a 30 minute observation, I produce 1196 DM trials, each comprising $N_s = 28.125 \times 10^6$ samples. I apply a set of 9 boxcar trials to each DM trial, implying that a total of $N_p = 1196 \times \sum_{i=0}^9 N_s/2^i \simeq 6.7 \times 10^{10}$ points are searched for pulses. Assuming each trial is independent, then the expected number of false alarms exceeding a threshold T during an observation is $N_p \operatorname{erfc}(T/(\sigma\sqrt{2}))$, where erfc is the complementary error function.

Applying this model to observations of GRB 101011A, I expect $N_p \simeq 130$ trials above 6σ and get 144 such detections, roughly tallying with my discussion to this point, but many detections are clearly associated with RFI and are not related to the Gaussian distribution. At the 6.6σ , significance of my detected pulse, I expect on average 2.75 false alarms per observation. Applying the Poisson distribution, I can compute the probability of finding exactly one false alarm during a 30 minute observation, at or above 6.6σ , of 17.5%. Alternatively, I can compute the probability of finding one or more pulses above 6.6σ of 76.0%.

These relatively high probabilities suggest that it is quite likely that I would find a 6.6σ candidate in a single observation, and that I would also expect between one and two candidates across my 9 observations, exactly




Figure 3.7: Plots of time-averaged long-time integrations for each GRB observation. For each GRB top panel: bandpass, 2nd panel: standard deviation, 3rd panel: kurtosis, 4th panel: channel weights. Solid lines are the mean values over 30 min, and dotted lines are the maximum values over 30 min.

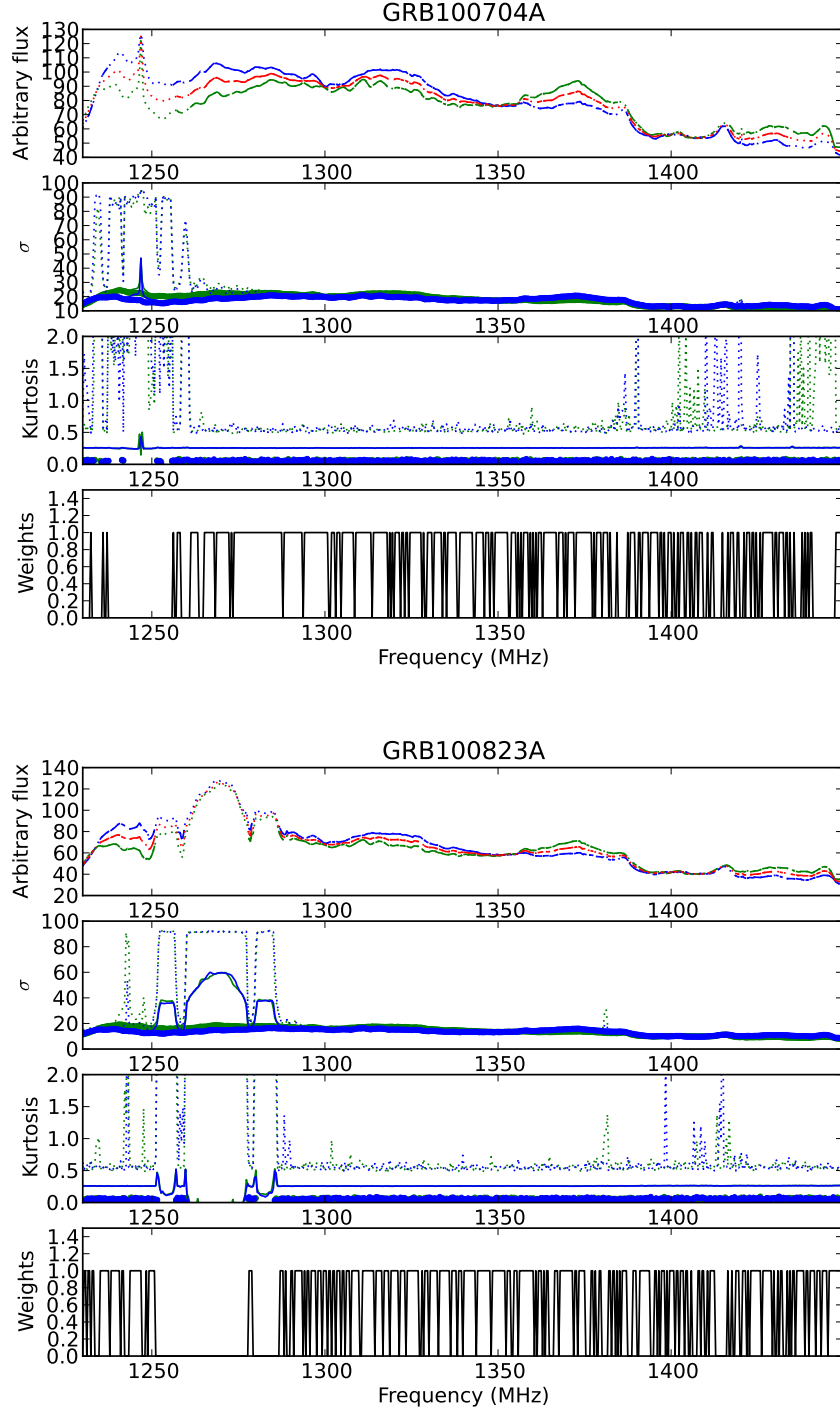


Fig. 3.7 continued.




Fig. 3.7 continued.




Fig. 3.7 continued.




Fig. 3.7 continued.

matching my results. However, I have neglected up to this point the ‘friend-of-friends’ algorithm employed by the single-pulse pipeline, which declares a candidate as viable only if a pulse is detected in 3 or more nearby DM and boxcar trials. According to Burke-Spolaor et al. (2011b), this algorithm reduces the number of candidates by a factor of $\simeq 100 - 1000$, implying that applying the algorithm reduces the probability of detecting exactly one candidate from 17.5% to less than 0.175%. Such an analysis suggests that the 6.6σ pulse I have detected is not a random noise fluctuation, with a confidence of greater than 99.8%.

Relaxing the assumption of independent trials implies that adjacent trials (either in DM or boxcar width) have overlapping time and frequency samples. Thus, if a given trial has a high value, an adjacent trial is also likely to have a high value, because many of the samples overlap. Therefore, relaxing the assumption of independence increases the false alarm probability, and I would expect to see clusters of false alarms in DM and boxcar space. The amount of overlapping samples between adjacent trials is of the order of 50% both in time and frequency. For a trial with a 6σ detection, which is already unlikely, the probability of the non-overlapping samples from an adjacent trial pushing the adjacent trial above 6σ is small. As the friend-of-friends algorithm requires three or more adjacent trials above 6σ to declare a detection, it effectively nullifies the correlation between adjacent trials, and the probabilities described




Figure 3.8: The raw receiver bandpass during observations of GRB 100823A, as a function of time and frequency. The gray scale is an arbitrary linear scale of the intensity. Long time integration measurements were badly affected by RFI from the *Beidou G1* satellite, despite zero weights being applied to the obvious interference. The interference was centred on 1270 MHz, the transmitting frequency of the satellite, and the interference was worst when the pointing direction passed closest to the satellite.




Figure 3.9: The baseline-subtracted receiver bandpass during observations of GRB 110412A, as a function of time and frequency. Long time integration measurements were affected by narrow band RFI at around 1230 MHz. The narrow-band interference was worst at around $t = 1300$ s, which corresponds to a bump in the broadband lightcurve.

above are still accurate.

For simplicity, I have neglected the affect of the additional weak RFI filter, which discards candidates with low DM pulses within 3 s (Section 3.2.3.3). Such a filter will reduce the number of candidates produced by random fluctuations even more, which in turn increases the confidence.

3.4.1.2 *Random noise fluctuation - A null trial*

The theoretical considerations in Section 3.4.1.1 assume Gaussian noise, and do not deal with the complexities of the RFI and detection algorithms very adequately. Perhaps a more reliable approach to determining the probability of such a pulse occurring by chance is to measure the number of false positives recorded when there is no GRB in the beam. One such method would be to observe a ‘blank’ patch of sky, not containing a GRB, for a period of time, and measure the number of candidates. Unfortunately, such observations were not made when the antenna and feed were available.

As an alternative, one can search for pulses in the data I already have, but randomising the channel ordering before dedispersion. Randomising the channels destroys the $t \propto \nu^{-2}$ dispersion relationship expected from an astronomical pulse, and almost guarantees that no genuinely astronomically dispersed pulses will be visible in the results, especially at large DMs. This technique, while inferior to taking additional data on a blank patch of sky, still has advantages over the theoretical approach. In particular, it uses the same detection algorithms as those used on the on-sky data, the noise statistics of the data are preserved, and the zero DM sequence (which is used for RFI flagging by the pipeline) is also preserved.

I conducted such a trial based on the data from GRB 101011A (which contained the brightest single pulse). A software-based pseudo-random number generator was seeded with random data from the Linux kernel entropy pool at the beginning of each observation, and the ordering of the channels was changed randomly every 25 s during the observation.

I simulated 211 30 minute observations with the randomised channel ordering, which yielded two pulse candidates. From this trial, I can conclude that the probability of detecting a $> 6\sigma$ candidate that passes all the additional criteria is approximately 1%, and I therefore rule out the possibility of my pulse being due to random fluctuations with a confidence of 99%.

For future studies, a blank sky observations should be performed after each GRB observation. The blank sky measurement should be as least as long as the GRB observation and immediately afterwards, which will enable an accurate characterisation of the background rate near the time of the GRB. Such a null trial will be extremely useful in being able to characterise determining the confidence of any GRB detections.

3.4.1.3 Impulsive RFI

Impulsive RFI has a number of attributes that make it often distinguishable from astronomical sources: it is usually most visible at low DMs, although it can extend to higher ones, depending on the exact waveform shape. Impulsive RFI can often be clustered in time, with a train of bursts appearing in a ~ 10 s interval, and it is often highly polarised.

The pulse seen just after GRB 101011A has a very high DM of 570 pc cm^{-3} , which argues against an RFI origin. The nearest low DM pulse appears to be some 120 s before the candidate pulse, which argues against the candidate being part of a train of RFI, or being a high-DM component of RFI which is also detected at low DMs (Figure 3.1). Finally, this pulse has approximately the same energy in both polarisations (Figure 3.3), arguing that it is either unpolarised, circularly polarised, or linearly polarised with an position angle approximately 45° to the feed angles. A detection in only one polarisation would imply a linearly polarised pulse and would almost certainly imply an RFI origin, so the polarisation properties of the pulse add credence to a non-RFI origin.

3.4.1.4 Lorimer burst events

When processing archival data from the Parkes 64 m telescope, Lorimer et al. (2007) detected a single 30 Jy burst with ~ 5 ms duration at a DM of 375 pc cm^{-3} . No other similar pulses have been reported, and some doubt has been cast on the astronomical interpretation (see Section 3.4.1.5). Lorimer et al. (2007) report an implied event rate of $3.8 \times 10^{-4} \text{ hr}^{-1} \text{ deg}^{-2}$ (based on a sample of 1) which is substantially larger than, but still potentially compatible with the GRB event rate, but note that no physical mechanism has been proposed for GRBs to emit such a pulse.

My pulse has similar properties to the Lorimer burst (LB). The DM of my pulse is higher than the LB, but this can be explained by differences in Galactic electron content, redshift and ionised material near the source. The Lorimer burst is shorter than my pulse (5 ms for the LB as opposed to 25 ms for my pulse). This shorter duration cannot be explained by reduced scatter broadening from a lower electron column, implying that the difference in pulse duration is intrinsic. The flux densities are roughly comparable (30 Jy for the LB, 7 Jy for ours). If the LB is associated with a GRB then it is likely that I have observed the same physical mechanism.

3.4.1.5 Atmospheric events “percytons”

In an archival search of Parkes Multibeam archive, Burke-Spolaor et al. (2011a) discovered 16 single pulses with similar spectral characteristics to those expected from a dispersed extragalactic pulse. These pulses, which Burke-Spolaor et al. (2011a) call ‘percytons’, have a number of unique properties.

In general, the perytons follow a $t \propto \nu^{-2}$ frequency sweep, although some pulses have significant departures from this sweep. The sweep times across the 288 MHz band centred at 1374 MHz are between 200 and 400 ms, but mostly clustered in the range 320–400 ms. The perytons are extremely bright (between 0.1–272 kJy), and have substantial amplitude modulations. All perytons were detected during the day, in general mid-morning, and mostly in the mid-winter months of June and July. The peryton event rate is $2.3 \times 10^{-7} \text{ deg}^{-2} \text{ hr}^{-1}$. Burke-Spolaor et al. (2011a) suggest that all these characteristics point to a terrestrial origin, and have additionally proposed an atmospheric Langmuir wave phenomenon, related to lightning (Burke-Spolaor, private communication).

The pulses I detect have somewhat different properties. Unfortunately, I do not have sufficient S/N to detect small departures from the cold plasma dispersion law, nor any substantial amplitude modulations. Figure 3.3 supports a standard cold plasma frequency sweep, and constant amplitude, although the significance is not very high. The flux densities of my pulses are approximately 7 Jy, within the range seen for the perytons. If I extrapolate the sweep time of my brightest burst for the configuration of Burke-Spolaor et al. (2011a) between 1230 and 1518 MHz, I obtain a sweep time of 536 ms. This sweep time is substantially higher than that of any of the known perytons, but within the general range, given the small number statistics. My brightest pulse was detected on 2011 Oct 11, well outside the usual Peryton season, and at 22:58 (local time) in the evening, outside the usual Peryton time. There had been no rainfall and no lightning detected within 50 km of the telescope within 24 hrs of the GRB time, disavouring the postulated lightning related cause. Finally, the event rates are vastly different, with my events being more common by over six orders of magnitude, although Burke-Spolaor et al. (2011a) note that the Peryton event rate is highly non-uniform, with 12 perytons detected within 0.5 hr in one instance.

From this analysis I conclude that the substantially higher event rate, longer sweep time, arrival date, time of day, and prevailing weather conditions of my pulse all make it unlikely to be a peryton.

3.4.1.6 *Gamma ray bursts*

Two out of nine (22%) of my observations contain a 1.4 GHz radio pulse, which begs the question, are the GRBs themselves the source of these pulses? Unfortunately, there is no known physical mechanism for producing such a burst within minutes of a GRB, in spite of the substantial theoretical work inspired by the detection of the LB.

Any proposed mechanism must explain a range of properties. First, my two detections were only present following long gamma ray bursts, so a mechanism must fit within the confines of the collapsar paradigm. The pulses are detected with flux densities of the order of 7 Jy, which requires a very pow-

erful mechanism if placed at the cosmological distances of GRBs. The short pulse width, of the order of 25 ms, requires a relatively compact causative region. My pulses are 524 and 1076 seconds after the gamma ray triggers, much later than the ~ 2 s dispersion delay, so any mechanism must explain this delay. Finally, any mechanism must also explain why both pulses are possibly coincident with a break in the X-ray light curve, and also must explain the complete lack of flaring at X-ray or gamma ray wavelengths simultaneous with the emission of the radio pulse.

Unfortunately, no host galaxies were observed for either GRB, so no redshift is available. For the remaining analysis, I assume a redshift of $z = 1$, which is approximately the mean GRB redshift (Savaglio et al., 2009).

I now try to interpret my pulses within the context of the collapsar model of MacFadyen & Woosley (1999), namely: the collapse of a massive ($> 10 M_{\odot}$) star producing a central engine (either black hole, or magnetar of mass $\sim 2 - 3 M_{\odot}$), which in turn accretes material from a disk which produces jets, the prompt gamma ray emission, and the shocks that pass through the surrounding material and into the interstellar medium.

Concentrating on the brightest of my radio pulses, the flux density of approximately 7 Jy implies a spectral luminosity of $2.8 \times 10^{25} \text{ erg s}^{-1} \text{ Hz}^{-1}$ at redshift one ($H_0 = 71 \text{ kms}^{-1} \text{ Mpc}^{-1}$, $\Omega_M = 0.27$, $\Omega_{\text{vac}} = 0.73$) assuming a spectral index of -0.7^{10} . This spectral luminosity is equivalent to a luminosity of $2.8 \times 10^{44} \text{ erg s}^{-1}$ over a bandwidth of 1 GHz. The intrinsic pulse width at redshift one is half the measured pulse width, equal to about 12 ms (neglecting intergalactic scattering), implying a total energy release in the radio burst of $3.3 \times 10^{42} \text{ erg}$ over a bandwidth of 1 GHz. The energy in the radio burst fits well within the canonical 10^{51} erg total energy release of the collapsar model, implying that the energy budget can support such a pulse. The pulse width of 12 ms implies a causative region of size $< 3.6 \times 10^8 \text{ cm}$, some 10-100 times larger than the innermost stable orbit of a black hole, depending on its mass and spin, and well within the magnetosphere of a magnetar. In either case, this suggests that the emission emerges from a region close to the central engine.

Perhaps the two most confounding problems are the time delay between the gamma ray trigger and the emission of the radio pulse, and that only one pulse is observed per GRB. If this is not an observational effect due to my detection threshold, the single detection implies that the radio pulse is related to a singular event that occurs 500-1000 s after the gamma rays are released. If the central engine is a promptly formed black hole, it forms just before the release of gamma rays. Therefore, there is no obvious mechanism related to the black hole, or GRB phenomenon, that can be invoked to explain the delay.

¹⁰In fact cannot accurately measure the spectral index of the pulse as the S/N is low. We assume -0.7 as it falls between the mean spectral indices for pulsars of -1.6 (Lorimer et al., 1995) and that of roughly 0 for magnetars (Camilo et al., 2008), and also corresponds to the spectral index of incoherent synchrotron emission.

On the other hand, for a magnetar central engine, Baumgarte et al. (2000) suggest that a super-Chandrasekhar ($> 1.4 M_{\odot}$) magnetar can be rotationally supported, and can avoid turning into a black hole as long as its angular momentum exceeds a threshold. If sufficient angular momentum is lost (e.g. due to winds), a black hole is formed after some delay. Metzger et al. (2011) suggest that the timescale for delayed black hole formation can be a few hundred seconds, but add that there is considerable uncertainty in this timescale. The exact mechanism for how this formation could release the radio pulse with the observed properties is still uncertain, but the delayed formation model has the desired characteristics: the delayed black hole formation happens only once, and occurs some time after the GRB trigger.

The alignment of the radio pulse with breaks in the X-ray light curve could also be related to the delayed formation of the black hole. The X-ray light curve is typically explained by a number of segments (Gehrels et al., 2009, and references therein): the initial rapidly falling afterglow is followed by a plateau lasting a few thousand seconds, which is in turn followed by a classical afterglow. The so-called plateau is usually explained as being due to additional energy being injected into the external shock, for example, from magnetar winds. The radio pulse after GRB 100704A occurs at the end of the X-ray plateau, which implies it is related to the end of the energy injection driving the plateau. One could imagine that the formation of the black hole event horizon encompassing the magnetar and its winds could be responsible both for the radio pulse and for the end of the X-ray plateau. On the other hand, the radio pulse following GRB 101011A occurs at the beginning of the X-ray plateau (which is not as clear as the plateau in GRB 100704A), casting some doubt on the delayed black hole formation as the source of this pulse.

The coincidence with the X-ray breaks could also be a statistical fluke. For GRB 101011A, the formal error in the break time is roughly 200 s, which is 10% of the observation duration. There are at least two breaks in the X-ray light curve predicted during the observation interval of 200-1800 s (Gehrels et al., 2009), suggesting a 20% chance of a random event falling within the statistical errors of a break in the X-ray light curve. The constraints on the break time in GRB 100710A are less strong, implying an even higher chance of a random event falling within the errors of an estimated X-ray break.

3.4.1.7 Upper limits

There remains a high probability that neither of my pulses are associated with an astronomical phenomenon. In spite of my best efforts to prove otherwise, there could still be some factor I have not considered in my telescope configuration that could cause short, dispersed pulse in my data. The fact that the pulses are detected at relatively low significance (the brightest is at 6.6σ), and that no theory has been proposed to explain the brightness or time delay of

the pulse, implies that further observational and theoretical work is desirable to confirm such pulses as having a GRB origin.

I can say with certainty that I did not detect any pulses at a significance $> 7\sigma$. Given my system equivalent flux density and processing method, I obtain an upper limit on the detection of single dispersed radio pulses at 1.4 GHz between 200 and 1800 s after a GRB of $1.27w^{-1/2}$ Jy, where $6.4 \times 10^{-5} \text{ s} < w < 32 \times 10^{-3} \text{ s}$ is the pulse width.

3.4.2 Pulsar detection

I have detected one repeating candidate (Section 3.3.3). In this section I discuss possible origins of this candidate.

3.4.2.1 Random noise or RFI

While this candidate appears to be astronomical in terms of its broadband nature, persistent emission, relatively narrow pulse profile and low RFI background, it is still only detected with low significance (6.1σ), implying that it could be RFI or a random noise fluctuation. A detailed theoretical study of the probability of detecting such a candidate by chance due to random noise fluctuations is extremely complex, and outside the scope of this thesis. Evidence from the ongoing HTRU pulsar survey with the Parkes 64 m telescope (Keith et al., 2010), with which I share the same site, backend and processing software, suggests that it is ‘quite common’ for 6σ candidates such as ours, to never again be detected in follow-up (Michael Keith, private communication). This implies that subtle RFI and systematic effects can increase the false detection rate at low levels of significance.

