

Mining MeerKAT data for minute to hour timescale transients and variable sources

Thesis submitted for the degree of

Master of Science

by


Sihle Gcilitshana

Supervised by

Prof Oleg Smirnov

&

Prof Ian Heywood



Department of Physics and Electronics

Rhodes University

08 February , 2024

Dedicated to my late Thozama

Declaration

I affirm that the research presented in this thesis is my work. In instances where I have incorporated the ideas and words of others, proper citation and referencing have been diligently applied to acknowledge the sources. I confirm my strict adherence to all academic honesty and integrity principles, ensuring that no misrepresentation, fabrication, or falsification of any idea, data, fact, or source has occurred in my submission. The timeframe for conducting this research spans from February 2022 to February 2024, coinciding with my enrolment as an M.Sc. student at Rhodes University.

.....

Sihle Gcilitshana

Date: February 2024

Place: Rhodes University

Acknowledgements

I want to express my sincere gratitude to my supervisors, Prof Oleg Smirnov and Prof Ian Heywood, for their invaluable guidance and constructive feedback throughout researching and writing this thesis. Their expertise and encouragement have been instrumental in shaping the direction of my work.

I am also thankful for the assistance provided by Dr Landman Bester, Dr Nadeem Oozer, Elizabeth Kamau, and Isaac Sihlangu, who offered valuable insights and thoughtful suggestions that significantly enhanced the quality of this thesis.

I extend my appreciation to the RATT team, whose contributions facilitated the smooth progress of my research. Special thanks to the SARA0 for providing access to essential equipment and facilities.

Finally, I want to acknowledge my family and friends' love, understanding, and patience. Their unwavering support has been my source of strength throughout this research.

This thesis would not have been possible without the collective support of these individuals. Thank you for being an integral part of my academic journey.

Sihle Gcilitshana

Abstract

In radio astronomy, minute-to-hour timescale transients and variable sources present an understudied population. We now have an unprecedented view of the transient radio sky due to the advent of new telescopes (such as the MeerKAT) with high instantaneous sensitivity, excellent snapshot imaging capabilities, and a large field of view. We can now probe into this population. This thesis presents the search for minute-to-hour timescale transient and variable sources in two of MeerKAT archival observations: the COSMOS and MACS J2140.2–2339 fields. The fields were observed for eight and five hours at the UHF band, respectively. We employed the PARROT transient and variable search pipeline currently being developed by the RATT group at Rhodes University. The pipeline’s input is a cross-calibrated measurement set, and its outputs are light curves extracted from all the sources in the restored image of the field. Using the light curves from the pipeline, we detected two variable sources in the COSMOS field, scintillating Active Galactic Nuclei (AGNs), which are most likely caused by the turbulent plasma in the interstellar medium. Due to persistent ionospheric diffraction, no variable sources were detected in the MACS J2140.2–2339 field, and no transients were detected in either field. The thesis also highlights areas where improvements to the PARROT pipeline can be implemented.

Contents

Declaration

Acknowledgements

Abstract

Contents

List of Figures

List of Tables

Abbreviations

Thesis Organisation

1	Introduction	1
1.1	Image plane transients and variable radio sources	2
1.1.1	AGN/ Blazars/ Quasars	3
1.1.2	Gamma-ray burst (GRBs) and Supernovae (SNs)	4
1.1.3	Intraday (IDV) and Intrahour (IHV) Variability	5
1.1.4	Pulsars	6
1.1.5	RRATs	6
1.1.6	Stellar flares	7
1.1.7	X-ray Binaries	7
1.2	Image plane transient and variable searches	8
1.3	MeerKAT telescope	12
1.4	Problem statement	13
1.5	Study objectives	15
2	Radio Interferometry & Image-plane transient detection	17
2.1	Radio telescope fundamentals	17

2.2	Radio Interferometry	19
2.3	Radio Interferometric Imaging	24
2.3.1	Imaging Software	25
2.4	Calibration	26
2.4.1	Radio Interferometer Measurement Equation (RIME)	26
2.4.2	Calibration Strategies	30
2.4.3	Calibration Software	34
2.5	Image-plane transient and variable searches	34
2.5.1	Variability Statistics	35
3	Methodology: Data Reduction, Light Curve Extraction, and Light Curve Analysis	39
3.1	Data Reduction: 1GC	41
3.2	PARROT Pipeline	42
3.2.1	PARROT pipeline automation	42
3.2.2	PARROT pipeline key steps	43
3.2.2.1	Self-calibration and peeling	43
3.2.2.2	Light curve extraction	45
3.3	Light Curve Analysis	49
3.3.1	A search for intrinsic variability	49
4	Mining MeerKAT Data: COSMOS & MACS J2140.2–2339 Fields	51
4.1	Observations	51
4.2	First generation calibration (1GC)	53
4.3	PARROT pipeline	53
4.3.1	Self-calibration and peeling	53
4.3.2	Light curve extraction	60
4.4	Light Curve Analysis	61
4.4.1	Systematic effects	62
4.4.1.1	Primary beam effects	62
4.4.1.2	Propagation effects	65
	Ionospheric scintillation,	65
	Interplanetary scintillation	67
	Ionospheric diffraction,	67
4.4.1.3	Imaging Artefacts	69
4.4.1.4	RFI and satellites	70
4.4.2	Light curve selection criteria	72
4.4.3	Searching for intrinsic variability in light curves	75
4.4.3.1	COSMOS	75
4.4.3.2	Analysis of the variable candidates	87
4.4.3.3	Discussion	89

4.4.3.4	MACS J2140.2–2339	90
5	Conclusions & Future Work	93
A	The PARROT Pipeline	97
B	A python code used to simulate fake variable sources	107
C	A Python code for calculating the variability metric (η, v) threshold	111
	Bibliography	113

List of Figures

1.1	The under construction MeerKAT radio telescope in Northern Cape, South Africa. Photo credit: SARAO	13
1.2	The figure shows a time-variable radio parameter space, which shows all the known transients and variable sources, adapted from (Cordes et al., 2004). Here, (L_v) represents luminosity (observed flux density multiplied by the squared distance of the source), while $(v.W)$ denotes the product of variability frequency and pulse width (variability timescale). It uses the Rayleigh-Jeans approximation to link these quantities to the brightness temperature, the minimum brightness temperature for coherent emission, with lines of constant brightness temperature marked on the figure. The highlighted pink region signifies an unpopulated area within this parameter space, particularly at frequencies around 1 GHz. This pink area pertains to variability occurring within a timeframe ranging from 1 second to 8 hours.	15
2.1	A two-element radio interferometer (Condon and Ransom, 2016).	21
2.2	Flow diagram outlining the procedure for self-calibration.	32
3.1	A workflow summarising the methodology.	40
4.1	This figure shows an improvement in the quality of the MFS restored image of the MACS J2140.2–2339 field, from 1GC (commonly known as crosscal) to delay self-calibrated visibilities. The upper inserted image represents a segment of the delay self-calibrated MFS restored image, while the lower inserted image represents the cross-calibrated MFS restored image before self-calibration.	56
4.2	This figure shows an improvement in the quality of the MFS restored image of the COSMOS field, from 1GC (commonly known as crosscal) to delay self-calibrated visibilities. The lower inserted image represents a segment of the delay self-calibrated MFS restored image, while the upper inserted image represents the cross-calibrated MFS restored image before self-calibration.	57

4.3	(a) This image shows the MFS restored image of COSMOS. The circular blue region corresponds to our RoI (one primary beam FWHM), within which we eventually search for intrinsic variability. The two red square regions highlight bright off-axis sources that impact the sources inside the circular blue region. In (b), the figure shows a model-subtracted high-time cadence snapshot image from the COSMOS. The presence of bright off-axis sources in the MFS-restored image introduces artefacts in the model-subtracted images, leading to temporal variations in sources within the circular region. Figures (c) and (d) correspond to the MACS J2140.2–2339 field.	59
4.4	This figure shows three problematic off-axis sources in the MACS J2140.2–2339 field before and after peeling using QUARTICAL	60
4.5	Two non-variable light curves from MACS J2140.2 – 2339 field.	61
4.6	Two non-variable light curves from COSMOS field	62
4.7	Three light curves from COSMOS affected by the primary rotation. Each source has an angular separation from phase centre (r) greater than 1. Source 4 has $r = 1.10$ deg, Source 31 has $r = 1.44$ deg, and Source 41 has $r = 1.50$ deg. Significant peak offsets in the light curves are plotted in red.	63
4.8	Three light curves from MACS J2140.2–2339 affected by the primary rotation. Each source has an angular separation from phase centre (r) greater than 1 degree: source 2 at $r = 1.18$ deg, source 22 at $r = 1.94$ deg, and source 50 at $r = 1.4$ deg. Significant variations in the light curves are plotted in red.	66
4.9	Two COSMOS light curves affected by ionospheric scintillation,	66
4.10	Two MACS J2140.2-2339 light curves affected by the ionospheric diffraction.	68
4.11	Three FITS cube (model subtracted) slices from MACS J2140.2-2339 displaying slight shifts in the apparent position of source one. The black spots inside the circular yellow region present an over-subtraction, and the white spots represent an under-subtraction	68
4.12	Image with strong artefacts around bright sources and the number of artefacts near the bright sources are incorrectly identified as real sources.	70
4.13	Three highly variable light curves of artefacts incorrectly identified as sources by PYBDSF.	71
4.14	Light curves of sources affected by satellite as they move through the side lobes of the primary beam. The light curves show synchronous short-timescale variations.	72
4.15	(a) and (c) show the selected sources (RoI) for intrinsic variability search in the COSMOS and MACS J2140.2–2339 respectively. (b) and (d) show a close up region of (a) and (c) of COSMOS and MACS J2140.2–2339 respectively	73

4.16	The plots show the light curves of five simulated variable sources as described in section 4.4.3.1	77
4.17	The histogram illustrates the log η -distribution of our light curves, with the curve (indicated in orange) representing the optimal fit normal distribution. Non-variable sources fall below the established threshold (indicated in a blue line), while sources with above η_{thresh} values are categorized as variable candidates.	78
4.18	This plot shows the reduced $\chi^2(\eta)$ plotted against the maximum and median flux density. The synthetic sources are highlighted by stars. The plots are in log-scale	79
4.19	This plot shows the reduced weighted $\chi^2(\eta)$ plotted against the variability index (ν). The synthetic sources are highlighted by stars.	80
4.20	(a) This plot shows the variability index ν plotted against the maximum and median flux density. The synthetic sources are highlighted by stars. The plot also shows the classification of the sources as either stable(non-varying), Low-moderate or Moderate-high variability	81
4.21	(a) This plot shows η plotted against the maximum flux density, with the colour bar indicating the phase distances from the centre (r). (b) This plot shows η plotted against the maximum flux density, with the colour bar indicating the observed flux density of each source reported by the source finder (PYBSDF). The maximum flux(μJy) in both plots are also in log scale.	82
4.22	The plot shows the distribution of η against the maximum flux density of the light curves obtained from the COSMOS field distribution without the effects of the primary beam (62 variable sources excluded). The numerical labels align with the 'sources ID' reported by the source finder (PYBDSF). PYBSDF). The maximum flux(μJy) is in log scale.	83
4.23	Four of the nine candidate variable light curves with $\eta > 1$ (continued next page)	85
4.24	Five other candidate variable light curves.	86
4.25	Restored MFS image of the COSMOS field showing source 25 and 28.	87
4.26	The plot shows the distribution of η against the maximum flux density of the light curves obtained from the MAC J2140.2–2339 field. The numerical labels align with the 'sources ID' reported by the source finder (PYBDSF).	91
4.27	The plots show five variable sources with a $\eta > 1.29$, after removing the sources affected by the PBR.	92

List of Tables

1.1	Transient or variable source searches and surveys (Mooley et al., 2016).	10
4.1	Characteristics of MeerKAT Observations for the COSMOS and MACS J2140.2 – 2339 fields.	52
4.2	Properties of the nine variable sources with $\eta > 1$	84
4.3	The Multiwavelength properties of our 2 potential variable sources . .	89

Abbreviations

1GC	1st Generation Calibration
2GC	2nd Generation Calibration
3GC	3rd Generation Calibration
AGN	Active Galactic Nuclei
ASKAP	Australian SKA Pathfinder
DIE	Direction Independent Effects
DDE	Direction Dependent Effects
FoV	Field of View
FRB	Fast Radio Burst
FT	Fourier Transform
HPBW	Half Power Beam Width
IDV	Intra Day Variability
IHV	Intra Hour Variability
LOFAR	Low Frequency Array
PSF	Point Spread Function
MS	Measurement Set
RATT	Radio Astronomy Techniques and Technologies
RRAT	Rotating Radio Transient
RIME	Radio Interferometric Measurement and Equation
RoI	Region of Interest
SKA	Square Kilometre Array
SNR	Signal to Noise Ratio
TraP	Transient Pipeline
VAST	Variables and Slow Transient
VLA	Very Large Array

Thesis Organisation

This thesis consists of five chapters, starting with Chapter 1, which introduces the transient and variable sources typically detected in the image plane, image-plane surveys and searches for transients and variable sources, the problem statement, and the research objective.

Chapter 2 offers an in-depth review of radio interferometry, Imaging, Calibration, and transients and variable detection in the image plane. Chapter 3 provides an overview of the methodology, including the PARROT pipeline, and an explanation of the analysis conducted on the extracted light curves.

Moving forward, Chapter 4 describes the application of the PARROT pipeline on two archival MeerKAT measurement sets and the results.

Lastly, Chapter 5 offers a summary and the research contributions. Furthermore, it discusses future works.

Chapter 1

Introduction

Transients are sources or events that suddenly appear and disappear during the observation period. On the other hand, *variable sources* are known sources that exhibit flux density variations over time (Rowlinson et al., 2016). These sources present an exciting branch of radio astronomy. The sky is dynamic, and almost all astrophysical objects, from stars to supermassive black holes in the nuclei of distant galaxies, exhibit unique variability (Macquart et al., 2010; Lazio et al., 2014)—their flux changes within a range of timescales. Studying the variability of sources is important to astronomers for several reasons: They provide insights into astrophysical processes that occur in the universe, enabling researchers to understand their underlying mechanisms and physical properties. For example, investigating fast transients can aid in understanding galaxy evolution (Abbott et al., 2017; Macquart et al., 2010). So within this context, the discovery of a fast radio burst (FRB) by Lorimer et al. (2007) opened up a lot of exciting possibilities for astronomers as a possible new way to look into energetic processes outside of our galaxy and as a cosmological probe (Macquart et al., 2010).

These variations can occur on various timescales, from fractions of a second to years; depending on their timescale of variation, they can be classified into two categories:

fast and *slow* transient and variable sources (Cordes et al., 2004, Bhat et al., 2013, Murphy et al., 2013). Though the separation between these two classes is ambiguous, fast transients or variable sources generally vary on a timescale of less than 1s. Slow transients or variable sources, on the other hand, vary from minutes to hours or days to years.

This introduction looks into transients and variable sources detected in the image plane. We also discuss surveys or studies that search for them using radio interferometers. These surveys highlight our gradual understanding and detection capabilities of these sources. Additionally, the studies highlight many unexplored areas (i.e., minute-to-hour timescale transients) waiting to be discovered or investigated. Ultimately, this introduction aims to emphasise the significance of this research, which is centred around exploring minute-to-hour timescale transients and variable sources utilising MeerKAT data.

1.1 Image plane transients and variable radio sources

There are various transient and variable radio sources or events, classified as either fast or slow sources or events. Fast transients or variable sources include rotating radio transients (RRATs), pulsars, FRBs, etc. On the other hand, slow transients or variable sources include active galactic nuclei (AGN), stellar flares, ultra-long neutron stars, X-ray binaries and supernovae/gamma-ray bursts (for an in-depth review of transients and variable sources, refer to Murphy et al., 2013). Fast transients or variable sources are usually detected using dedicated high-speed hardware and software that searches beamformed data with high-time resolution. This allows for the precise measurement of events or changes that occur rapidly (Serylak et al., 2012). Slow transients or variables are typically detected in the image plane. Within this

context, minute-to-hour timescale transients or variable sources are classified as slow transients and variable sources. They are more likely to be detected in the image plane.

For this research, we will discuss the types of transients or variable sources typically detected in the image plane (note that some sources described below can be detected using beamformer techniques).

1.1.1 AGN/ Blazars/ Quasars

First on the list of variable sources that can be detected using time series imaging are AGNs. They are highly energetic and compact regions at the centres of galaxies exhibiting activity and variability. The variability of these compact sources was initially discovered by [Dent \(1965\)](#) and has remained a subject of keen interest to astronomers. Notably, [Fiedler et al. \(1987\)](#) presented the discovery of Extreme Scattering Events (ESEs) in AGNs, sampling 33 AGNs with flux densities of a few Janskys daily for over six years. All sources displayed variability over this period, characterised by generally smooth flux variations with quasi-periods ranging from months to years. AGNs exhibit variability on timescales ranging from days to several years, often attributed to intrinsic phenomena, ESEs and interstellar scintillation (ISS). Intrinsic variability can arise from processes such as accretion onto supermassive black holes situated at the core of galaxies ([Smith et al., 2010](#); [Komossa, 2006](#)) or relativistic jets. Variations in flux density within the jets manifest as flares, with durations spanning from days to months ([Hovatta et al., 2008](#)). When observed from the side, galaxies hosting AGNs are classified as radio galaxies. Alternatively, if the radio galaxy is observed directly down the jet, it is termed a blazar.

1.1.2 Gamma-ray burst (GRBs) and Supernovae (SNs)

GRBs and SNs are a second important phenomena that could be detected in the image plane. They are among the most energetic transient events in the universe and are sometimes believed to be related (Kulkarni et al., 1998). The supernova explosion is typically brief, lasts for milliseconds, and leaves behind an afterglow detectable across various wavelengths, including radio waves. The radio afterglow exhibits a gradual increase in flux, reaching a peak persisting for weeks to months and eventually fading away. Supernova explosions are classified into three main subgroups: Type Ia, Type Ib/Ic, and Type II (Weiler et al., 2004; Filippenko, 1997).

Type Ia supernovae occur in binary star systems where one star is a white dwarf, and the other is either a main-sequence star or a giant star (Livio, 2001). The white dwarf accretes matter from its companion until it reaches a critical mass, triggering a runaway nuclear fusion reaction that completely disrupts the white dwarf, resulting in a highly luminous explosion (Panagia et al., 2006).

Type Ib/Ic supernovae also involve the explosion of a massive star, but they lack hydrogen in their spectra, indicating that the progenitor star had shed its outer hydrogen layers before the explosion. This shedding of hydrogen typically occurs due to stellar winds or interactions with a binary companion (Yoon et al., 2010).

Type II results from the core collapse of massive stars. The explosion of Type II supernovae can leave behind a neutron star or, in some cases, if the core collapse is particularly massive, a black hole (Woosley and Weaver, 1986).

GRBs are sudden, intense bursts of gamma-ray radiation caused by the launch of relativistic jets in cataclysmic events (Piran, 2005). The GRB population can be categorised into two sub-groups (Kouveliotou, 1997): short GRBs, with durations less than 2 seconds, believed to originate from the merger of two neutron stars

([Eichler and Levinson, 1999](#)); and long GRBs, lasting longer than 2 seconds, likely produced during the collapse of massive stars ([Woosley and Weaver, 1986](#)). Similar to supernovae, gamma-ray bursts are often followed by a broad-band afterglow, observable across various wavelengths from high-energy gamma-rays to radio waves, sometimes persisting for years.

1.1.3 Intraday (IDV) and Intrahour (IHV) Variability

Compact sources with angular sizes smaller than one milliarcsecond (mas), such as pulsars and AGNs, are susceptible to propagation effects like interstellar and interplanetary scintillation ([Heeschen et al., 1987](#), [Wang et al., 2021](#)). Interstellar scintillation occurs when radio waves from compact sources propagate through the turbulent interstellar plasma, leading to extreme random variations. The variability induced by interstellar scintillation spans from minutes to months. Interstellar scintillation can be classified into intra-day variability (IDV) or intra-hour variability (IHV), depending on the timescale of variation. IHV exhibits variations on timescales from minutes to hours, while IDV varies on timescales of hours to days. The phenomena of IDV and IHV have been extensively investigated (see [Heywood, 2023](#); [Heeschen et al., 1987](#); [Roland et al., 2009](#); [Simonetti et al., 1985](#); [Jauncey et al., 2020](#), [Wang et al., 2021](#), [Oosterloo et al., 2020](#)) and have been identified across a wide frequency range, from gamma-ray to radio frequencies ([Wagner and Witzel, 1995](#)). Interplanetary scintillation refers to the rapid scintillation (1 second or less) of very compact radio sources as their signal propagates through turbulent interplanetary plasma caused by the solar wind ([Hewish et al., 1964](#)).

1.1.4 Pulsars

Pulsars are rapidly rotating neutron stars that emit focused radio (and high-energy) radiation from their magnetic poles. Typically, pulsars exhibit periods ranging from milliseconds to seconds (Lorimer et al., 2007; Hessels et al., 2006; Tan et al., 2018; Wang et al., 2023). However, recent discoveries have revealed rotating neutron stars with variability occurring within minutes (see Caleb et al., 2022; Hurley-Walker et al., 2022). These ultra-long variable sources have been detected and characterised using image-domain products. The variability of pulsars on timescales of minutes to hours can arise from various phenomena. These include the self-gravitational lensing effect proposed in pulsar-black hole systems (Xiao and Shen, 2024), instances of deep nulling where pulse energy suddenly diminishes to zero and then returns to its typical state, a phenomenon known as 'moding' (Wang et al., 2007), interstellar scattering, which changes the mean flux of the pulsar over timescales of minutes to months (Cordes, 2005), and eclipses within binary systems (Camilo et al., 2000).

1.1.5 RRATs

Rotating radio transients (RRATs) are a subset of pulsars believed to be caused by an extreme form of moding (Backer, 1970). They were first discovered in search of archival data from the Parkes Multi-Beam Pulsar Survey (PMPS) by McLaughlin et al. (2006). They are sporadic, isolated radio pulses with short durations ranging from milliseconds to seconds. They are characterised by their irregular and infrequent bursts of radio emissions. The pulses have peak flux densities (at 1.4 GHz) ranging from ~ 100 mJy up to 10 Jy (Keane et al., 2010). The underlying physical mechanisms governing the sporadic emission of RRATs remain unclear, and several hypotheses have been proposed to account for the unpredictable nature of their radio

bursts. One hypothesis posits that RRATs may be conventional pulsars, exhibiting exceptionally powerful radio emission pulses (Brylyakova and Tyul'bashev, 2021; Weltevrede et al., 2006), or they could be pulsars with a high proportion of null pulses (Knight et al., 2006).

1.1.6 Stellar flares

Stellar flares are phenomena characterised by sudden eruptions of energy caused by the explosive release of magnetic energy (Bastian, 1994). The types of stars observed to flare up include young, active stars (such as T Tauri stars), classic close binaries (such as RS Canum Venaticorum (RS CVn) Binaries), M-dwarfs, Brown Dwarfs, and White Dwarfs (Pettersen, 1989). The flares in these stars are attributed to the high rotation rates and convective cores of these stars, which generate powerful global magnetic fields (Rigney et al., 2022). The variability of stellar flares ranges from minutes to days, depending on whether the flares are coherent or incoherent (Toet et al., 2021). Coherent flares typically exhibit variability spanning from minutes to days, while incoherent flaring shows variations from milliseconds to hours.

1.1.7 X-ray Binaries

X-ray binaries are binary systems consisting of a companion star and compact objects such as black holes or neutron stars. In these systems, the compact objects accrete material from the companion, leading to the heating of the accreted material and the emission of X-rays due to strong gravitational pull (Schatz and Rehm, 2006, Fender et al., 1997). These systems also exhibit strong, highly variable radio emissions, believed to result from the acceleration of relativistic electrons within the X-ray binaries. These electrons emit radio waves through the synchronous emission

mechanism when interacting with magnetic fields (Mirabel and Rodriguez, 1999, Gregory and Seaquist, 1974). X-ray binaries are categorised into two groups: High-mass X-ray binaries (HMXBs) and Low-mass X-ray binaries (LMXBs). HMXBs consist of massive, luminous stars and compact objects, with accretion facilitated by strong stellar winds (Liu et al., 2006). LMXBs, on the other hand, involve low-mass stars transferring material to compact objects either through Roche lobe overflow or stellar winds. The variability period of X-ray emissions in these binaries typically ranges from days to months (Liu et al., 2007).