3.4.2.2 A new Galactic pulsar

The DM, at 50.7 pc cm^{-3} is in agreement with the predicted maximum Galactic value from NE2001 of 50 pc cm^{-3} , with a corresponding DM distance of 4.6 kpc. The NE2001 model values have considerable uncertainty, implying that it is entirely possible for this DM to be explained by the Galactic disk, albeit with a considerably larger distance than 4.6 kpc (Gaensler et al., 2008). There is no known pulsar within 5 degrees in the ATNF pulsar catalogue (Manchester et al., 2005), which is expected given the high latitude ($b = +33^\circ$), and that most previous surveys have concentrated mainly at low latitudes.

If this candidate is not RFI or random noise, it is in all likelihood, the chance detection of a new high-latitude pulsar.

3.4.2.3 *A gamma ray burst*

Given the vast majority, if not all, of the DM of this candidate can be explained by the Galactic distribution of free electrons, the possibility that it is associated with an extragalactic GRB is remote. The Galactic contribution to the DM would have to be substantially underestimated, and the source extremely nearby, to maintain such a low DM. I find no optical counterpart in SuperCOSMOS (Hambly et al., 2001) down to $B < 22$, implying that this source is unlikely to be associated with a nearby galaxy. In light of the low DM and lack of optical counterpart, I consider the GRB origin of this source a remote possibility.

3.4.3 Long time integrations

Ultimately, my experiment was not designed to detect flux changes over long time intervals. My feed was not equipped with a Dicke radiometer, which would have enabled us to calibrate the gain fluctuations (Section 3.3.4) over the 1 s integration time, but would also have contaminated the high-time resolution studies. In addition, I did not regularly observe flux calibrators during my observations, as this would have ruined the periodicity searches (which require continuous data streams). My detection limits are therefore completely limited by the temperature-related gain variations in my system, of approximately $5\% \text{ hr}^{-1}$. I would consider change of $20\% \text{ hr}^{-1}$ as a candidate, which implies that a change in source flux density of 760 Jy would have constituted a candidate.

I did not detect any variation that was not immediately attributable to interference, either from satellites or the sun. I set an upper limit of a change of 760 Jy on any long-duration emission ($> 1 \text{ s}$) from a GRBs within 200-1800 s, at 1.4 GHz.

3.5 CONCLUSION

I have searched for prompt radio emission from gamma ray bursts at 1.4 GHz, using a robotic telescope and a pulsar backend. My telescope was typically on source within 200 s of the gamma ray trigger.

I detected single dispersed pulses following two GRBs at significances $> 6\sigma$. Simple statistical arguments, and a null trial based on randomising channels on existing data, rule out random fluctuations as the origin of these pulses at $> 99.8\%$ and $\sim 99\%$, respectively. Weak impulsive RFI and atmospheric origins of these pulses remain a possibility. If the radio pulses are associated with the corresponding GRBs, they could be related to changes in the central engine, in particular the delayed formation of a black hole due to spin-down of a rotationally-supported magnetar. If the single pulse is not related to the GRB, I set an upper limit for radio pulses from GRBs between 200 to 1800 s

of $1.27w^{-1/2}$ Jy, where $6.4 \times 10^{-5} \text{ s} < w < 32 \times 10^{-3} \text{ s}$ is the pulse width. This limit is substantially better than the limit of Katz et al. (2003) at 611 MHz.

I have detected one candidate repeating, dispersed source and conclude that, if it is not related to RFI or random fluctuations, it is most likely a new high-latitude Galactic pulsar. I have detected no candidates on timescales $> 1 \text{ s}$, but my experiment was not primarily designed for such detections. Nonetheless, I set an upper limit of a change of 760 Jy on any long-duration emission ($> 1 \text{ s}$) between 200 to 1800 s from my GRB triggers, in agreement with Dessenne et al. (1996) at 151 MHz.

The detection of single dispersed pulses in this experiment is intriguing. Clearly the next step is to determine whether these pulses are related to their GRBs, for which the key problem is ruling out RFI, statistical fluctuations and other equipment-related sources as the origin of these pulses. The simplest future experiment to rule out these origins is to use a coincidence detection, by employing the same telescope setup at two widely separated sites. The fact that some 20% of GRBs seem to be accompanied by a radio pulse, even with my relatively poor sensitivity, suggests that sensitivity is not the key factor in this experiment. Therefore, similar dishes, feeds and backends can be used. More important parameters of this experiment are the short on-source time (preferably $< 200 \text{ s}$), and a wide separation between antennas. A simultaneous detection of a single pulse at two widely separated sites, even at 6σ , would almost certainly rule out RFI and statistical fluctuations, and render atmospheric effects a very remote possibility. Once such a detection is made, the future is wide open to probe the astrophysical and cosmological implications of these phenomena.

Two efficient new techniques for detecting dispersed radio pulses with interferometers: the Chirpolator and the Chimageator

This Chapter along with appendices C and D have been published as Bannister & Cornwell (2011). It is reproduced here with minor formatting changes.

4.1 INTRODUCTION

4.1.1 Scientific Motivation

THE study of the high time resolution radio sky has illuminated the physics of our Galaxy, enabled exquisite measurements of physics at the extremes of gravity, density and magnetic field, and uncovered a plethora of exotic objects. The main class of object enabling these measurements are radio pulsars: rapidly rotating magnetic neutron stars that emit periodic, short pulses at radio frequencies. While pulsars are interesting astrophysical laboratories in their own right, they can also be used to test predictions of general relativity through observations of binary pulsar systems (e.g., Lyne et al. 2004), search for gravitational waves with ensembles of widely-spaced pulsars (Yardley et al., 2010, and references therein), and to test theories of matter at the most extreme densities (e.g., Demorest et al. 2010). The short pulses emitted by pulsars undergo propagation effects during their passage through the Galactic ISM, which enables measurements of the Galactic magnetic field structure (e.g., Van Eck et al. 2011) and free electron density (e.g. Cordes & Lazio 2002). There are also many more pulsars to be found: barely 2000 of the estimated 30000 potentially detectable pulsars (Lorimer et al., 2006) have been discovered.

Radio pulsars are chiefly identified by searching for periodic, dispersed radio emission. By contrast, searches for single pulses of radio emission have uncovered other types of objects, of which the most notable are the so-called rotating radio transients (RRATs; McLaughlin et al. 2006). RRATs, like pulsars, are rotating neutron stars but emit only sporadically and are of interest because they may hold the key to the so-called missing supernova problem (Keane & Kramer, 2008). Searches for single pulses have also yielded a number of intriguing short-duration radio transients, which do not fit the classical models of pulsars or RRATs (e.g., Lorimer et al., 2007, Burke-Spolaor et al., 2011a, Keane et al., 2011).

In spite of these discoveries, there is still much to do, as the parameter space of radio transients is relatively poorly explored (Cordes et al., 2004). Exploration of this parameter space opens the potential for discovering new objects and new physics. These motivations are behind at least eight ongoing pulsar and single-pulse surveys (McLaughlin, 2011), and more surveys are being planned (e.g., Macquart et al. 2010).

4.1.2 Improving pulsar and single-pulse surveys

When surveying for pulsars and single-pulse sources, a desirable figure of merit is the product of instantaneous sensitivity and field of view, known as “survey speed.” Improving survey speed has three important consequences: (1) it reduces the integration time required to reach a flux density limit for periodic

sources; (2) it reduces the computational requirements to search for pulsars in tight binary systems;¹¹ (3) it enables a deeper search of the parameter space for single-pulse transients, in terms of rarity or faintness.

In recent years, the optimal approach for maximising survey speed has been to use large steerable and fixed single-dish telescopes feeding multi-beam receivers and wide-band electronics. This approach, however, is probably nearing its limit. Steerable single-dish engineering has reached its practical limit at diameters of about 100 m and there are limited sites available for large fixed reflectors, which also suffer from a limited sky coverage. Similarly, modern digital electronics have improved to the state where 2 GHz bandwidth is readily achievable, but the physics of pulsar emission and of the ISM limit the majority of the pulsar radio emission to the range 0.1–10 GHz, so an increase in processing bandwidth is unable to produce large improvements in sensitivity. Finally, the field of view of multi-beam receivers cannot grow indefinitely, as a large multi-beam receiver simply blocks too much of the dish aperture to be efficient.

The likely way forward for improvements in survey speed, therefore, is to use arrays of antennas. Using a large number of small antennas achieves simultaneously large field of view and good sensitivity, and therefore give a good survey speed. Using an array does have a dramatic cost, however: for configurations of interest to future surveys, the computational requirements increase to the point where the data processing becomes almost infeasible (Smits et al., 2009).

4.1.3 The problem: processing requirements

The desire to use arrays of antennas to simultaneously obtain large field of view and high sensitivity presents major data processing challenges, both in terms of the required data and the implied operation rates. These challenges have inspired a number of novel approaches. For example, Daishido et al. (2000) proposed the so-called fast Fourier transform telescope (FFTT) for a pulsar survey, which used FFT beamforming, first proposed by Williams (1968), and a square array geometry. This approach, also known as direct imaging, has attractive properties in terms of operation rates and has been extended, with particular emphasis on 21 cm tomography, to arrays of regular, arbitrary hierarchies of grids by Tegmark & Zaldarriaga (2010) and to arbitrary array geometries by Morales (2008). In a novel experiment, Janssen et al. (2009) enhanced the field of view over standard techniques, by employing a uniform linear geometry and phased array beamforming, which introduced deliberate ambiguities into the synthesised beam (i.e., grating lobes). Recently, Trott et al. (2011) proposed a method for searching for transient sources directly in visibility space, rather than image space. The visibility space approach

¹¹Pulsars in very small orbits provide the most powerful tests of general relativity.

is promising for arrays with sparse, arbitrary geometries, as the number of visibilities is much smaller than the number of pixels. But, it is not yet clear whether this approach will achieve substantial computational savings.

Some of the above approaches rely on having control over the array geometry. In many cases, controlling the geometry is not possible because it is driven by other requirements e.g., the $u - v$ coverage. The SKA and its pathfinders fall into this category. In such cases, one can reduce the processing requirements by falling back to the reduced sensitivity of incoherent processing, analysing a smaller field of view than available from the primary beam (D’Addario, 2010), or pointing all antennas at different parts of the sky in a “Fly’s Eye” mode (Macquart, 2011). Nonetheless, a fully coherent, wide field and computationally tractable system is a desirable goal.

We also note that the processing requirements are not limited to the number of arithmetic operations. In fact, in modern supercomputing problems, the processing bottleneck is not the number of arithmetic operations, but rather the data bandwidth into the processor (Leback et al., 2008). The bandwidth bottleneck has been identified as a key problem for correlators for large interferometers, with a number of proposed solutions (e.g., Lutomirski et al. 2011; Carlson 2010). To our knowledge, the data requirements of classical fast transient detection techniques have not been discussed explicitly in the literature, so we consider them in this Chapter.

4.1.4 Two new techniques

It is in this context that we propose two new techniques, which we have named after the frequency swept signals on which they operate (“chirps”). The first technique, which we call “The Chirpolator”,¹² operates by correlating the chirps received by pairs of antennas. The second technique, which we call “The Chimageator”,¹³ operates by gridding the cross-multiplied voltages from all telescopes to form an image at every sampling time. Both techniques are applicable to arbitrary array configurations, exploit the full sensitivity of the telescope, and have substantially lower data rate requirements than classical coherent techniques. Thus, these new techniques may be favoured over classical techniques in many regimes of computer economics and array geometry.

This Chapter is organised as follows: in Section 4.2, we provide background on the problem of searching for dispersion emission in interferometric data and on the classical solutions to this problem. In Section 4.3, we describe The Chirpolator and in Section 4.4 we describe The Chimageator. We describe a simple model to compare our techniques with classical results in Section 4.5 and present the results of this comparison in Section 4.6. We discuss the

¹²A portmanteau of “chirp” and “correlator”.

¹³A portmanteau of “chirp,” “image,” and “correlator”. The “imaging” part is inspired by the direct imaging approach of Daishido et al. (2000).

implications of these two algorithms on future telescope design and science outcomes in Section 4.7 and draw our conclusions in Section 4.8. In Appendix C and Appendix D we present detailed analysis of the algorithms, a number of novel optimisations, and a discussion of implementation considerations.

4.2 BACKGROUND

4.2.1 Dispersion in the Interstellar Medium

Before being received by telescopes on Earth, electromagnetic waves from an astronomical source must pass through the ISM, a plasma containing non-relativistic unbound electrons. As they travel through the ISM, the waves undergo dispersion, or frequency-dependent delay between two frequencies ν_1 and ν_2 according to the following formula:

$$t = \frac{e^2}{2\pi m_e c} \text{DM} (\nu_2^{-2} - \nu_1^{-2}), \quad (4.1)$$

where t is the time from the beginning of the pulse, $\nu_1 > \nu_2$, and the physical constant is:

$$\mu = \frac{e^2}{2\pi m_e c} \simeq 4.15 \text{ms}, \quad (4.2)$$

for frequencies in GHz. The DM describes the number of electrons between the observer and the emitting source, defined as:

$$\text{DM} = \int_0^d n_e dl, \quad (4.3)$$

where n_e is the electron density and d is the distance to the source. DM is usually quoted in units of cm^{-3}pc .

If a narrow pulse is emitted by a source, the frequency of the signal received at Earth be in the form of the dispersion. For our analysis, we will consider the form of this dispersed pulse as a complex voltage time series. To form this time series, we will compute the instantaneous phase of the signal, by integrating the instantaneous frequency of the signal. The instantaneous frequency can be found by rearranging Equation 4.1. We approximate this equation with a Taylor series around $t = T/2$, which yields

$$\nu_2(t) = \left(\nu_1^{-2} + \frac{t}{\mu \text{DM}} \right)^{-1/2} \quad (4.4)$$

$$= \sum_{i=0}^{\infty} a_i (t - T/2)^i, \quad (4.5)$$

where $\nu_2(t)$ is the instantaneous frequency of the signal, t is the time from the beginning of the pulse, and T is the time taken for the pulse to traverse the bandwidth of interest $B = \nu_1 - \nu_2$. The first three Taylor coefficients are

$$a_0 = \alpha^{-1/2} \quad (4.6)$$

$$a_1 = \frac{-1}{2} \frac{1}{\mu\text{DM}} \alpha^{-3/2} \quad (4.7)$$

$$a_2 = \frac{3}{8} \frac{1}{(\mu\text{DM})^2} \alpha^{-5/2}, \quad (4.8)$$

where

$$\alpha = \nu_1^{-2} + \frac{T}{2\mu\text{DM}} \quad (4.9)$$

$$= \frac{1}{2} ((\nu_1 - B)^{-2} + \nu_1^{-2}). \quad (4.10)$$

Plots of the true dispersion law and its linear and second-order approximations are shown in Figure 4.1.

To obtain the formula for the voltage time series of the signal received on Earth, we first write the phase of the signal, as given by

$$\phi(t) = 2\pi \int_0^t \nu_2(t') dt' \quad (4.11)$$

$$= 4\pi\mu\text{DM} \left(\nu_1^{-2} + \frac{t}{\mu\text{DM}} \right)^{1/2} \quad (4.12)$$

$$= 2\pi \sum_{i=0}^{\infty} \frac{a_i}{i+1} (t - T/2)^{i+1} \quad (4.13)$$

$$\simeq 2\pi \left[\frac{T}{2} (-a_0 + a_1/2) + t(a_0 - a_1 T/4) + t^2 a_1/2 \right], \quad (4.14)$$

where we have expanded out the phase to the $i = 1$ Taylor term. Finally, we can write the complex voltage as

$$s(t) = \exp(j\phi(t)). \quad (4.15)$$

In a typical telescope system, the absolute phase of the voltage (i.e., the constant term in Equation 4.14) is not important, and the fixed frequency (t term) is removed by down-conversion, so that only the t^2 and higher terms are relevant. Therefore, in the main text we approximate the dispersion with a signal of the form




Figure 4.1: Delay vs. frequency for a pulse with DM of $100 \text{ cm}^{-3} \text{ pc}$, over a frequency range of 400 MHz centred at 1.4 GHz. The true dispersion (Equation 4.1) and the first- (linear chirp), second- and third-order Taylor series approximations around the delay midpoint are shown. The third-order approximation is barely visible as it very closely matches the true dispersion.

$$s(t) = \exp(\pi j \dot{f} t^2), \quad (4.16)$$

where $\dot{f} = a_1 \simeq -B/T$ is the gradient of the linear frequency trajectory as shown in Figure 4.1. Signals of this form are known as complex linear chirps. In Appendix C and Appendix D we consider the higher order terms.

4.2.2 A taxonomy of methods

An astronomer wishing to search for, or study short-duration radio pulses may employ any one of a wide range of dedispersion and array processing techniques, as shown in Figure 4.2. General properties of these methods are shown in Table 4.1. We give more detailed descriptions of the classical methods in the following sections.

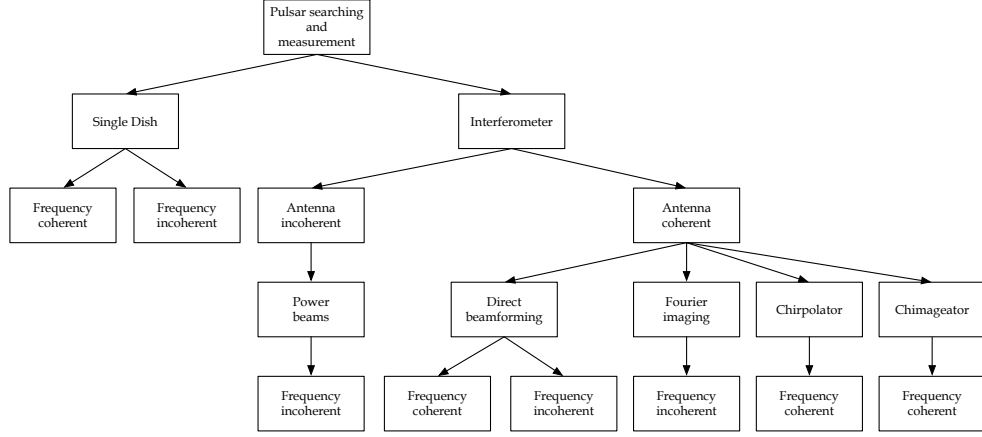


Figure 4.2: Taxonomy of approaches to pulsar searching and measurement.

Table 4.1: Methods of searching for pulsars and scaling relations for various figures of merit. The columns are the name of the method, the resolution in radians, the number of pixels in an image, and the scaling relation for the sensitivity with number of antennas M . Parabolic dishes and equal maximum baselines in the u - and v -directions have been assumed.

Method	$\Delta\theta$ (rad)	N_{pix}	Sensitivity Scaling
Power beams	$1.17 \frac{\lambda}{D}$ (primary beam)	1	$\propto M^{1/2}$
Fourier imaging / Direct beamforming	$\frac{\lambda D}{b_{\text{max}}}$	$\left(\frac{b_{\text{max}}}{D}\right)^2$	$\propto M$
Chirpolator / Chimageator	$0.844 \frac{c}{B b_{\text{max}}}$	$1.92 \left(\frac{B b_{\text{max}} \lambda}{c D}\right)^2$	$\propto M$

4.2.3 Dedispersion

The effect of dispersion on a short-duration pulse is to smear it out in time. In order to determine the emitted pulse shape or to detect the pulse with maximum S/N, the effect of the dispersion must be undone, in a process known as dedispersion. Two methods of dedispersion can be employed, as described below.

4.2.3.1 Incoherent dedispersion

Incoherent dedispersion involves two steps. First, the raw telescope voltages are passed through an analysis filterbank, such as an analogue filterbank, FFT, or polyphase filterbank to form a set of channelised outputs. Each filter output is then squared, and integrated over an interval (typically 0.1–10 ms) to form a spectrogram, or time-frequency plane. In the second step, a range of trial DMs are searched by summing across frequency channels, after delaying each frequency channel according to the trial DM of interest. We call this method “frequency incoherent”, because only the filterbank amplitudes are summed and the phase information is discarded. Incoherent dedispersion is typically used in pulsar surveys because the filtering only needs to be computed once and a range of DM trials can be performed on the same filtered output relatively cheaply.

While early workers used analogue filterbanks, more recent projects make use of digitally sampled baseband voltage signals, and perform the digital filtering and dedispersion in hardware and software. Taylor (1974) proposed a computationally efficient method of performing incoherent dedispersion known as the “tree” method, which requires fewer additions than a naive implementation, but assumes linear dispersion. The linear assumption can be relaxed by adding padding channels, which marginally increases the computational cost.

4.2.3.2 Coherent dedispersion

Coherent dedispersion operates on raw telescope voltages, and involves convolving the telescope voltages with the impulse response corresponding to the inverse of the ISM (i.e., the inverse of Equation 4.1), thereby forming the maximum S/N filter or “matched” filter. We call this method “frequency coherent” as the data are processed without discarding the phase. Coherent dedispersion is used during pulsar monitoring, when the DM is approximately known, as the inverse filtering preserves the emitted pulse shape more faithfully than incoherent dedispersion. Coherent dedispersion is not used for pulsar surveys, as the computational cost of performing multiple DM trials is prohibitive. Coherent dedispersion is most often performed on digitally sampled complex baseband signals.

4.2.4 Array processing

When using an array, the question arises of how to best combine the signals from two or more antennas. In this section we describe three common approaches.

4.2.4.1 Power beams

Power beams are formed by envelope detecting the output of each antenna, and summing the resulting powers across antennas. We call this method ‘antenna incoherent’ as the envelope detection removes the phase information before the sum across antennas. The power beam is sensitive to the entire sky, as long as the integration time of the envelope detector is longer than the largest geometric delay, and is usually limited by the primary beam of the telescope antennas. The penalty for power beams is that the sensitivity is poor, as it scales as $M^{1/2}$, where M is the number of antennas. The output of the power beams can only be incoherently dedispersed, as the phase information is discarded by the envelope detector at the antenna.

4.2.4.2 Direct beamforming (Tied Array Beams)

Direct beamforming involves delaying the voltage signal from each telescope to compensate for the array geometry and summing the resulting voltages. This technique is also known as “tied array beam forming.” We refer to this method as “antenna coherent” as the phase information is preserved. The resulting beam has the size of a synthesised beam, which is much smaller than the telescope primary beam. Unlike power beams, the full array sensitivity is preserved as it scales with M . As a tied array beam provides a voltage stream, either coherent or incoherent dedispersion can be used. Multiple tied array beams can be deployed to increase the field of view.

4.2.4.3 Fourier imaging

Fourier imaging involves cross-correlating the telescope voltages with one another to form a set of complex “visibilities”, which are Fourier transformed to form an image. Cross-correlation can be performed either by an initial filtering step followed by cross-multiplication (so-called FX correlation), or cross-correlation followed by a Fourier transform (XF correlation). Each pixel of the image must be separately incoherently dedispersed, as the pixels are spatially independent. Coherent dedispersion cannot be used because each pixel contains only amplitude information.

Fourier imaging achieves the full array sensitivity over the full primary beam of the individual antennas. As only incoherent dedispersion can be used, Fourier imaging is most suited to surveys. Fourier imaging requires a

so-called “corner-turn”, or matrix transpose between the imaging and dedispersion stages, which can result in very high data rates between the two steps. Fourier imaging has recently been used by Law et al. (2011) and Wayth et al. (2011) to detect giant pulses from the Crab pulsar.

4.3 THE CHIRPOLATOR

In this section we provide an intuitive description of the Chirpolator and provide a derivation of the equations, beginning with the simplest two antenna case. We then extend the results to multiple antennas in one dimension (1D). Extensions to three-dimensional (3D) telescope geometries, the effects of non-linear dispersion delay and novel techniques for efficiently implementing the Chirpolator are all described in Appendix C.

4.3.1 Intuitive Description

Here, we describe an overview of the Chirpolator to aid the intuition of the reader. Put simply, the Chirpolator exploits the observation that when a linear chirp received by one antenna is multiplied by a delayed linear chirp received at another antenna, the result is a fixed-frequency tone whose frequency is proportional to the geometric delay. The discrete Fourier transform (DFT) of these tones can be coherently combined across all antenna pairs to form a detection metric.

A more rigorous mathematical description is described in Section 4.3.2 and in the following.

1. We model the dispersed pulse from an astronomical source as a finite-duration linear chirp ($s(t)$), i.e., a signal whose frequency sweeps linearly across the bandwidth (B) in a time T (see Figure 4.3, top panel). Such a signal has a constant frequency gradient $\dot{f} = B/T$.
2. This signal is received by two antennas, indexed p and q . The signal ($s_p(t)$) at antenna p is delayed with respect to the signal ($s_q(t)$) at antenna q by an unknown geometric delay (τ).
3. The difference in frequency between the two signals is constant for the duration of the pulse and is equal to $\dot{f}\tau$ (Figure 4.3, top panel).
4. If we multiply the signal from antenna p by the conjugate of the signal from antenna q , the result ($x_{pq}(t)$) is a tone at fixed frequency (Figure 4.3, bottom-left panel). This multiplication is equivalent to “down-conversion” (also known as “mixing”), which shifts the frequency of a signal in a radio frequency system. In the mixing case, an incoming signal is multiplied by a fixed-frequency Local Oscillator (LO), and the result has a centre frequency which is the *difference* between the centre

frequency of the incoming signal and the LO frequency. In our case, both the “LO” and the incoming signal are sweeping at the same rate (\dot{f}) but the frequency difference remains fixed. Thus, the signal at antenna p is effectively ‘down-converted’ by an “LO” (provided by antenna q) that is perfectly matched in frequency, yielding a fixed-frequency tone.

5. We have a fixed frequency tone, with unknown frequency ($\dot{f}\tau$) and duration T . In practise, this tone will also be contaminated by noise. By taking the DFT of this signal ($X_{pq}[k]$), all the energy of the sinusoid is coherently added into a small number of DFT bins, while the noise adds incoherently. Therefore, taking the DFT increases the S/N by approximately the square root of the DFT length. Also, for most arrays of interest, the signal can be more compactly expressed in a DFT, as the range of possible frequencies is much smaller than the number of samples (see Appendix C.3.1), which reduces the downstream data and processing rates.
6. We repeat the above two steps for each antenna pair, and produce a DFT spectrum for each (Figure 4.3, bottom-right panel). The spectrum has a peak at frequency k_0 . This frequency is proportional to the geometric delay (τ), which is in turn proportional to the baseline length (b_{pq}) and angle of arrival (θ) (see Figure 4.4). The value of the peak of the spectrum ($X_{pq}[k_0]$) is a complex number whose phase (Φ_{pq}) is also a function of the geometric delay (τ).
7. Finally, we form an image, which is a detection metric ($P(\theta)$) for a range of trial directions of interest. To produce the detection metric for a given direction of interest, we compute the expected arrival frequency (k'_0) and DFT phase correction (Φ_{pq}^*) for a given antenna pair. We then pick out the DFT bin at the expected frequency ($X_{pq}[k'_0]$) and multiply by the phase correction (Φ_{pq}^*) so that the bins for all pairs have the same absolute phase (see Figure 4.7). A vector sum of the phase-corrected DFT bins over all antenna pairs is a coherent sum across all antennas, and yields a detection metric in the direction of interest.
8. In practise, both the time of arrival and actual DM (equivalent to T) are not known in advance. Therefore, we repeat the above procedure in a sliding window fashion and assuming a range of DMs. This repetition can be efficiently implemented using a number novel of techniques (see Appendix C.3).

4.3.2 Two-antenna case

In this section we develop a more rigorous description of the Chirpolator. We begin by considering a single pulse which has been dispersed by the ISM, which




Figure 4.3: Schematic illustrating the “the Chirpolator” operating with two antennas. Top left: two linear chirps are received by antennas p and q , with one delayed by τ . Bottom left: after taking the product of the two voltage time series, the result x_{pq} has constant frequency over most of the duration of the chirp. Bottom right: the DFT of x_{pq} yields a peak at k_0 .

we approximate by a linear chirp impinging on an ideal (perfectly calibrated) two-antenna array. A schematic of the scheme is shown in Figure 4.3.

This technique was described by Gershman et al. (2001) as the maximum likelihood detector for a single chirp, which they termed a “chirp beamformer”.

As described in Section 4.2.1, the voltage waveform received by an antenna can be written as a complex linear chirp with unit amplitude:

$$s(t) = \exp\left(\pi j \dot{f} t^2\right), \quad (4.17)$$

where $\dot{f} = -B/T$ is the chirp rate, B is the system bandwidth, and T is the time taken for the chirp to cross the bandwidth. Assume this signal is received by two antennas, with the signal delayed at antenna q by τ seconds with respect to the arrival at antenna p . The product of the chirp received by antenna p , with its delayed and conjugated counterpart from antenna q is

$$x_{pq}(t) = s_p(t)s_q^*(t) \quad (4.18)$$

$$= s(t)s^*(t - \tau) \quad (4.19)$$

$$= \exp\left(\pi j \dot{f} t^2\right) \exp\left(-\pi j \dot{f} (t - \tau)^2\right) \quad (4.20)$$

$$= \exp\left(\pi j \dot{f} (2t\tau - \tau^2)\right), \quad (4.21)$$

which is a complex sinusoid of frequency $\dot{f}\tau$ and phase $-\pi j \dot{f} \tau^2$. Taking the product in this way is also termed “mixing”. We have assumed here that $\tau \ll T$, which implies that the signals received by both antennas substantially overlap in time. If we have discrete-time sampling, we simply replace t by n/f_s , where n is the sample number and f_s is the sampling frequency. For complex Nyquist sampling, $f_s = B$.

The sampled version of $x_{pq}(t)$ is therefore

$$x_{pq}[n] = \exp\left(\pi j \dot{f} \left(\frac{2n\tau}{f_s} - \tau^2\right)\right). \quad (4.22)$$

If we take the DFT of $x_{pq}[n]$ over N samples, where $N = f_s T = BT$, and by using the standard result of the DFT of a complex sinusoid of finite duration,¹⁴ we obtain

$$X_{pq}[k] = \text{DFT}\{x_{pq}[n]\} \quad (4.23)$$

$$= \sum_{n=0}^{N-1} \exp\left(\frac{-2\pi j n k}{N}\right) \exp\left(\pi j \dot{f} \left(\frac{2n\tau}{f_s} - \tau^2\right)\right) \quad (4.24)$$

$$= \sum_{n=0}^{N-1} \exp\left(\frac{-2\pi j n k}{N} + \frac{2\pi j n \dot{f} \tau}{f_s}\right) \exp\left(-\pi j \dot{f} \tau^2\right) \quad (4.25)$$

$$= \exp\left(-\pi j \dot{f} \tau^2\right) \sum_{n=0}^{N-1} \exp\left(\frac{-2\pi j n}{N} \left(k - \frac{\dot{f} \tau N}{f_s}\right)\right) \quad (4.26)$$

$$= \Phi_{pq}(k - k_0) D_N(k - k_0), \quad (4.27)$$

where k_0 is the frequency of $x_{pq}[k]$ (in units of DFT bins), given by

$$k_0 = \frac{\dot{f} \tau N}{f_s} \quad (4.28)$$

$$= \frac{(B/T)\tau(f_s T)}{f_s} \quad (4.29)$$

$$= B\tau, \quad (4.30)$$

¹⁴Using the shift theorem, the Fourier transform of a delta function and the similarity theorem.

where $D_N(x)$ is a real-valued amplitude term, whose shape is the Dirichlet kernel, defined as

$$D_N(x) = \begin{cases} N & x = 0 \\ \frac{\sin(\pi x)}{\sin(\pi x/N)} & x \neq 0, -N < x < N \end{cases} \quad (4.31)$$

and $\Phi_{pq}(x)$ is a unit-amplitude complex phase term given by

$$\Phi_{pq}(x) = \exp\left(-\pi j \dot{f} \tau^2\right) \exp\left(-\pi j x \frac{N-1}{N}\right) \quad (4.32)$$

$$\simeq \exp(-\pi j x). \quad (4.33)$$

For the two-antenna case, we can write the geometric delay τ simply as

$$\tau = \frac{b_{pq} \sin \theta}{c}, \quad (4.34)$$

where b_{pq} is the distance between the antennas, θ is the angular separation of the source from the phase centre, and c is the speed of light in the medium. The discrete frequency of x_{pq} , k_0 , corresponds to the position of the peak in the spectrum $X_{pq}[k]$ and is related to the baseline length, angle of arrival, and bandwidth by

$$k_0 = B\tau \quad (4.35)$$

$$= B \frac{b_{pq} \sin \theta}{c}. \quad (4.36)$$

Thus, the frequency of the mixed signal is linearly related to the sine of the angle arrival and the baseline length, as sketched in Figure 4.4.

To get an idea of the important factors in the above expressions, we can substitute typical values for current medium-sized dish-based radio telescopes. For a baseline of 1 km, bandwidth of 400 MHz, dispersion delay of 1 s and beamwidth of 1° , $\tau^2 \ll 1$ and the first exponential term of Equation 4.33 approaches 1. In the same regime, the number of samples, N , is large, of the order $N = f_s T = BT > 10^5$. The large number of samples has two consequences. First, in Equation 4.33, $(N-1)/N$ approaches 1, and most importantly the amplitude term, $D_N(x)$ in Equation 4.31, has only a small region of support around k_0 . This allows computational savings by allowing us to truncate the computation of DFT bins to a few bins centred around k_0 , meaning that calculation of the full DFT spectrum is not required, and downstream processing is also considerably reduced.




Figure 4.4: Response of the Chirpolator to a source offset from the phase centre. As the baseline length (b_{pq}) between antennas increases, the position of the peak in the DFT (k_0) increases linearly, with a gradient given by $B \sin(\theta)/c$ (see Equation 4.36). The amplitude of the DFT, $D_N(k - k_0)$, is shown to illustrate that there is some smearing of the signal around the expected frequency k_0 .

4.3.3 Multiple telescopes in 1D

To detect a chirp with a given \dot{f} coming from an unknown direction, we form a detection metric, or intensity image, over a range of directions of interest. The detection metric is formed by phasing up results from all pairs of antennas. For simplicity, we assume that the array is perfectly calibrated, and we ignore the smearing from the $D_N(x)$ term. Assuming that the DFT spectrum is a single delta function, with all energy in the bin:

$$k'_0 = \text{round}(k_0), \quad (4.37)$$

then for a particular direction of interest θ , we compute k'_0 for each antenna pair using Equations 4.36 and 4.37, and then compute the value of a single DFT bin $X_{pq}[k'_0]$. We can essentially phase up each antenna pair by multiplying by the conjugate of the known phase term in Equation 4.33 and a detection metric can be formed by performing a vector sum across the phased-up antenna pairs according to

$$P(\theta) = \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} \Phi_{pq}^*(k'_0 - k) X_{pq}[k'_0]. \quad (4.38)$$

The procedure can be repeated for a range of θ . If the direction of interest, θ , and the actual angle of arrival coincide, the DFT will have a peak at k'_0 with a value of $X_{pq}[k'_0]$. By substituting Equation 4.27, $P(\theta)$ reduces to

$$P(\theta) = \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} \Phi_{pq}^*(k'_0 - k_0) \Phi_{pq}(k'_0 - k_0) D_N(k'_0 - k_0) \quad (4.39)$$

$$= \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} D_N(k'_0 - k_0) \quad (4.40)$$

$$\simeq \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} D_N(0) \quad (4.41)$$

$$= N \frac{M(M-1)}{2}. \quad (4.42)$$

In general, the quantity $P(\theta)$ will be complex-valued. We are not interested in the absolute phase of the signal, so a more useful metric for thresholding is

$$E(\theta) = |P(\theta)|^2. \quad (4.43)$$

The fact that the sum over antenna pairs is a vector sum means that the resulting S/N scales with M rather than the $M^{1/2}$ scaling for non-coherent addition.

The time sequence of $E(\theta)$ for a particular value of θ can be considered a typical time sequence of power measurements, and can be subjected to the usual pulsar detection methods such as periodicity and acceleration searches.

4.3.4 Compensating for the smearing in $D_N(x)$

Equations 4.37 and 4.39 assume that all the energy is concentrated in a single bin. For an arbitrary angle of arrival, this is not the case, and in the worst case, the energy can be spread over all the bins in the DFT (see Figure 4.3). For large N , $D_N(x)$ has relatively compact support, so we can truncate the number of DFT bins we compute, as well as the number of bins that need to be summed for a given direction of arrival and antenna pair. We can choose, therefore, to truncate the computations to $2F + 1$ bins centred around k'_0 . In practise, one can choose a value of F that provides the best trade between computational cost and S/N.

To capture the energy with support $[-F, F]$ around k_0 , we perform a matched filter operation against the expected amplitude response function, which is the shifted $D_N(x)$. Therefore, Equation 4.39 can be trivially generalised to

$$P(\theta) = \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} \sum_{k=-F}^{k=+F} \Phi_{pq}^*(k + k'_0 - k_0) D_N(k - k_0) X_{pq}[k + k'_0] \quad (4.44)$$

4.4 THE CHIMAGEATOR

We now describe an alternative method for combining signals from multiple telescopes based on gridding cross-multiplied voltages. Once again, we begin with an intuitive description and provide more mathematical rigour in later sections.

4.4.1 Intuitive Description

Here, we describe an overview of the Chimageator to aid the intuition of the reader. The first three steps of the Chimageator are exactly the same as the Chirpolator (Section 4.3.1), i.e., the Chimageator exploits the observation that when a linear chirp received by one antenna is multiplied by a delayed linear chirp received at another antenna, the result is a fixed-frequency tone whose frequency is proportional to the geometric delay. The difference between the two techniques is how the cross-multiplied data are combined: the Chimageator exploits an efficient spatial FFT at each sample time. The result is a dynamic spectrum in which the chirp deposits energy along a linear trajectory in the time-spatial frequency plane. The gradient of the trajectory is proportional to the geometric delay. We sum DFT bins along a range of trial trajectories to form a detection metric.

1. We begin as with the Chirpolator, by assuming a linear chirp which sweeps across the bandwidth (B) in time T , with gradient $\dot{f} = B/T$.
2. As with the Chirpolator, the chirp is received by two antennas and the resulting signals from these antennas are multiplied together (mixed).
3. The resulting mixed signal ($x_{pq}(t)$) has constant frequency. Once again the frequency is proportional to the distance between antennas and to the angle of arrival.
4. We would like to take a spatial FFT of the mixed signals over all antennas at each sample time. Much like the Fourier transform in regular interferometry, this spatial FFT requires the signals to be sampled on a regular grid. To form a regular grid ($x'_l[n]$), we take the sampled, mixed signal from each pair of antennas ($x_{pq}[n]$) and average those products which have the same inter-antenna spacing (l), and therefore the same (and therefore redundant) geometric delays. This process is known as “gridding”. Gridding can also be used to interpolate a non-uniform array geometry onto a uniform grid so that the FFT can be used.

5. The gridded signals from all antennas comprise a “space-time tone,” i.e., for a given sample number (n), the spatial frequency of the tone is proportional to the angle of arrival (θ). Similarly, for a given inter-antenna spacing (l), the temporal frequency is proportional to θ .
6. For each sample number n , we take the DFT of the gridded signals over the spatial dimension (which can be implemented as an FFT). The result is the DFT of a single tone ($X_k[n]$), which has a peak at the bin k_0 .
7. Unlike with the Chirpolator, the peak in the DFT (k_0) is not a constant. In fact the peak increases linearly with the sample number n and is proportional to the arrival direction θ . As a result, a pulse of duration T arriving from a direction θ will trace out a linear trajectory in time where it will cross a number of spatial DFT bins (Figure 4.5) during its duration.
8. The angle of arrival and DM (equivalent to pulse duration T) are unknown. Therefore, at each sample time, we assume a set of trial angles and durations, which correspond to a set of trial trajectories. To form a detection metric $P(\theta)$ for each angle and duration, we sum along the trial trajectory (applying a phase correction Φ^* as we go).
9. Additional optimisations are possible. For example, the shorter trajectories can be calculated as partial sums along longer trajectories with the same gradient, and the spatial FFTs can be averaged before performing the trajectory sums. These optimisations are discussed in Appendix D.2.

4.4.2 Formulation for a uniform linear array

In this section, we develop a more rigorous description of the Chimageator.

Consider a linear, perfectly calibrated array of M antennas, uniformly spaced with inter-element spacing L . If a linear chirp impinges on the array, the product of the signals from two antennas, indexed p , and q is, therefore, given by

$$x_{pq}(t) = s_p(t)s_q^*(t) \quad (4.45)$$

$$= s(t - p\tau)s^*(t - q\tau) \quad (4.46)$$

$$= \exp\left(\pi j \dot{f}(t - p\tau)^2\right) \exp\left(-\pi j \dot{f}(t - q\tau)^2\right) \quad (4.47)$$

$$= \exp\left(-\pi j \dot{f}(2t\tau(p - q)) + \tau^2(p^2 - q^2)\right) \quad (4.48)$$

$$\simeq \exp\left(-2\pi j \dot{f}t\tau(p - q)\right), \quad (4.49)$$

where

$$\tau = \frac{L}{c} \sin(\theta). \quad (4.50)$$

As before, we can form the sampled signal by replacing t with n/f_s . Next, we combine the values of $x_{pq}[n]$ for all baselines with the same spacing l and the sample number n in a process known as gridding. The uniform linear array has redundant spacings which can be combined and weighted according to

$$x'_l[n] = \sum_{p=0}^{M-l-1} w_i x_{p,p+l}[n], \quad (4.51)$$

where l runs from 1 to $M - 1$ (the auto-correlations are ignored). $x'_l[n]$ corresponds to the visibility measured by combining all baselines with spacing $(l + 1)D$ and w_i are weights. $w_i = 1$ corresponds to “uniform” weighting, yielding the maximum S/N but reduced resolution. $w_i = 1/(N - l)$ corresponds to natural weighting, and yields maximal resolution but reduced S/N. Through a suitable choice of weights, an arbitrary array be interpolated onto a regular grid as required for the spatial FFT, by gridding with a spatially varying set of weights. The interested reader is referred to Taylor et al. (1999, Chapter 7, Section 3) for details.