1.2 Image plane transient and variable searches

This section reviews searches of image plane transients and variable sources such as stellar flares, X-ray binaries, GRBs, AGNs, or pulsars. These searches are typically divided into targeted and untargeted searches/surveys depending on the scientific objectives. Targeted searches are defined as focused investigations aimed at specific objects, phenomena, or regions based on prior knowledge or hypotheses. Examples of targeted searches include works by Bignall et al. (2003), Heeschen et al. (1987), Oosterloo et al. (2020), and Heywood (2023). Untargeted searches, on the other hand, are broad and exploratory surveys of the sky that do not focus on specific targets but aim to discover new phenomena or objects. Examples of these include works by Wang et al. (2023), Wang et al. (2021), and Rowlinson et al. (2016). In this research, we will discuss searches/surveys that are untargeted.

Bower et al. (2007) pioneered the first untargeted large multi-epoch search of transient and variable sources using the very large array (VLA) archival surveys. Subsequently, numerous surveys for radio transients or variable sources have been carried

out using archival data and preliminary datasets from next-generation radio telescopes. The main objectives of these surveys have typically revolved around two key goals: (a) the identification and in-depth investigation of specific objects of interest and (b) the analysis of population statistics related to radio transients, anticipating upcoming surveys. An updated overview of these searches for transient and variable sources, initially compiled by Ofek et al. (2011) and later revised by Murphy et al. (2013), Mooley et al. (2016), and Laura N. Driessen¹ is provided in Table 1.1. We have also performed additional updates to the table, incorporating some further searches.

As stated at the outset of this introduction, we aim to utilise Table 1.1 to emphasise two critical aspects pertinent to our research. Firstly, showcase the gradual improvement of our knowledge regarding transients and variable searches and how these insights can be used to advance future variable and transient searches. Secondly, the table allows us to draw attention to the less explored area of intraday variability searches. In Table 1.1, the surveys or searches primarily target identifying very slow, transient, and variable sources, with timescales between observations spanning days, months, and even years. Notable contributions include works by Bannister et al. (2011); Frail et al. (1994); Carilli et al. (2003); de Vries et al. (2004); Aoki et al. (2013); Bell et al. (2011); Levinson et al. (2002); Ofek et al. (2011); Thyagarajan et al. (2011); Bower and Saul (2011); Bhandari et al. (2018); Mooley et al. (2013); Croft et al. (2010), and Croft et al. (2011). Furthermore, only Rowlinson et al. (2022); Hobbs et al. (2016); Heywood et al. (2016); Wang et al. (2021), Aoki et al. (2013), and Wang et al. (2023) explored variability on timescales less than a day (intraday variability). For this research, we only briefly summarised the intraday variability searches:

¹https://pure.manchester.ac.uk/ws/portalfiles/portal/248796605/FULL_TEXT.PDF

TABLE 1.1: Transient or variable source searches and surveys (Mooley et al., 2016).

Frequency (GHZ)	Area (deg ²)	Epochs	Timescale	RMS (mJy)	Sources	Transient	Variable	References
1.4	2.6	16	1-12 days, 1-3 months	0.7	250	0	1%	Frail et al. (1994)
1.4	2500	2	~ <i>years</i>	0.45	7181	1	...	Levinson et al. (2002)
1.4	0.22	3	19 days, 17 months	0.015	...	0	2%	Carilli et al. (2003)
1.4	120	2	7 yrs	0.15	9086	0	1.4 %	de Vries et al. (2004)
5 or 8.4	0.07 < R < 40	944	22 yrs			10	...	Bower et al. (2007)
1.4	690	2	15yrs	3.94	4408	0	~ 0.1%	Croft et al. (2010)
0.84	2776	3011	1days - 20 yrs	2.8	29730	2	~ 100	Bannister et al. (2011)
1.4	0.2	~ 1852	1 days - 23 yrs	2	10	0	...	Bower and Saul (2011)
1.4	690	12	days-months	38	4408	0	~ 0.5%	Croft et al. (2011)
1.4	690	151	days-years	~ 1	...	0	...	Bell et al. (2011)
1.4	0.2	~ 55000	~ <i>days</i> years	~ 1	279407	...	1627	Thyagarajan et al. (2011)
327 & 321.6	6.5	..	1-3 months	0.17	..	1	2	Jaeger et al. (2012)
1.4	2870	~ 1000	1 days	300	...	1	...	Aoki et al. (2013)
1.5	0.44	49	1days - 3 months	0.003	736	0	~ 1%	Mooley et al. (2013)
1.4	0.5	3	7yrs	0.15	1436	0	89	Hodge et al. (2013)
5.5	0.3	3	~ 2.4yrs		142	0	~ 3 - 25%	Bell et al. (2015)
1.4	60	390	2min-13hrs	21	...	0	1	Hobbs et al. (2016)
0.86	56	6	8days-8yrs	~ 1	~ 500	1	8	Hancock et al. (2016)
1.4	5.97	3	12h-1 week	1	3722	0	1	Heywood et al. (2016)
115 & 190	2275	151	~ <i>months</i>	0	0	Carbone et al. (2016)
60	175	2149	~ 4 <i>months</i>	1	0	Stewart et al. (2016)
0.86	150	8	~ <i>days</i>	0.3	1653	0	9	Bhandari et al. (2018)
1.4	...	48	~ <i>weeks</i>	0.317	..	1	0	Drissen et al. (2020)
1.4	8444	172	days-6yrs	0.001	2713	...	58	Sarbadhicary et al. (2021)
1.4	30	slow	minutes-1 yrs	0.035	40859	0	6	Wang et al. (2021)
1.4	155	62	1 week	0.42	...	0	6	Rowlinson et al. (2022)
~ 1.28	..	6 - 167	3 months - 3 yrs	~ 22 - 136	~ 89000	...	168	Andersson et al. (2023)
1	1476	...	15min	30	38	Wang et al. (2023)

- Heywood et al. (2016) conducted a study on wide-field broad-band radio imaging using the Australian SKA Pathfinder (ASKAP-BETA) telescope. They used a 1-hour by 12-point pointing survey with nine beams per pointing, using 12 hours of observing time to cover roughly 150 square degrees in the frequency range of 711 to 1015 MHz. Over a week, this procedure was performed three times. A thorough examination of these three observational epochs for variability and transient events and a detailed comparison to the SUMSS survey revealed one notably variable source: a potential quasar candidate.

- [Hobbs et al. \(2016\)](#) conducted a study that utilised observations from the ASKAP-BETA telescope to investigate transient radio signals near the intermittent pulsar (PSR J1007-5907). They employed three separate transient data processing pipelines (MATCH-FILTER pipeline, LOFAR transient pipeline, and VAST pipeline) on short, two-minute snapshot images extracted from 13 hours of BETA observations. Although the presence of the pulsar (PSR J1007-5907) was confirmed, no additional transient radio phenomena were identified from the compact stellar remnants left behind.
- [Wang et al. \(2021\)](#) conducted a study using ASKAP. They used the observation to investigate and detect radio variables on hour-long timescales. They conducted an untargeted survey spanning a 30-square-degree area, involving multiple 10-hour observations spread over days to months, focusing on a central frequency of 945 MHz. Through this effort, they identified six rapid scintillators in 15-minute model-subtracted images with a sensitivity of approximately 200 mJy. Notably, two of these scintillators demonstrated extreme intra-hour variability, with variability indices reaching around 40% and timescales as brief as tens of minutes. The variability index is a variability statistic used to quantify the source's variability by comparing the light curve's standard deviation to its mean. Additionally, five variables displayed a linear arrangement in the sky, with an angular width of about one arcminute and a length of approximately 2 degrees, indicating the presence of a substantial plasma filament in their foreground.
- [Rowlinson et al. \(2022\)](#) conducted a search and identification of transients and variable sources using the MeerKAT. They utilised observations obtained through the ThunderKAT programme, which involves weekly monitoring of XRBs. They employed the TRAP pipeline to scan for transients and variable

sources. The study identified four likely active galactic nuclei (AGNs), one source with potential as either a galactic source (such as a pulsar or quiescent XRB) or an AGN, and one variable pulsar. However, no transient sources were discovered—defined as sources within the same field that had not been previously detected in prior images.

- [Wang et al. \(2023\)](#) conducted a study on transient and variable sources with a timescale of 15 minutes using the ASKAP pilot surveys. The survey employed 36 beams per field. Each field was monitored for 8 to 10 hours at a frequency of 1GHz. The cumulative duration of the observation period was 505 hours, resulting in a total of 1476 square degrees across the sky. The data were then transferred to the China SKA regional centre for processing and searching for transients and variables, and then the final data products (variable and transient candidate catalogues, image cutouts and model-subtracted snapshot images) were downloaded to a local machine for further analysis. As a result, 38 variable and transient sources were identified. There are a total of seven known pulsars, with one of them being identified explicitly as an eclipsing millisecond pulsar named PSR J2039 – 5617. In addition, there are eight sources associated with stars, with only one previously recognised as a radio star. Of the 23 remaining objects, 22 are associated with active galactic nuclei or galaxies.

1.3 MeerKAT telescope

The emerging wide-field radio telescopes such as the planned Square Kilometre Array (SKA, [Dewdney et al., 2009](#)) and (MeerKAT, [Jonas, 2009](#)) promise an unprecedented view of the radio transient sky to improve the detection rates of transients and

variable sources. The MeerKAT telescope, shown in Figure 1.1, in South Africa, comprises 64 individual dish antennas, each with a diameter of 13.5 meters. This telescope is managed by the South African Radio Astronomy Observatory (SARAO) and serves as a precursor to the SKA. MeerKAT has three receivers: an L-band receiver operating at 856–1712 MHz, a UHF-band receiver operating at 580–1015 MHz, and an S-band receiver operating at 1750–3500 MHz. MeerKAT’s noteworthy feature is its high instantaneous sensitivity and strong snapshot imaging capabilities, allowing for the detection of variability on sub-hourly timescales down to levels as low as 50–100 $\mu\text{Jy}/\text{beam}$ in snapshot images, making the instrument an excellent tool for probing into faint and minute-to-hour timescale variable sources and transients.



FIGURE 1.1: The under construction MeerKAT radio telescope in Northern Cape, South Africa. Photo credit: SARAO

1.4 Problem statement

The success of the surveys and studies highlighted by Table 1.1, coupled with their detailed methodologies and advanced instrumentation, provides a valuable framework for designing and executing future studies in transients and variable

source detections. Although these studies have yielded useful results, minute-to-hour timescale transients and variable sources remain noticeably under-researched compared to the investigation of very slow transients or variables. The intraday survey results from Table 1.1 also point to a limited number of variable source discoveries, with very few detections of transient events. This could be attributed to several reasons, such as observational challenges, instrument limitations, and data volume and processing.

Furthermore, considering the parameter space illustrated in Figure 1.2, it shows all known transient and variable sources based on their luminosity, observing frequency, and variability period. Specifically, within frequencies around 1 GHz, the pink region corresponds to variability occurring within 1 second to 8 hours. Notably, there have been very few detections, particularly in the coherent luminosity region. This area remains largely unexplored, presenting a possibility of an untouched population of transient and variable sources.

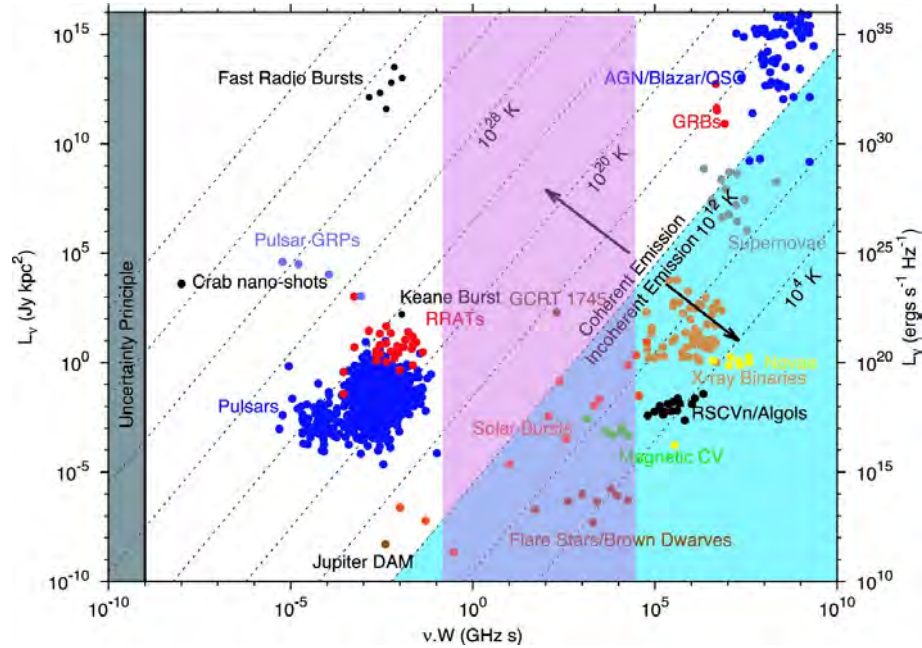


FIGURE 1.2: The figure shows a time-variable radio parameter space, which shows all the known transients and variable sources, adapted from (Cordes et al., 2004). Here, (L_v) represents luminosity (observed flux density multiplied by the squared distance of the source), while ($v.W$) denotes the product of variability frequency and pulse width (variability timescale). It uses the Rayleigh-Jeans approximation to link these quantities to the brightness temperature, the minimum brightness temperature for coherent emission, with lines of constant brightness temperature marked on the figure. The highlighted pink region signifies an unpopulated area within this parameter space, particularly at frequencies around 1 GHz. This pink area pertains to variability occurring within a timeframe ranging from 1 second to 8 hours.

1.5 Study objectives

In response to the problem statement outlined in the discussion above, we now have instruments like MeerKAT at our disposal. It's high sensitivity and wide field of view present an ideal means to explore faint transient and variable source signals while surveying extensive regions of the sky to maximise our search potential.

The RATT² team at Rhodes University is developing a software pipeline (the PAR-ROT pipeline, refer [Smirnov et al., 2023](#)) to probe minute-to-hour timescale transient and variable sources. This pipeline's feasibility is that it will be used daily to detect these sources for follow-ups. The objective of this research is to test the pipeline being currently built on different MeerKAT archival data with two objectives:

- Detect transient and variable sources using the pipeline.
- Highlight areas where we can further improve the pipeline.

Given the success of the overall project by the RATT team, there is a possibility of uncovering a new population of transients or variable sources.

²<https://ratt.center/>

Chapter 2

Radio Interferometry & Image-plane transient detection

Our ability to detect transient and variable sources in the image plane depends on a thorough understanding of the fundamental concepts of radio astronomy. These include how an interferometer observes the data, known as visibilities, calibration, which corrects for instrumental and atmospheric effects introduced into visibilities, and imaging of visibilities. In this chapter, we discuss the mathematical description of these concepts. We also discuss how transients and variable sources are typically detected in the image plane.

2.1 Radio telescope fundamentals

For the work discussed in this thesis, there are three important radio telescope characteristics to consider: sensitivity, resolution and field of view (FoV).

- **Sensitivity**

Sensitivity is expressed in terms of minimum detectable flux density, or signal-to-noise ratio, and is calculated using the radiometer equation (Condon and Ransom, 2016):

$$S_{min} = \frac{2K.T_{sys}}{A_e\sqrt{\Delta\nu\Delta t}} \quad (2.1)$$

where K is the Boltzmann constant, $\Delta\nu$ is the bandwidth of a receiver in Hz, A_e is the effective area of the telescope (A_e) in square metres, and T_{sys} is the telescope system temperature. A lower sensitivity value indicates a more sensitive telescope capable of detecting weaker signals. (Lorimer and Kramer, 2005).

- **Resolution**

The resolution of a radio telescope is determined by its ability to distinguish between two closely spaced sources in the sky. It is often quantified by a parameter called the angular resolution. For a single-dish telescope, the angular resolution is inversely proportional to the diameter of the telescope's aperture and the wavelength of the observed radiation. For a filled circular aperture, this is expressed using the Rayleigh equation (Condon and Ransom, 2016):

$$\theta = 1.22 \frac{\lambda}{D} \quad (2.2)$$

For an interferometer, the resolution is proportional to the longest baseline (the separation between the antennas).

- **Field of view (FoV)**

The FoV (Ω) of an instrument is often defined as a function of the telescope's half-power beam width (HPBW). It is the point in the beam at which the

signal power is half that of the peak. It's typically measured in degrees or arcseconds.

$$\Omega \propto \theta^2 \propto \text{HPBW}^2 \quad (2.3)$$

For an interferometer composed of dishes, the FoV (Ω) is effectively the same thing as the *resolution* of a single dish, and it can be increased with the use of phased-array feeds or multi-beam receivers. A larger field of view in a radio telescope is beneficial for detecting transient events and monitoring variable sources. It allows for broader coverage of the sky, increasing the likelihood of capturing time-sensitive or short-lived transient events ([Condon and Ransom, 2016](#)).

2.2 Radio Interferometry

Let us now consider the sensitivity and resolution of the radio telescope. Both of these characteristics are improved by increasing the aperture of the telescope. If we increase the collecting area, the minimum detectable source brightness (S_{min}) described in Equation 2.1 decreases, which means that the telescope can detect very weak signals. Again, if we increase the diameter of the telescope, the angular resolution increases; to obtain one arcsec resolution at ~ 1 GHz, a dish with a diameter of 250 km will be required. Building such a large dish is not practical. Therefore, the concept of radio interferometer was introduced (e.g., Michelson stellar interferometer). Aperture synthesis interferometry is a technique that combines a number of separated telescopes to simulate the effect of a single virtual dish with a diameter comparable to the longest separation between the telescopes.

Figure 2.1 shows an example of a simple interferometer where antennas 1 and 2 observe the sky in the direction \hat{s} , and the distance between the two antennas (1 and 2) is the baseline (b). The wavefront of the observed sources does not arrive simultaneously in both antennas; for example, the wavefront arrives at τ_g later in antenna 1. This time lag between antennas is called the geometric delay (τ_g). Therefore, to correct it, some delay τ_g is normally added to antenna 2.

Below, we summarise how an interferometer measures complex visibilities. For simplicity, we consider a quasi-monochromatic interferometer. The voltage response of antenna 1 and antenna 2 at time t is given by:

$$\begin{aligned} V_1 &= V \cos[\omega(t - \tau_g)] \\ V_2 &= V \cos[\omega(t)] \end{aligned} \tag{2.4}$$

V_1 and V_2 voltages represent amplified versions of the antenna input voltages. V denotes the amplitude of the received signal, ω denotes the angular frequency of the signal, and time t indicates the time at which the signal is observed. Finally, τ_g represents the time delay associated with signals reaching the two telescopes. In the correlator, the voltages are multiplied, and the output voltage is determined by:

$$V_1 V_2 = V^2 \cos[\omega(t)] \cos[\omega(t - \tau_g)] \tag{2.5}$$

The correlator then averages in time, which averages out the high-frequency (t -dependent) term in the expression, leaving the τ -dependent term. Now, the response of the interferometer can be written as:

$$\begin{aligned} R &= \langle V_A V_B \rangle \\ &= \frac{V^2}{2} \cos[\omega(\tau_g)] \end{aligned} \tag{2.6}$$




FIGURE 2.1: A two-element radio interferometer (Condon and Ransom, 2016).

The correlator's output voltage R exhibits sinusoidal variation as the Earth's rotation changes the direction of the source in relation to the baseline vector. These sinusoidal patterns are called fringes. An extended source (slightly spread out) in the sky influences Equation 2.6 differently from point sources. To understand its effect on the interferometer, we consider it a collection of many individual point sources. This approximation helps us analyse and understand how the interferometer responds to extended sources by breaking them down into smaller components. Equation 2.6 can be expressed in terms of a brightness distribution denoted by $I(s)$ integrated across the sky. The equation is given as follows:

$$R = \int I(s) \cos(\omega \tau_g) d\Omega \quad (2.7)$$

Equation 2.7 contains an even cosine function that is only sensitive to the even portion of any given brightness distribution. (This distribution can be represented as a combination of two components: the even part (I_{even}) and the odd part (I_{odd}), which adds up to the entire distribution ($I = I_{\text{even}} + I_{\text{odd}}$.) To complete the measurement, a correlator also computes a ‘sine’ interferometer response: a phase shift is used to convert the cosine response to a sine. Hence, Equation 2.7 for a sine interferometer can be written as:

$$R = \int I(s) \sin\left(\frac{2\pi \cdot b \cdot s}{\lambda}\right) d\Omega \quad (2.8)$$

note that we have substituted $\omega\tau_g = -2\pi\left(\frac{b \cdot s}{\lambda}\right)$.

A complex visibility combines the cosine and sine responses into a complex number. It is mathematically convenient to interpret Equation 2.8 as complex exponentials using Euler’s formula:

$$V = \int I(s) \exp\left(\frac{i2\pi \cdot b \cdot s}{\lambda}\right) d\Omega \quad (2.9)$$

Equation 2.9 represents the response to an extended source with a brightness distribution $I(s)$ of the two-element quasi-monochromatic interferometer with complex visibility.

To establish the connection between the visibility V , which represents spatial information on the radio brightness distribution in the sky, and the brightness $I(s)$, a set of baseline coordinates and their corresponding sky coordinates are introduced (Thompson et al., 2017; Condon and Ransom, 2016). The observation coordinate system for visibility is defined using spatial frequency (u, v, w) , which is measured

in units of wavelength. In this system, the eastern direction is defined by the coordinate u , the northern direction is defined by the coordinate v , and the reference direction along vector \hat{s} (refer to Figure 2.1) is defined by the coordinate w . The (l, m, n) represents the source coordinate system, which corresponds to the directional cosines of the vector \hat{s} . Given that \hat{s} is a unit vector, we may determine that $n = \sqrt{1 - l^2 - m^2}$. Therefore, we define the relationship in the coordinate systems as:

$$V(u, v, w) = \int \int \frac{I_v(l, m)}{\sqrt{1 - l^2 - m^2}} \exp[-2\pi i(ul + vm + w\sqrt{1 - l^2 - m^2})] dl dm \quad (2.10)$$

In some cases, the additional integral element can be excluded to get the Fourier transform representation of the sky brightness distribution $I(l, m)$ and the interferometer response, the complex visibility $V(u, v)$. For example, if the observation area centred on the reference direction is extremely small, then $\sqrt{1 - l^2 - m^2} \approx 1$ (Condon and Ransom, 2016), and Equation 2.10 simplifies to:

$$V(u, v, w)e^{2\pi i w} = \int \int I_v(l, m) \exp[-2\pi i(ul + vm)] dl dm \quad (2.11)$$

The phase factor on the right-hand side of the equation corrects the observed phase of the visibilities as if the antennas were perfectly located in the uv plane (Rohlfs and Wilson, 2013). This suggests:

$$V(u, v, w)e^{2\pi i w} \approx V(u, v, 0) \quad (2.12)$$

Hence, Equation 2.11 can be expressed as the result of the Van Cittert–Zernike theorem. (Thompson et al., 2017):

$$V(u, v) = \int \int I_v(l, m) \exp[-2\pi i(ul + vm)] dl dm \quad (2.13)$$

2.3 Radio Interferometric Imaging

One key insight drawn from Equation 2.13 is the Fourier transform relationship between the visibilities observed by the interferometer, and the distribution of the sky brightness. The process of performing the inverse Fourier transform to estimate the sky brightness distribution of visibilities is called “radio interferometric imaging” (Thompson et al., 2017; Condon and Ransom, 2016; Perley et al., 1989). There are commonly two methods used to approximate the Fourier transform: direct Fourier transform (DFT) and fast Fourier transform (FFT) (for a comprehensive review, see Perley et al., 1989). Let’s consider Equation 2.13, which corresponds to continuous complex visibilities. However, sampling of a radio interferometer is incomplete on the uv plane, so a sampling function is introduced that assigns zero where data was not sampled. The sampling function is given by:

$$S(u, v) = \sum_{k=0} \delta(u - u_k, v - v_k) \quad (2.14)$$

where each 2D delta function represents a different sampled visibilities (u_k, v_k) of their actual measured visibilities. We then substitute equations 2.14 to 2.13 to obtain a “dirty image” of the sky:

$$\begin{aligned} I^D(l, m) &= \int \int V(u, v) S(u, v) \exp[-2\pi i(ul + vm)] du dv \\ &= B(l, m) * I_v(l, m) \end{aligned} \quad (2.15)$$

where $B(l, m)$ is called the dirty beam or point spread function (PSF), a Fourier transform of the sampling function.

In order to retrieve the true sky brightness $I(l, m)$ from the dirty image $I^D(l, m)$, a process called deconvolution is performed. A common approach is based on CLEAN family of algorithms originally introduced by (Högbom, 1974).

The CLEAN algorithm begins by utilising the dirty beam and dirty image. The process involves iteratively subtracting the dirty beam from the dirty image until either a specified noise level is achieved or the maximum iterations are reached. The extracted CLEAN components are stored in a model image, with each source being represented by a basic delta function. The resulting image is formed by convolving the CLEAN model of the field with the so-called CLEAN beam (or restoring beam), which is then added back onto the final residual image. The CLEAN beam is usually formed as a 2D Gaussian fit to the main lobe of the dirty beam.

2.3.1 Imaging Software

In this project, we utilised the software called WSCLEAN (Offringa et al., 2014). WSCLEAN employs w-stacking, where the visibilities are gridded onto different w-layers and w-projection corrections are applied after performing inverse Fourier transforms. Additionally, WSCLEAN offers useful functionalities such as multi-scale CLEAN (Offringa and Smirnov, 2017), multi-frequency synthesis (MFS) imaging (Conway et al., 1990), and Briggs robust weighting (Briggs, 1995).

Additionally, algorithms like WSCLEAN may leave behind residuals with flux density levels equal to the stopping threshold. Consequently, the flux from sources can be systematically underestimated in the recovered model image, leading to biases in self-calibration procedures (Offringa and Smirnov, 2017). To mitigate this issue,

WSCLEAN utilises an imaging mask. This mask can be generated automatically, such as using the auto-mask parameter in WSCLEAN, or by incorporating an external mask, such as BREIZORRO, which creates a binary mask ([Ramaila et al., 2023](#)).

2.4 Calibration

In order to improve the image quality from the imaging process, calibration is performed. Calibration is a process that corrects for instrumental and atmospheric effects within radio interferometric data. These effects are categorized into two types: direction-independent effects (DIEs), which remain constant across the sky (e.g., antenna gains), and direction-dependent effects (DDEs), which vary based on the position of each source in the sky (e.g., primary beam or ionosphere). This section introduces the radio interferometer measurement equation (RIME), which provides a mathematical framework to represent the systematic errors affecting visibility ([Hamaker et al., 1996](#); [Smirnov, 2011](#)). Furthermore, we give an overview of calibration strategies (1GC, 2GC, and 3GC). This section concludes with a brief summary of the software used in this research.

2.4.1 Radio Interferometer Measurement Equation (RIME)

In mathematical terms, calibration involves determining gains that minimise the error between the measured visibilities taken by an interferometer and the predicted visibilities. The predicted visibilities are calculated using the radio interferometer measurement equation (RIME), which accounts for all instrumental and atmospheric factors. Below, we outline the mathematical derivation of the RIME.

Assuming a quasi-monochromatic plane wave, we can represent its state at any point using a complex vector. If we utilise a perpendicular coordinate system (xyz) , where the z -axis points from the antenna to the source, the signal vector originating from the source can be expressed as follows:

$$\mathbf{e} = \begin{bmatrix} e_x \\ e_y \end{bmatrix} \quad (2.16)$$

When the signal from the source arrives at the (usually dual) antenna feeds, it is converted into two complex voltages denoted by ‘ \mathbf{v} ’. These voltages are expressed as follows:

$$\mathbf{v} = \begin{bmatrix} v^X \\ v^Y \end{bmatrix} = \mathbf{J}\mathbf{e} \quad (2.17)$$

The “Jones matrix”, denoted as \mathbf{J} , is a 2×2 matrix representing all atmospheric and instrumental effects. In the case of a two-element interferometer, introduced in Section 2.2, each antenna is labelled as ‘p’ and ‘q’. Each antenna generates a voltage vector, which we’ll denote as v_p and v_q . These voltage vectors are correlated to produce four pairwise correlations, which can be organised into a visibility matrix (\mathbf{V}_{pq}) as shown below:

$$\mathbf{V}_{pq} = \begin{bmatrix} \langle v_p^X \bar{v}_q^X \rangle & \langle v_p^X \bar{v}_q^Y \rangle \\ \langle v_p^Y \bar{v}_q^X \rangle & \langle v_p^Y \bar{v}_q^Y \rangle \end{bmatrix} \quad (2.18)$$

Here, the notation $\langle \cdot \rangle$ signifies the average over the correlator channel and integration time intervals, and $(\bar{\cdot})$ denotes the conjugate. To simplify the expression, Equation

2.18 can be represented as the product of two vectors:

$$\mathbf{V}_{pq} = 2 \left\langle \begin{bmatrix} v_p^X \\ v_p^Y \end{bmatrix} \begin{bmatrix} \bar{v}_q^X & \bar{v}_q^Y \end{bmatrix} \right\rangle = 2 \langle v_p v_q^H \rangle \quad (2.19)$$

The symbol $\langle \cdot \rangle^H$ refers to the Hermitian transpose. We then rewrite Equation 2.19 in terms of Equation 2.17, we can now express the complex output (the visibility matrix) of the interferometer as follows:

$$\mathbf{V}_{pq} = 2 \langle \mathbf{J}_p \mathbf{e} (\mathbf{J}_q \mathbf{e})^H \rangle = 2 \langle \mathbf{J}_p (\mathbf{e} \mathbf{e}^H) \mathbf{J}_q^H \rangle \quad (2.20)$$

As we conduct averaging across both frequency and time, we operate under the assumption that the Jones matrices remain constant throughout these averaging periods. This allows us to factor them out of the averaging operation $\langle \cdot \rangle$, and then we can express the visibility matrix as follows:

$$\mathbf{V}_{pq} = 2 \mathbf{J}_p \langle \mathbf{e} \mathbf{e}^H \rangle \mathbf{J}_q^H = \mathbf{J}_p \begin{bmatrix} \langle e_x \bar{e}_x \rangle & \langle e_x \bar{e}_y \rangle \\ \langle e_y \bar{e}_x \rangle & \langle e_y \bar{e}_y \rangle \end{bmatrix} \mathbf{J}_q^H \quad (2.21)$$

The elements within the matrix in Equation 2.21 have a simple relationship to the Stokes parameters, which characterise the polarisation properties of an electromagnetic signal. We can denote them collectively as \mathbf{B} , representing the brightness of the source.

$$2 \begin{bmatrix} \langle e_x \bar{e}_x \rangle & \langle e_x \bar{e}_y \rangle \\ \langle e_y \bar{e}_x \rangle & \langle e_y \bar{e}_y \rangle \end{bmatrix} = \begin{bmatrix} I + Q & U + iV \\ U - iV & I - Q \end{bmatrix} = \mathbf{B} \quad (2.22)$$

This leads to the most fundamental version of the RIME below.

$$\mathbf{V}_{pq} = \mathbf{J}_p \mathbf{B} \mathbf{J}_q^H \quad (2.23)$$

The quantity \mathbf{V}_{pq} can now be expressed as the product of three 2×2 matrices. \mathbf{J}_p and \mathbf{J}_q represent the instrumental response, while \mathbf{B} characterises the properties of the source. We can also rewrite the RIME in onion form, which is expressed as:

$$\mathbf{V}_{pq} = \mathbf{J}_{pn} (\dots (\mathbf{J}_{p2} \mathbf{J}_{p1} \mathbf{B} \mathbf{J}_{q1}^H \mathbf{J}_{q2}^H) \dots) \mathbf{J}_{qm}^H \quad (2.24)$$

Because matrix multiplication doesn't generally commute, it's important to maintain the order of matrices to reflect the sequence of propagation effects. The equation can be reformulated for N sources as the following summation:

$$\mathbf{V}_{pq} = \sum_s \mathbf{J}_{sp} \mathbf{B}_s \mathbf{J}_{sq}^H \quad (2.25)$$

We then present a complete set of Equations 2.25, highlighting the distinction between terms on the left that relate to the antenna's perspective along the propagation path and those on the right that originate from the source's perspective.

$$\mathbf{J}_{sp} = \mathbf{G}_p \mathbf{E}_{sp} \mathbf{K}_{sp} \quad (2.26)$$

In Equation 2.26, G_p presents the collective impact of all terms independent of the source (direction), while E_{ps} serves as its counterpart for direction-dependent factors. K_{sp} is a scalar phase term dependent on the source, originating from geometric delays. By integrating Equation 2.26 into Equation 2.25, we arrive at the following

result:

$$\begin{aligned} \mathbf{V}_{pq} &= \mathbf{G}_p \sum_s (\mathbf{E}_{sp} \mathbf{K}_{sp} \mathbf{B}_s \mathbf{K}_{sq}^H \mathbf{E}_{sq}^H) \mathbf{G}_q^H \\ &= \mathbf{G}_p \left(\sum_s \mathbf{E}_{sp} \mathbf{X}_{spq} \mathbf{E}_{qs}^H \right) \mathbf{G}_q^H \end{aligned} \quad (2.27)$$

This version of the RIME provides a foundation to start discussing the calibration. If we recall, calibration involves finding a Jones matrix (\mathbf{J}) that minimises the difference between some observed and predicted visibilities. In mathematical terms, we can represent calibration as a process that minimizes:

$$\chi^2 = \sum_{pq} \left\| D_{pq} - \mathbf{G}_p \left(\sum_s \mathbf{E}_{sp} \mathbf{X}_{spq} \mathbf{E}_{qs}^H \right) \mathbf{G}_q^H \right\| \quad (2.28)$$

Where D_{pq} is the observed data. Once a suitable \mathbf{J} is calculated, we can then use them to recover the corrected visibilities from the observed visibilities. Since calibration is an optimisation problem, various optimisation algorithms exist. In this research, we will not discuss the calibration algorithms in detail. While there are a lot of optimization algorithms, the two commonly implemented in radio astronomy are the Levenberg Marquardt (LM) (Levenberg, 1944; Marquardt, 1963) and statistically efficient and fast calibration (STEFICAL) (Mitchell et al., 2008; Smirnov and Tasse, 2015). LM splits the complex visibilities and gains into real and imaginary parts before solving for gains. (STEFICAL) solves the gains directly in complex numbers (complex optimisation).

2.4.2 Calibration Strategies

Over time, many calibration approaches have been developed to address issues arising from the advancement of new radio telescopes with a wide field of view and

advanced technology (Fomalont and Perley, 1999). Typically, calibration is divided into three categories (Noordam and Smirnov, 2010): first-generation calibration (1GC), second-generation calibration (2GC) (Cornwell and Wilkinson, 1981), and third-generation calibration (3GC) (Noordam and Smirnov, 2010).

- **First Generation Calibration**

The first-generation calibration techniques involve employing a calibrator field to correct propagation and instrumental effects, subsequently applying these corrections to the target (Fomalont and Perley, 1999). In order to accomplish this, the target and calibrator fields must be observed intermittently. There are generally two types of calibrators: primary and secondary calibrators. The primary calibrator is a very strong source with a stable and well-known flux density at the observing frequency, usually a bright quasar. Secondary calibrators are moderately bright sources that are known to be compact (e.g., extragalactic point sources). There are very few primary calibrators (only two for MeerKAT), so for any given observation, they are likely to be far away from the target source and, therefore, not suitable to calibrate atmospheric effects as they will be seen through different parts of the atmosphere. A secondary calibrator (closer to the target) is therefore employed. Below, we describe the systematic effects addressed by 1GC (a.k.a. cross-calibration):

- **Bandpass and fluxscale calibration (B)**, primary calibrators are observed to correct for the antenna and receiver’s frequency response and overall flux scale.
- **Delay calibration (K)**, primary or secondary calibrators are observed to correct phase delays that appear as a frequency slope. These calibrators are used to equalise the propagation difference of each signal at the correlators and remove the large phase slopes across the observing band.

- **Gain calibration (G)**, Secondary calibrators are employed to determine the complex-valued gains specific to each antenna. These gains are interpolated over time and applied to the target field.

- **Second Generation Calibration**

After executing 1GC and applying the antenna gains derived from the calibrators to the target field, an initial image of the target can be obtained. The image quality can be improved further using the self-calibration framework (see Figure 2.2). Self-calibration is the technique that uses the target data itself to improve the antenna-based gain corrections further. Furthermore, self-calibration interchanges between the image and visibility domains. Below, we will describe the self-calibration framework.

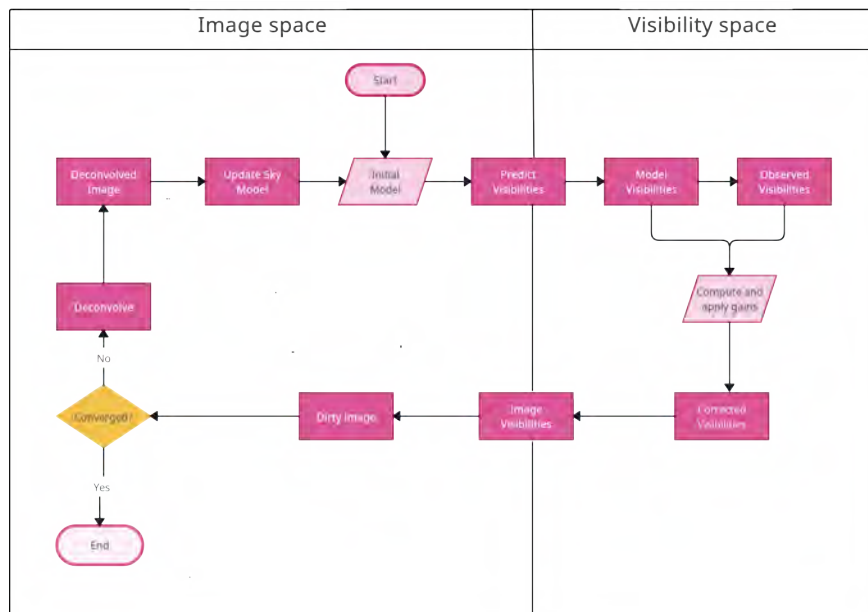


FIGURE 2.2: Flow diagram outlining the procedure for self-calibration.

- **Initial model**, obtain an initial model of the sky. The method used to obtain this initial model can vary. A common approach is to use an

imager such as WSCLEAN, and utilise the resulting model image as the initial model.

- **Predict visibilities**, convert the model image back to visibility via Fourier transform; this is normally an additional function of the imager. This is referred to as the prediction step.
- **Compute and apply gain**, the model visibilities are used with the observed visibilities to solve for complex gains. The output is the corrected visibilities.
- **Dirty image and Deconvolve**, dirty image is formed by taking the inverse Fourier transform of the corrected visibilities. The dirty image is then deconvolved to make a restored image and a new model image. This updated model is then used as the input to step 1, and the entire process is repeated until convergence (or a subjectively satisfactory result).

- **Third Generation Calibration**

The advancement of new telescopes with increased field-of-view and sensitivity causes direction-dependent effects to become prominent. These include frequency attenuation due to the decreasing primary beam gain at higher frequencies (frequency-dependent gain) and the rotation of the beam through parallactic angle, leading to a gradual, time-variable primary beam gain on off-axis sources. These effects result in PSF-like error patterns in images that change over time, creating challenges when identifying variable sources in the image plane. Addressing these issues requires direction-dependent calibration. Various methods for this exist, such as the differential gains approach proposed by [Smirnov \(2011\)](#) and the peeling approach introduced by [Noordam \(2004\)](#), and [Noordam and Smirnov \(2010\)](#). Differential gains prevent gain contamination by solving for the DIES and DDEs simultaneously. The peeling process, in

most cases, proceeds by subtracting the strongest source in the field and then the second strongest, etc. Below, we give a summary of the peeling strategy normally performed by calibration software:

- Select the brightest sources.
- Calculate self-calibration solutions (i.e., delays) for the chosen sources.
- Subtract (or peel) the contribution of the selected source’s sky model multiplied by the relevant gains.
- Repeat the process, starting again with step one, using the updated visibilities.

2.4.3 Calibration Software

In this project, we made use of QUARTICAL software ([Kenyon et al., 2023](#)) to perform 2GC and 3GC. QUARTICAL is the successor to CUBICAL ([Kenyon et al., 2018](#)). It implements a suite of fast radio interferometric calibration routines exploiting complex optimization.

2.5 Image-plane transient and variable searches

This section provides an overview of techniques for identifying transients and variable sources within images. Typically, detecting these sources involves the construction of light curves and employing statistical analysis to detect significant variability. There are two standard methods used to detect variables and transients: difference imaging ([Alard and Lupton, 1998](#)) and catalogue comparison ([Hopkins et al., 2002](#)).

Difference imaging involves the creation of a reference image (based on the full-time integrated synthesis image), which is then subtracted from individual snapshots taken at different epochs. The resulting differential images have positive or negative excursions, indicating whether a source has brightened or dimmed compared to the reference image. Additionally, light curves can be generated using e.g. aperture photometry to show the variation in the differential images. This approach has proven effective in optical data analysis, especially in densely populated regions where catalogue-based surveys may encounter difficulties due to source overlap (Alard and Lupton, 1998).

The catalogue comparison approach varies depending on the timescale of interest. Generally, catalogues are compared directly. This approach involves first using a source finder to establish a reference catalogue and then measuring the flux independently for each epoch, followed by matching the catalogues to create light curves.

In this thesis, we employed a somewhat hybrid approach. We begin with a catalogue approach, utilising a source finder to identify and characterise sources in the full synthesis image. For every catalogue coordinate, we then convert them into pixel positions, and measure the flux within a time-axis cube to create high time-resolution light curves. The light curves are then analysed for variability. One drawback of the catalogue approach is that if a source hasn't been identified beforehand in the full image, it won't be measured. This means that truly short transient events, occurring only once and thus too weak to show up in the full image, will go unnoticed.

2.5.1 Variability Statistics

Astronomical objects display a wide range of patterns in their light curves. For instance, interstellar scintillation manifests as random fluctuations around an average

value, while flares occur intermittently, with brief active periods and much higher brightness compared to their usual state. To identify if a light curve is more or less variable compared to other light curves, we calculate two variability statistics, reduced chi-square (χ_v^2 : (Bridle et al., 1977)) and variability index (ν : (Bridle et al., 1977)).

- **Reduced χ^2 (η)**

η is a statistical test used to assess the goodness of fit of a model to a set of data points. In light curves, it can be used to search for significant variations. η is calculated by comparing the observed flux values in the light curve to the expected flux values under the assumption of constant flux. The formula is:

$$\eta = \frac{1}{N - p} \sum_{i=1}^N \frac{(S_i - \bar{S})^2}{\sigma_i^2} \quad (2.29)$$

$$\bar{S} = \frac{\sum_{i=1}^N S_i / \sigma_i^2}{\sum_{i=1}^N 1 / \sigma_i^2} \quad (2.30)$$

\bar{S} is the weighted mean flux density, σ_i is the estimated flux density error, S_i is the flux density measured. p is the number of parameters in the model. $v = N - p$ are the degrees of freedom. A smaller η_v indicates a better fit to the data if η_v , while bigger η_v indicates a poor fit to the data. A η_v threshold is established to distinguish between variable and non-variable light curves. Light curves with η_v values greater than the established threshold are considered variable and η_v less than the threshold are considered non-variable

- **Variability index(ν)**

The variability index is defined as a ratio of the standard deviation of the flux to the mean:

$$\nu = \frac{\sigma}{\bar{S}} \quad (2.31)$$

where σ is the standard deviation of the light curve and \bar{S} is the weighted mean. The variability index (ν) is useful for quantifying the variation of a source. Additionally, the η statistic takes into account signal-to-noise (via σ_i), the variability index does not, hence it is generally useful as a measure of variability in the case of high SNR.

Chapter 3

Methodology: Data Reduction, Light Curve Extraction, and Light Curve Analysis

The main aim of this project is to mine MeerKAT data for minute-to-hour timescale transients and variable sources. This section reviews the methodology used to mine for these sources. Figure 3.1 provides the summary of the methodology workflow. Most importantly, we review the pipeline currently being developed for minute-to-hour timescale transient and variable source detection by the RATT team at Rhodes University. We call this the PARROT pipeline, based on the name of the project (and object) for which it was initially developed (Smirnov et al., 2023). Briefly, the PARROT pipeline proceeds by self-calibrating or peeling the data and constructing a deep sky model (based on the full time integrated synthesis image), then subtracting the corresponding model visibilities from the corrected visibilities to emphasise any temporal variations in the sources. The model-subtracted visibilities are imaged into high-time-cadence (up to the native correlator resolution) snapshots

and stacked into a time-axis cube. Finally, the cube is used to extract light curves of sources.

As we recall, interferometric data reduction can be implemented in three stages, or “generations” (1GC, 2GC, 3GC). In our project, 1GC is executed prior to running the PARROT pipeline using CARAcacal (Józsa et al., 2020), while light curve analysis is conducted post-running the pipeline. Consequently, we structure the methodological explanation into three parts: (i) data reduction: 1GC, (ii) PARROT pipeline: data reduction (2GC and 3GC) and light curve extraction, and (iii) light curve analysis.

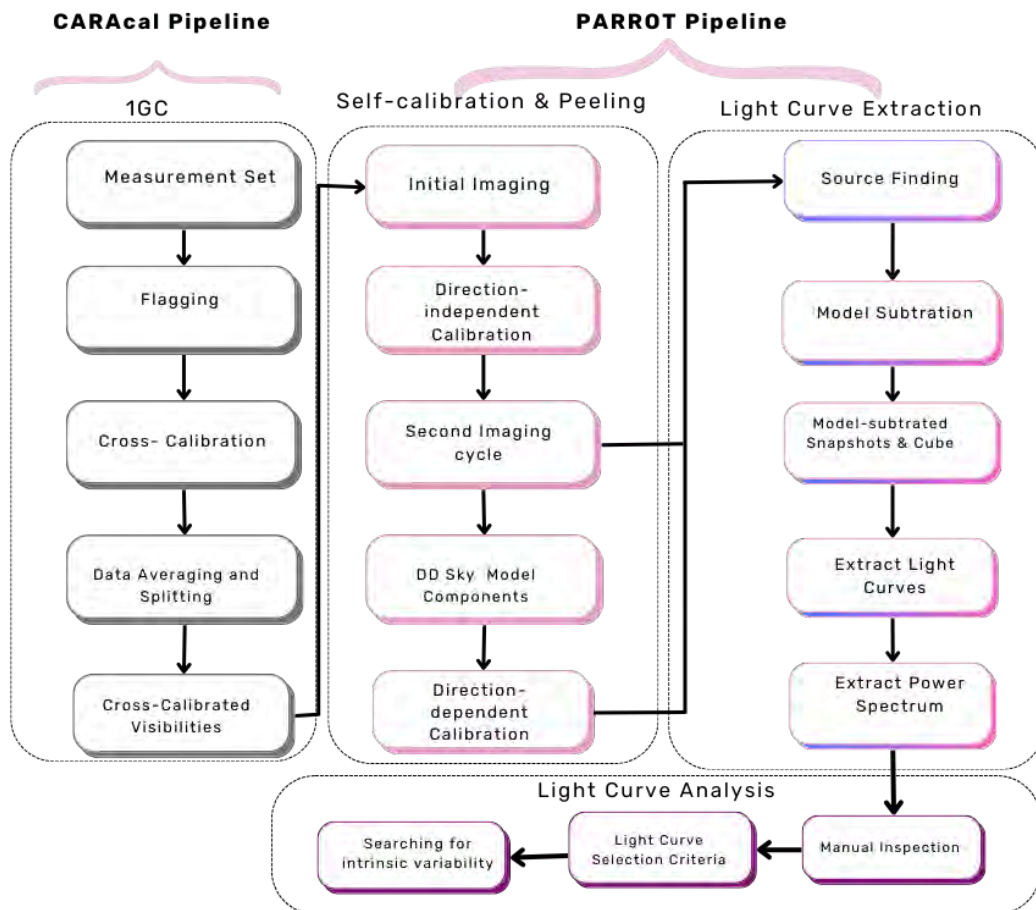


FIGURE 3.1: A workflow summarising the methodology.