For a single chirp, we can substitute Equation 4.49 and assuming natural weighting and a uniform linear array, the gridded voltages simplify to

$$x'_l[n] = \sum_{p=0}^{M-l-1} w_i \exp\left(-2\pi j \dot{f} n \tau (p - (p + l)) / f_s\right) \quad (4.52)$$

$$= \exp\left(2\pi j \dot{f} n \tau l / f_s\right) \sum_{p=0}^{M-l-1} w_i \quad (4.53)$$

$$= \exp\left(2\pi j \dot{f} n \tau l / f_s\right). \quad (4.54)$$

$$(4.55)$$

A spatial DFT of the gridded voltages yields

$$X_k[n] = DFT \{x'_l[n]\} \quad (4.56)$$

$$= \sum_{l=0}^{M-1} \exp\left(\frac{-2\pi j k l}{N}\right) \exp\left(2\pi j \dot{f} \tau l \frac{n}{f_s}\right) \quad (4.57)$$

$$= \sum_{l=0}^{M-1} 2\pi j l \left(\frac{-k}{N} + \tau \dot{f} \frac{n}{f_s}\right) \quad (4.58)$$

$$= D_M(k - k_0) \Phi(k - k_0). \quad (4.59)$$

where $D_M(f)$ is the Dirichlet kernel defined earlier and Φ is a unit-amplitude phase term. This spatial DFT can be efficiently implemented as a Fast Fourier transform.

A chirp crossing a bandwidth $B = f_s$ in time T , arriving from angle θ , puts power in the DFT bin given by

$$k_0(n, T, \theta) = \dot{f} \tau M \frac{n}{f_s} \quad (4.60)$$

$$= \frac{n}{T} \tau M \quad (4.61)$$

$$= \frac{n}{T} \frac{ML}{c} \sin(\theta) \quad (4.62)$$

$$\simeq \frac{n}{T} \frac{b_{\max}}{c} \sin(\theta), \quad (4.63)$$

which is very similar to Equation 4.36 for the Chirpolator described previously, with the key difference that in this case, the k_0 term now depends linearly on the sample number n rather than on the baseline length b_{pq} . Thus, a chirp will appear as power along a diagonal trajectory in n - k space, as shown in Figure 4.5.

For a chirp beginning at $n = 0$, the trajectory ends at the DFT bin given by

$$k_{0,\text{end}} = k_0(T f_s, T, \theta) \quad (4.64)$$

$$= f_s \frac{b_{\max}}{c} \sin(\theta). \quad (4.65)$$

To form an image, we can sum across the diagonal trajectory in DFT bins and time, applying the inverse of the phase term to produce an intensity image for a given dispersion delay, given by:

$$P_T(\theta) = \sum_{n=0}^{f_s T} \Phi^*(k_0(n, T, \theta)) X_k(k_0(n, T, \theta)), \quad (4.66)$$

with the scalar energy computed as in Equation 4.43.

4.5 METHOD OF COMPARISON

We have described two new antenna-coherent techniques for detecting dispersed pulses with interferometers. In this section, we describe our method of comparing our techniques to two existing classical techniques with roughly equivalent sensitivity: Fourier imaging, and direct beamforming with frequency incoherent processing (see Figure 4.2). These classical techniques are described in Sections 4.2.4.2, 4.2.4.3, and 4.2.3.1.




Figure 4.5: Trajectories of linear chirps with varying durations (T) and angles of arrival (θ) after gridding and Fourier transforming with the Chimageator. As the sample number (n) increases, the peak of the DFT (k_0) increases linearly (see Equation 4.63). Two types of trajectory are shown with dashed lines: the θ_{\max} case, which corresponds a range of dispersion delays, and a single arrival angle at the edge of the field of view; and the T_{\max} case, where each trajectory corresponds to the longest dispersion of interest and a range of arrival angles. The thick line is the trajectory corresponding to $T_{\max}/2$ and $\theta_{\max}/2$, which lies along the T_{\max}, θ_{\max} trajectory and can be therefore be computed from the partial sums along the T_{\max}, θ_{\max} trajectory.

Ideally, we would like to compare the techniques in terms of the true implementation costs. But, evaluating the true implementation cost is complicated by a number of considerations:

- The choice of survey parameters (e.g., minimum and maximum DMs, centre frequency).
- The telescope parameters (e.g., number of antennas, system bandwidth, baseline distribution).
- The economics of available technologies.
- The details of the implementation on a given technology. For example, how an algorithm is parallelised over a number of processors.
- The techniques do not yield equivalent sensitivities in certain situations (e.g., Section 4.7.1.2).
- The parametrisation of the algorithms themselves.

To explain the final point further: an implementation of an algorithm requires a set of parameters that affects both the cost and the sensitivity of that implementation (e.g., number of channels for interferometric imaging or F for the Chirpolator). For each algorithm, the relationship between the parameters and sensitivity is complicated, and there is no straightforward way to choose realisations that yield equivalent sensitivities for all techniques so that their costs can be compared fairly.

4.5.1 A simple model for evaluating algorithm cost

To help illustrate, in very approximate terms, the differences in operations and data rates required by different methods, we propose a simple model. In this model, we split each algorithm up into two basic functional blocks: the processing required before an integrate-and-dump step and the processing required after it. We also consider the data rate required between the two blocks, i.e., immediately after the integrate-and-dump step. We acknowledge that this model does not consider very important details of how data are transported within each block, and we also acknowledge that the bandwidth bottlenecks may indeed be within each block, rather than between the two. But, the bandwidth requirements inside each block are a strong function of the way the processing is parallelised inside each block, and quantifying the many different methods for doing this parallelisation are outside the scope of this Chapter.

This functional breakdown applies to the techniques as follows:

The Chirpolator. The pre-integrator step is the sliding-DFT (Appendix C.3.2). The integrated output is a set of DFT results per DM trial. The post-integrator step is the imaging per DM trial. Details of the data and operation rates are described in Appendix C.

The Chimageator. The pre-integrator steps include gridding and integration to the shortest sampling interval. The integrated output is a sequence of partially averaged images. The post-integrator steps include the remaining integration for the full range of DM trials and the imaging. Details of the data and operation rates are described in Appendix D.

Fourier Imaging. The pre-integrator steps include cross-correlation and integration. The integrated output is the visibilities. The post-integrator steps include gridding, FFT, and tree incoherent dedispersion (Taylor, 1974). For the bandwidth requirement, we sum both the requirements for both the visibilities, and the “corner turn” required for dedispersion. Operation rates are described by Cordes (1997)

Direct Beamforming. We form as many tied array beams as required to cover the entire primary beam. The pre-integrator steps include the beam forming and integration. The integrated output is a power spectrum per beam. The post-integrator step is tree incoherent dedispersion. Operation rates are described by Cordes (1997).

4.5.2 Array, survey, and algorithm parameters

To arrive at concrete values of bandwidth and operation rate, we must define a full set of parameters for an array, survey and each algorithm. To motivate our example, we choose a set of parameters based on the SKA from Cordes (1997), as shown in Table 4.2. Clearly, evaluating the performance of all techniques as a function of all parameters results in a highly multidimensional data set. For the sake of simplicity, we leave only one free parameter: the number of antennas in the array (M). We let M go from 2 antennas up to 2000 antennas, which covers the range of values for SKA and its pathfinders.

4.6 RESULTS

For all but the largest arrays, the Chimageator and the Chirpolator have substantially superior bandwidth requirements than the classical techniques (Figure 4.6). The superior bandwidth requirements are achieved because of a difference in timescale that needs to be sampled by the integrate-and-dump step: our techniques sample the shortest dispersion delay, while the classical techniques sample the shortest dedispersed pulse duration. As a dedispersed

Table 4.2: Parameters used in our example model. The parameters for the Chirpolator and the Chimageator are defined in Appendices C and D, respectively.

Parameter	Value
Array parameters	
System bandwidth (MHz)	400
Antenna size (m)	12
Maximum baseline (m)	1000
Centre frequency (GHz)	1.4
Number of polarisations	2
Survey parameters	
Minimum DM (cm^{-3}pc)	10
Maximum DM (cm^{-3}pc)	1000
Fourier imaging and direct beamforming	
Number of frequency channels	1000
Number of DM trials	1000
Integration time (seconds)	10^{-4}
Bytes per visibility (post correlator)	2
Bytes per image pixel	1
Chirpolator specific	
DM step (ϵ)	0.1
Smearing support size (F)	1
Time oversampling (κ_t)	4
Spatial oversampling (κ_s)	1
Bytes per DFT bin	2
Chimageator specific	
DM step (ϵ)	0.1
Smearing support size (F)	1
Time oversampling ($\kappa_{t,0}$)	4
Spatial oversampling (κ_s)	1
Operations per grid point	50
Bytes per FFT bin	2

pulse can be substantially shorter than the dispersion delay, the classical techniques must dump their integrators at a much higher rate, therefore requiring larger bandwidth between the functional blocks. One additional factor worsens the bandwidth requirements for Fourier imaging in particular: below about $\simeq 100$ antennas, the bandwidth is dominated by the dedispersion “corner turn”.

In terms of post-integrator operation rate, the Chirpolator betters all other techniques up to $\simeq 200$ antennas. This low rate for small arrays is a consequence of both the low input bandwidth and the fact that the imaging operates on a per-baseline basis. Above $\simeq 200$ antennas, the Direct beamforming method is the clear winner, as the dedispersion cost is fixed by the longest baseline, rather than by the number of antennas.

In terms of pre-integrator operation rate, Fourier imaging is clearly the most efficient approach for all array sizes of interest, with our two techniques presented here requiring between 2 and 4 orders of magnitude more operations

for equivalent array sizes.

4.7 DISCUSSION

For any array size, there is no clearly superior algorithm in all measures. The Chirpolator has a high pre-integrator operation rate, but has good post-integrator and data requirements for small to medium arrays. The Chimageator has a consistently high post-integrator operation rate. Fourier imaging is computationally attractive but has very high data rate requirements, either due to the corner turn in small arrays or the visibility data rate in large arrays. Direct beamforming has very high operations and data rate requirements for small arrays, but becomes somewhat competitive for larger ones. The preferred algorithm, therefore, will depend on the details of the array, survey and algorithm parameters, and on the economics of available computing technologies.

The economics of supercomputing are changing rapidly. The increase in arithmetic capability of processors has been well described by Moore's Law, that is, the number of transistors (and by inference, arithmetic capability) on a chip doubles every 18 months. While this prodigious improvement is very welcome for the arithmetic part of the problem, it does not hold for data rate, which has traditionally grown much more slowly. We propose that, because the arithmetic capability of processors is outstripping the bandwidth capability, our techniques with their superior data rate performance, will become more and more favourable as technology progresses, in spite of their requirements for higher operation rate. Therefore, on the timescale of the SKA and its pathfinders, our techniques may be preferred over the classical ones.

4.7.1 Further work

4.7.1.1 *Effect of calibration errors*

In our analysis we have assumed an ideal, perfectly calibrated array, in which all the antenna gains are equal and have zero relative phase. In practise, each antenna will have uncalibrated errors in gain and phase that will affect the performance of our algorithms. While a detailed discussion of the effect of calibration errors is outside the scope of this Chapter, we present here a simple proof that phase errors (which we model as delay errors) in the Chirpolator case will result in decoherence across the array and reduced S/N.

If we assume that the uncalibrated delay error between two antennas is τ_{err} then we can make the substitution $\tau \rightarrow \tau + \tau_{\text{err}}$ into Equation 4.35 to obtain the frequency of the tone after mixing:




Figure 4.6: Data and operation rates for a dispersed pulse survey as a function of algorithm and number of antennas. Survey parameters are given in Table 4.2. Top panel: data rate between pre-integrator and post-integrator steps. Middle panel: operation rate before the integrator. Bottom panel: operation rate after the integrator. “Fourier + Tree” and “Direct + Tree” signify Fourier imaging and Direct beamforming, respectively, to form beams, and using tree dedispersion described by Taylor (1974).

$$\hat{k}_0 = B(\tau + \tau_{\text{err}}) \quad (4.67)$$

$$= k_0 + k_{\text{err}}, \quad (4.68)$$

Therefore, a delay error changes the frequency of the mixed signal, and shifts the entire DFT spectrum from $X_{pq}[k]$ to $X_{pq}[k + k_{\text{err}}]$. The shift in the DFTs reduces the amplitude of the detection metric, which is formed by a vector sum of the phase-corrected DFT bins from each antenna pair. The detection metric has a maximum value when all the phase-corrected DFT bins have the same absolute phase. If an antenna pair contains a delay error, each phase-corrected DFT bin will not have the same absolute phase, and the vector sum will not be over a straight line (Figure 4.7), resulting in reduced amplitude of the sum. This process can be quantified by substituting \hat{k}_0 into Equation 4.38:

$$P_{\text{err}}(\theta) = \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} \Phi_{pq}^*(k'_0 - k) X_{pq}[\hat{k}_0] \quad (4.69)$$

$$= \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} \Phi_{pq}^*(k'_0 - k) \Phi_{pq}(k'_0 + k_{\text{err}} - k_0) D_N(k'_0 + k_{\text{err}} - k_0) \quad (4.70)$$

$$= \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} \exp(j\pi(k'_0 - k)) \exp(-j\pi(k'_0 + k_{\text{err}} - k_0)) D_N(B\tau_{\text{err}}) \quad (4.71)$$

$$= \sum_{p=0}^{M-1} \sum_{q=p+1}^{M-1} \exp(-j\pi B\tau_{\text{err}}) D_N(B\tau_{\text{err}}) \quad (4.72)$$

$$\leq P_{\text{ideal}}(\theta). \quad (4.73)$$

The inequality in Equation 4.73 is a result of the triangle inequality for vector addition (see Figure 4.7), and the fact that $D_N(x) \leq D_N(0)$.

It is clear from this argument that delay errors will result in a reduced detection metric, resulting in a drop in S/N. We leave a quantitative analysis of this effect, and of other calibration effects, for future work.

4.7.1.2 Extension to millisecond pulsars

Both our methods have assumed that a chirp is received in isolation, meaning that during the duration T of a chirp, no other chirps are received. This condition is violated for millisecond pulsars, which have short periods and can have large DMs. The combination of short period and large DM means that a chirp will not have finished traversing the system bandwidth B before a subsequent chirp is received.

We can write the isolated chirp condition for a pulsar with period P as:




Figure 4.7: Errors in delay calibration reduce the amplitude of the detection metric. Here we plot the formation of the detection metric $P(\theta)$ as the vector (i.e., complex) sum of the phase-corrected DFT results from three antenna pairs. In the ideal case (red), the phase correction (Φ_{pq}^*) perfectly corrects for the known phase in the DFT bins (X_{pq}), and each result has the same absolute phase. The resulting detection metric ($P_{\text{ideal}}(\theta)$) is fully coherent. If delay errors are present, each DFT bin has a residual phase that is different for each antenna pair. The resulting detection metric ($P_{\text{err}}(\theta)$) has a smaller amplitude, because the vectors do not add into a straight line.

$$P > T \quad (4.74)$$

$$> \mu\text{DM}(\nu_1^{-2} - \nu_2^{-2}). \quad (4.75)$$

If the isolated chirp condition is not satisfied, there are multiple chirps occupying the bandwidth at any one time. These additional chirps produce additional mixing products at the multiplication steps (i.e., in Equations 4.18 and 4.45) that appear at frequencies that are outside the frequencies searched in the isolated chirp case. If only isolated chirp processing is performed, the energy in the additional mixing products is effectively lost, with a resulting loss in S/N. Our techniques will still operate effectively, but the S/N achieved will not be as high as when processed by other methods. Quantifying loss of




Figure 4.8: Fraction of known pulsars that violate the isolated chirp condition, as a function of system bandwidth and centre frequency (f_c). Known pulsars are taken from the ATNF pulsar catalogue.

energy to mixing products, and the resulting loss in S/N, is outside the scope of this Chapter.

To determine what fraction of pulsars violate this condition, we use the ATNF pulsar catalogue (Manchester et al., 2005).¹⁵ This catalogue contains the DMs and periods for all known pulsars. At a bandwidth of 400 MHz at 1.4 GHz, 30% of known pulsars have periods that are too high to satisfy the isolated chirp condition for their DMs (Figure 4.8), indicating that this effect is important.

4.8 SUMMARY

We have described two new techniques for detecting dispersed pulses with radio interferometers, which we call the Chirpolator and the Chimageator. These techniques have antenna-coherent sensitivities in the isolated chirp case, and have substantially lower data rate requirements than other coherent meth-

¹⁵<http://www.atnf.csiro.au/research/pulsar/psrcat>. Version 1.43. Accessed 2011 March 29.

ods for realistic array configurations (including the configurations of the SKA and its pathfinders). For small to medium array sizes, the Chirpolator is also more efficient than classical techniques in terms of post-integrator operation rate. While the pre-integrator operation rates our methods high in some cases, the changing economies of computer design may favour the lower bandwidth requirements of our new techniques in spite of their high operation rate requirements.

Conclusion

THE new age of wide field radio astronomy is fast approaching. Upcoming telescopes will provide an unprecedented view of the dynamic universe, probing new parts of the transient parameter space and potentially discovering new classes of sources. This thesis is a contribution to the preparations for this new age, presenting two pilot surveys, and a range of new processing techniques.

5.1 PILOT SURVEYS

The pilot surveys described in this thesis provide important statistical distributions for future survey planning, as well as experience and data to aid algorithm development. This thesis contains the results of two pilot surveys: a blind survey and a GRB follow-up survey.

Chapter 2 described a blind, archival survey for transient and variable radio sources using the Molonglo Observatory Synthesis Telescope (MOST) at 843 MHz. It discovered 53 highly variable sources and 15 transient sources. This survey determined surface densities of 0.268 deg^{-2} for variable sources and $1.3 \times 10^{-2} \text{ deg}^{-2}$ for transient sources, with timescales between one day and three years. The brightness distribution and cadences of the transient and variable sources were calculated and discussed. The light curves for the transient and variable sources have all been published and can be used to train future classification algorithms. The lessons learned during the survey processing have also necessitated the development of a number of new algorithms.

Chapter 3 described an automatic follow-up survey in which I searched for prompt 1.4 GHz radio emission from gamma ray bursts at high time resolution. This survey discovered single, dispersed pulses following two long GRBs. An event related to the delayed formation of a black hole at the centre of the GRB is a possible origin of these pulses. If the single pulses are not related to the GRB, this survey sets an upper limit for radio pulses from GRBs in the time between 200 and 1800 s of the gamma ray trigger, of $1.27w^{-1/2} \text{ Jy}$, where $6.4 \times 10^{-5} \text{ s} < w < 32 \times 10^{-3} \text{ s}$ is the pulse width. This survey sets a limit of 760 Jy on any long-duration emission ($> 1 \text{ s}$) between 200 and 1800 s from the GRB trigger. The high time resolution limits are some of the strongest limits on prompt radio emission from gamma ray bursts in the time between 200 and 1800 s of the burst.

5.2 TECHNIQUES

The new techniques developed in this thesis will improve the efficiency and robustness of upcoming surveys for radio transients.

Chapter 2 described two techniques for application to image-based radio variability surveys. The first technique corrects for an overall gain term in

an image, by comparing the measured flux densities with the catalogued flux densities. This technique flags poor quality images, in which the flux of the sources cannot be reliably determined. This Chapter also describes a statistical technique to determine whether the flux density errors have been correctly estimated, based on a population of light curves with different numbers of measurements. This statistical technique also provides a direct measurement of the number of variable sources in a population.

Chapter 4 presented two efficient new algorithms for detecting dispersed radio emission in interferometric data: the Chirpolator and the Chimageator. When compared with classical methods, these techniques excel in the regime of sparse arrays, where they both require substantially lower data rates, and where the Chirpolator requires a much lower post-integrator operation rate. Compared with classical techniques, these techniques are better matched to future supercomputing architectures for which the arithmetic capability is outstripping the bandwidth capability, and are therefore suitable for use by interferometer-based fast transient surveys.

5.3 THE FUTURE

The transient and variable sources described in Chapter 2 are, for the most part, unidentified. To aid identification, 96 hr of follow-up time with the Australia Telescope Compact Array (ATCA) at 5 and 8 GHz have been obtained¹⁶, as well as 5 nights of optical follow-up with the WiFeS instrument on the ANU 2.3 m telescope. These observations should uncover the identity of most of these transient and variable sources. Of particular interest are the transient and variable sources coincident with resolved optical counterparts. These sources do not appear to have the properties of standard source classes, so these observations should shed light on what these sources actually are. The statistics measured by this survey have also found scientific application, with my surface density measurement already being used to constrain the neutron star interpretation of the of the Bower et al. (2007) transients (Ofek et al., 2011), as well as to constrain the rate of tidal disruption flares (Bower, 2011; van Velzen et al., 2011).

The survey results and techniques described in this thesis will find almost immediate application in the planning of surveys with new wide field radio telescopes such as ASKAP, MWA and LOFAR. The surface density, flux distribution and cadence measurements will have a direct influence on planning VAST and other transient surveys, and the light curves can be used to train automatic classification algorithms. The image gain calibration and flagging techniques described in Chapter 2 have already been implemented as part of a prototype processing pipeline for the VAST survey, and the statistical technique may soon be incorporated in VAST also. The prototype pipeline

¹⁶Project code C2252

will form the basis of the production software, which will be running on the ASKAP engineering telescope (known as BETA) as early as 2012.

The single pulses following GRBs, described in Chapter 3, are intriguing. The key next step is to confirm that such pulses are related to GRBs, by taking additional observations that do not suffer from the same uncertainties as the results presented here. One method is to use a coincidence detection, with two widely separated observations observing the same GRB, with a detection by both telescopes comprising a convincing confirmation. Another opportunity is to detect the pulse with an interferometer, perhaps deploying the Chirpolator or the Chimageator. Such an approach has the advantage of also increasing the sensitivity of the experiment and the significance of the detection.

The Chirpolator and the Chimageator, described in Chapter 4, can be considered for upcoming interferometer-based fast transients surveys. A software implementation of the Chirpolator has been developed, which operates on voltage data from the Giant Metrewave Radio Telescope (GMRT). This implementation could be used, almost without modification, for a fast transients survey with the GMRT. Both techniques could also find application in the ASKAP fast transients survey, CRAFT. Implementing these techniques on ASKAP would increase the sensitivity by a factor of six over the previously proposed incoherent approach. Alternatively, a six antenna subset of ASKAP (such as BETA) could obtain the same sensitivity as the 36-antenna ASKAP summed incoherently. Cost trade-offs of these two approaches are required to decide on the best approach for such a survey.

5.4 SUMMARY

This thesis has made important contributions to the new, wide-field era of transient radio astronomy. The two pilot surveys have provided transients statistics both for blind fields at low time resolution, and for GRBs at high time resolution. The transient and variable sources discovered in this thesis await follow-up and confirmation and could soon see exciting new discoveries. This thesis describes a range of new techniques that will improve the efficiency and robustness of upcoming radio surveys and provide new insights in this exciting new wide-field era. Ultimately, these surveys and techniques could find application in the Square Kilometre Array, which will revolutionise our understanding of radio transients, and provide humanity with another reason to be fascinated with the dynamic sky.

Bibliography

- Alard, C. Lupton, R. H. 1998, *ApJ*, 503, 325
- Alcock, C., et al. 1997, *ApJ*, 486, 697
- Amy, S. W., Large, M. I., Vaughan, A. E. 1989, *PASA*, 8, 172
- Andrew, B. H., MacLeod, J. M., Harvey, G. A., Medd, W. J. 1978, *AJ*, 83, 863
- Backer, D. C. 1970, *Nature*, 228, 42
- Ball, L., Crawford, D. F., Hunstead, R. W., Klammer, I., McIntyre, V. J. 2001, *ApJ*, 549, 599
- Balsano, R. J., et al. 1998, in *American Institute of Physics Conference Series*, Vol. 428, *Gamma-Ray Bursts, 4th Huntsville Symposium*, ed. C. A. Meegan, R. D. Preece, & T. M. Koshut (Melville, New York, USA: AIP), 585–589
- Bannister, K. W. Cornwell, T. J. 2011, *ApJS*, 196, 16
- Bannister, K. W., Murphy, T., Gaensler, B. M., Hunstead, R. W., Chatterjee, S. 2011, *MNRAS*, 412, 634
- Barthelmy, S. D., Butterworth, P., Cline, T. L., Gehrels, N., Fishman, G. J., Kouveliotou, C., Meegan, C. A. 1995, *Ap&SS*, 231, 235
- Bastian, T. S., Dulk, G. A., Leblanc, Y. 2000, *ApJ*, 545, 1058
- Baumgarte, T. W., Shapiro, S. L., Shibata, M. 2000, *ApJ*, 528, L29
- Becker, R. H., Helfand, D. J., White, R. L., Proctor, D. D. 2010, *AJ*, 140, 157
- Becker, R. H., White, R. L., Helfand, D. J. 1995, *ApJ*, 450, 559
- Bell, M. E., et al. 2011, *MNRAS*, 415, 2

- Bell Burnell, S. J. 1977, in *Annals of the New York Academy of Sciences*, Vol. 302, Eighth Texas Symposium on Relativistic Astrophysics, ed. M. D. Pagiannis, 685
- Bignall, H. E., et al. 2003, *ApJ*, 585, 653
- Błażejowski, M., et al. 2005, *ApJ*, 630, 130
- Bock, D. C.-J., Large, M. I., Sadler, E. M. 1999, *AJ*, 117, 1578
- Bode, M. F., et al. 2006, *ApJ*, 652, 629
- Bonnarel, F., et al. 2000, *A&AS*, 143, 33
- Bower, G., Saul, D., Bloom, J., Bolatto, A., Filippenko, A., Foley, R., Perley, D. 2007, *ApJ*, 666, 346
- Bower, G. C. 2011, *ApJ*, 732, L12
- Bower, G. C. Saul, D. 2011, *ApJ*, 728, L14
- 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
- Bower, G. C., et al. 2010, *ApJ*, 725, 1792
- Briskin, W. F., Macquart, J., Gao, J. J., Rickett, B. J., Coles, W. A., Deller, A. T., Tingay, S. J., West, C. J. 2010, *ApJ*, 708, 232
- Burke-Spolaor, S. Bailes, M. 2010, *MNRAS*, 402, 855
- Burke-Spolaor, S., Bailes, M., Ekers, R., Macquart, J.-P., Crawford, III, F. 2011a, *ApJ*, 727, 18
- Burke-Spolaor, S., et al. 2011b, *MNRAS*, 416, 2465
- Caldwell, N. Rose, J. A. 1997, *AJ*, 113, 492
- Cameron, P. B., et al. 2005, *Nature*, 434, 1112
- Camilo, F., Ransom, S. M., Halpern, J. P., Reynolds, J., Helfand, D. J., Zimmerman, N., Sarkissian, J. 2006, *Nature*, 442, 892
- Camilo, F., Reynolds, J., Johnston, S., Halpern, J. P., Ransom, S. M. 2008, *ApJ*, 679, 681
- Campbell-Wilson, D. Hunstead, R. 1994a, *IAU Circ.*, 6052
- Campbell-Wilson, D. Hunstead, R. W. 1994b, *PASA*, 11, 33

- Cannizzo, J. K., Markwardt, C. B., Ukwatta, T. N., Rowlinson, A., de Pasquale, M. 2010a, GCN Report, 305
- Cannizzo, J. K., et al. 2010b, GCN Circular, 11331
- Carilli, C. L., Ivison, R. J., Frail, D. A. 2003, ApJ, 590, 192
- Carlson, B. 2010, SKA Memo Series, 127
- Chatterjee, S. Murphy, T. 2010, in BAAS, Vol. 41, 515
- Cimò, G., Beckert, T., Krichbaum, T. P., Fuhrmann, L., Kraus, A., Witzel, A., Zensus, J. A. 2002, PASA, 19, 10
- Clark, D. H. Stephenson, F. R. 1977, The historical supernovae (Oxford: Pergamon Press)
- Cline, D. B. 1996, Nuclear Physics A, 610, 500
- Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., Broderick, J. J. 1998, AJ, 115, 1693
- Cordes, J., Lazio, T., McLaughlin, M. 2004, NewAR, 48, 1459
- Cordes, J. M. 1986, ApJ, 311, 183
- Cordes, J. M. 1997, SKA Memo Series, 97
- Cordes, J. M. 2009, SKA Memo Series, 109
- Cordes, J. M. Lazio, T. J. W. 2002, arXiv:astro-ph/0207156
- Cordes, J. M. McLaughlin, M. A. 2003, ApJ, 596, 1142
- Cornwell, T. J. Narayan, R. 1993, ApJ, 408, L69
- Cortiglioni, S., Mandolesi, N., Morigi, G., Ciapi, A., Inzani, P., Sironi, G. 1981, Ap&SS, 75, 153
- Croft, S., Bower, G. C., Keating, G., Law, C., Whysong, D., Williams, P. K. G., Wright, M. 2011a, ApJ, 731, 34
- Croft, S., Tomsick, J. A., Bower, G. C. 2011b, ApJ, 740, 87
- Croft, S., et al. 2010, ApJ, 719, 45
- Croom, S. M., Smith, R. J., Boyle, B. J., Shanks, T., Miller, L., Outram, P. J., Loaring, N. S. 2004, MNRAS, 349, 1397
- D’Addario, L. 2010, SKA Memo Series, 123

- Daishido, T., et al. 2000, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. H. R. Butcher, Vol. 4015, 73–85
- Dale, D. A., Giovanelli, R., Haynes, M. P., Scodeggio, M., Hardy, E., Campusano, L. E. 1998, *AJ*, 115, 418
- D’Elia, V., et al. 2010, *GCN Circular*, 11310
- Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., Hessels, J. W. T. 2010, *Nature*, 467, 1081
- Dessenne, C. A.-C., et al. 1996, *MNRAS*, 281, 977
- Dressler, A. Shectman, S. A. 1988, *AJ*, 95, 284
- Eck, C. R., Cowan, J. J., Branch, D. 2002, *ApJ*, 573, 306
- Fender, R. 2004, *NewAR*, 48, 1399
- Fender, R., et al. 1999, *ApJ*, 519, L165
- Fender, R. P. 2007, in *Bursts, Pulses and Flickering: Wide-Field Monitoring of the Dynamic Radio Sky (PoS(Dynamic2007)30)*
- Fiedler, R. L., et al. 1987, *ApJS*, 65, 319
- Frail, D. A., Kulkarni, S. R., Berger, E., Wieringa, M. H. 2003, *AJ*, 125, 2299
- Frail, D. A., Kulkarni, S. R., Bloom, J. S. 1999, *Nature*, 398, 127
- Frail, D. A., Kulkarni, S. R., Nicastro, L., Feroci, M., Taylor, G. B. 1997, *Nature*, 389, 261
- Gaensler, B. M. Hunstead, R. W. 2000, *PASA*, 17, 72
- Gaensler, B. M., Madsen, G. J., Chatterjee, S., Mao, S. A. 2008, *PASA*, 25, 184
- Gaensler, B. M., et al. 2005, *Nature*, 434, 1104
- Gal-Yam, A., et al. 2006, *ApJ*, 639, 331
- Gehrels, N., Ramirez-Ruiz, E., Fox, D. B. 2009, *ARA&A*, 47, 567
- Gelbord, J. M., et al. 2011, *GCN Circular*, 11922
- Gershman, A. B., Pesavento, M., Amin, M. G. 2001, *IEEE Transactions on Signal Processing*, 49, 2924
- Gezari, S., et al. 2006, *ApJ*, 653, L25

- Ginzburg, V. L. 1973, *Nature*, 246, 415
- Gould, A. 2005, *New Ast Rev*, 49, 424
- Green, A. J., Cram, L. E., Large, M. I., Ye, T. 1999, *ApJS*, 122, 207
- Gregory, P. C. Kronberg, P. P. 1972, *Nature*, 239, 440
- Gregory, P. C. Taylor, A. R. 1986, *AJ*, 92, 371
- Grimm, H.-J., McDowell, J., Zezas, A., Kim, D.-W., Fabbiano, G. 2007, *ApJS*, 173, 70
- Grupe, D., et al. 2010a, *GCN Circular*, 10929
- Grupe, D., et al. 2010b, *GCN Report*, 293
- Güdel, M. 2002, *ARA&A*, 40, 217
- Hambly, N. C., et al. 2001, *MNRAS*, 326, 1279
- Han, J. L., Manchester, R. N., Berkhuijsen, E. M., Beck, R. 1997, *A&A*, 322, 98
- Hancock, P. P., Gaensler, B. M., Murphy, T. 2011, *ApJ*, 735, L35
- Hankins, T. H. Eilek, J. A. 2007, *ApJ*, 670, 693
- Hankins, T. H., Kern, J. S., Weatherall, J. C., Eilek, J. A. 2003, *Nature*, 422, 141
- Hankins, T. H., Rankin, J. M., Eilek, J. A. 2009, in *astro2010: The Astronomy and Astrophysics Decadal Survey* (<http://mingus.as.arizona.edu/~bjw/astro2010/>), 112
- Hannikainen, D., Hunstead, R., Campbell-Wilson, D. 1998, *NewAR*, 42, 601
- Hansen, B. M. S. Lyutikov, M. 2001, *MNRAS*, 322, 695
- Hardy, G. H. Wright, E. M. 2008, *An introduction to the theory of numbers*, 6th edn. (Oxford University Press)
- Hester, J. J. 2008, *ARA&A*, 46, 127
- Hewish, A. 1975, *Science*, 188, 1079
- Hewish, A., Bell, S., Pilkington, J., Scott, P., Collins, R. 1968, *Nature*, 217, 709
- Hey, J. S., Parsons, S. J., Phillips, J. W. 1946, *Nature*, 158, 234

- Holland, S. T., et al. 2010, GCN Circular, 10884
- Hopkins, A. M., Afonso, J., Chan, B., Cram, L. E., Georgakakis, A., Mobasher, B. 2003, AJ, 125, 465
- Hopkins, A. M., Miller, C. J., Connolly, A. J., Genovese, C., Nichol, R. C., Wasserman, L. 2002, AJ, 123, 1086
- Hotan, A. W., van Straten, W., Manchester, R. N. 2004, PASA, 21, 302
- Hunstead, R. W. 1972, ApL, 12, 193
- Hurley, K. 2011, Advances in Space Research, 47, 1337
- 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., Roy, S., Pal, S., Lazio, T. J. W., Ray, P. S., Kassim, N. E., Bhatnagar, S. 2007, ApJ, 660, L121
- Immler, S., et al. 2010, GCN Circular, 10895
- Inoue, S. 2004, MNRAS, 348, 999
- Inzani, P., Sironi, G., Mandolesi, N., Morigi, G. 1982, in American Institute of Physics Conference Series, Vol. 77, Gamma Ray Transients and Related Astrophysical Phenomena, ed. R. E. Lingenfelter, H. S. Hudson, & D. M. Worrall (Melville, New York, USA: AIP), 79–84
- Jacobsen, E. Lyons, R. 2003, Signal Processing Magazine, IEEE, 20, 74
- Jakobsson, P., et al. 2006, A&A, 447, 897
- Jansky, K. 1933, Proceedings of the Institute of Radio Engineers, 21, 1387
- Janssen, G. H., Stappers, B. W., Braun, R., van Straten, W., Edwards, R. T., Rubio-Herrera, E., van Leeuwen, J., Weltevrede, P. 2009, A&A, 498, 223
- Johnston, S., et al. 2008, Experimental Astronomy, 22, 151
- Jones, D. H., et al. 2004, MNRAS, 355, 747
- Kaplan, D. L., Hyman, S. D., Roy, S., Bandyopadhyay, R. M., Chakrabarty, D., Kassim, N. E., Lazio, T. J. W., Ray, P. S. 2008, ApJ, 687, 262
- Katz, C. A., Hewitt, J. N., Corey, B. E., Moore, C. B. 2003, PASP, 115, 675
- Keane, E. F. Kramer, M. 2008, MNRAS, 391, 2009
- Keane, E. F., Kramer, M., Lyne, A. G., Stappers, B. W., McLaughlin, M. A. 2011, MNRAS, 415, 3065

- Keith, M. J., et al. 2010, MNRAS, 409, 619
- Kellermann, K. I. Pauliny-Toth, I. I. K. 1968, ARA&A, 6, 417
- Kellermann, K. I. Pauliny-Toth, I. I. K. 1969, ApJ, 155, L71
- Kesteven, M. J. Turtle, A. J. 1991, IAU Circ., 5181
- Kesteven, M. J. L., Bridle, A. H., Brandie, G. W. 1977, AJ, 82, 541
- Kirk, J. G. Wassmann, M. 1992, A&A, 254, 167
- Koay, J. Y., et al. 2011, AJ, 142, 108
- Koposov, S. Bartunov, O. 2006, in Astronomical Society of the Pacific Conference Series, Vol. 351, Astronomical Data Analysis Software and Systems XV, ed. C. Gabriel, C. Arviset, D. Ponz, & S. Enrique (Orem, Utah, USA: ASP), 735
- Koranyi, D. M., Green, D. A., Warner, P. J., Waldram, E. M., Palmer, D. M. 1995, MNRAS, 276, L13
- Kouveliotou, C., Meegan, C. A., Fishman, G. J., Bhat, N. P., Briggs, M. S., Koshut, T. M., Paciesas, W. S., Pendleton, G. N. 1993, ApJ, 413, L101
- 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
- Large, M. I., Campbell-Wilson, D., Cram, L. E., Davison, R. G., Robertson, J. G. 1994, PASA, 11, 44
- Law, C. J., Jones, G., Backer, D. C., Barott, W. C., Bower, G. C., Gutierrez-Kraybill, C., Williams, P. K. G., Werthimer, D. 2011, ApJ, in press
- Lazio, T. J. W., Cordes, J. M., de Bruyn, A. G., Macquart, J.-P. 2004, NewAR, 48, 1439
- Lazio, T. J. W., Fey, A. L., Gaume, R. A. 2001, Ap&SS, 278, 155
- Lazio, T. J. W., et al. 2010, AJ, 140, 1995
- Leback, B., Doerfler, D., Heroux, M. 2008, in CUG 2008: Crossing the Boundaries, Cray User Group Conference series (Philomath, Oregon, USA: Cray User Group)
- Lenc, E., Garrett, M. A., Wucknitz, O., Anderson, J. M., Tingay, S. J. 2008, ApJ, 673, 78

- Levan, A. J., Wynn, G. A., Chapman, R., Davies, M. B., King, A. R., Priddey, R. S., Tanvir, N. R. 2006, MNRAS, 368, L1
- Levin, L., et al. 2010, ApJ, 721, L33
- Levinson, A., Ofek, E. O., Waxman, E., Gal-Yam, A. 2002, ApJ, 576, 923
- Lipunov, V. M. Panchenko, I. E. 1996, A&A, 312, 937
- Lorimer, D., Bailes, M., McLaughlin, M., Narkevic, D., Crawford, F. 2007, Science, 318, 777
- Lorimer, D. R., Yates, J. A., Lyne, A. G., Gould, D. M. 1995, MNRAS, 273, 411
- Lorimer, D. R., et al. 2006, MNRAS, 372, 777
- Lovell, J. E. J., et al. 2008, ApJ, 689, 108
- Lutomirski, A., Tegmark, M., Sanchez, N. J., Stein, L. C., Urry, W. L., Zaldarriaga, M. 2011, MNRAS, 410, 2075
- Lyne, A. G., et al. 2004, Science, 303, 1153
- Lyubarsky, Y. 2008, ApJ, 682, 1443
- M. Hoskin, ed. 1997, Cambridge Illustrated History of Astronomy (Cambridge University Press)
- MacFadyen, A. I. Woosley, S. E. 1999, ApJ, 524, 262
- Macquart, J. 2007, ApJ, 658, L1
- Macquart, J., et al. 2010, PASA, 27, 272
- Macquart, J.-P. 2011, ApJ, 734, 20
- Maitia, V., Lestrade, J.-F., Cognard, I. 2003, ApJ, 582, 972
- Manchester, R. N. 1985, PASA, 6, 89
- Manchester, R. N., Hobbs, G. B., Teoh, A., Hobbs, M. 2005, AJ, 129, 1993
- Manchester, R. N., et al. 2001, MNRAS, 328, 17
- Mangano, V., et al. 2010, GCN Circular, 11135
- Matsumura, N., et al. 2009, AJ, 138, 787
- Mauch, T., Murphy, T., Buttery, H. J., Curran, J., Hunstead, R. W., Piestrzynski, B., Robertson, J. G., Sadler, E. M. 2003, MNRAS, 342, 1117

- McLaughlin, M. A., et al. 2006, *Nature*, 439, 817
- McLaughlin, M. 2011, in *American Institute of Physics Conference Series*, Vol. 1379, *American Institute of Physics Conference Series*, ed. E. Göğüş, T. Belloni, Uuml. Ertan (Melville, New York, USA: AIP), 48–55
- McMahon, P. 2008, Master's thesis, Univ. Cape Town
- Metzger, B. D., Giannios, D., Thompson, T. A., Bucciantini, N., Quataert, E. 2011, *MNRAS*, 413, 2031
- Metzger, B. D., Quataert, E., Thompson, T. A. 2008, *MNRAS*, 385, 1455
- Mills, B. Y. 1981, *PASA*, 4, 156
- Moortgat, J. Kuijpers, J. 2005, in *NATO ASIB Proc. 210: The Electromagnetic Spectrum of Neutron Stars*, ed. A. Baykal, S. K. Yerli, S. C. Inam, & S. Grebenev (Dordrecht, The Netherlands: Springer), 103
- Morales, M. F. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 345, *From Clark Lake to the Long Wavelength Array: Bill Erickson's Radio Science*, ed. N. Kassim, M. Perez, W. Junor, & P. Henning, 452
- Morales, M. F. 2008, arXiv:0812.3669
- Morales, M. F., Hewitt, J. N., Kasper, J. C., Lane, B., Bowman, J., Ray, P. S., Cappallo, R. J. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 345, *From Clark Lake to the Long Wavelength Array: Bill Erickson's Radio Science*, ed. N. Kassim, M. Perez, W. Junor, & P. Henning (Orem, Utah, USA: ASP), 512
- Moshir, M., Kopman, G., Conrow, T. A. O. 1992, *IRAS Faint Source Survey*, Explanatory supplement version 2 (ESO)
- Moshir, M. e. 1990, in *IRAS Faint Source Catalogue*, version 2.0
- Murphy, T., Gaensler, B. M., Chatterjee, S. 2008, *MNRAS*, 389, L23
- Murphy, T., Mauch, T., Green, A., Hunstead, R., Piestrzynska, B., Kels, A., Sztajer, P. 2007, *MNRAS*, 382, 382
- Muxlow, T. W. B., et al. 2010, *MNRAS*, 404, L109
- Nakar, E. 2007, *Phys. Rep.*, 442, 166
- Narayan, R. 1992, *Royal Society of London Philosophical Transactions Series A*, 341, 151