3.1 Data Reduction: 1GC

CARACAL is a dedicated pipeline developed for reducing HI and continuum data from the MeerKAT and other radio interferometric telescopes (Józsa et al., 2020). It is built upon STIMELA (Makhathini, 2023), a platform-independent radio interferometric pipelining framework that uses containerisation technologies like Docker and Singularity. STIMELA facilitates the seamless integration of various software packages, making it well-suited for pipeline creation. The CARACAL pipeline outputs are calibrated visibility data, continuum and spectral images, and diagnostic plots for an MS. The configuration of the CARACAL pipeline includes flagging, cross-calibration (a.k.a. 1GC), and self-calibration. However, this project only utilises CARACAL up to and including the 1GC. Below, we give a more detailed description of 1GC via CARACAL:

- Initial step involves splitting calibrator scans from the measurement set (MS).
- Subsequently, calibrator scans undergo flagging for radio frequency interference (RFI) using either AOFLAGGER (Offringa, 2010) or TRICOLOUR (Hugo et al., 2022).
- Solutions for bandpass, delays, flux scale, and gains are obtained.
- Likewise, target data is then split out, and the previously derived solutions are applied to it.
- The target data is (optionally) averaged down in frequency, and additional flagging for RFI is performed.

3.2 PARROT Pipeline

This section describes the PARROT pipeline, which performs 2GC (a.k.a. self-calibration) and 3GC (peeling) and ultimately extracts the light curves. The PARROT pipeline is shown in the Appendix A. Note that the parameters used in the pipeline are specific to the COSMOS MS (refer to Chapter 4). In this section, we will explain a general overview of the pipeline.

3.2.1 PARROT pipeline automation

The pipeline is developed within the framework of STIMELA2¹ (the successor of STIMELA), a workflow management system for developing radio interferometry data pipelines (Smirnov et al., in prep). STIMELA2 revolves around two essential components: a `recipe` and `cabs`. A `recipe` is a YAML document that defines the sequence of tasks to be executed within the pipeline. In contrast, `cabs` are YAML documents that define specific tasks, which can be run in the operating system natively, via Singularity containers, or within a Python virtual environment. The parameters of every pipeline step are validated against the underlying schema of the `cab` that executes it. A `cab` can take various forms – e.g. a binary executable, a Python function, a CASA task.

The pipeline offers the flexibility of running each step independently (e.g., `$ stimela run recipe.yml -s step-1`). Alternatively, (a subset of) the pipeline steps can be executed in sequence (e.g., `$ stimela run recipe.yml -s from_step:to_step`). The pipeline’s core input is a measurement set that has been calibrated with 1GC, and its primary output consists of light curves extracted from all the sources in the restored image of the field. Below, we will describe the pipeline key steps.

¹<https://github.com/caracal-pipeline/stimela>

3.2.2 PARROT pipeline key steps

3.2.2.1 Self-calibration and peeling

(I) Initial Imaging, is performed in three steps, `image-1`, `mask-1`, and `predict-1` respectively.

- `image-1` uses WSCLEAN to create a deconvolved image (a.k.a restored image), dirty image, sky model image, PSF image and residual image. Preferred imaging parameters such as the number of channels, loop gain, weighting scheme, the number of clean iterations, pixel scale, threshold, auto-masking or an external mask, and multi-frequency synthesis (MFS) imaging options are specified within the recipe or externally.
- `mask-1` uses the deconvolved image from the previous step to make an imaging mask using the BREIZORRO tool (Ramaila et al., 2023), parameters such as the threshold are specified. A typical starting threshold at this stage is 10σ , we then lower the threshold at subsequent stages.
- `predict-1`: uses WSCLEAN to convert the sky model image from `image-1` back to model visibilities. This step utilises similar parameters as `image-1`.

(II) Direction-independent calibration, is executed in a single step, `DI_cal-2` using QUARTICAL (Kenyon et al., 2023). Its purpose is to address direction-independent effects such as delays (K). QUARTICAL utilises the sky model from `predict-1` and observed visibilities to calculate gains. To mitigate for potential absorption of artefacts from the model into the calibration solution, QUARTICAL uses solution intervals (time and frequency), which reduces overfitting. One may refer to the guidelines provided in Sob et al. (2021) to make an optimal selection of these intervals.

(III) **The second imaging cycle**, is executed similarly to the initial imaging process. This cycle involves the same three steps: `image-2`, `predict-2`, and `mask-2`.

- `image-2` generates a deep restored image from DI-calibrated visibilities with similar imaging parameters as `image-1`. However, this step uses an external mask (from `mask-1`), a lower threshold, and a higher number of minor cycles are allowed.
- `mask-2` uses the deconvolved image from the previous step to create a mask with a more stringent threshold compared to `mask-1`.
- Finally, `predict-2` is employed to convert the new sky model image back to model visibilities. One might conclude self-calibration at this stage if the field exhibits no DDEs. However, if DDEs which manifest as PSF-like artefacts are visible in the restored image, direction-dependent calibration is performed as described below.

(IV) **Predicting direction-dependent sky model components**, is executed at step `DD_predict`, using `CRYSTALBALL`². A sky model list from the previous imaging process and DS9-format region files identifying one or more problematic bright sources are specified (prior to running the pipeline). `CRYSTALBALL` is then used to predict the model visibilities associated with these sources into separate model column(s).

(V) **Direction-dependent calibration**, is executed at the `DD_cal-2` using `QUARTICAL` to address direction-dependent effects. Each problematic off-axis source specified at `DD_predict` is peeled, which results in a reduction of PSF-like artefacts. The

²<https://github.com/caracal-pipeline/crystalball>

peeling process is described in Section 2.4.2. The final output from this step are corrected visibilities with off-axis sources peeled off (DD-corrected visibilities).

(VII) The final imaging cycle, is executed similarly to the second imaging cycle. `image-3` is used to make the final restored and sky model images from DD-calibrated visibilities. `Predict-3` is used to convert the final sky model image back to visibilities.

3.2.2.2 Light curve extraction

(VIII) Source Finding, is performed in three separate steps (of which only the first one is strictly required):

- `source-catalog` makes use of the PYBDSF (Mohan and Rafferty, 2015) source finder to identify and characterise sources in the final restored image (from the final imaging cycle). The parameters consist of the `rms_box`, for example `[40, 20]`. The first integer (40) is a 2-D sliding box utilised to compute the rms and mean of the image, while the second integer (20) indicates the desired number of pixels in each box. Additional parameters, such as `thresh_pix` and `thresh_isl`, are employed to identify pixels in the image with brightness surpassing the local background noise by a specified factor (e.g., 4). Subsequently, the algorithm identifies islands of high pixel values exceeding the noise threshold (e.g., 6). These islands undergo iterative fitting processes into point and 2D Gaussian components. The final process creates a catalogue detailing the properties of the sources (such as flux density, right ascension (RA), declination (DEC), etc).

- `source-catalog-spi` is an optional step (not run by default). This uses a per-band image cube, corrected for the nominal primary beam, to get spectral index estimates for the detected sources.³
- An optional `make-master-catalog` step merges multiple source catalogues into one master catalogue. This is not necessary for single-epoch observations, but when dealing with multi-epoch observations, a master catalogue is useful for consistent labelling of the extracted light curves.

(IX) Subtract sky model, is executed by the `subtract_model` step. This uses the standard `taql` utility to subtract the model data column generated by the last `predict` step from the DD-corrected data, yielding residual or model-subtracted visibilities. This model subtracted data serves as the basis for detecting significant genuine variations in sources. Note that this step is conceptually similar to the continuum subtraction step used in spectral line studies.

(X) Model-subtracted snapshots and cube stacking, this is performed in two steps: `image-htc`, and `cube-htc`:

- `image-htc` employs `WSCLEAN` to generate high-time cadence snapshot images, utilising the “`intervals-out`” parameter. Depending on the specific timescale of interest (and available disk space), the snapshot images may be created at cadences up to the native correlator resolution (2 or 8s for MeerKAT).
- `cube-htc` stacks the individual snapshot images into a single FITS cube using the `fitstools.py` utility.

³In-band spectral indices (particularly of fainter sources) determined by such source finders are unreliable and not treated as a science product here. The only reason the PARROT pipeline includes this step is to detect “interesting” i.e. flat- or inverted-spectrum sources, which tend to be prime candidates for rapid variability and scintillation.

(XI) Light curve extraction, is performed by a Python function, `image_utils.py`. For every source in the catalogue created earlier, global coordinates are converted into pixel position (x, y) . Light curves are extracted as follows:

- For each pixel position (x, y) , the flux is obtained by taking a slice through the cube, $S_i(x, y)$, where i represents the time axis of the cube.
- For each pixel position (x, y) , the flux error (at time i) is estimated as the rms of pixel values centred around the pixel position (x, y) .
- Finally, a light curve is plotted utilising the fluxes and flux errors, and the corresponding timeslots.
- Variability metrics (discussed below) are also calculated for each plotted light curve.

Note that due to the procedure outlined above, the resulting light curves are effectively mean-subtracted, a.k.a. differential. The time-averaged flux of each source (as captured by the deep sky model) has been subtracted from the visibilities from which the snapshots are made. The extracted light curves, therefore have a mean of zero, and any variability in the source will manifest itself as positive or negative excursions with respect to zero.

For each light curve extracted, some statistics are calculated in order to quantify its variability. The first variability metric is the variability index (v). This offers a dimensionless value that characterizes the degree of variability within a light curve in relation to the mean flux. Higher v values indicate a heightened level of variability relative to the mean flux, while lower values signify less variability. The v metric is

computed as:

$$v = \frac{\sigma}{\bar{S}}, \quad (3.1)$$

$$\sigma = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (S_i - \bar{S})^2} \quad (3.2)$$

Where σ represents the standard deviation of the light curve, and \bar{S} is the flux of the source based on the full synthesis image (as reported by the source finder). The second metric we compute is the variability maximum of a source, determined by calculating the ratio between the maximum absolute value of the (differential) light curve and the observed flux of a source. The formula is expressed as follows:

$$V_{\max} = \frac{|L_{\max}|}{\bar{S}} \quad (3.3)$$

Another useful variability metric is the standard deviations ratio denoted by vm , where we calculate the ratio between the standard deviation and median standard deviation of each light curve. The standard deviation is sensitive to outliers, meaning a few extreme values can significantly impact its value. In contrast, the median standard deviation is less affected by outliers since it is based on the median, which is less influenced by extreme values. Therefore, a high ratio suggests that the light curve may exhibit outliers or extreme values affecting the standard deviation. The formula for calculating the vm ratio is:

$$vm = \frac{\sigma}{\sigma_{\text{median}}}, \quad (3.4)$$

$$\sigma_{\text{median}} = \begin{cases} \sigma_{(\frac{N+1}{2})} & \text{if } N \text{ is odd} \\ \frac{\sigma_{(\frac{N}{2})} + \sigma_{(\frac{N}{2}+1)}}{2} & \text{if } N \text{ is even} \end{cases} \quad (3.5)$$

The final variability metric calculated is the reduced χ^2 statistic, denoted as η (refer to section 2.5.1 for a detailed explanation). This metric is useful in assessing the goodness of fit between the constant light curve represented by the mean flux and the actual light curve (\bar{S}). A lower value indicates a better fit and, in our case, a non-varying light curve. A higher value indicates a poor fit, meaning the light curve varies significantly. The formula is given by:

$$\eta = \frac{1}{N-p} \sum_{i=1}^N \frac{(S_i - \bar{S})^2}{\sigma_i^2} \quad (3.6)$$

$$\bar{S} = \frac{\sum_{i=1}^N S_i / \sigma_i^2}{\sum_{i=1}^N 1 / \sigma_i^2} \quad (3.7)$$

3.3 Light Curve Analysis

3.3.1 A search for intrinsic variability

The variability metrics (η, v) for each light curve from the pipeline are used to search for intrinsic variability⁴. The reduced χ^2 statistic (η) indicates the significance of variation in a light curve. We establish a threshold for distinguishing between variable and non-variable sources by calculating a $\log(\eta)$ distribution fitted with a normal distribution characterised by parameters ($\mu = \mu_{\log \eta}$) and ($\sigma = \sigma_{\log \eta}$). The

⁴Through the rest of this work, we use the term ‘‘intrinsic variability’’ in the sense of astrophysical variability (i.e. intrinsic to the source itself, or due to scattering effects), as opposed to variability induced by instrumental or atmospheric effects.

threshold (η_{thresh}) is determined by $\eta_{\text{thresh}} = \mu + n\sigma$, where n is a tunable parameter, a good starting value is $n = 2$. Any source with $\eta > \eta_{\text{thresh}}$ is considered significantly variable.

To estimate the mean (μ) and standard deviation (σ) of the η distribution, we adopt a strategy outlined by Rowlinson et al. (2016), Sarbadhicary et al. (2021) and Wang et al. (2023). We initially construct a histogram of $\log(\eta)$ using a Bayesian block for binning. Bayesian Blocks belong to a category of algorithms from ASTROPY⁵ that select the most advantageous binning strategy, aiming to minimise the error in the histogram’s representation of the data. Lastly, we apply a non-linear least squares method to fit a normal distribution to the histogram. From the fitted distribution, we can then calculate the mean and standard deviation. A Python code for calculating the threshold can be found in Appendix C.

The variability index (v) serves to measure the extent of variation in light curves. In this project, we will not employ the variability index as a criterion for intrinsic variability detection, as demonstrated in Wang et al. (2023). Instead, it will be used solely to quantify the variability of variable sources. For instance, a source with $\eta > \eta_{\text{thresh}}$ and $v > v_{\text{thresh}}$ will be categorised as having “moderate-high” variability, while sources with $\eta > \eta_{\text{thresh}}$ and $v < v_{\text{thresh}}$ will be categorised as “low-moderate” variability (Sarbadhicary et al., 2021).

⁵<https://www.astropy.org/>

Chapter 4

Mining MeerKAT Data: COSMOS & MACS J2140.2–2339 Fields

In this section, we discuss the application of the PARROT pipeline to two archival measurement sets from MeerKAT. Firstly, we provide a comprehensive overview of the observation and a detailed explanation of the parameters utilised in the PARROT pipeline. Additionally, we discuss the systematic effects of the extracted light curves, using them as a basis to establish our selection criteria. The criteria are designed to exclude numerous sources affected by systematic effects. Lastly, we outline the analysis of light curves given by the selection criteria to detect intrinsic variability.

4.1 Observations

In line with the primary goal of this project, we implemented the PARROT pipeline on archival measurement sets (MSs) from MeerKAT. We selected the COSMOS

and MACS J2140.2–2339 fields. These measurement sets were chosen without prior knowledge or bias regarding the intrinsic variability of the sources. MACS J2140.2–2339 fields (proposal ID: SCI-20190418-KA-01) is one of the samples of relaxed galaxy clusters selected to detect and characterise the central diffuse radio emission (Trehaeven et al., 2023). Additionally, the COSMOS field serves as one of the four extragalactic fields of the MIGHTEE Large Survey Projects, a galaxy evolution survey aimed at investigating the formation and evolution of galaxies over cosmic time (Heywood et al., 2022).

Table 4.1 provides an overview of the observations. The COSMOS observation spanned 8 hours, with a dedicated on-source time of 6 hours. During these observations, the primary calibrator (J0408–6545) was revisited twice, each scan lasting approximately 10 minutes. The secondary calibrator (J1008 + 0730) was visited ten times for 2–3 min in between the target scans. The MACS J2140.2 – 2339 field was observed for a total of 5 hours. The primary calibrator (J1939 – 6342) was visited once for 10 min. The secondary calibrator (J2152–2828) and the target source (J2140.2–2339) were then visited alternately for 2 and 15 min. Both observations utilised the UHF-band receiver, covering a frequency range from 580 to 1015 MHz. The correlator operated in 4096 channel mode, and the data integration time was set to 8s.

TABLE 4.1: Characteristics of MeerKAT Observations for the COSMOS and MACS J2140.2 – 2339 fields.

Field	R.A	Dec	Track(sec)	N_{channels}	N_{antennas}	Primary Calibrator	Secondary Calibrator
COSMOS	10:00:28.6	+02.12.21.0	28875.5	4096	63	J0408 – 6545	J1008 + 0730
J2140.2 – 2339	21:40:15.180	–23.39.39.90	20648.2	4096	64	J1939 – 6342	J2152 – 2828

4.2 First generation calibration (1GC)

For each MS, we utilised the CARACAL pipeline to perform 1GC. Each MeerKAT observation included scans of targets, as well as primary and secondary calibrators. Using calibrator sources, CARACAL derived corrections for instrument gains and propagation effects (such as instrumental gain fluctuations and atmospheric conditions), which were then applied to the target. Refer to Chapter 3 for a full review of the steps.

4.3 PARROT pipeline

4.3.1 Self-calibration and peeling

After a successful 1GC run, 2GC and 3GC are performed using the PARROT pipeline detailed in this thesis. 2GC uses the target data to correct the previously calculated instrumental complex gains, which improves the image quality. 3GC is employed to address direction dependent effects.

(I) Initial imaging: We created shallow MFS restored images and MFS model images for each measurement set using WSCLEAN (step `image-1`). The parameters used included Briggs weighting, with a value of -0.3 (the value was selected following a thorough evaluation of various weighing schemes and values). This choice was made to minimise potential image artefacts around the affected sources. Additionally, six subbands were used, and a deconvolution threshold of 10σ was set, which mitigated artefact absorption in the sky models. For the COSMOS field, we created a 10125-pixel square image with a scale of 1.7 arcsec; for MACS J2140.2–2339, we

created an 8000-pixel square image with a scale of 1.7 arcsec ¹. Notably, we chose to generate the images in both fields without auto-masking because it was observed that auto-masking led to artefact absorption by the sky models. The resulting restored images were then used to create initial masks using BREIZORRO (`mask-1`), with a threshold set at 6. Finally, initial sky model images were converted into visibilities (`predict-1`).

(II) Direction-independent calibration: For each MS, we used the K-Jones solver of QUARTICAL to correct for delay. The time solution interval was set to 1 integration (i.e., 8s), and the frequency solution interval was set to 0 (i.e., full band), effectively establishing intervals within which gains were constrained to remain constant. This approach helps mitigate the overfitting of the gains. The delay corrected visibilities were stored.

(III) Second imaging cycle: The delay corrected visibilities of each MS were subjected to imaging using the same parameters as the `image-1` step. However, we employed an external mask generated by `mask-1`. In addition, we set the deconvolution threshold to 6σ , and the max minor cycle iterations were set to 1 million. Subsequently, MFS-restored images were obtained. We generated updated external masks (`mask-2`) from these images for the final imaging process. The `predict-2` step used WSCLEAN to convert the sky models created at this stage back into visibilities.

Figures 4.1 and 4.2 show delay corrected MFS-restored images generated for each MS. Each figure shows the reduction of image artefacts from cross calibration (1GC)

¹The size of the image is chosen to ensure that all sufficiently bright off-axis are included in the image and therefore deconvolved, hence the difference between the fields.

to self-calibration (2GC). Additionally, persistent direction-dependent effects (DDEs) in the restored images manifesting as PSF-like artefacts around moderately bright sources were present. Peeling (3GC) is employed to address this issue, and striking the right balance between the number of sources to be peeled and minimising flux suppression in our field is important. A study by [Sob et al. \(2021\)](#) highlights that excessive freedom in Direction Dependent solutions can lead to unintended suppression of unmodeled sources. Since our sky model is constructed based on an 8-hour synthesis image, any inherent temporal variability remains unmodelled by definition, and could be suppressed.




FIGURE 4.1: This figure shows an improvement in the quality of the MFS restored image of the MACS J2140.2 – 2339 field, from 1GC (commonly known as crosscal) to delay self-calibrated visibilities. The upper inserted image represents a segment of the delay self-calibrated MFS restored image, while the lower inserted image represents the cross-calibrated MFS restored image before self-calibration.




FIGURE 4.2: This figure shows an improvement in the quality of the MFS restored image of the COSMOS field, from 1GC (commonly known as crosscal) to delay self-calibrated visibilities. The lower inserted image represents a segment of the delay self-calibrated MFS restored image, while the upper inserted image represents the cross-calibrated MFS restored image before self-calibration.

In both fields, prominent direction-dependent effects (DDEs) are primarily caused by the primary beam. We decided to restrict ourselves to peeling no more than three sources per field to keep unmodelled flux (and variability) suppression down. Therefore, we preferred to peel DDE-affected sources only within the primary beam FWHM ($r \sim 1^\circ$), which we designate to be our Region of Interest (RoI) within which we eventually search for variable sources. This RoI is illustrated by the blue circle in Figure 4.3, while the red squares indicate image artefacts affecting the RoI.

Figures 4.3(a) and 4.3(b) display the COSMOS MFS-restored and model-subtracted images, respectively. The problematic sources in COSMOS also affect the model-subtracted high-cadence images, causing significant image artefacts across the RoI. Figures 4.3(c) and 4.3(d) show the MACS J2140.2–2339 MFS-restored image and the model-subtracted image. The RoI is defined the same as in the previous field, and we also aimed to peel only inside the RoI. Although the artefacts from problematic sources are not directly apparent in the model-subtracted image, they are still visible in the eventual per-snapshot images, inducing apparent temporal variation on nearby sources.

(IV) Direction-dependent sky model components: For each MS, we created circular regions (in DS9 format) corresponding to the bright off-axis sources (in our case, all the sources within the red squares in Figure 4.3). We used CRYSTALBALL to convert the corresponding CLEAN components back to visibilities.

(V) Direction-dependent calibration: We employed QUARTICAL to peel the problematic sources for each MS using the overall sky model, and the direction-dependent sky models from the previous step. We also solved for delays with the same solution interval as direction-independent calibration. QUARTICAL then produced corrected visibilities, with the designated sources peeled off.

(VI) Final Imaging: We then generated final deep MFS images, using the peeled visibilities. The MSs are subjected to imaging using the same parameters as in the second imaging cycle. A deeper external mask generated by `mask-2` is employed. In addition, we set the deconvolution threshold to 3σ . Subsequently, MFS-restored images were obtained. The `predict-3` step was used to convert the sky model




FIGURE 4.3: (a) This image shows the MFS restored image of COSMOS. The circular blue region corresponds to our RoI (one primary beam FWHM), within which we eventually search for intrinsic variability. The two red square regions highlight bright off-axis sources that impact the sources inside the circular blue region. In (b), the figure shows a model-subtracted high-time cadence snapshot image from the COSMOS. The presence of bright off-axis sources in the MFS-restored image introduces artefacts in the model-subtracted images, leading to temporal variations in sources within the circular region. Figures (c) and (d) correspond to the MACS J2140.2–2339 field.

created at this stage back into visibilities. Figure 4.4 shows the reduction of image artefacts in the MACS J2140.2 – 2339 field.




FIGURE 4.4: This figure shows three problematic off-axis sources in the MACS J2140.2–2339 field before and after peeling using QUARTICAL

4.3.2 Light curve extraction

(VII) Source Finding: We created a catalogue of the sources in our fields using the method outlined in Chapter 3.

(VIII) Model subtraction: We subtracted the sky model visibilities (obtained from predict-3) from the corrected visibilities for each MS.

(IX) Model-subtracted snapshots and cube stacking: Utilising the model-subtracted visibilities, we created snapshot images at high time cadence (8 seconds) and stacked them into a time-axis FITS cube.

(X) Light curve extraction: The source catalogues and model-subtracted FITS cubes were used to extract light curves for each field. We extracted light curves for

10241 sources in the COSMOS field, and 13978 sources in the MACS J2140.2 – 2339 field.

Figure 4.5 and 4.6 show examples of extracted light curves (for non-varying sources, in this case) from COSMOS and MACS J2140.2 – 2339.




FIGURE 4.5: Two non-variable light curves from MACS J2140.2 – 2339 field.

4.4 Light Curve Analysis

In this section, we describe our analysis of the light curves extracted by the pipeline. A suspicious number of light curves from both the COSMOS and MACS J2140.2 – 2339 fields are highly variable, suggesting that the variability may not just be a result of intrinsic astronomical phenomena. Systematic errors introduced by instrumental and environmental effects can lead to apparent fluctuations over and above the systematic effects already mitigated by the pipeline’s data processing (calibration). These errors include **primary beam effects**, **propagation effects**, **imaging artefacts**, and **RFI/satellite effects**.