- Nice, D. 2007, Pulsar Observations with the GBT: A retrospective, and speculations on future discoveries, ed. Lockman, F. J., Ghigo, F. D., & Balser, D. S. (Charlottesville, Virginia, USA: The National Radio Astronomy Observatory), 435
- Niinuma, K., et al. 2007, *ApJ*, 657, L37
- Noutsos, A., Johnston, S., Kramer, M., Karastergiou, A. 2008, *MNRAS*, 386, 1881
- Nyquist, H. 1928, Transactions of the American Institute of Electrical Engineers, 47, 617
- 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
- O’Sullivan, J. D., Ekers, R. D., Shaver, P. A. 1978, *Nature*, 276, 590
- Palmer, D. M. 1993, *ApJ*, 417, L25
- Perlmutter, S., et al. 1999, *ApJ*, 517, 565
- Pshirkov, M. S. Postnov, K. A. 2010, *Ap&SS*, 330, 13
- Qin, B., Wu, X.-P., Chu, M.-C., Fang, L.-Z., Hu, J.-Y. 1998, *ApJ*, 494, L57
- Rankin, J. M. 1986, *ApJ*, 301, 901
- Rees, M. J. 1977, *Nature*, 266, 333
- Richards, J. W., et al. 2011, *ApJ*, 733, 10
- Rickett, B. J. 1986, *ApJ*, 307, 564
- Riess, A. G., et al. 1998, *AJ*, 116, 1009
- Robertson, J. G. 1991, *Aust. J. Phys.*, 44, 729
- Röttgering, H. 2003, *NewAR*, 47, 405
- Roy, S., Hyman, S. D., Pal, S., Lazio, T. J. W., Ray, P. S., Kassim, N. E. 2010, *ApJ*, 712, L5
- Ryder, S. D., Sadler, E. M., Subrahmanyan, R., Weiler, K. W., Panagia, N., Stockdale, C. 2004, *MNRAS*, 349, 1093
- Sagiv, A. Waxman, E. 2002, *ApJ*, 574, 861

- Sagiv, A., Waxman, E., Loeb, A. 2004, *ApJ*, 615, 366
- Sault, R. J. Killeen, N. 2009, *Miriad User Guide*, Australia Telescope National Facility, Sydney, Australia
- Savaglio, S., Glazebrook, K., Le Borgne, D. 2009, *ApJ*, 691, 182
- Saxton, C. J., et al. 2010, *GCN Circular*, 11357
- Seielstad, G. A., Pearson, T. J., Readhead, A. C. S. 1983, *PASP*, 95, 842
- Shannon, C. 1949, *Proceedings of the IRE*, 37, 10
- Shapirovskaia, N. Y. 1985, *Soviet Astronomy Letters*, 11, 289
- Shibata, M., Suwa, Y., Kiuchi, K., Ioka, K. 2011, *ApJ*, 734, L36
- Smits, R., Kramer, M., Stappers, B., Lorimer, D. R., Cordes, J., Faulkner, A. 2009, *A&A*, 493, 1161
- Soderberg, A. M., Chevalier, R. A., Kulkarni, S. R., Frail, D. A. 2006, *ApJ*, 651, 1005
- Soderberg, A. M., et al. 2010, *Nature*, 463, 513
- Sokoloski, J. L., Luna, G. J. M., Mukai, K., Kenyon, S. J. 2006, *Nature*, 442, 276
- Staelin, D. H. Reifstein, III, E. C. 1968, *Science*, 162, 1481
- Starling, R. L. C., et al. 2011, *GCN Circular*, 11931
- Stern, B. E., Tikhomirova, Y., Kompaneets, D., Svensson, R., Poutanen, J. 2001, *ApJ*, 563, 80
- Strauss, M. A., Huchra, J. P., Davis, M., Yahil, A., Fisher, K. B., Tonry, J. 1992, *ApJS*, 83, 29
- Taylor, G. B., Carilli, C. L., Perley, R. A., eds. 1999, *Astronomical Society of the Pacific Conference Series*, Vol. 180, *Synthesis Imaging in Radio Astronomy II* (Orem, Utah, USA: ASP)
- Taylor, J. H. 1974, *A&AS*, 15, 367
- Tegmark, M. Zaldarriaga, M. 2010, *Phys. Rev. D*, 82, 103501
- Teräsranta, H., et al. 2004, *A&A*, 427, 769
- Trott, C. M., Wayth, R. B., Macquart, J.-P. R., Tingay, S. J. 2011, *ApJ*, 731, 81

- Turtle, A. J., Campbell-Wilson, D., Bunton, J. D., Jauncey, D. L., Kesteven, M. J. 1987, *Nature*, 327, 38
- Usov, V. V. Katz, J. I. 2000, *A&A*, 364, 655
- Van Eck, C. L., et al. 2011, *ApJ*, 728, 97
- van Velzen, S., Körding, E., Falcke, H. 2011, *MNRAS*, L310
- Verheijen, M. A. W., Oosterloo, T. A., van Cappellen, W. A., Bakker, L., Ivashina, M. V., van der Hulst, J. M. 2008, in *American Institute of Physics Conference Series*, Vol. 1035, *The Evolution of Galaxies Through the Neutral Hydrogen Window*, ed. R. Minchin & E. Momjian, 265–271
- Voges, W., et al. 1999, *A&A*, 349, 389
- Walker, M. A. 2007, in *Astronomical Society of the Pacific Conference Series*, Vol. 365, *SINS - Small Ionized and Neutral Structures in the Diffuse Interstellar Medium*, ed. M. Haverkorn & W. M. Goss, 299
- Wayth, R. B., Briske, W. F., Deller, A. T., Majid, W. A., Thompson, D. R., Tingay, S. J., Wagstaff, K. L. 2011, *ApJ*, 735, 97
- Weiler, K. W., Panagia, N., Montes, M. J., Sramek, R. A. 2002, *ARA&A*, 40, 387
- Weiler, K. W., van Dyk, S. D., Sramek, R. A., Panagia, N. 2004, *NewAR*, 48, 1377
- Williams, J. R. 1968, *The Journal of the Acoustical Society of America*, 44, 1454
- Wilson, A. S. Colbert, E. J. M. 1995, *ApJ*, 438, 62
- Wright, A. E., Griffith, M. R., Burke, B. F., Ekers, R. D. 1994, *ApJS*, 91, 111
- Wright, E. L. 2006, *PASP*, 118, 1711
- Yardley, D. R. B., et al. 2010, *MNRAS*, 407, 669
- Zanardo, G., et al. 2010, *ApJ*, 710, 1515

Variables from the MOST Survey

This appendix contains radio light curves of all variable sources shown in Table 2.5, in RA order (Fig. A.1). 843 MHz flux density measurements from MOST are shown as black squares. Flux density measurements or limits from NVSS, where available, are shown as open circles with x errors indicating the observation duration of the NVSS survey. For SUMSS J011019–455112, flux density measurements at 1.4 GHz are shown as open triangles. For SUMSS J223225–615308 the Parkes-MIT-NRAO (PMN) flux density measured at 4.85 GHz (Wright et al., 1994) is shown as an open triangle.

The NVSS flux densities published in Bannister et al. (2011) (Fig. A1) had not been bias corrected. Fig. A.1 contains the bias corrected NVSS flux densities.

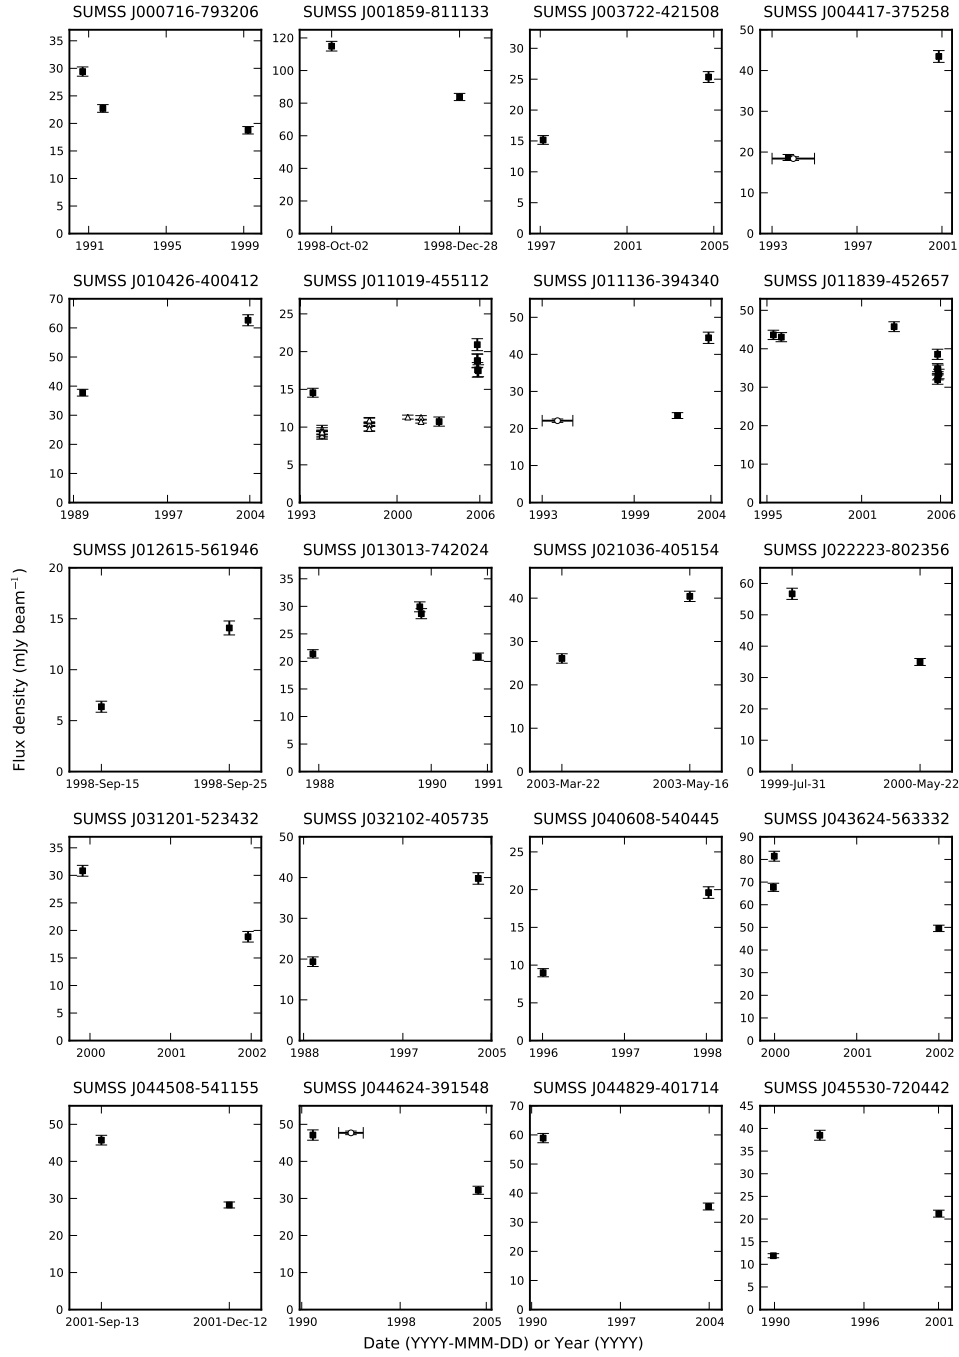


Figure A.1: Light curves of variable sources

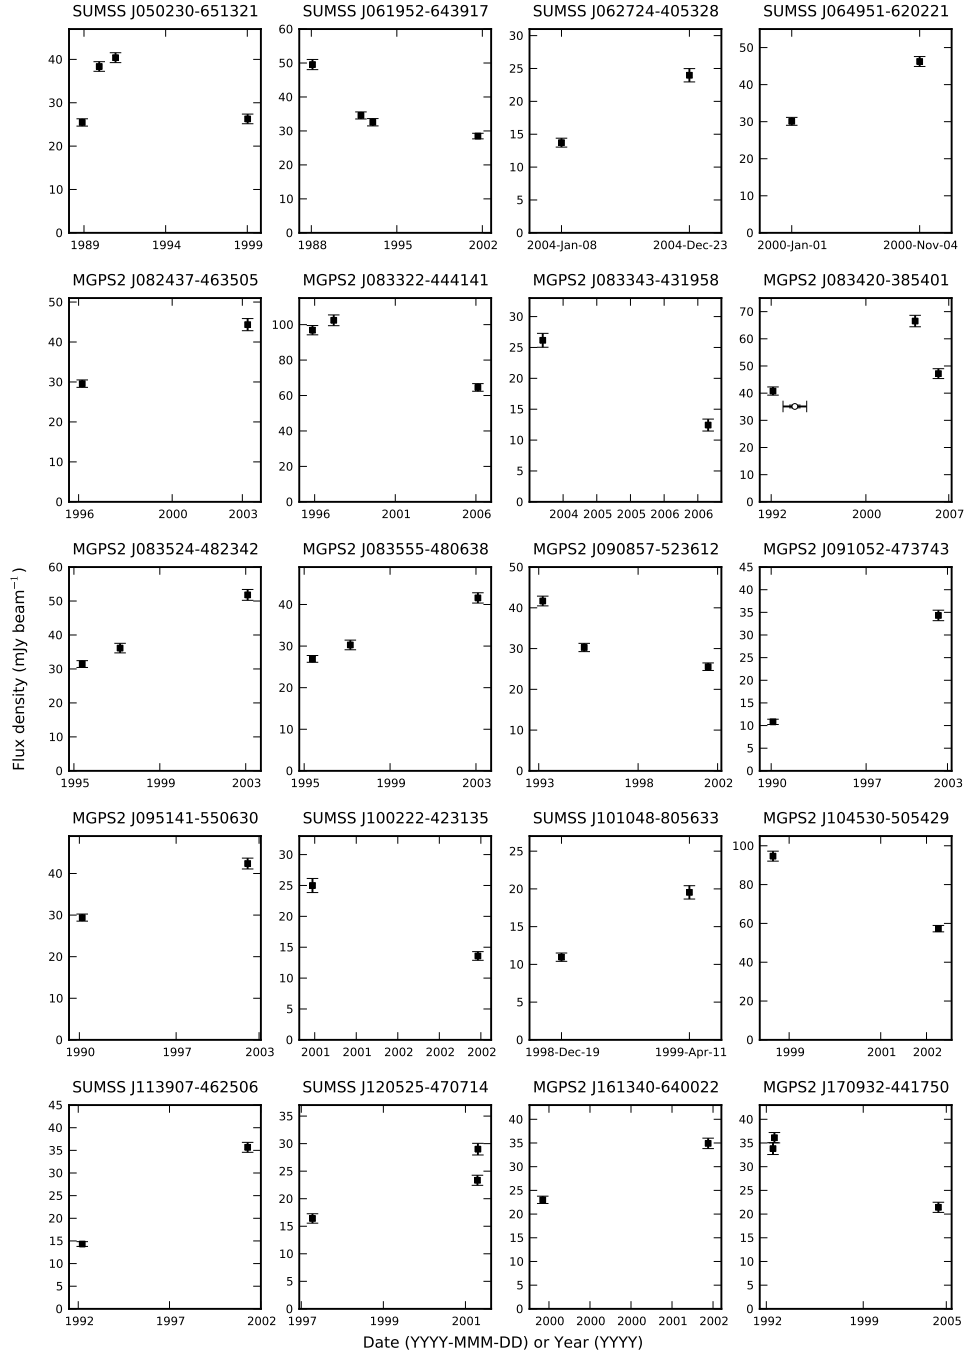


Figure A.1 (continued)

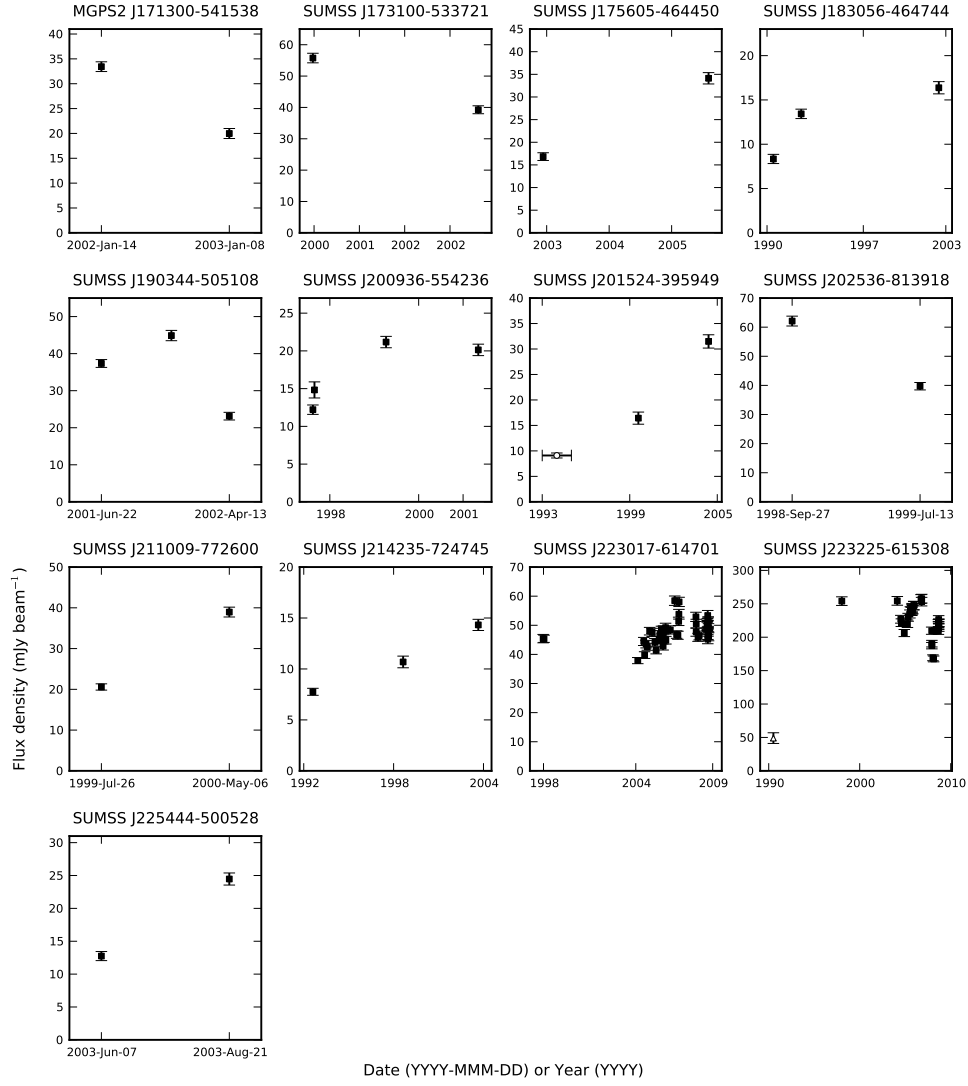


Figure A.1 (continued)

Transients from the MOST survey

This appendix contains radio light curves of all transient sources shown in Table 2.6, in RA order (Fig. B.1). Each source is shown on a single row. The leftmost panel is the radio light curve. 843 MHz flux density measurements from MOST are shown as black squares. Flux density measurements or limits from NVSS, where available, are shown as open circles with x errors indicating the observation duration of the NVSS survey. The centre and right panels show the images of the source at the minimum and maximum detection significance, in time order from left to right, with the position of the source encircled. The grey scale of all images is linear from -5 to 20 mJy beam^{-1} .

Due to an incorrect matching radius used to associate MOST sources with NVSS, some transient sources were described in Bannister et al. (2011) (Fig. B1) as having non-detections in NVSS, and the NVSS flux densities had not been bias corrected. Fig. B.1 contains the bias corrected NVSS flux densities and the correct NVSS identifications.

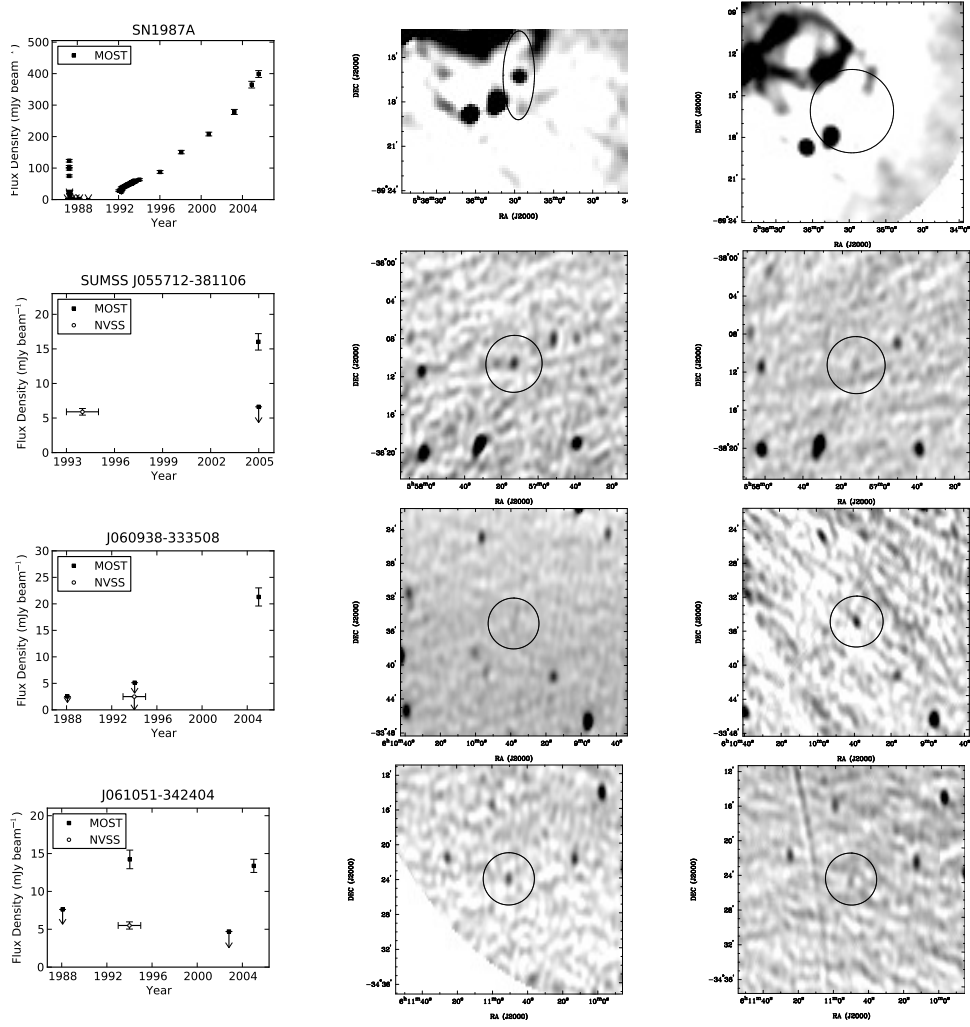


Figure B.1: Transient sources. Left panel: radio light curve. The centre and right panels show the images of the source at the minimum and maximum detection significance, in time order from left to right, with the position of the source encircled

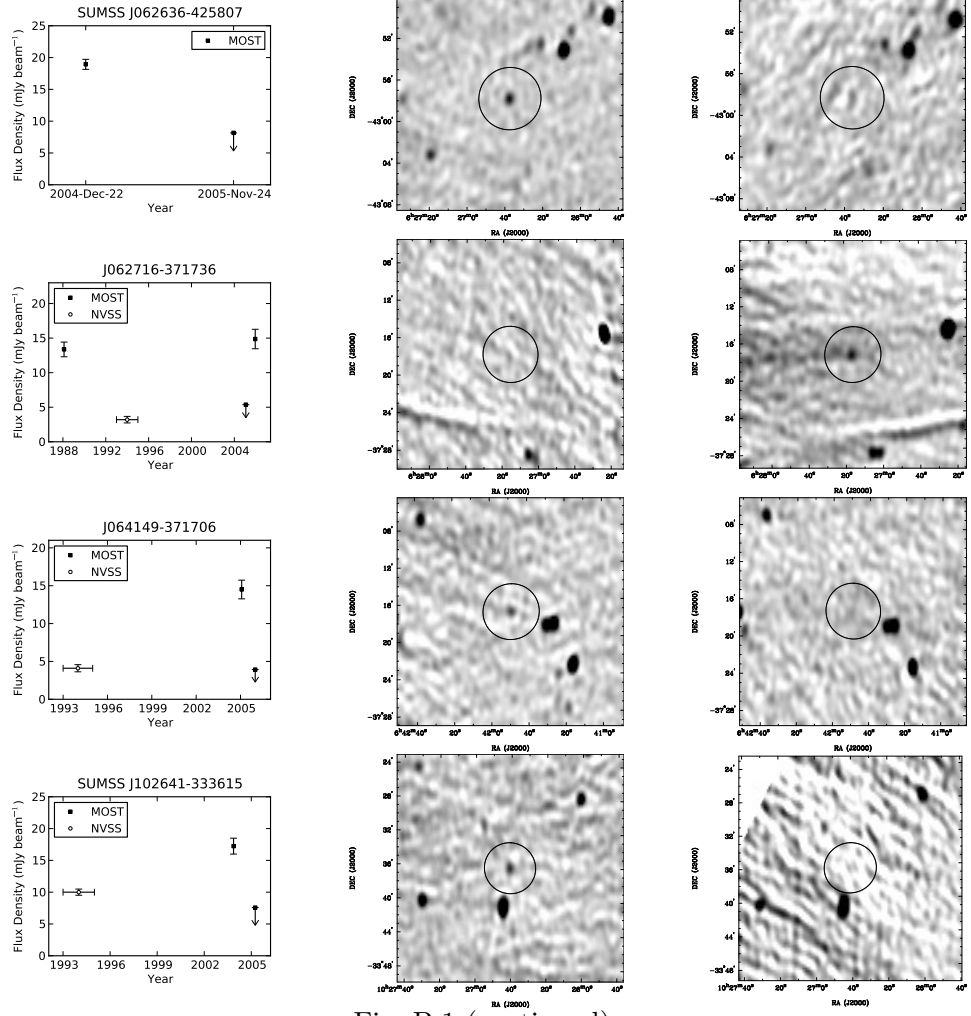


Fig. B.1 (continued).




Fig. B.1 (continued).




Fig. B.1 (continued).

The Chirpolator: analysis and implementation

In this appendix, we describe additional extensions to the Chirpolator to include multiple telescopes in 3D and non linear dispersion. We describe novel methods for efficiently implementing the Chirpolator and also derive equations for the resolution and data and operation rate requirements.

C.1 MULTIPLE TELESCOPES IN 3D

The generalisation to arbitrary arrays of elements in three dimensions is most easily done in the notation of interferometry (Taylor et al., 1999, Chapter 2).

If we measure the $[u, v, w]^T$ baseline vector in units of distance (not wavelength), then the geometric delay for a 3D array is

$$\tau = \frac{ul + vm + w \left(\sqrt{1 - l^2 - m^2} - 1 \right)}{c}, \quad (\text{C.1})$$

where l and m are the direction cosines in the u - and v - directions, respectively. l and m define the angle of interest analogous to θ in the 1D case.

The method of computing the intensity image then proceeds in much the same manner, with k_0 computed with Equations 4.35 and C.1, and with $P(\theta)$ evaluated over two angular dimensions instead of one.

C.2 NON LINEAR DISPERSION

In the main text, beginning at Equation 4.17, we have assumed a linear chirp. In fact, at most frequencies and bandwidths of interest (i.e., below 10 GHz, and bandwidths >100 MHz), the cold plasma dispersion law is much more accurately modelled as $\propto \nu^{-2}$ as shown in Equation 4.1 and Figure 4.1. In this section we describe the effect of the true dispersion law on Chirpolator processing (decoherence), and propose a solution (oversampling).

C.2.1 The Problem: Decoherence in the DFT Bins

To determine the effect of the higher order terms on Chirpolator processing, we begin by considering the frequency of the mixed signal x_{pq} (Equation 4.18), which is the difference between the instantaneous frequencies of the signals from the two antennas. Assuming a delay $\tau \ll T$, the instantaneous frequency difference between the two chirps is given by:

$$\nu_{\text{mix}}(t) = \nu_2(t - \tau) - \nu_2(t) \quad (\text{C.2})$$

$$\simeq -a_1\tau - 2a_2\tau(t - T/2), \quad (\text{C.3})$$

where we have used the Taylor expansion described in Equation 4.5. From Equation C.3 we can see that the effect of nonlinear dispersion on the Chirpolator processing is to smear out the signal across a wider range of frequencies

after mixing the two antenna signals (Figure C.1). The departure of the frequency from the linear assumption is significant for typical array configurations and dispersion (Figure C.2), and is worst far from the phase centre, on the long baselines, and at $t = 0$, where it can be approximated as the difference between the linear approximation and the third-order Taylor series (Equation C.3):

$$\delta_{\text{mix}} = \tau T^2 \left(a_2 - \frac{3}{4} a_3 \tau T \right). \quad (\text{C.4})$$

When a signal with non-constant frequency is passed through a DFT, the amplitude of the DFT output is reduced, which we call decoherence.

We identify three regimes in which the system operates.

- The smearing is $\ll 1$ bin, in which case the decoherence is small and can be ignored.
- The smearing is ~ 1 bin, in which case the signal still occupies only one bin, but the decoherence within that bin is significant. In this case, the DFT must be broken into a number of sub-integrations, with each sub-integration requiring a complex phase rotation to recover the coherence.
- The smearing is > 1 bin, in which case there is energy in multiple bins. The DFTs must be broken into a number of sub-integrations. The final output must be formed with a complex phase rotation of a range of sub-integrations of *different* DFT bins.

If the smearing is > 1 bin, (e.g., Figure C.1), the true dispersion occupies a higher DFT bin than the linear assumption for approximately half the pulse duration. To capture energy from the higher frequencies, additional DFTs must be computed that would not be required under the linear assumption. In the nonlinear case, the maximum number of DFT bins increases from $k_{0,\text{max}}$ to $k_{0,\text{max}} + \delta_{\text{mix}}$, which increases the operations and data rate requirements for the DFT step. At the worst case longest baseline of 1 km, at 1.4 GHz, 400 MHz and 0.5° from the phase centre, $\delta_{\text{mix}} = 7$, and the number of DFT bins required increases by a factor of $\delta_{\text{mix}}/k_{0,\text{max}} = 63\%$. The additional DFT bins increases non-linearly as a function of baseline, so accurately estimating the total increase over the whole array requires a knowledge of the exact baseline distribution. To obtain an approximate figure, assuming a baseline distribution where the mean baseline length is half the maximum baseline length, we propose that the total increase is approximately half the worst case figure, i.e., 32%.




Figure C.1: Simulated mixing frequency as a function of time for a single antenna pair of the Chirpolator (see Equation eq:approxerror1). A range of approximations are shown. The parameters for this simulation were: a DM of $100 \text{ cm}^{-3} \text{ pc}$ and a bandwidth of 400 MHz centred at 1.4 GHz, $\theta = 0.5^\circ$ and a baseline of 1 km.

C.2.2 The solution: Oversampling

The key to handling the nonlinear dispersion, therefore, is to dump the integrator more often than required for the nonlinear case (oversample), and phase correct the results to obtain the coherence again. To quantify the amount of oversampling required where the smearing is > 1 bin we need to quantify the response of the DFT to the mixed, non linearly dispersed signal. As shown in Figure C.1, the frequency of the mixed signal is well approximated by the third-order Taylor expansion described in Equation 4.5. The phase of the mixed signal is, therefore, given by the integral

$$\phi_{\text{mix}}(t) = \int_0^t 2\pi\nu_{\text{mix}}(t')dt' \quad (\text{C.5})$$

$$= 2\pi(-a_1\tau t - a_2\tau t(t-T)) \quad (\text{C.6})$$

$$= 2\pi(t(-a_1\tau + a_2\tau T) - t^2a_2\tau). \quad (\text{C.7})$$




Figure C.2: Error in mixing frequency δ_{mix} as a function of centre frequency and bandwidth (B), assuming DM of $100 \text{ cm}^{-3} \text{ pc}$, 1 km baseline and $\theta = 0.5^\circ$.

The sampled, mixed signal can then be expressed as

$$m[n] = \exp(j\phi_{\text{mix}}[n]) \quad (\text{C.8})$$

$$= \exp \left[2\pi j \left(\frac{n}{f_s} (-a_1\tau + a_2\tau T) - \left(\frac{n}{f_s} \right)^2 a_2\tau \right) \right] \quad (\text{C.9})$$

and we take the DFT over $N = f_s T$ samples to obtain

$$X_m[k] = \sum_{n=0}^{N-1} \exp(-2\pi j k n / N) m[n] \quad (\text{C.10})$$

$$= \sum_{n=0}^{N-1} \exp 2\pi j \left[n(-a_1\tau + a_2\tau T - k/N) + n^2(-a_2\tau/f_s^2) \right] \quad (\text{C.11})$$

The term that is linear with n has already been dealt with in Equation 4.27 and is simply the DFT of a single tone, so we turn our attention to the n^2 term and define the sum:

$$G(a, L) = \sum_{n=0}^{L-1} \exp(2\pi j a n^2), \quad (\text{C.12})$$

where:

$$a = -a_2 \tau / f_s^2. \quad (\text{C.13})$$

From Equation C.12 it is clear that $G(0, L) = L$, and that for non-zero values of a , the $a n^2$ term introduces oscillations, effectively moving the instantaneous frequency into the adjacent DFT bins, so that $|G(a, L)| < L$ for non-zero a .

We want to determine how large L can be made before some fraction of the energy will be lost to adjacent DFT bins. Equation C.12 defines the result of summing a chirp with an initial instantaneous frequency of zero, which is essentially the centre of the DFT bin. We define the coherence loss or loss in amplitude as

$$\eta = |G(a, L)| / L. \quad (\text{C.14})$$

Therefore, the value of L that maintains a required η is the number of samples to traverse half the DFT bin and maintain a given loss. To calculate the required oversampling for chirp that crosses an entire DFT bin, we can pose the question: what oversampling factor $\kappa_t = N/2L$ is required to maintain $\eta(a, L)$ above a specified threshold?

To get an approximation of the required oversampling factor, we have simulated a typical case for the SKA, with a one-sided frequency smearing of the order of seven DFT bins, which is reasonably large in the context of Figure C.2. We conclude that four times oversampling yields $\eta \simeq 99\%$ (Figure C.3). We have found empirically that the required oversampling is independent of DM.

Because the time of arrival is not known, a one would typically require $\simeq 4$ times oversampling to obtain a sample which is integrated over a large fraction of the incoming signal. The equivalence of the oversampling rates required for time oversampling, and nonlinear dispersion correction, implies that nonlinear dispersion does not substantially drive the oversampling in this instance.

C.3 IMPLEMENTATION OPTIMISATIONS

In the main text we assumed a single value of \dot{f} (equivalently a single value of the DM), and that the DFT window is exactly time-aligned with the chirp. In practise, neither the time of arrival for the chirp nor the \dot{f} are known in




Figure C.3: One method for handling nonlinear dispersion with the Chirpulator is to increase the oversampling rate. Above is the required oversampling rate (κ_t) for a nonlinear chirp vs. the geometric delay (τ) for single a baseline operating at $f_c = 1.4$ GHz, $B = 400$ MHz, and a range of different coherence losses (η). The vertical dashed line is the geometric delay for $\theta = 0.5^\circ$ and a baseline of 1 km. The required oversampling rate is independent of DM.

advance and we would like to maximise our chances of finding the signal. The maximum likelihood approach to the problem of maximising the detection probability when the waveform parameters are unknown, is to pass the signal through many different matched filters, each with a particular realisation of the unknown parameters. In our case, we would evaluate $P(\theta)$ and $X_{pq}[k]$ independently on a range of values of \hat{f} and on a set of overlapping windows in time.

Significant computational savings can be made as described in the following sections.

C.3.1 Compute only DFTs Required on a baseline basis

We do not have to compute the same number of DFT bins for each pair of antennas. In fact, for a given pair of antennas, we only have to compute the DFT for values of k up to approximately $k_{0,\max}$, as illustrated in Figure

4.4. With values from typical radio telescopes, $k_{0,max} \simeq 100$ for the longest baselines and $k_{0,max} \simeq 5$ for the shortest baselines. If the baseline distribution is such that the average baseline is half the maximum baseline, this strategy saves a factor of two in DFT operations and data rate.

C.3.2 Efficient calculation of $X_{pq}[k]$ with sliding DFTs

We consider problem of computing DFT values for overlapping time windows. For typical array configurations and DMs, the number of samples in the DFT (N) is of order 10^5 , whereas the number of usable DFTs is of the order $k_{0,max} \simeq 10^2$, meaning that computing a full FFT would result in a very large number of unused bins. In addition, we do not require a DFT result every sample, which means a sliding window DFT result every $L < N$ samples is adequate.

A naive method to computing the sliding window DFT is to (1) compute the dot product of N input samples with a complex sinusoid of appropriate frequency, then (2) shift the input sequence by $L < N$ samples, and (3) compute the dot product on the shifted samples, with the same complex sinusoid. This naive method requires N complex multiplications per L samples, per DFT bin, and corresponds to an operation rate of roughly $f_s N/L$ per DFT bin.

Jacobsen & Lyons (2003) describe a “the sliding DFT,” a more efficient method for computing a small number of DFT bins in a sliding window manner. The sliding DFT is a recursive filter that produces a sliding window DFT output according to

$$S_k[n] = S_k[n-1] \exp(-2\pi j k/N) - x[n-N] + x[n] \quad (\text{C.15})$$

where $S_k[n]$ is the sliding window DFT output for sample n and bin k , and $x[n]$ is the sampled input sequence. Equation C.15 is effectively a moving average filter implemented as a Cascaded Integrator Comb (CIC), with a complex resonator embedded in the integrator feedback path. The sliding DFT has an operation rate of only $\simeq 3f_s$ per DFT bin, which is significantly less than that required for the naive method.

In practise, we do not require an output every sample, so the operation rate can be further reduced by computing a block-based sliding DFT. In this case, we compute the partial DFTs, time-indexed by m in blocks of L samples:

$$V_k[m] = \sum_{n=0}^{L-1} \exp(-2\pi j(n+mL)k/N) x[n+mL] \quad (\text{C.16})$$

and form the DFT over the full number of samples N by applying a moving average filter on the partial DFTs:

$$S_k[m] = S_k[m-1] + V_k[m] - V_k[m-N/L]. \quad (\text{C.17})$$




Figure C.4: The DFT can be efficiently computed in a sliding window manner (see Equation C.17). For a DFT bin number k , the current value of the DFT bin ($S_k[m]$) is formed by taking the previous value of the DFT bin ($S_k[m-1]$), adding the most recent partial DFT ($V_k[m]$) and subtracting the oldest partial DFT ($V_k[m - N/L]$).

This method is illustrated in Figure C.4.

The block-based sliding DFT has a lower operation rate than the sliding DFT, because the moving average (CIC) stage (Equation C.17) operates at the block rate, rather than the sample rate, resulting in an operation rate of $\simeq f_s + 2f_s/L = f_s(1 + 2/L) \simeq f_s$.

C.3.3 Factorising the DFTs

As the DM and therefore the value of \dot{f} is unknown, a search in \dot{f} is required to maximise S/N of pulse at unknown DM. This search through \dot{f} is equivalent to varying the size of the DFT: N . One might choose to use a bank of DFTs, each with a length of $N = N_0 d$, where N_0 is the length of the DFT corresponding to the shortest DM of interest and d is a positive integer. If we have a bank of DFTs, each starting at the most recent sample and extending back in time by N samples, then we can factorise some of the DFTs by noting that some of the basis functions for the long DFTs can be formed by concatenating the basis functions for the short DFTs. By way of example, the result of the S_2 bin for the length $2N_0$ window can be trivially computed by summing the adjacent, non-overlapping results of S_1 over the length N_0 window, as illustrated in Figure C.5.

If we write $S_k^N[n]$ as the DFT result for bin k at sample time¹⁷ n for a length N DFT, we can say that a DFT of length $N_0 d$ can be computed from the sum of D shorter length $N_0 d/D$ DFTs, if it can be written as:

$$S_k^{N_0 d}[n] = \sum_{d'=0}^{D-1} S_{k/D}^{N_0 d/D}[n - d' N_0] \quad (\text{C.18})$$

where k and d are integers.

The bin $S_k^{N_0 d}$ can be factorised if and only if d/D and k/D are integers, that is d and k must have a common, non-unity factor D which implies that d and k cannot be co-prime. The probability of two random integers being co-prime is approximately 61% (Hardy & Wright, 2008), which implies that approximately 39% of DFT bins can be factorised. If the shorter DFT results are already available, computing the factorised DFT requires $D-1$ operations, which is trivial in comparison with N operations to compute the full DFT. As a result, DFT factorisation saves approximately 39% in complex operations.

C.3.4 Efficiently computing negative DFT bins

Computing the DFT output requires the multiplication of the complex input sample with the complex exponential. On a standard computer, the complex numbers are stored as real and imaginary parts, and the complex multiplication is performed in the following way:

$$r_k[n] = x[n] \exp(-2\pi j k n / N) \quad (\text{C.19})$$

$$= (a + jb)(c - jd) \quad (\text{C.20})$$

$$= ac + bd + j(-ad + bc) \quad (\text{C.21})$$

¹⁷In practise, one would compute the factorised DFTs on the sliding DFT block outputs indexed by m . We have kept the full sample rate n here for clarity.