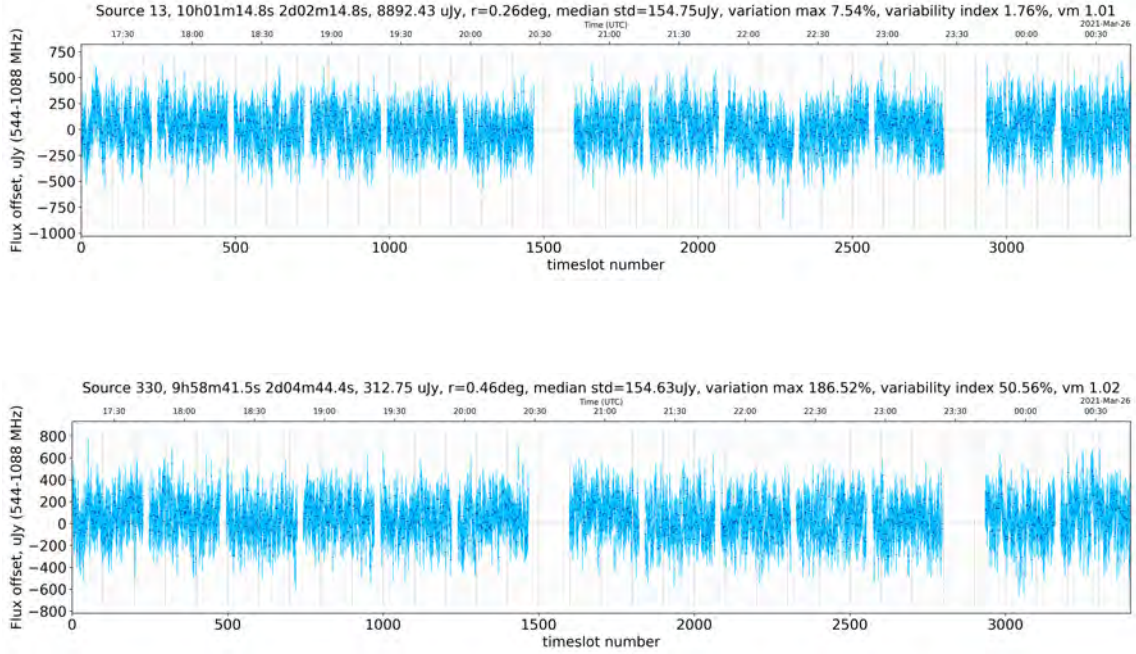


FIGURE 4.6: Two non-variable light curves from COSMOS field

4.4.1 Systematic effects

4.4.1.1 Primary beam effects

Primary beam rotation is the most prominent source of systematic variability in light curves in the COSMOS and MACS J140.2–2339 fields. MeerKAT’s primary beam is slightly elongated, which, coupled with parallactic angle rotation, imprints a slow time-variable primary beam gain onto sources located on the flank of the beam, as each source traces an arc through the primary beam pattern. Figures 4.7 and 4.8 show the primary beam rotation effect on light curves in the COSMOS and MACS J2140.2 – 2339 fields, respectively.

In addition to the primary beam rotation, the primary beam also introduces frequency-dependent effects. The primary beam scales with frequency. The effective observing frequency changes as we move away from the phase centre. While brightness is

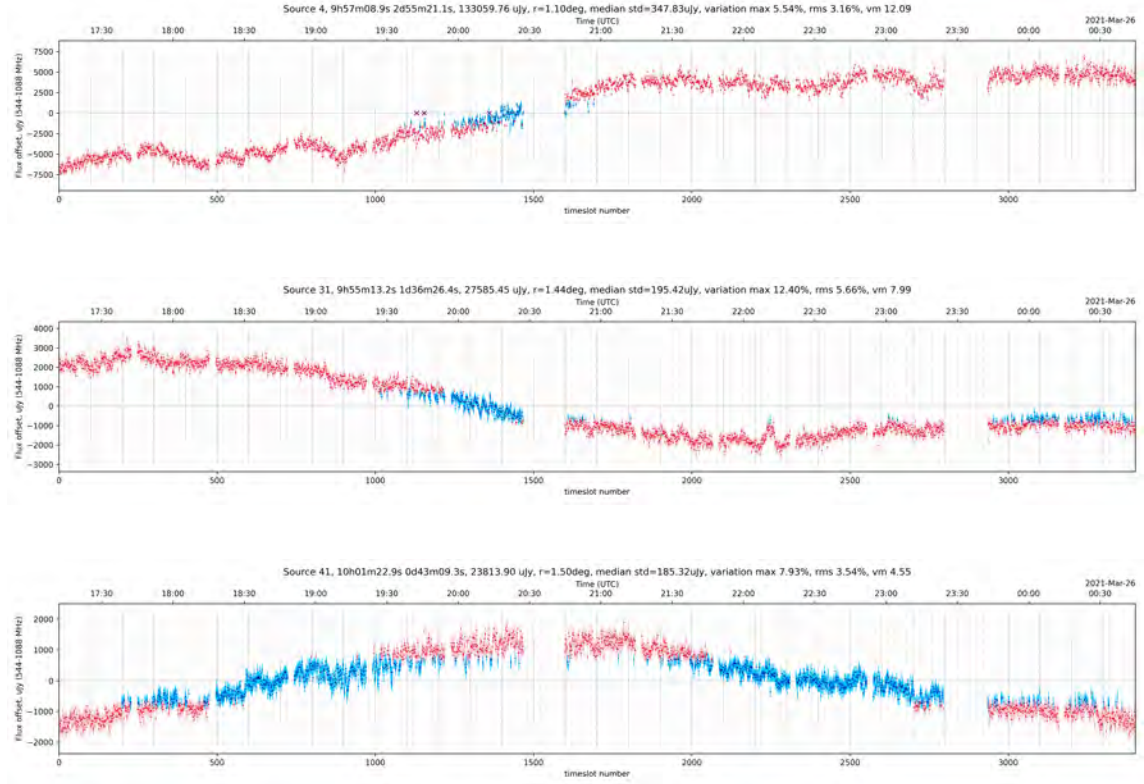


FIGURE 4.7: Three light curves from COSMOS affected by the primary rotation. Each source has an angular separation from phase centre (r) greater than 1. Source 4 has $r = 1.10$ deg, Source 31 has $r = 1.44$ deg, and Source 41 has $r = 1.50$ deg. Significant peak offsets in the light curves are plotted in red.

often quoted at the band centre frequency, this may not accurately represent all sources. The brightness can differ significantly with the range of effective observing frequency. To address this issue, the effective observing frequency is typically calculated as a function of position. This improves the accuracy and constancy of the flux density. However, the corrections for this effect are small, especially for sources with smooth spectra that are not extremely steep, which is the case for most of the sources (Heywood et al., 2022). Therefore, we can ignore frequency attenuations for our transient and variable detection analysis. However, we note most frequency-dependent primary beam effects can probably be effectively addressed by a future frequency-dependent version of the PARROT pipeline, which is outside the scope of

the present work. Below, we only discuss methods for correcting the primary beam rotation.

There are three ways one might try to mitigate the primary beam rotation :

- Model and correct for the rotation: given a primary beam model, one could, in principle, predict the time-variable beam gain in the direction of each source. Detailed holography measurements (de Villiers, 2023) of the MeerKAT beam became available midway through this project. However, even with an accurate beam model, predicting the effect on a *wideband* light curve is not entirely trivial: the primary beam scales with frequency, so the temporal variation is different at different frequencies, and thus the cumulative effect on the wide-band flux measurement is dependent on the spectral shape of the source (steep spectrum sources would appear less variable, while inverted spectrum sources more variable), which is not known a priori.

Therefore, we leave a more detailed investigation of this approach to future work and only sketch two ways of approaching the problem. Firstly, given a light curve (measured along a known arc through the primary beam pattern), we could attempt to fit a simple spectral index model for each source. Secondly, we could introduce primary beam correction into the imaging process, using the IDG gridder of WSCLEAN, or the DDFACET imager (Tasse et al., 2018). Both approaches require additional computational complexity but certainly merit further investigation.

- Compute per-channel light curves in frequency without the need to model the source spectral index by correcting the PB rotation on finer frequency intervals. This approach can improve the accuracy and reduce concerns regarding intrinsic source spectra. However, it's important to note that there may be

a SNR penalty associated with this method, limiting its applicability to all sources.

- Fit and subtract a smooth “baseline” function from each light curve. This has the drawback of absorbing any slow intrinsic variability but would retain rapid variability. Note also that near transit, the primary beam can rotate quite rapidly (particularly at higher elevations), so this approach will have limited utility in those regimes. We also leave this to future investigation.
- Ignore the affected light curves and restrict our analysis to sources within the central region of the field, where the effect is minimal – this motivates our choice of the RoI above.

4.4.1.2 Propagation effects

Ionospheric scintillation, caused by turbulence in the ionosphere ([Thompson et al., 2017](#)), produces small stochastic variations in the light curves. Being stochastic, this effect cannot be directly mitigated. The angular scale of ionospheric scintillation is on the order of arcminutes or more. Thus, we can expect it to affect any source that appears point-like at MeerKAT’s resolution. We do not attempt to quantify the effect further in this work but only note that it should be possible, for any given observation and ensemble of light curves, to estimate the overall effect of ionospheric scintillation on the field as a whole ([Loi et al., 2015](#)). Figure 4.9 shows examples of what we presume to be ionospheric scintillation in the COSMOS field.

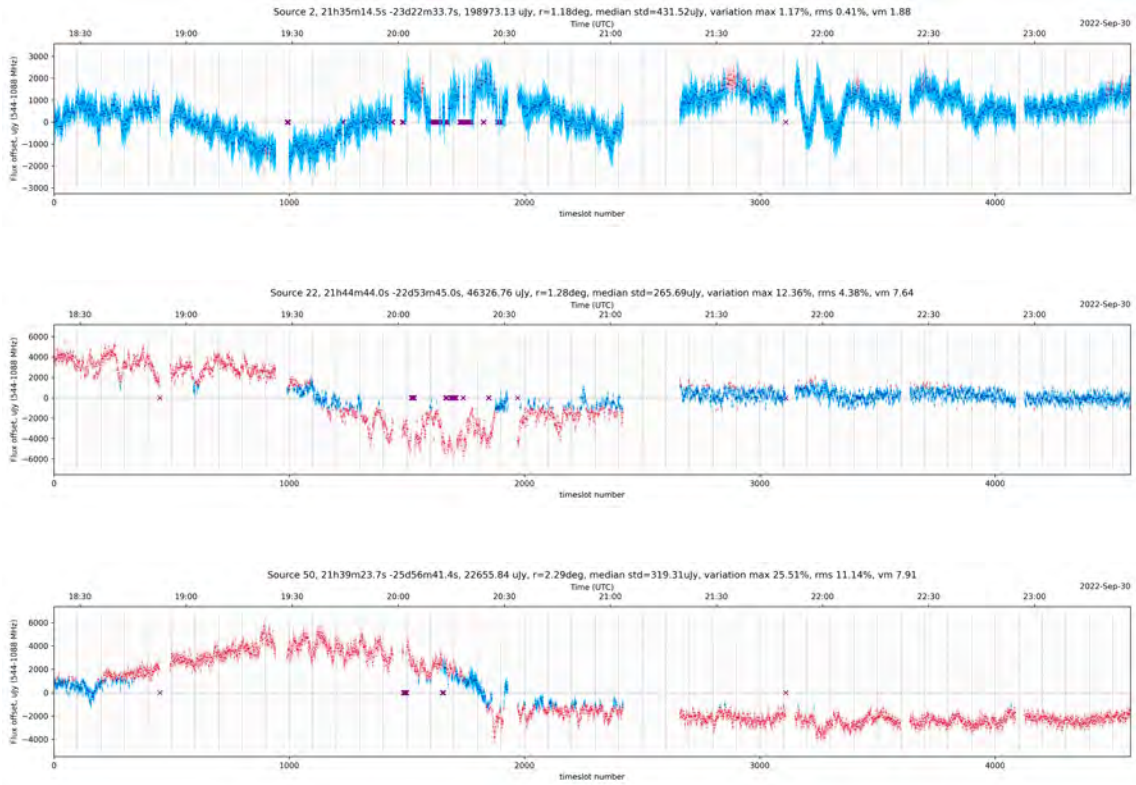


FIGURE 4.8: Three light curves from MACS J2140.2–2339 affected by the primary rotation. Each source has an angular separation from phase centre (r) greater than 1 degree: source 2 at $r = 1.18$ deg, source 22 at $r = 1.94$ deg, and source 50 at $r = 1.4$ deg. Significant variations in the light curves are plotted in red.

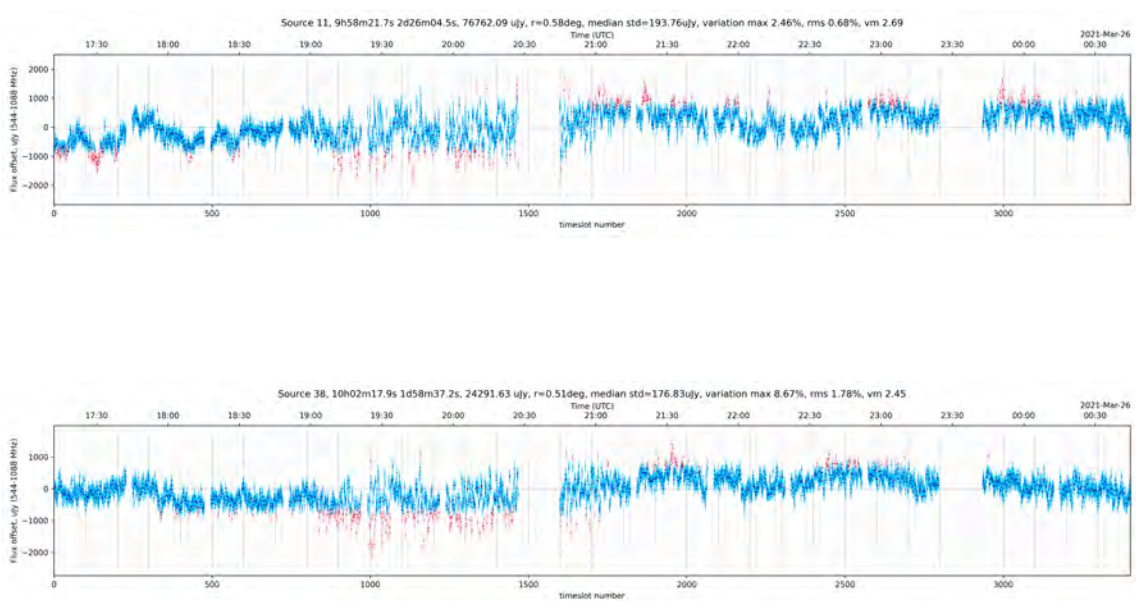


FIGURE 4.9: Two COSMOS light curves affected by ionospheric scintillation,

Interplanetary scintillation (IPS) (Thompson et al., 2017) is a similar effect caused by propagation through solar wind particles. The timescale of IPS being ~ 1 s, i.e. less than our correlator dump time, it manifests itself as a stochastic variation in the timeslot-to-timeslot flux measurements. The angular scale of IPS is arcseconds or less. Thus, only truly compact sources should be affected. We do not study IPS further in this work, but we do note that MeerKAT light curves such as those produced by our pipeline could be a scientifically useful probe of IPS per se, and also provide the potential for measuring source extents well below the resolution of the array.

Ionospheric diffraction, as a first-order effect, results in a slight variation in the apparent position of a source, while higher-order effects produce distortions in the source shape (effectively, a “defocusing”). Given the simplistic way in which we measure the light curves in our pipeline (i.e. pixel value at a fixed position in the image plane), these effects will imprint themselves as different signatures in the light curves. For example, Figure 4.10 shows distinct sinusoidal-like fluctuations in two light curves from the MACS J2140.2–2339 field. Visual inspection of the image cubes confirms that these correspond to small shifts in the apparent source positions. Figure 4.11 shows the three slices (selected randomly) in the cube, where we can see the shifts in source positions. We speculate this is consistent with ionospheric diffraction, particularly travelling ionospheric disturbances (TIDs) passing over the array. Dealing with this effect is left to future work. However, we note ionospheric diffraction is more prevalent at lower frequencies, such as the UHF band. Therefore, extracting light curves at different ends of the band can be used to confirm that, indeed, these are ionospheric effects and obtain cleaner light curves at the expense

of SNR. Additionally, we can also explore more sophisticated approaches to extracting flux measurements from snapshots (for example, advanced aperture photometry methods (Vanderburg and Johnson, 2014) could mitigate this effect.

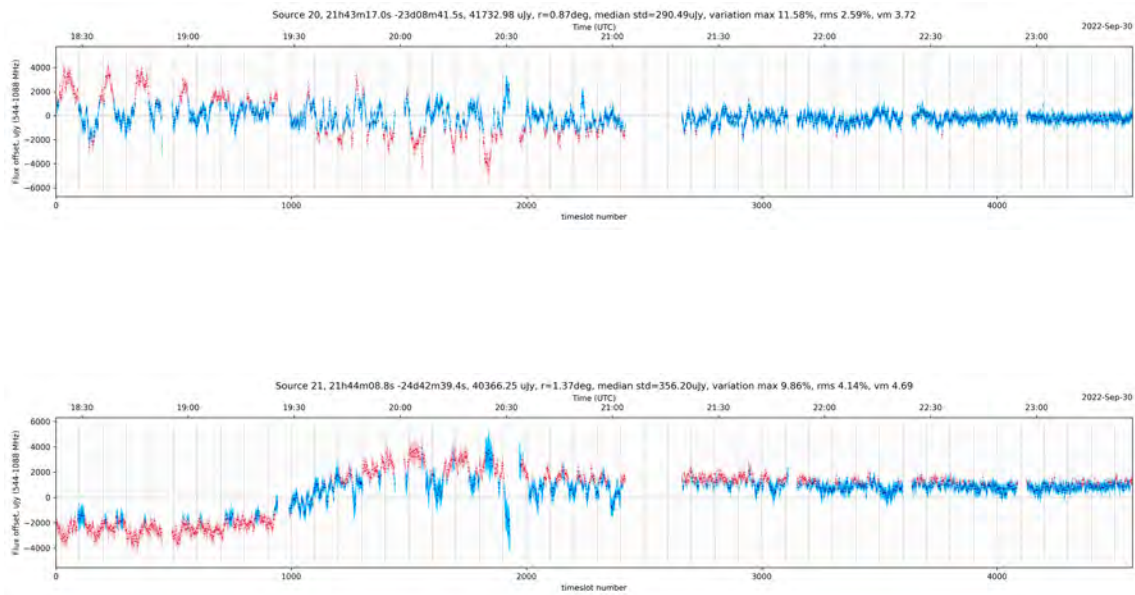


FIGURE 4.10: Two MACS J2140.2-2339 light curves affected by the ionospheric diffraction.




FIGURE 4.11: Three FITS cube (model subtracted) slices from MACS J2140.2-2339 displaying slight shifts in the apparent position of source one. The black spots inside the circular yellow region present an over-subtraction, and the white spots represent an under-subtraction

4.4.1.3 Imaging Artefacts

Besides directly affecting the light curves as discussed above, DDEs (primary beam, ionosphere) indirectly contaminate our measurements in two other ways. DDEs primarily manifest themselves as PSF-like error patterns centred on bright off-axis sources. This results in (i) false source detections during source finding, which yield fake light curves with very high temporal variability (see Figure 4.13), and (ii) contamination, manifesting itself as false variability in the light curves of genuine nearby sources. Deconvolution artefacts produced by bright extended – and particularly, slightly resolved – sources can have a similar impact (although this is much smaller for the two fields in question). These effects can be mitigated as follows:

- A few of the brightest DDE-affected sources can be peeled (see Figure 4.4), subject to flux suppression considerations discussed above.
- Better deconvolution and DD-calibration techniques will improve the situation in the future.
- False detections can be minimised by carefully tuning PYBDSF parameters. The default settings tend to produce many false detections in the presence of DDEs (see Figure 4.12). Adjusting the `rms-box` parameter, which determines the size of the region over which PYBDSF computes the local noise statistics, is crucial. This should be aligned with the typical scale at which DDE artefacts exhibit significant variation. For example, if DDE-affected regions have a typical span of 20 pixels, the default PYBDSF setting of 63 results in an over-smoothed rms map. Setting a smaller `rms-box=(20,10)` effectively mitigates for artefacts being identified as sources (Mohan and Rafferty, 2015).

- If the artefacts themselves cannot be sufficiently mitigated, we exclude all “contaminated” sources, that is, sources within a certain distance of DDE-contributing sources, from further analysis.




FIGURE 4.12: Image with strong artefacts around bright sources and the number of artefacts near the bright sources are incorrectly identified as real sources.

We acknowledge that in future investigations, employing spectral domain analysis to identify intrinsic variability can help minimize the impact of image artefacts. This is attributed to the fact that image artefacts, as shown in Figure 4.12, are influenced by telescope resolution (refer to Equation 2.2). Consequently, these artefacts exhibit unphysical spectra as their positions shift due to the λ/D scaling.

4.4.1.4 RFI and satellites

RFI is pervasive in radio observations. While our pipeline performs conventional flagging during the data reduction process (see above) to exclude RFI-contaminated




FIGURE 4.13: Three highly variable light curves of artefacts incorrectly identified as sources by PYBDSF.

visibilities from imaging and calibration, some weaker RFI invariably remains undetected at that processing stage. We note two particular effects in our light curves.

Firstly, some image snapshots exhibit high image rms due to unflagged RFI. This is obvious from visual inspection of the snapshots. During light curve extraction, we flag timeslots where the local rms around a source exceeds the median rms around that source by some factor (e.g. 3).

Secondly, we note that (particularly, satellite-induced) RFI can manifest itself as temporal gain variations, due to the receiving system being driven into non-linearity (B. Hugo, priv. comm.) This results in synchronous short-timescale changes in the

flux scale across the entire field (refer to Figure 4.14), as the satellite moves through the sidelobes of the primary beam. Since this affects all sources simultaneously, the effect can, in principle, be detected. It is probably prudent to flag the affected time intervals entirely since any flux measurements at these points cannot be reliable.

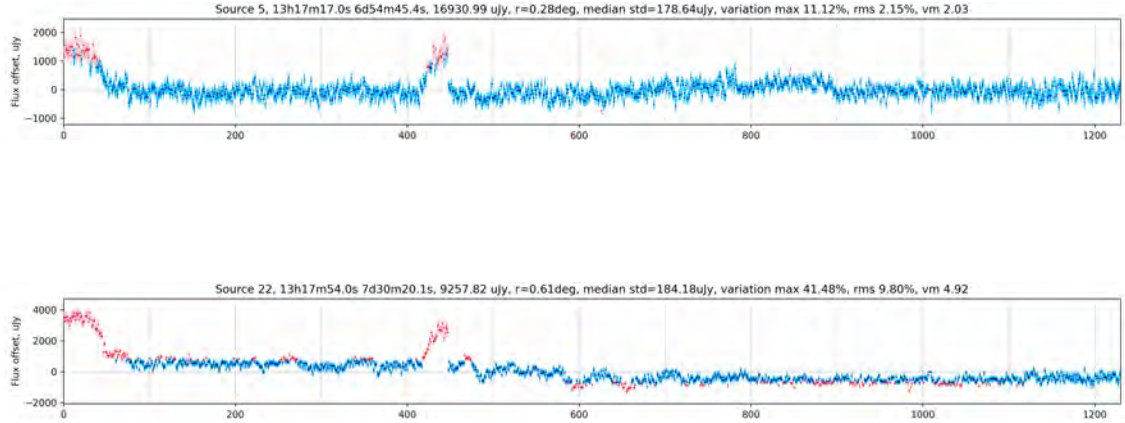


FIGURE 4.14: Light curves of sources affected by satellite as they move through the side lobes of the primary beam. The light curves show synchronous short-timescale variations.

4.4.2 Light curve selection criteria

We obtained a total of 10241 sources from the COSMOS and a total of 13978 sources from the MACS J2140.2–2339 field. To ensure the reliability and accuracy of our analysis, we employ a careful selection process to choose light curves less affected by the systematic effects described above. The following criteria have been established;

1. **Primary Beam Effects:** Bright sources on the flank of the beam are more affected by PBEs. This motivates our choice of RoI above. Figure 4.15 visually represents the RoI and selected sources for the analysis.
2. **Ionospheric Effects:** Both MS are subject to ionospheric scintillation and diffraction. The COSMOS field experiences scintillation between 19:00 and




FIGURE 4.15: (a) and (c) show the selected sources (RoI) for intrinsic variability search in the COSMOS and MACS J2140.2–2339 respectively. (b) and (d) show a close up region of (a) and (c) of COSMOS and MACS J2140.2–2339 respectively

21:00 UTC, as illustrated in Figure 4.9. MACS J140.2–2339 field undergoes ionospheric diffraction from the commencement of the observation at 18:30 until 21:00 (refer to Figure 4.10). We removed most of the light curves affected on both MSs.

3. **Imaging Artefacts:** As mentioned earlier, two types of image artefacts impact the light curves. Firstly, PSF-sidelobe-like spikes are erroneously identified as sources by PYBDSF. For both MSs, we applied the strategy outlined

in section 4.4.1.3, which effectively prevented these spikes from being misidentified as sources. This correction was implemented in the source finding step in the PARROT pipeline. Concerning sources near the bright off-axis source in the COSMOS, the artefacts are mainly attributed to the primary beam. The selected RoI excluded numerous affected sources and those in their vicinity. In the MACS J140.2–2339 field, after peeling the three problematic cases (as shown in Figure 4.4), we were left with minor image artefacts within the RoI. We determined that these artefacts do not significantly impact the sources, so we deemed them to be of no major concern.

4. **RFI and satellites:** In both MSs, our flagging during 1GC effectively flagged RFI, leading to light curves with minimal to no RFI. Light curves affected by weak RFI were appropriately flagged using the PARROT pipeline. Consequently, no light curves were excluded from our analysis.
5. **Extended Sources:** We excluded all light curves from extended sources from the analysis since these cannot be expected to show intrinsic variability based on fundamental physical considerations.
6. **Faint Sources :** We excluded all light curves from sources with flux below $50\mu\text{Jy}$, since their per-snapshot SNR was too low.

After the selection process, we were left with **3,241** light curves for COSMOS and **6,428** for MACS J2140.2–233.

4.4.3 Searching for intrinsic variability in light curves

4.4.3.1 COSMOS

This section reviews the analysis of 3,241 light curves from the COSMOS field. For a source to be deemed a potential variable source, it needs to be a compact or point source with $\eta > \eta_{\text{thresh}}$ and the light curve must be distinctly variable compared to other light curves. This is because the intrinsic variability of celestial sources is unlikely to manifest similarly across different sources. If this occurs, it suggests potential influence from systematic effects. Given that some light curves retained systematic effects despite the data reduction and selection criteria, we injected five simulated variable sources into our dataset (refer to Figure 4.16). This was done to check if the remaining systematic effects could compromise our ability to detect intrinsic variability. Additionally, it provided insights into where we are more likely to find variable sources on the $(\eta; v)$ distribution plots if they exist.

We simulated the light curves by introducing Gaussian profiles with random widths and peaks into the high-time-cadence cube used to extract light curves. These Gaussian profiles were added to existing non-variable source regions (see Appendix B for the Python code to simulate the sources). The simulated sources and their characteristics are as follows:

- **Source 101** displays seven distinct peaks, each characterised by a Gaussian profile with peak amplitude between 500 and 1000 μJy . The width of each Gaussian peak denotes the distribution spread, which varies between 3 and 7 data points, equivalent to a duration of 24 seconds to 80 seconds. These Gaussian peaks represent a temporal increase in flux. Significant flux increase within light curves is plotted in red.

- **Source 230** displays a single distinct Gaussian peak between timeslot 2000 and 2500. The Gaussian peak is 1300 μJy and has a half width lasting 1 minute.
- **Source 464** is similar to source 101.
- **Source 700** displays seven distinct peaks similar to source 101. These peaks share the same amplitudes but differ in their widths. In this plot, the widths are noticeably broader, spanning a duration from 1 minute to 5 minutes.
- **Source 1086** is similar to source 700.

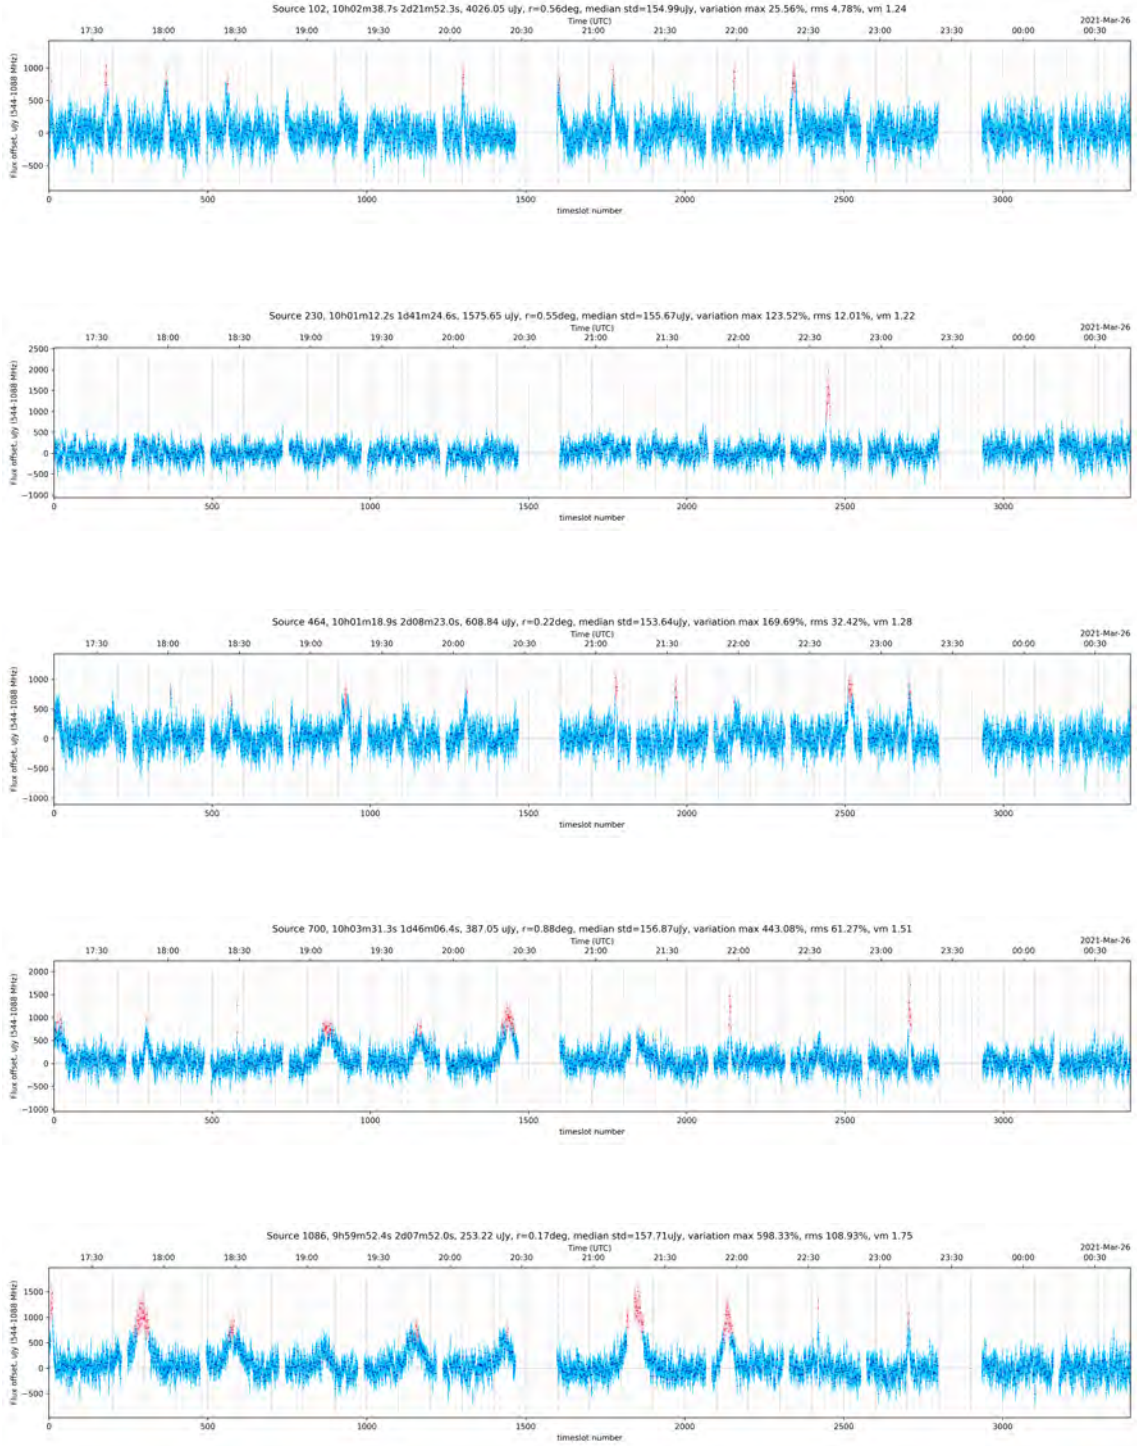


FIGURE 4.16: The plots show the light curves of five simulated variable sources as described in section 4.4.3.1

We utilised the variability metrics (η ; ν) from all 3241 light curves in the COSMOS field to search for genuine significant variability. We calculated the threshold using the strategy outlined in section 3.3. The threshold (η_{thresh}) is determined by $\eta_{\text{thresh}} = \mu + n\sigma$, where any source with $\eta > \eta_{\text{thresh}}$ is considered significantly variable compared to other light curves. In this field, we calculated $\mu = 0.94$ and $\sigma = 0.03$, Therefore, we chose a threshold $\eta_{\text{thresh}} = \mu + 2\sigma = 0.94 + 2(0.03) = 1$ (see Figure 4.17).

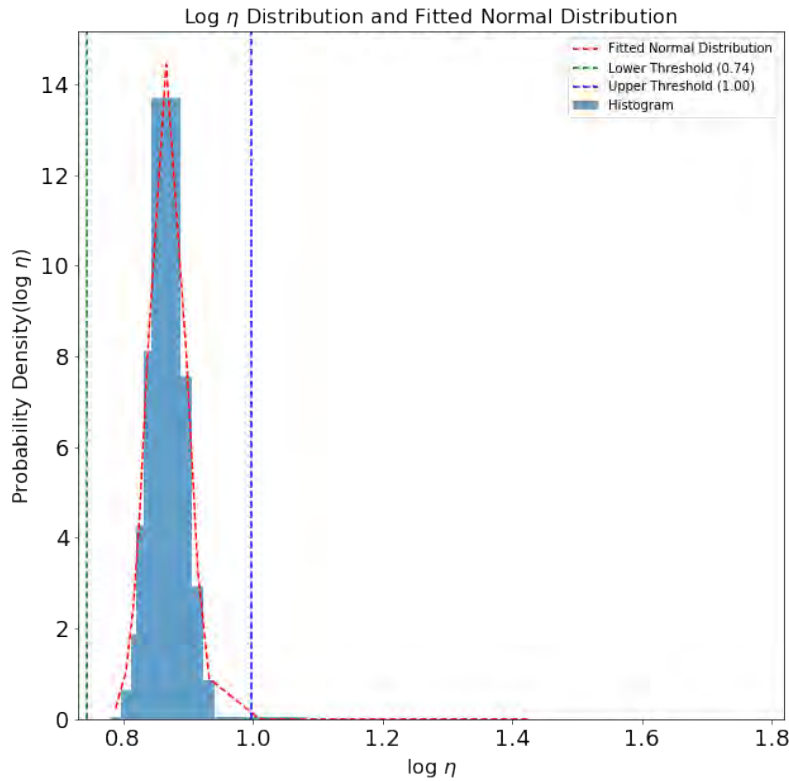


FIGURE 4.17: The histogram illustrates the $\log \eta$ -distribution of our light curves, with the curve (indicated in orange) representing the optimal fit normal distribution. Non-variable sources fall below the established threshold (indicated in a blue line), while sources with above η_{thresh} values are categorized as variable candidates.

Figure 4.18 shows the distribution of the variability metrics (η) against maximum and median flux. The thresholds for each variability metric are shown using the dotted line. After examining the plots, we identified **76** sources with $\eta > 1$. The plots also highlighted significant variability in the simulated variable sources, indicated by stars. The simulated variable light curves demonstrated that sources with relatively

fast variations tend to exhibit a low η value compared to slowly varying sources. In our case, slow variation is introduced by all sources affected by the primary beam and simulated variable light curves with broader Gaussian widths. For example, Source 230, with a single bright Gaussian peak lasting 1 minute (once off), exhibits a low η compared to Source 700 and 1086. This indicates that the η metric is less sensitive to fast variations. Therefore, any source with variability shorter than 1 minute will not have a significantly high η value and is more likely not detected in this analysis.




FIGURE 4.18: This plot shows the reduced $\chi^2(\eta)$ plotted against the maximum and median flux density. The synthetic sources are highlighted by stars. The plots are in log-scale

Moreover, we utilised the variability index (ν) to quantify the variability of our light curves. We first plotted the η - ν plot, setting both the thresholds for η and ν to 0. Figure 4.19 shows ν plotted against η . Light curves with $\eta < 0$ are classified as non-varying (stable light curves), those with $\eta > 0$ and $\nu > 0$ are classified as moderate to highly variable light curves, and finally, those with $\eta > 0$ and $\nu < 0$ are classified as low to moderately variable sources. For completeness, we then plotted

ν against the maximum flux and median flux, with each classification plotted in different colors (see Figure 4.20).

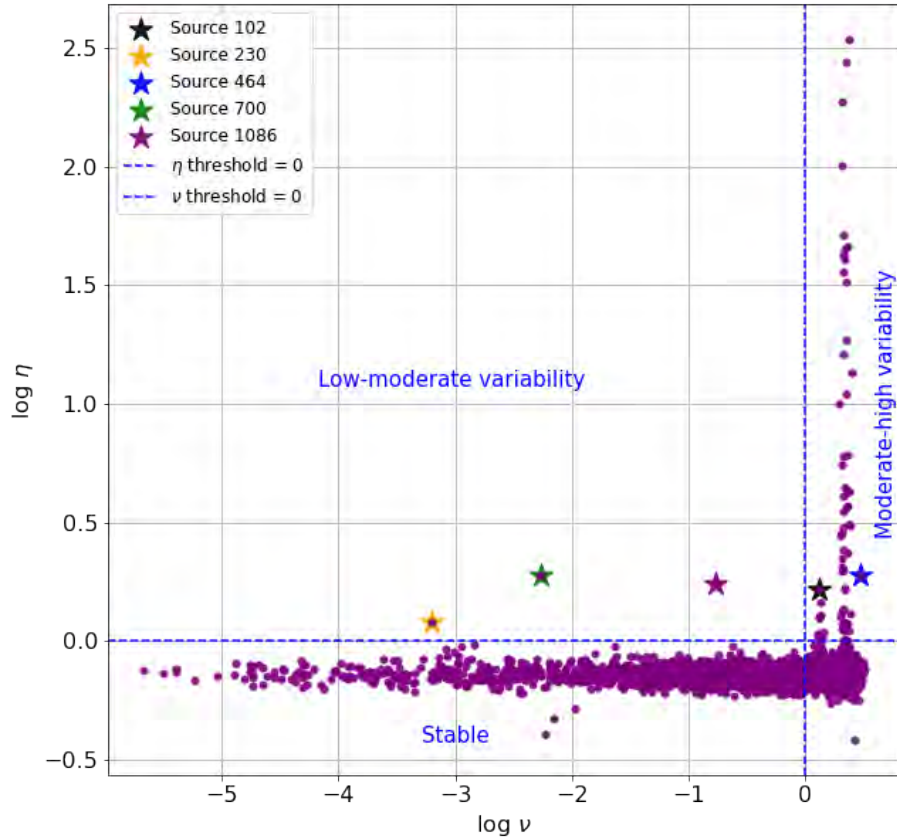


FIGURE 4.19: This plot shows the reduced weighted $\chi^2(\eta)$ plotted against the variability index (ν). The synthetic sources are highlighted by stars.

After obtaining η and ν variability plots and identifying 76 candidate variable sources, we excluded the light curves of the injected simulated sources. Since all injected sources were detected, we conclude that if any significantly variable real sources are present among our light curves, they were likely to be among the 76 identified as variable candidates. Moving forward, our focus lay in understanding the origin of variability that accounted for the 71 candidates. We aimed to discern whether this variability was attributed to systematic effects or intrinsic variability. To achieve this, we eliminated sources where we suspected the variability was

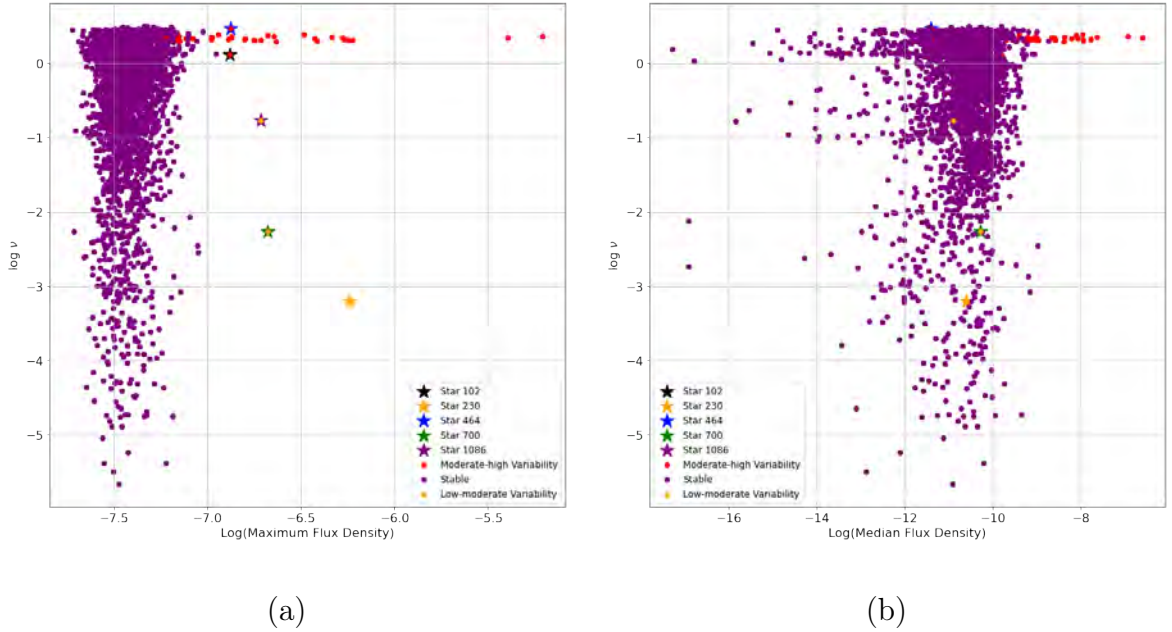


FIGURE 4.20: (a) This plot shows the variability index ν plotted against the maximum and median flux density. The synthetic sources are highlighted by stars. The plot also shows the classification of the sources as either stable(non-varying), Low-moderate or Moderate-high variability

induced by systematics. Those for which we could not make a definitive determination were considered potential variable sources, requiring further investigation to establish their variable nature conclusively. We began by examining two potential contributing factors:

- To get a handle on primary beam effects, we produced plots illustrating the distribution of η against maximum flux and distance from the pointing centre for each source. PB effects are more likely to affect sources with a distance greater than $r > 0.5^\circ$, inducing an artificial temporal variation and, subsequently, a high η value.
- Furthermore, we investigate the connection between flux and a higher η values by plotting η against the maximum flux as a function of observed flux as reported by PYBDSF. This stemmed from the observation that all detected

variable candidates were on the brighter end (10 to 200 mJy) of the source flux distribution, suggesting that faint sources are less sensitive to the η metric. Intuitively, this also makes sense from simple SNR considerations. Both of these plots are shown in Figure 4.21.

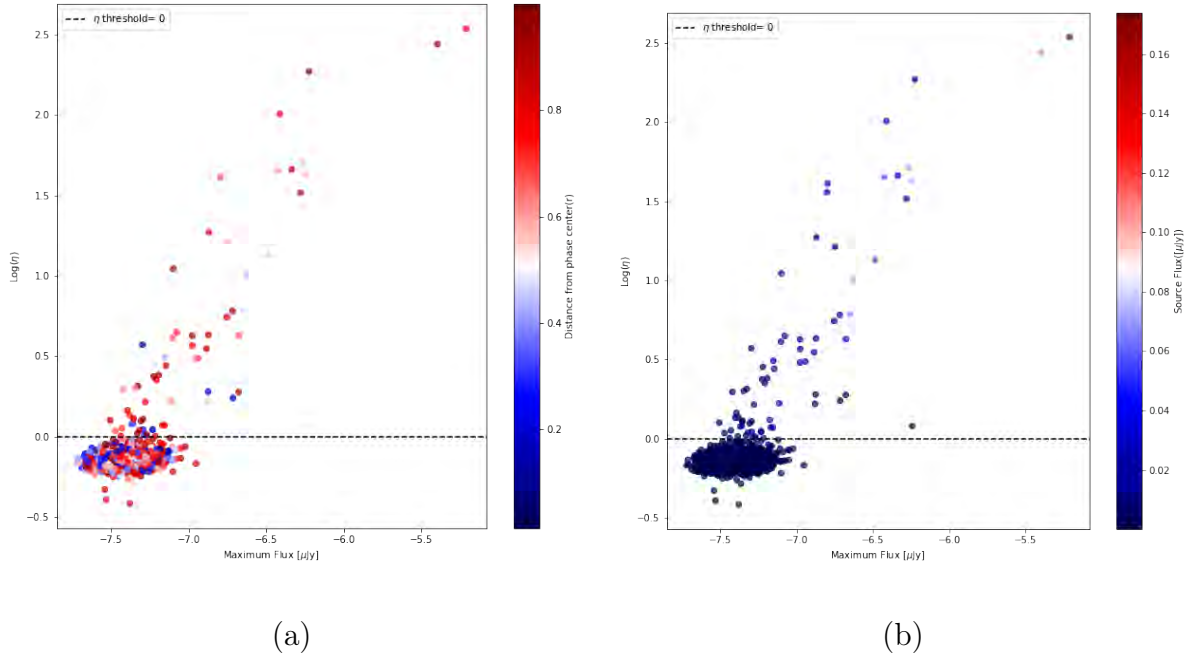


FIGURE 4.21: (a) This plot shows η plotted against the maximum flux density, with the colour bar indicating the phase distances from the centre (r). (b) This plot shows η plotted against the maximum flux density, with the colour bar indicating the observed flux density of each source reported by the source finder (PYBSDF).

The maximum flux(μ Jy) in both plots are also in log scale.

From Figure 4.21(a), it is evident that a significant portion of our variable candidates are positioned at $r > 0.5^\circ$, suggesting that their variability can be attributed to the primary beam. While the variability induced by the primary beam rotation in the $0.5^\circ < r < 1^\circ$ range may not be as pronounced as in sources with $r > 1^\circ$, they still display variability compared to sources within $r = 0.5^\circ$. Therefore, we excluded all sources with $r > 0.5^\circ$ from the 71 candidates, resulting in the exclusion of 62 sources.