Figure C.5: Calculation of some DFTs can be factorised into the sum of two adjacent DFT results. In this example, we illustrate how to calculate the $k = 2$ bin of the length $2N_0$ DFT ($S_2^{2N_0}[n]$) by adding the results of two, adjacent, non-overlapping $k = 1$ bins of the length N_0 DFT ($S_1^{N_0}[n] + S_1^{N_0}[n - N_0]$). In the notation of Appendix C.3.3, this example corresponds to $k = d = D = 2$.

where $r_f[k]$ is the result of the multiplication of the input sample with the sinusoid of frequency k , a and b are the real and imaginary parts of the complex input sample, and c and d are the real and imaginary parts of the complex sinusoid.

Assuming the phase centre is set to the centre of the primary beam, the DFTs must be computed for frequencies over the range $[-k_{0,\max}, k_{0,\max}]$ to cover the full field of view. Therefore, each positive bin has a negative counterpart. To compute the negative frequency, we could also separately calculate

$$r_{-k}[n] = x[n] \exp(2\pi jkn/N) \quad (\text{C.22})$$

$$= (a + jb)(c + jd) \quad (\text{C.23})$$

$$= ac - bd + j(ad + bc) \quad (\text{C.24})$$

The calculation of $r_k[n]$ and $r_{-k}[n]$ naively requires 12 operations (eight multiplications and four additions). But the multiplications are common between the two results (Equations C.21 and C.24), which means both results can be computed with only eight operations (four multiplications and four additions). This results in a saving of 33% over the naive implementation.

C.4 PERFORMANCE

C.4.1 Resolution

From the definitions of Equations 4.38 and 4.43, the amplitude of the response to a chirp on a given baseline is $|P(\theta)|^2 = D_N^2(x)$. Therefore we approximate the spatial resolution of the ‘‘The Chirpolator’’ as the full width half maximum (FWHM) of $D_N^2(x)$ on the longest baseline. The FWHM of $D_N^2(x)$ is defined by

$$D_N^2(2x_{\text{FWHM}}) = \frac{1}{2} D^2(0) \quad (\text{C.25})$$

$$\implies \frac{\sin(2\pi x_{\text{FWHM}})}{\sin(2\pi x_{\text{FWHM}}/N)} = \frac{1}{\sqrt{2}} N \quad (\text{C.26})$$

Taking the third-order Taylor expansion of the sin terms, and solving for the non-trivial solutions of x , we obtain

$$x_{\text{FWHM}} = \frac{1}{\pi} \sqrt{\frac{6(1 - 1/\sqrt{2})}{1 - 1/N^2\sqrt{2}}} \quad (\text{C.27})$$

$$\simeq \frac{1}{\pi} \sqrt{6(1 - 1/\sqrt{2})} \quad (\text{C.28})$$

$$\simeq 0.844. \quad (\text{C.29})$$

To convert Equation C.29, which is the width of the main lobe in units of DFT bins, to an angle, we rearrange Equation 4.36, which yields

$$\sin \theta = \frac{ck_0}{b_{pq}B} \quad (\text{C.30})$$

and by applying the small angle formula and substituting Equation C.29 as the DFT bin resolution (i.e., $\Delta k_0 = x_{\text{FWHM}}$), we obtain the spatial resolution of the Chirpolator:

$$\Delta \theta = \Delta k_0 \frac{c}{Bb_{\max}} \quad (\text{C.31})$$

$$= 0.844 \frac{c}{Bb_{\max}}. \quad (\text{C.32})$$

C.4.2 Operation rates

We compute the number of operations required to form images of the full field of view of a telescope comprised of parabolic dishes. We keep to the convention of Cordes (1997) of counting complex operations, where a complex multiplication and accumulation is considered a single operation. As such, we have not accounted for the 33% saving in floating point operations for the DFT as described in Appendix C.3.4. Also, for clarity, we have not included the additional DFTs required to support nonlinear dispersion (δ_{mix} , see Appendix C.2), as this substantially complicates the analysis, is only significant on the longest baselines and roughly balances the 33% saving described above.

To begin, we assume the half-width beam of a parabolic reflector, at 25% of the peak amplitude is (Cordes, 1997)

$$\theta_{\max} = 0.585 \frac{\lambda}{D}. \quad (\text{C.33})$$

The full width of the beam at 25% amplitude is $2\theta_{\max}$ and we set the phase centre to the centre of the primary beam. If we compute only the required DFT bins (as described in Appendix C.3.1), the number of DFT bins that must be computed for a single DM over all baselines and the full primary beam is given by

$$N_{\text{DFT}} = \sum_{p=0}^{M-1} \sum_{q=0, q \neq p}^{M-1} \frac{B}{c} b_{pq} 2 \sin \theta_{\max} \quad (\text{C.34})$$

$$= 2 \frac{B}{c} \sin \theta_{\max} \sum_{p=0}^{M-1} \sum_{q=0, q \neq p}^{M-1} b_{pq} \quad (\text{C.35})$$

$$= 2 \frac{B}{c} \sin \theta_{\max} \frac{M}{2} (M-1) \bar{b} \quad (\text{C.36})$$

$$\simeq \frac{B}{c} \theta_{\max} M^2 b_{\max}, \quad (\text{C.37})$$

where we have employed the small angle formula for \sin and assumed a distribution of baselines such that the average baseline length is approximately half of the maximum baseline length. Each of these DFT bins requires $\simeq f_s = B$ operations per second (using the block-based sliding DFT, see Appendix C.3.2), and assuming we measure N_{DM} dispersion measures then we require N_{DM} different DFT banks. The operation rate for the DFT step is, therefore, given by

$$\dot{C}_{\text{DFT}} = BN_{\text{DFT}}N_{\text{DM}}\rho_f N_{\text{pol}} \quad (\text{C.38})$$

$$\simeq 0.585 \frac{\lambda}{D} \frac{B^2}{c} M^2 b_{\text{max}} N_{\text{DM}} \rho_f N_{\text{pol}} \quad (\text{C.39})$$

where $\rho_f = 0.61$ is a factor to account for factorising the DFTs across the DM banks as described in Appendix C.3.3 and N_{pol} is the number of polarisations.

To form an image for a given DM, a dot product with the truncated amplitude response function, across all baselines must be performed for every pixel. The number of pixels in an image is

$$N_{\text{pix}} = \left(\kappa_s \frac{2\theta_{\text{max}}}{\Delta\theta} \right)^2 \quad (\text{C.40})$$

$$= 1.92 \left(\kappa_s \frac{Bb_{\text{max}}\lambda}{cD} \right)^2 \quad (\text{C.41})$$

For each pixel, we require a dot product with the response function per baseline, implying

$$N_{\text{ops-per-pixel}} = (2F + 1)M^2. \quad (\text{C.42})$$

An image is formed per DM bank at a rate given by $T_i \kappa_t$ where T_i is the dispersion delay associated with the i th DM of interest, and $\kappa_t \geq 1$ is the time over sampling factor. We choose set of DM banks that is a geometric progression¹⁸ according to

$$T_i = T_0(1 + \epsilon)^i \quad 0 \leq i < N_{\text{DM}} \quad (\text{C.43})$$

with $\epsilon < 1$ an overlap factor which can be chosen by a trade off between computation and S/N. The number of DM banks required to cover the range of DMs from T_0 to T_{max} is given by

$$N_{\text{DM}} \simeq \frac{\log(T_{\text{max}}/T_0)}{\log(1 + \epsilon)} \quad (\text{C.44})$$

¹⁸We can choose a geometric progression for the DM bank lengths, and, when $N_i > N_0$, round N_i to an integer multiple of N_0/L to take maximum advantage from factorisation as required in Appendix C.3.3.

and images are produced at a rate

$$\dot{N}_{\text{image}} = \sum_{i=0}^{N_{\text{DM}}-1} \frac{\kappa_t}{T_i} \quad (\text{C.45})$$

$$= \sum_{i=0}^{N_{\text{DM}}-1} \frac{\kappa_t}{T_0(1+\epsilon)^i} \quad (\text{C.46})$$

$$= \frac{\kappa_t}{T_0} \sigma_{\text{DM}} \quad (\text{C.47})$$

where

$$\sigma_{\text{DM}} = \frac{1 - (1+\epsilon)^{-N_{\text{DM}}+1}}{1 - (1+\epsilon)^{-1}} \quad (\text{C.48})$$

Finally, the operation rate for the imaging step is

$$\begin{aligned} \dot{C}_{\text{img}} &= N_{\text{pix}} N_{\text{ops-per-pixel}} \dot{N}_{\text{image}} N_{\text{pol}} \\ &\simeq 1.92 \left(\kappa_s \frac{B b_{\text{max}} \lambda}{c D} \right)^2 (2F+1) M^2 \frac{\kappa_t}{T_0} \sigma_{\text{DM}} N_{\text{pol}} \end{aligned}$$

C.4.3 Data rates

The DFT step takes in voltages for all antennas and produces N_{DFT} outputs sufficient to produce images at a rate of \dot{N}_{image} . The data rate between the DFT step and the imaging step is therefore

$$R_{\text{DFT-out}} = N_{\text{DFT}} \rho_f \dot{N}_{\text{image}} N_{\text{pol}} N_{\text{bytes-per-DFT-bin}}. \quad (\text{C.49})$$

The data rate at the output of the imaging is

$$R_{\text{img-out}} = \dot{N}_{\text{image}} N_{\text{pix}} N_{\text{pol}} N_{\text{bytes-per-pixel}}. \quad (\text{C.50})$$

The Chimageator: analysis and implementation

In this appendix we describe methods for efficiently implementing the Chimageator. We also derive equations for the resolution and operations, and data rate requirements.

D.1 NON LINEAR DISPERSION

The Chimageator is also affected by non linear dispersion. As with The Chirpolator (see Appendix C.2), the mixing frequency in the non linear case is no longer constant resulting non linear trajectories in n - k space and higher $k_{0,\max}$. The solutions may be to increase the time oversampling κ_t and spatial oversampling κ_s with additional phase correction to recover coherence.

As the Chimageator is not really very competitive in the near term (Figure 4.6) we leave a detailed treatment for a future paper.

D.2 IMPLEMENTATION OPTIMISATIONS

D.2.1 Optimising operations in the imaging step

In order to form an image with the Chimageator, we need not sum across all possible trajectories. From Figure 4.5, one can see that there is a family of arrival angles and dispersions with the same gradient, but whose durations, T , differ. For trajectories on the same gradient, the results for all chirp durations can be computed by cumulative sum, i.e., the $P_{T_i}(\theta_i)$ can be computed recursively from the result $P_{T_{i-1}}(\theta_j)$ where T_i , T_{i-1} , θ_i and θ_j are chosen to have the same gradient in n - k space. The gradient of the trajectory is given by

$$\dot{k}_0(T, \theta) = \frac{1}{T} \frac{b_{\max}}{c} \sin(\theta). \quad (\text{D.1})$$

The requirement for a shorter trajectory to be calculated from a longer trajectory implies

$$\dot{k}_0(T_i, \theta_i) \simeq \dot{k}_0(T_{i-1}, \theta_j) \quad (\text{D.2})$$

$$\implies T_i \simeq T_{i-1} \frac{\sin \theta_i}{\sin \theta_j} \quad (\text{D.3})$$

where the equivalence of the gradient can be traded depending on S/N and computational requirements.

Because of this recursive property, there are a many fewer independent calculations required to search through the DM (equivalent to the T_i) and θ space than one might naively expect. We require only enough operations to calculate the trajectories that end on the rectangle bounded by $N_{\max} = f_s T_{\max}$ in the n axis, and $k_{0,\max} = f_s b_{\max} \sin \theta_{\max} / c$ in the k axis, as shown in Figure

4.5. All shorter trajectories can be obtained as partial sums of calculation of the longer trajectory with the same (or similar) gradient.

This optimisation works because all trajectories are linear, which means all the shorter trajectories can be constructed from partial sums of a single longer trajectory with the same gradient. When considering the non linear dispersion, the trajectories are no longer linear and a short trajectory does not lie along the path of a single long trajectory. One possible solution is to consider the trajectories as piece-wise linear. The shorter trajectories can then be constructed from the piece-wise partial sums over a *number* of long trajectories. We leave a detailed treatment of this approach to a future paper.

D.2.2 Sampling

For typical interferometers and dispersions, the gradient of the trajectory is reasonably small, which implies an integrate-and-dump operation after the gridding and FFT step can reduce the required data volumes and downstream processing requirements. If we assume that a trajectory of duration T is sampled κ_t times, and assume the oversampling is proportional to the final DFT bin,

$$\kappa_t = k_{0,\text{end}}\kappa_{t,0} \quad (\text{D.4})$$

$$= f_s \frac{b_{\text{max}}}{c} \sin(\theta) \kappa_{t,0} \quad (\text{D.5})$$

then we can produce a sequence of integrated samples indexed by m :

$$Y_k[m] = \sum_{n=0}^{f_s T / \kappa_t} X_k[n + m f_s T / \kappa_t], \quad (\text{D.6})$$

from which we can form the intensity image in Equation 4.66 in much the same way, but at a reduced rate.

A long integration time will smear out the signal and result in a loss of coherence, which results in a practical limit for how small $\kappa_{t,0}$ can be made. Figure D.1 illustrates the amount of coherence loss which is achieved for a given value of $\kappa_{t,0}$. This plot suggests that oversampling the trajectory over five times during a its duration is sufficient to keep the coherence above 95%.

D.2.3 Spectral smearing

As with The Chirpolator, when the instantaneous frequency of the chirp is between discrete DFT bins, the energy is spread out along all DFT bins. To recover some S/N in this case, we sum along $2F$ additional terms in the frequency direction (effectively widening the trajectory) to improve the S/N, in much the same way as described in Section 4.3.4.




Figure D.1: The Chimageator can operate with non linear dispersion as long as a sufficiently high oversampling rate is chosen. Above is the required oversampling factor ($\kappa_{t,0} = \kappa_t/k_{0,\text{end}}$) as a function of trajectory gradient (\dot{k}), to maintain a range of coherence loss levels. The simulated array had centre frequency 1.4 GHz, bandwidth 3 MHz, 4 antennas and 500 m spacing. The input signal was a linear chirp with dispersion delay corresponding to $20 \text{ cm}^{-3} \text{ pc}$. The trajectory gradient was calculated for 100 arrival angles from 0 to 3° . The increase at small gradients is due to $k_{0,\text{end}} \simeq 0$ for small gradients but $\kappa_t = 1$.

D.3 PERFORMANCE

D.3.1 Resolution

The Chirpolator and Chimageator have the same resolution characteristics. This is demonstrated by considering $k_{0,\text{max}}$ for the two methods. With the Chimageator, $k_{0,\text{max}}$ occurs when the sample number is the final sample of the chirp, i.e., $n = f_s T = BT$. For The Chirpolator, $k_{0,\text{max}}$ occurs on the maximum baseline, b_{max} . In either case, it has a value

$$k_{0,\text{max}} = B \frac{b_{\text{max}}}{c} \sin(\theta_{\text{max}}) \quad (\text{D.7})$$

and the resolution is given by Equation C.32

D.3.2 Operation rates

The first steps in Chimageator processing are the multiplication and gridding stages. Until now, we have assumed a uniform linear array, which makes gridding reasonably straightforward. For more complex geometries, a larger gridding kernel is required. A trade between the size of gridding support and the quality of the images is outside the scope of this paper¹⁹, but for dimensioning purposes, one can consider a 7×7 pixel grid kernel, resulting in $N_{\text{ops-per-grid-point}} \simeq 50$ including the multiplication of the voltages from the two antennas.

A grid point must be formed from each pair of antennas at the sampling rate, resulting in an operation rate for gridding of:

$$\dot{C}_{\text{gridding}} = f_s \frac{M}{2} (M - 1) N_{\text{ops-per-grid-point}} \quad (\text{D.8})$$

Assuming a spatial oversampling of κ_s the number of pixels in the grid plane is:

$$N_{\text{pix}} = (2k_{0,\text{max}}\kappa_s)^2 \quad (\text{D.9})$$

The operation rate for the spatial FFT step is

$$\dot{C}_{\text{FFT}} = f_s N_{\text{pix}} \log_2 N_{\text{pix}} \quad (\text{D.10})$$

and the operation rate for the integration step is

$$\dot{C}_{\text{int}} = f_s N_{\text{pix}}. \quad (\text{D.11})$$

The total pre-integrator operation rate is therefore

$$\dot{C}_{\text{pre-int-total}} = \dot{C}_{\text{gridding}} + \dot{C}_{\text{FFT}} + \dot{C}_{\text{int}} \quad (\text{D.12})$$

The total pre-integrator operation rate is dominated by the FFT for sparse arrays, while for dense arrays, it is dominated by the gridding.

To compute the data and operation rates of imaging, we begin by assuming that we use N_{DM} logarithmically spaced set of trial DMs as described in Equations C.43 and C.44, and that each trajectory is sampled at the rate

$$R_i = \frac{\kappa_t}{T_i} \quad (\text{D.13})$$

The computations are then broken into the two types of trajectory shown in Figure 4.5: the trajectories with fixed angle θ_{max} and variable T_i , and the trajectories with fixed time T_{max} and variable angle, θ_i . In each case, a single

¹⁹For detecting a single source in the field, which is a likely case in our example, gridding is not required at all. We include it this section for completeness

trajectory requires $\kappa_t(2F + 1)$ operations per integration step. Therefore, the computation rate for a single trajectory is

$$\dot{C}_{\text{traj}} = R_i \kappa_t(2F + 1) \quad (\text{D.14})$$

$$= \frac{\kappa_t^2(2F + 1)}{T_i} \quad (\text{D.15})$$

$$= \frac{k_{0,\text{end}}^2 \kappa_{t,0}^2(2F + 1)}{T_i} \quad (\text{D.16})$$

We now consider the total operation rate for the 1D case and assuming half the beamwidth.

The fixed angle trajectories have fixed $k_{0,\text{end}} = k_{0,\text{max}}$, and variable T_i , resulting in an operation rate of

$$\dot{C}_{\theta_{\text{max}}} = \sum_{i=1}^{N_{\text{DM}}} \dot{C}_{\text{traj}} \quad (\text{D.17})$$

$$= \sum_{i=1}^{N_{\text{DM}}} \frac{k_{0,\text{end}}^2 \kappa_{t,0}^2(2F + 1)}{T_i} \quad (\text{D.18})$$

$$= k_{0,\text{max}}^2 \kappa_{t,0}^2(2F + 1) \sum_{i=1}^{N_{\text{DM}}} \frac{1}{T_i} \quad (\text{D.19})$$

$$= k_{0,\text{max}}^2 \kappa_{t,0}^2(2F + 1) \sigma_{\text{DM}} \quad (\text{D.20})$$

where σ_{DM} is defined in Equation C.48.

The fixed time trajectories, have fixed $T_i = T_{\text{max}}$ and variable $k_{0,\text{end}}$, resulting in an operation rate of:

$$\dot{C}_{T_{\max}} = \sum_{k=0}^{k_{0,\max}\kappa_s} \dot{C}_{\text{traj}} \quad (\text{D.21})$$

$$= \sum_{k=0}^{k_{0,\max}\kappa_s} \frac{k_{0,\text{end}}^2 \kappa_{t,0}^2 (2F+1)}{T_i} \quad (\text{D.22})$$

$$= \sum_{k=0}^{k_{0,\max}\kappa_s} \frac{(k/\kappa_s)^2 \kappa_{t,0}^2 (2F+1)}{T_i} \quad (\text{D.23})$$

$$= \frac{1}{T_{\max}} \frac{\kappa_{t,0}^2}{\kappa_s^2} (2F+1) \sum_{k=0}^{k_{0,\max}\kappa_s} k^2 \quad (\text{D.24})$$

$$= \frac{1}{6T_{\max}} \frac{\kappa_{t,0}^2}{\kappa_s} (2F+1) k_{0,\max} (\kappa_s k_{0,\max} + 1) (2\kappa_s k_{0,\max} + 1) \quad (\text{D.25})$$

$$\simeq \frac{1}{3T_{\max}} \kappa_{t,0}^2 \kappa_s (2F+1) k_{0,\max}^3 \quad (\text{D.26})$$

where κ_s is the desired spatial oversampling. The total operation rate is the sum of the two sets in the 1D case. In the 2D case, the computation is squared, so that the total operation rate in two dimensions, for the full beamwidth is

$$\dot{C}_{\text{Total}} = \left(2(\dot{C}_{\theta_{\max}} + \dot{C}_{T_{\max}}) \right)^2 \quad (\text{D.27})$$

D.3.3 Data rates

Assuming that this integrate-and-dump operates at the highest rate $R_0 = \kappa_t/T_0$ and the longer integrations can be formed from the short integrations in the imaging step, then the data rate at the output of the integrate and dump step is

$$R = N_{\text{pix}} N_{\text{bytes-per-pix}} R_0 \quad (\text{D.28})$$

$$= (2k_{0,\max}\kappa_s)^2 N_{\text{bytes-per-pix}} k_{0,\max} \kappa_{t,0} / T_0 \quad (\text{D.29})$$

$$= 4\kappa_s^2 \kappa_{t,0} k_{0,\max}^3 N_{\text{bytes-per-pix}} / T_0 \quad (\text{D.30})$$