TABLE 4.2: Properties of the nine variable sources with $\eta > 1$

Source ID	R.A	Dec	Flux(μ Jy)	η	r	$v(\%)$	Var max(%)	vm
Source 4	9:59:46.23	2:36:03.15	21 290.26	1.64	0.43	1.06	3.68	1.40
Source 5	10:01:22.31	2:01:12.58	20 682.30	1.08	0.29	0.88	3.07	1.14
Source 6	10:01:31.29	2:26:40.13	20 518.44	1.08	0.35	0.88	3.07	1.13
Source 7	10:00:57.97	1:51:34.41	18 264.42	1.01	0.37	0.96	3.80	1.10
Source 8	9:59:59.08	1:48:38.81	15 696.58	1.05	0.41	1.13	4.49	1.11
Source 11	9:59:57.91	2:18:10.62	11 245.68	1.04	0.16	1.58	6.42	1.13
Source 24	10:00:05.35	2:30:30.27	5 828.65	1.02	0.35	2.93	11.83	1.09
Source 25	10:00:00.53	2:15:32.02	5 654.38	1.77	0.13	4.10	14.54	1.47
Source 28	10:02:24.06	2:16:22.37	5 269.23	1.57	0.49	4.13	13.99	1.38

We speculate that the COSMOS field is affected by an additional systematic, namely a slow drift in the flux scale, which imparts a multiplicative effect on the light curves. Visual inspection of Figure 4.23 shows a smooth longer-term variation that appears to be correlated between sources 4, 5, 6, 7, 11, 24. This is more exaggerated in source 4 (21.2 mJy) than in source 11 (11.2 mJy). Such a multiplicative effect would artificially increase η for brighter sources. It is suggested that future versions of the light curve pipeline include computations of correlations to reveal such effects explicitly.

We note that this effect could possibly be mitigated by an additional round of amplitude self-calibration in the pipeline, and/or by correlating over an ensemble of light curves across the field. We leave this investigation for future work. In the meantime, we only consider the three candidates with the highest η values, namely sources 4, 25, and 28. We note that source 4 is the brightest of these, and that its lightcurve is visually correlated with those of the weaker candidates, suggesting that its high η value is due to the same systematic. On the other hand, sources 25 and 28 show a distinctive variability pattern that is dissimilar to the other light curves. These are our final variable candidates.

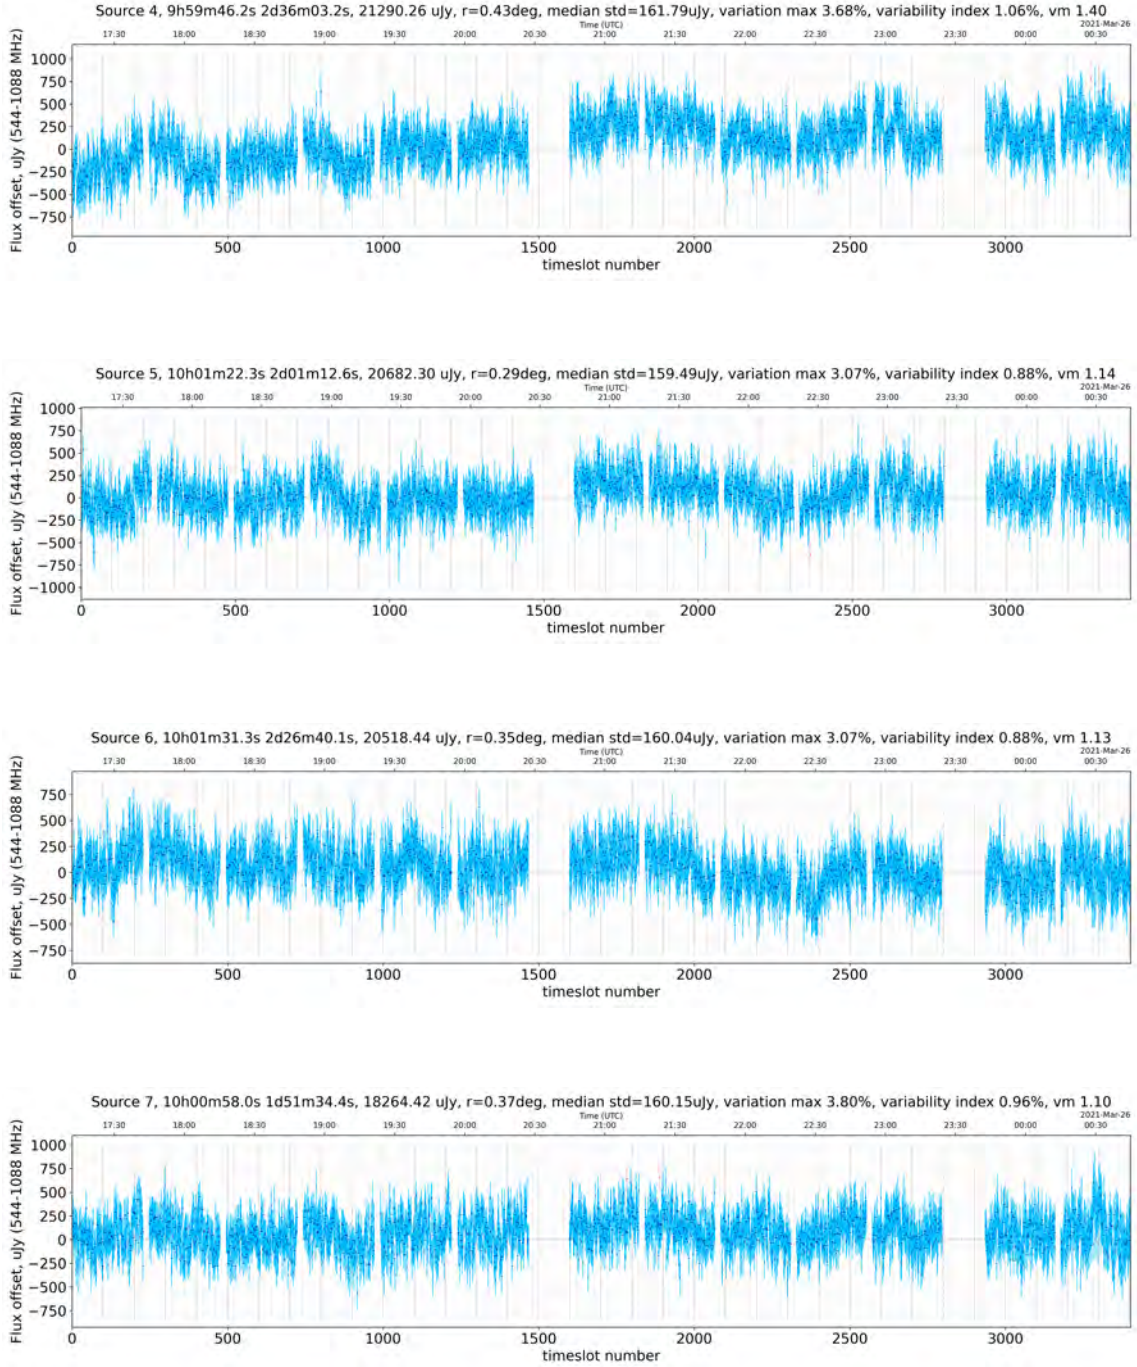


FIGURE 4.23: Four of the nine candidate variable light curves with $\eta > 1$ (continued next page)

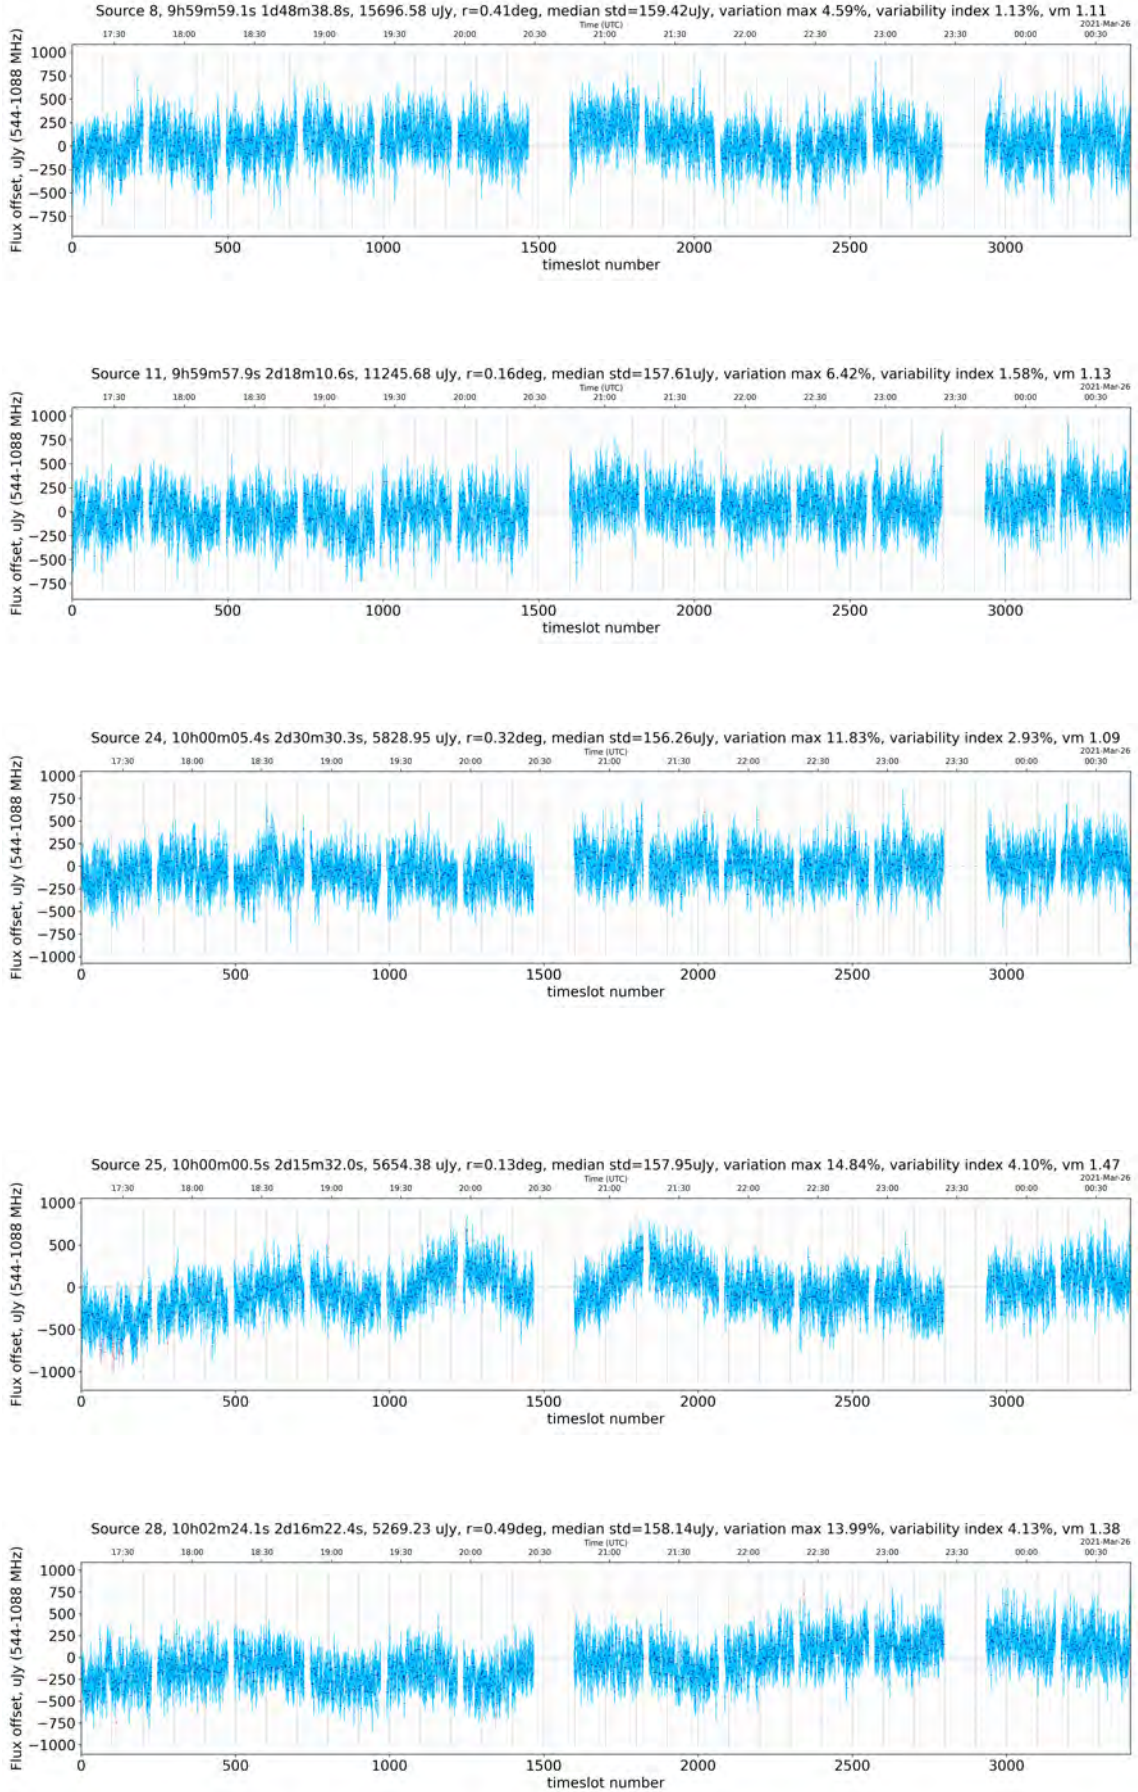


FIGURE 4.24: Five other candidate variable light curves.

4.4.3.2 Analysis of the variable candidates

This section examines the characteristics of sources 25 and source 28. Figure 4.25 below shows the sources in the MFS restored image. Source 25 has $\eta = 1.77$




FIGURE 4.25: Restored MFS image of the COSMOS field showing source 25 and 28.

and $v = 4.10\%$ (classified as a Moderate-high variability source). This source holds the highest η value among the variable candidates. What sets this source apart is its distinct variability pattern when compared to every other light curve in the field, especially those nearby that have similar fluxes. The light curve of the sources has two gradually increasing peaks between 19:30 and 22:00 UTC. Source 28 has a $\eta = 1.57$ and $v = 4.13\%$ (classified as a Moderate-high variability source). This source has also been identified as a variable source due to having a high (η) value.

Notably, this source also displays a distinct pattern of variability that sets it apart from other sources. Its flux gradually fluctuates through the timestamps.

To determine the physical characteristics of our two variable sources, we utilised multi-wavelength datasets from the COSMOS field, referencing the VLA-COSMOS 3 GHz large project survey conducted by [Smolčič et al. \(2017\)](#). This survey spanned 2.6 square degrees in the COSMOS field and catalogued 10,830 sources with a significance level ($\sigma > 5$). They cross-matched the radio data to identify counterparts with optical, near-infrared, mid-infrared, and X-ray surveys. They successfully identified counterparts for 93% of the sources with redshifts ($z < 6$). Each source was classified as a star-forming galaxy or active galactic nucleus (AGN) based on criteria such as X-ray luminosity, observed mid-infrared colour, and UV-FIR spectral energy distribution. They further categorised AGNs into low-moderate and moderate-high luminosity AGNs.

In our study, we cross-matched our two variable sources with the catalogue from [Smolčič et al. \(2017\)](#). We found counterparts for one source: Source 25, which is classified as a high luminosity AGN. The VLA COSMOS 3 GHz catalogue found no counterpart for source 28. Consequently, we cross-matched source 28 with another catalogue from [Sarbadhicary et al. \(2021\)](#), which conducted a radio variability search in the COSMOS field using VLA with B-configuration at 1.4 GHz. Although Source 28 was matched with a CHILES VEDERS CV17 AGN, the study did not specify its luminosity classification. Below, we compiled a table of the two variable sources and their classification (Table 4.3)

TABLE 4.3: The Multiwavelength properties of our 2 potential variable sources

Source ID	R.A	Dec	Variability	Classification	reference
Source 25	10:00:00.53	2:15:32.02	Moderate-high	High Luminosity AGN	Smolčić et al. (2017)
Source 28	10:02:24.06	2:16:22.37	Moderate-high	AGN	Sarbadhicary et al. (2021)

4.4.3.3 Discussion

As outlined in the introduction, AGNs are known to be intrinsically variable over periods ranging from several days to years. This implies that the variability observed in sources 25 and 28 is more likely extrinsic. Consequently, we naïvely propose that these AGNs are exhibiting intra-hour variability, characterised by quasi-random variations in intensity within a single hour. This is often associated with the effects of turbulent plasma in the interstellar medium, and the nature of the plasma clouds responsible for the rapid variations remains unclear ([Wang et al., 2021](#), [Oosterloo et al., 2020](#), [Rickett et al., 1984](#)).

Following a note from ([Oosterloo et al., 2020](#)) to emphasise the importance of detecting IHV, they noted a recent development involving the statistical association of three known IHVs with nearby hot stars, which implies that the scattering is caused by the ionized sheath of tiny self-gravitating molecular clumps, resembling the cometary knots observed in the Helix Nebula ([Walker et al., 2017](#)). They further noted that this discovery would be highly intriguing if validated, as these clouds could represent a new category of objects within the interstellar medium. Given the limited number of currently known IHVs, discovering more IHV objects is essential to understanding their statistical properties better and establishing more secure empirical associations.

In conclusion, intrinsic variability was not identified in the COSMOS field, but we

identified potential IHVs which could be useful in interstellar scintillation studies. This outcome underscores the PARROT pipeline’s capability to explore intrinsic variability and extrinsic variability (i.e., IHV and IDV). Note that a detailed follow-up of these sources exceeds the scope of this project. Nevertheless, the prospect of conducting a follow-up study on these sources remains open for consideration in future research.

4.4.3.4 MACS J2140.2–2339

In this section, we now consider the analysis of the MACS J2140.2–2339 field. The search for intrinsic variability in these sources is similar to that of the COSMOS field, and we apply the lessons learned above. Namely, we restrict our RoI to the inner part of the primary beam, $r = 0.5^\circ$. We use the reduced chi-square (η) of all the light curves to identify significant variations. The threshold for η is set at 1.29. Light curves below this threshold are classified as non-variable, while those surpassing it are considered variable sources. Figure 4.26 presents the distribution of η plotted against the maximum flux of the light curve. A total of 5 sources exhibit $\eta > 1.29$.

Figure 4.27 illustrates the light curves of the five variable candidates. The variability observed in all these sources can very likely be attributed to ionospheric diffraction, detailed in section 4.4.1. While this ionospheric influence is less noticeable in more moderate and fainter sources, as shown in Figure 4.5, it is likely also present. This introduces uncertainty regarding the observed variability, and any fluctuations in these light curves are deemed unreliable due to the consistent impact of ionospheric diffraction.

In summary, the results from MACS J2140.2–2339 highlight the limitations of the current version of the PARROT pipeline. As detailed in section 4.4.1, applying a more sophisticated approach to extract light curves may offer a potential solution to this issue. This is an obvious avenue for future investigation.




FIGURE 4.26: The plot shows the distribution of η against the maximum flux density of the light curves obtained from the MAC J2140.2–2339 field. The numerical labels align with the 'sources ID' reported by the source finder (PYBDSF).




FIGURE 4.27: The plots show five variable sources with a $\eta > 1.29$, after removing the sources affected by the PBR.

Chapter 5

Conclusions & Future Work

This thesis presents the PARROT pipeline used to mine for minute-to-hour timescale transients and variable sources. As aligned with our thesis objectives, (i) we applied and tested the pipeline to different MeerKAT measurement sets to detect transients and variable sources, and (ii) highlighted areas where we could improve the current pipeline for future work.

(i) PARROT pipeline application

The PARROT pipeline was applied and tested on two MeerKAT measurement sets, both observed in the UHF band: the COSMOS and MACS J2140.2–2339 fields. The COSMOS field was observed for 8 hours, while MACS J2140.2 – 2339 was observed for 5 hours. We extracted 10241 light curves from COSMOS and 13978 from MACS J2140.2 – 2339. We manually inspected the light curves and identified systematic effects affecting the observations, and we applied selection criteria to exclude most of the sources affected by systematic effects. We obtained 3,211 light curves for COSMOS and 6,428 MACS J2140.2 – 233. Furthermore, we then applied the variability

statistic (η, v) to identify significant variability and quantify the variable source sources. The analysis revealed the presence of two variable sources in the COSMOS field, indicative of potential scintillating AGNs likely caused by turbulent plasma in the interstellar medium. However, a detailed follow-up on these sources falls beyond this project's scope. On the other hand, no intrinsically variable sources were detected in the MACS J2140.2 – 2339 field due to the significant impact of ionospheric diffraction in this field. No transients were detected on either observation.

(ii) Future work

The application of the pipeline to different MeerKAT observations highlighted areas where we can further improve the current version of the PARROT pipeline. Below, we will give a summary of the potential improvements.

- **Systematic effects**

In Section 4.4.1, we discussed systemic affecting COSMOS and MAC J2140.2–2339 at UHF band, which are also likely to affect other MeerKAT observations. Primary beam and ionospheric effects emerged as the most significant effects. Typically, the PBE affects sources within $1.5 \times$ primary beam full width at half maximum (FWHM) at UHF. However, in this thesis, we observed the primary beam effect influencing sources from $0.5 \times$ the primary beam FWHM. Rather than directly addressing this effect, we opted to exclude sources greater than $0.5 \times$ the primary beam FWHM. Unfortunately, this exclusionary criterion resulted in many sources within our field being omitted from the analysis for intrinsic variability. Consequently, the probability of detected variable sources and transients decreased. For future work, we aim to implement the mitigation strategies outlined in section 4.4.1 to tackle primary beam effects in the PARROT pipeline. This strategy aims to expand our search area from 0.5

to 1.5 times the primary beam FWHM, increasing our chances of discovering variable sources.

Ionospheric diffraction is another significant systematic effect with the potential to impact additional observations at lower frequencies. This effect manifests as different signatures in different observations, posing immediate identification challenges. In future investigations, our objective is to explore methods for accurately identifying ionospheric diffraction and to develop approaches for extracting light curves that account for the observational impact of these effects. This, in return, will significantly reduce the spurious variability in light curves and improve our search for intrinsic variability.

- **Search for intrinsic variability (transients and variable sources)**

The current method employed for detecting significant variation in light curves computes the reduced χ^2 , denoted as (η) . As previously stated, η is not sensitive to fast variations; rather, it is sensitive to slower variations. In the future, we intend to employ a sliding window to determine η , with the η window varying with SNR. This will ensure that fast variations occurring only once will not be missed

Furthermore, no transients were detected in our observation. This may be attributed to the nature of using photometry for extracting light curves, which relies on the prior identification of sources by the source finder. Any source not detected by the source finder prior remains unmonitored. Transients, being brief and abrupt events emerging and disappearing during observations, are more likely to remain undetected. For future work, we aim to incorporate a transient detection strategy into our PARROT pipeline. For instance, we may consider implementing strategies similar to those discussed in [Fijma et al. \(2023\)](#) and [Wang et al. \(2023\)](#).

Appendix A

The PARROT Pipeline

```
1
2 #!/usr/bin/env -S stimela run
3
4 _include:
5   - rrat-cabs.yml
6
7
8 ## This augments the standard 'opts' config section to tweak
9   logging settings
10
11 opts:
12   log:
13     dir: logs/log-{{config.run.datetime}}
14     name: log-{{info.fqname}}
15     nest: 2
16     symlink: log
17   rlimits:
18     NOFILE: 100000
19
20 ratt-rrat:
21   name: ratt-rrat
22   info: "The parrot pipeline application: COSMOS"
23   assign:
24     dir-out: '{{recipe.dirs.base}}/{{recipe.dirs.sub}}{{recipe.
25       output-suffix}}'
26     image-prefix: '{{recipe.dir-out}}/im{{info.suffix}}{{recipe.
27       variant}}/im{{info.suffix}}{{recipe.variant}}'
28     log.dir: '{{recipe.dir-out}}/logs/log-{{config.run.datetime}}'
29
30 ## Some more directory assignments
31
32 dirs:
```

```
30     ms: ../msdir
31     temp: "${config.run.env.HOME}/tmp"
32     base: .
33
34     casa-spw: 0
35
36 ## Extra lightcurve suffix, added if convolving
37
38     lc-suffix: =IF(recipe.convolve.arcsec or recipe.convolve.
39     time, "-{recipe.convolve.time:.0f}s-{recipe.convolve.arcsec
40     :.0f}asec", "")
41     lc-suffix-sub: =IF(recipe.convolve-sub.arcsec or recipe.
42     convolve-sub.time, "-{recipe.convolve-sub-time:.0f}s-{recipe
43     .convolve-sub.arcsec:.0f}asec", "")
44
45 ## Yaml file containing the measurement sets
46
47 assign_based_on:
48   _include: observation-sets.yml
49
50 inputs:
51   obs:
52     choices: [L1, L2, L3, U0, U1, U2, U3, U3b, U3c]
53     info: "Selects observation"
54     default: L1
55   output-suffix:
56     dtype: str
57     default: '-qc2'
58   variant:
59     dtype: str
60     default: ''
61   dir.out:
62     dtype: str
63   htc_cadence:
64     dtype: int
65     default: 1
66   htc_size:
67     dtype: int
68     default: 512
69   htc_scale:
70     dtype: str
71     default: '{recipe.pixel_scale}'
72   ampcal:
73     info: do a step of amplitude calibration
74     dtype: bool
75     default: false
76   convolve:
77     arcsec:
78       dtype: float
79       default: 0
80     time:
81       dtype: float
```

```
78         default: 120
79     convolve-sub:
80         arcsec:
81             dtype: float
82
83 ## Self-calibration and peeling
84
85 steps:
86     image-1:
87         info: "Shallow I stoke clean"
88         _use: lib.steps.wsclean.rrat
89         params:
90             column: DATA
91             nchan: 8
92             gain: 0.15
93             weight: "briggs -0.5"
94             niter: 150000
95             scale: 1.7asec
96             auto-threshold: 10
97             size: 10125
98
99     mask-1:
100         info: "Make an external mask"
101         recipe: make_masks
102         params:
103             restored_image: "{previous.restored-mfs}"
104             prefix: "{previous.prefix}"
105             threshold: 6
106
107     predict-1:
108         info: "Convert model image back to visibility."
109         _use: lib.steps.wsclean.predict
110         params:
111             ms: '{recipe.ms}'
112             scale: 1.7asec
113             size: 10125
114             prefix: '{steps.image-1.prefix}'
115             nchan: '{steps.image-1.nchan}'
116
117     DI_cal-2:
118         info: "Delay calibration."
119         _use: lib.steps.quartical.k
120
121     image-2:
122         info: "Deep I stoke clean"
123         _use: lib.steps.wsclean.rrat
124         params:
125             scale: 1.7asec
126             niter: 1000000
127             size: 10125
128             fits-mask: '{steps.mask-1.mask}'
129             auto-threshold: 2
```

```
130
131 mask-2:
132   info: "Make another mask"
133   recipe: make_masks
134   params:
135     restored_image: "{previous.restored-mfs}"
136     prefix: "{previous.prefix}"
137     threshold: 3
138
139 predict-2:
140   _use: lib.steps.wsclean.predict
141   params:
142     ms: '{recipe.ms}'
143     scale: 1.7asec
144     size: 10125
145     prefix: '{steps.image-2.prefix}'
146     nchan: '{steps.image-2.nchan}'
147
148 predict-copycol-3:
149   info: "Copies the model column to MODEL3_DATA."
150   skip: =not recipe.copy-model-columns
151   cab: msutils.copycol
152   params:
153     msname: =recipe.ms
154     fromcol: MODEL_DATA
155     tocol: MODEL{info.suffix}_DATA
156
157 DI_cal-3:
158   info: 'This is the second round of delay calibration.'
159   _use: lib.steps.quartical.k
160
161 image-3:
162   _use: lib.steps.wsclean.rrat
163   params:
164     scale: 1.7asec
165     niter: 1000000
166     size: 10125
167     fits-mask: '{steps.mask-2.mask}'
168     auto-threshold: 0.5
169
170 mask-3:
171   recipe: make_masks
172   params:
173     restored_image: "{previous.restored-mfs}"
174     prefix: "{previous.prefix}"
175     threshold: 3
176
177 predict-3:
178   _use: lib.steps.wsclean.predict
179   params:
180     ms: '{recipe.ms}'
181     scale: 1.7asec
```

```
182     size: 10125
183     prefix: '{steps.image-3.prefix}'
184     nchan: '{steps.image-3.nchan}'
185
186 predict-copycol-3a:
187     skip: =not recipe.copy-model-columns
188     cab: msutils.copycol
189     params:
190         msname: =recipe.ms
191         fromcol: MODEL_DATA
192         tocol: MODEL{info.suffix}_DATA
193
194 dd_predict-1:
195     cab: crystalball
196     info: "Predicts off-axis into separate column"
197     params:
198         ms: '{recipe.ms}'
199         within: ./obs-u0-qc2/im2/cosmos_region_1.reg
200         sky_model: "{steps.image-3.source-list}"
201
202 dd_predict-2:
203     cab: crystalball
204     info: "Predicts off-axis into separate column"
205     params:
206         ms: '{recipe.ms}'
207         within: ./obs-u0-qc2/im2/cosmos_region_2.reg
208         output_column: DIR2_COSMOS_DATA
209         sky_model: "{steps.image-3.source-list}"
210
211 dd_predict-3:
212     cab: crystalball
213     info: "Predicts off-axis into separate column"
214     params:
215         ms: '{recipe.ms}'
216         within: ./obs-u0-qc2/im2/cosmos_region_3.reg
217         output_column: DIR3_COSMOS_DATA
218         sky_model: "{steps.image-3.source-list}"
219
220 DD_cal-3a:
221     info: "This is the peeling process"
222     _use: lib.steps.quartical.k-de
223     params:
224         output:
225             subtract_directions: [1,2,3]
226
227 image-3a:
228     _use: lib.steps.wsclean.rrat
229     params:
230         size: 10125
231         scale: 1.7asec
232         weight: "briggs -0.3"
233         nchan: 8
```

```
234     fits-mask: '{steps.mask-3.mask}'
235     auto-threshold: 0.5
236
237 mask-3a:
238     recipe: make_masks
239     params:
240         restored_image: "{previous.restored-mfs}"
241         prefix: "{previous.prefix}"
242         threshold: 3
243
244 predict-3a:
245     _use: lib.steps.wsclean.predict
246     params:
247         ms: '{recipe.ms}'
248         scale: 1.7asec
249         size: 10125
250         prefix: '{steps.image-3a.prefix}'
251         nchan: '{steps.image-3a.nchan}'
252
253 cube-1:
254     cab: stack_freq_cube
255     params:
256         images: =steps.image-*.restored
257         cube: "{recipe.image-prefix}-cube.fits"
258
259 compute-power-beam:
260     cab: mdv-beams-to-power-beam
261     params:
262         mdv_beams: =recipe.mdv-beams
263         power_beam: =recipe.power-beam
264         skip_if_outputs: fresh
265
266 derive-obs-specific-power-beam:
267     cab: derive-power-beam
268     params:
269         cube: =steps.cube-1.cube
270         images: =steps.image-1.restored
271         outcube: =STRIPEXT(current.cube) + ".pbcorr.fits"
272         power_beam: =recipe.power-beam
273         beaminfo: "{steps.image-1.prefix}-powerbeam.p"
274         nband: 6
275         skip_if_outputs: fresh
276
277 save-ms-info:
278     cab: save-ms-info
279     params:
280         ms: =recipe.ms
281         ms-info-file: '{recipe.ms}.yaml'
282         skip_if_outputs: exist
283
284 ms-info:
285     cab: load-ms-info
```

```
286     params:
287         ms-info-file: =steps.save-ms-info.ms-info-file
288
289 ## Light Curve Extraction
290
291 source-catalog:
292     cab: pybdsm
293     params:
294         image: =steps.image-1.restored-mfs
295         thresh_pix: 6
296         thresh_isl: 3
297         rms_box: [40,20]
298         rms_map: true
299         flagging_opts: true
300         flag_maxsize_bm: 10
301         catalog_format: ascii
302         outfile_srl: =STRIPEXT(current.image) + '.srl'
303         outfile_gaul: =STRIPEXT(current.image) + '.gaul'
304
305 source-catalog-spi:
306     _use: ratt-rrat.steps.source-catalog
307     params:
308         image: =steps.derive-obs-specific-power-beam.outcube
309         detection_image: =steps.image-1.restored-mfs
310         spectralindex_do: true
311         outfile_srl: =STRIPEXT(current.image) + '.spi.gaul'
312         outfile_gaul: =STRIPEXT(current.image) + '.spi.srl'
313
314 make-master-catalog:
315     cab: make-master-catalog
316     params:
317         master_catalog: ./im1-MFS-image.ecsv
318         ra0: 150.1191666666667deg
319         dec0: 2.205833333333333deg
320         max_radius_deg: 0.5
321         catalogs:
322             L: [./im1-MFS-image.srl, 0, main]
323             spi_L: [./im1-cube.pbcorr.spi.srl, 6, srl]
324         search_box_radec: [303.19deg, -20.492deg, 15arcsec, 40
arcmin]
325         search_box_frame: geocentricmeanecliptic
326         search_box_minflux: 50uJy
327
328 add-corrected-data:
329     cab: msutils.addcol
330     params:
331         msname: =recipe.ms
332         colname: RESIDUAL2_DATA
333
334 subtract_model:
335     cab: taql_update
336     params:
```

```
337     ms: "{recipe.ms}"
338     commands: =LIST("set", "RESIDUAL2_DATA={steps.image-2.
column}-MODEL_DATA")
339
340     image-3ahtc:
341     _use: lib.steps.wsclean.dirty
342     params:
343         column: RESIDUAL3a_DATA
344         prefix: =DIRNAME(recipe.image-prefix) + '-cad{recipe.
htc_cadence}/images/cad{recipe.htc_cadence}'
345         intervals-out: =steps.ms-info.num_intervals / recipe.
htc_cadence
346         nchan: 1
347         size: 6000
348         scale: 1.7asec
349         make-psf: true
350         scale: '{recipe.htc_scale}'
351
352     cube-3ahtc:
353     cab: stack_time_cube
354     params:
355         images: =steps.image-*htc.restored-timeint
356         cube: =DIRNAME(recipe.image-prefix) + '-cad{recipe.
htc_cadence}/cube-{recipe.htc_cadence}.fits'
357         ms: =recipe.ms
358         skip_if_outputs: fresh
359         tags: [lightcurves]
360
361     extract-metadata:
362     cab: extract_fits_metadata
363     params:
364         images: =previous.images
365         timestamps_file: =STRIPEXT(previous.cube) + '.
timestamps.p'
366         beams_file: =STRIPEXT(previous.cube) + '.beams.p'
367         skip_if_outputs: fresh
368         tags: [lightcurves]
369
370     cube-convolve:
371     info: 'Optional step , not included in this project'
372     skip: true
373     cab: convolve_image
374     params:
375         image: =steps.cube-3ahtc.cube
376         size_arcsec: =recipe.convolve.arcsec
377         size_sec: =recipe.convolve.time
378         outimage: =STRIPEXT(current.image) + '{recipe.lc-
suffix}.fits'
379         skip: =not recipe.lc-suffix
380         skip_if_outputs: fresh
381         tags: [lightcurves]
382
```

```

383 cube-convolve-sub:
384     skip: true
385     cab: convolve_image
386     params:
387         image: =steps.cube-3ahtc.cube
388         size_arcsec: =recipe.convolve-sub.arcsec
389         size_sec: =recipe.convolve-sub.time
390         outimage: =STRIPEXT(current.image) + '{recipe.lc-
suffix-sub}.fits'
391     skip: =not recipe.lc-suffix-sub
392     skip_if_outputs: fresh
393     tags: [lightcurves]
394
395 extract-lightcurves:
396     info: extracts lightcurves from convolved cube
397     skip: =not recipe.lc-suffix
398     cab: extract_lightcurves
399     params:
400         cube: =IF(recipe.lc-suffix, steps.cube-convolve.
outimage, steps.cube-2htc.cube)
401         outdir: '{recipe.dir-out}/lc{recipe.htc_cadence}{
recipe.lc-suffix}'
402         catalog: =steps.make-master-catalog.master_catalog
403         regfile: '{current.outdir}/lc{recipe.htc_cadence}-{
recipe.obs}.reg'
404         statsfile: '{current.outdir}/lc{recipe.htc_cadence}-{
recipe.obs}.stats.p'
405         output_file_label: "-{recipe.obs}"
406         srctype: p
407         minflux: 50uJy
408         flag_excess_std: 6
409         within: =recipe.lightcurves-within
410         nsrc: 10000
411         ncpu: =root.ncpu
412         beaminfo: =steps.derive-obs-specific-power-beam.
beaminfo
413         fluxcols: =recipe.catalog-flux-columns
414         skip_if_outputs: fresh
415         tags: [lightcurves]

```

LISTING A.1: An example of the PARROT pipeline applied to the COSMOS field.

Appendix B

A python code used to simulate fake variable sources

```
1 import glob
2 import numpy as np
3 from astropy.io import fits
4 from astropy.time import Time
5
6 def natural_sort(l):
7     convert = lambda text: int(text) if text.isdigit() else
8     text.lower()
9     alphanum_key = lambda key: [convert(c) for c in re.split('
10 ([0-9]+)', key)]
11     return sorted(l, key=alphanum_key)
12
13 def gaussian(x, amp, mean, std_dev):
14     return amp * np.exp(-((x - mean)**2) / (2 * std_dev**2))
15
16 def give_edges(locations, nx, ny, nx_psf, ny_psf):
17     nx0 = nx_psf // 2
18     ny0 = ny_psf // 2
19     edges = []
20
21     for location in locations:
22         p, q = location
23         dxl = p - nx0
24         xl = np.maximum(dxl, 0)
25         dxu = p + nx0
26         xu = np.minimum(dxu, nx)
27         dyl = q - ny0
28         yl = np.maximum(dyl, 0)
29         dyu = q + ny0
30         yu = np.minimum(dyu, ny)
31         xlpfsf = np.maximum(nx0 - p, 0)
32         xupfsf = np.minimum(nx0 + nx - p, nx_psf)
33         ylpfsf = np.maximum(ny0 - q, 0)
34         yupfsf = np.minimum(ny0 + ny - q, ny_psf)
```

```

33     edges.append((slice(xl, xu), slice(y1, yu), slice(
xlpsf, xupsf), slice(y1psf, yupsf)))
34
35     edges = np.array(edges)
36     return edges
37
38 def process_images():
39     # File paths
40     Residual_fits = glob.glob("./madroon-1/obs-u0-qc2/im2htc-
cad1/images_original/cad1-t*-image.fits")
41     PSF_fits = glob.glob("./madroon-1/obs-u0-qc2/im2htc-cad1/
images_original/cad1-t*-psf.fits")
42     Residual_fits = natural_sort(Residual_fits)
43     PSF_fits = natural_sort(PSF_fits)
44
45     # Extract timestamps
46     timestamps = [Time(fits.open(image)[0].header.get('DATE-
OBS'), format='isot', scale='utc').mjd for image in
Residual_fits]
47     timestamps = np.array(timestamps)
48
49     # Generate Gaussian profiles
50     np.random.seed(42)
51     num_profiles = 10
52     gaussian_profiles = []
53
54     for _ in range(num_profiles):
55         t = np.linspace(0, 3000, 3000)
56         num_peaks = 7
57         amp = np.random.uniform(0.0005, 0.0008, num_peaks)
58         mean = np.linspace(2156, 2169, num_peaks)
59         std_dev = np.random.uniform(0.5, 1.5, num_peaks)
60         profile = np.sum([gaussian(t, amp[i], mean[i], std_dev
[i]) for i in range(num_peaks)], axis=0)
61         gaussian_profiles.append(profile)
62
63     gaussian_profiles = np.array(gaussian_profiles)
64
65     # Image dimensions and time
66     locations = [(1853, 2665), (2615, 4094), (2557, 3140),
(1389, 3926), (3319, 3158)]
67     nx = ny = nx_psf = ny_psf = 6000
68     ntime = 2942
69
70     # Generate edges
71     edges = give_edges(locations, nx, ny, nx_psf, ny_psf)
72
73     # Process images
74     for i, (res, psf) in enumerate(zip(Residual_fits, PSF_fits
)):
75         ff1 = fits.open(psf)
76         hdr1 = ff1[0].header

```

```
77     img1 = ff1[0].data[0, 0][::-1, :].T
78     ff = fits.open(res)
79     hdr = ff[0].header
80     img = ff[0].data[0, 0][::-1, :].T
81
82     for j in range(5):
83         Ix, Iy, Ix_psf, Iy_psf = edges[j]
84         img[Ix, Iy] += img1[Ix_psf, Iy_psf] *
gaussian_profiles[j, i]
85
86     hdu = fits.PrimaryHDU(header=hdr1)
87     img = img.T[::-1, :][None, None]
88     img = np.require(img, dtype=img1.dtype, requirements='
F')
89     hdu.data = img
90     hdu.writeto(f'./madroon-1/obs-u0-qc2/im2htc-cad1/
images/cad1-t{str(i).zfill(4)}-image.fits', overwrite=True)
91
92 process_images()
```

LISTING B.1: A python code used to simulate fake variable sources

Appendix C

A Python code for calculating the variability metric (η, v) threshold

```
1 from scipy.optimize import curve_fit
2 from astropy.stats import bayesian_blocks
3 import matplotlib.pyplot as plt
4
5 # define Chi square values
6
7 chi = []
8
9 def fit_log_h_distribution(log_h_data, lower_threshold=None,
10 upper_threshold=None):
11     # Create a histogram with Bayesian Blocks binning
12     bins = bayesian_blocks(log_h_data)
13     hist, bin_edges = np.histogram(log_h_data, bins=bins,
14 density=True)
15
16     # Find bin centers
17     bin_centers = 0.5 * (bin_edges[1:] + bin_edges[:-1])
18
19     # Define a normal distribution function to fit
20     def normal_distribution(x, mean, std_dev):
21         return (1 / (std_dev * np.sqrt(2 * np.pi))) * np.exp
22 (-((x - mean) ** 2) / (2 * std_dev ** 2))
23
24     # Initial guesses for mean and standard deviation
25     initial_guess = (np.mean(log_h_data), np.std(log_h_data))
26
27     # Fit the histogram data to the normal distribution
28     function
29     params, covariance = curve_fit(normal_distribution,
30 bin_centers, hist, p0=initial_guess)
31
32     # Extract the best-fit parameters
33     mean_log_h, std_dev_log_h = params
```

```
30 # Create a plot of the histogram and the fitted normal
distribution
31 plt.figure(figsize=(20,20))
32 plt.hist(log_h_data, bins=bins, density=True, alpha=0.7,
label='Histogram')
33 plt.plot(bin_centers, normal_distribution(bin_centers, *
params), 'r--', label='Fitted Normal Distribution')
34
35 # Add lower threshold line if specified
36 if lower_threshold is not None:
37     plt.axvline(x=lower_threshold, color='g', linestyle='
--', label=f'Lower Threshold ({lower_threshold:.2f})')
38
39 # Add upper threshold line if specified
40 if upper_threshold is not None:
41     plt.axvline(x=upper_threshold, color='b', linestyle='
--', label=f'Upper Threshold ({upper_threshold:.2f})')
42
43 plt.xlabel('log  $\eta$ ')
44 plt.ylabel('Probability Density(log  $\eta$ )')
45 plt.legend()
46 plt.title('Log  $\eta$  Distribution and Fitted Normal
Distribution')
47
48 # Show the plot
49 plt.show()
50
51 return mean_log_h, std_dev_log_h
52
53 # usage:
54 if __name__ == "__main__":
55     log_h_data = chi
56     lower_threshold = np.mean(log_h_data) - 1 * np.std(
log_h_data)
57     upper_threshold = np.mean(log_h_data) + 1 * np.std(
log_h_data)
58     mean_log_h, std_dev_log_h = fit_log_h_distribution(
log_h_data, lower_threshold, upper_threshold)
59     print(f"Mean log h: {mean_log_h:.2f}")
60     print(f"Standard Deviation log h: {std_dev_log_h:.2f}")
61     print(f"Lower Threshold: {lower_threshold:.2f}")
62     print(f"Upper Threshold: {upper_threshold:.2f}")
```

LISTING C.1: A Python code for calculating the variability metric (η, v) threshold

References

- B. P. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, V. B. Adya, et al. Gw170817: observation of gravitational waves from a binary neutron star inspiral. *Physical review letters*, 119(16):161101, 2017.
- C. Alard and R. H. Lupton. A method for optimal image subtraction. *The Astrophysical Journal*, 503(1):325, 1998.
- A. Andersson, C. Lintott, R. Fender, J. Bright, F. Carotenuto, L. Driessen, M. Espinasse, K. Gasealahwe, I. Heywood, A. J. van der Horst, et al. Bursts from space: Meerkat—the first citizen science project dedicated to commensal radio transients. *Monthly Notices of the Royal Astronomical Society*, 523(2):2219–2235, 2023.
- T. Aoki, T. Tanaka, K. Niinuma, K. Takefuji, S. Kida, R. Nakao, N. Nomura, K. Sugisawa, and T. Daishido. Reliability of radio transients detected in the nasu sky survey. *The Astrophysical Journal*, 781(1):10, 2013.
- D. Backer. Pulsar nulling phenomena. *Nature*, 228(5266):42–43, 1970.
- K. Bannister, T. Murphy, B. M. Gaensler, R. Hunstead, and S. Chatterjee. A 22-yr southern sky survey for transient and variable radio sources using the molonglo observatory synthesis telescope. *Monthly Notices of the Royal Astronomical Society*, 412(1):634–664, 2011.
- T. Bastian. Stellar flares. *Fragmented Energy Release in Sun and Stars: The Interface between MHD and Plasma Physics*, pages 261–274, 1994.
- M. Bell, M. Huynh, P. Hancock, T. Murphy, B. Gaensler, D. Burlon, C. Trott, and K. Bannister. A search for variable and transient radio sources in the extended chandra deep field south at 5.5 ghz. *Monthly Notices of the Royal Astronomical Society*, 450(4):4221–4232, 2015. URL <https://doi.org/10.1093/mnras/stv882>.
- M. E. Bell, R. P. Fender, J. Swinbank, J. C. A. Miller-Jones, C. J. Law, B. Scheers, H. Spreeuw, M. W. Wise, B. W. Stappers, R. A. M. J. Wijers, J. W. T. Hessels, and J. Masters. *Monthly Notices of the Royal Astronomical Society*, 415(1):2–10, 07 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.18631.x. URL <https://doi.org/10.1111/j.1365-2966.2011.18631.x>.
- S. Bhandari, K. W. Bannister, T. Murphy, M. Bell, W. Raja, J. Marvil, P. J. Hancock, M. Whiting, C. M. Flynn, J. D. Collier, et al. A pilot survey for transients

- and variables with the australian square kilometre array pathfinder. *Monthly Notices of the Royal Astronomical Society*, 478(2):1784–1794, 2018.
- N. D. R. Bhat, J. N. Chengalur, P. J. Cox, Y. Gupta, J. Prasad, J. Roy, M. Bailes, S. Burke-Spolaor, S. S. Kudale, and W. van Straten. Detection of fast transients with radio interferometric arrays. *The Astrophysical Journal Supplement Series*, 206(1):2, apr 2013. doi: 10.1088/0067-0049/206/1/2. URL <https://dx.doi.org/10.1088/0067-0049/206/1/2>.
- H. E. Bignall, D. L. Jauncey, J. E. J. Lovell, A. K. Tzioumis, L. Kedziora-Chudczer, J.-P. Macquart, S. J. Tingay, D. P. Rayner, and R. W. Clay. Rapid variability and annual cycles in the characteristic timescale of the scintillating source pks 1257–326. *The Astrophysical Journal*, 585(2):653, mar 2003. doi: 10.1086/346180. URL <https://dx.doi.org/10.1086/346180>.
- G. C. Bower and D. Saul. *The Astrophysical Journal Letters*, 728(1):L14, 2011. URL <https://doi.org/10.48550/arXiv.1101.0121>.
- G. C. Bower, D. Saul, J. S. Bloom, A. Bolatto, A. V. Filippenko, R. J. Foley, and D. Perley. Submillijansky transients in archival radio observations. *The Astrophysical Journal*, 666(1):346, 2007.
- A. Bridle, M. Kesteven, and G. Brandie. Variability of extragalactic sources at 2.7 ghz. ii. flux densities of 550 sources and further evidence for variations. *Astronomical Journal*, vol. 82, Jan. 1977, p. 21-31., 82:21–31, 1977.
- D. Briggs. High fidelity interferometric imaging: Robust weighting and nmls deconvolution. In *American astronomical society meeting abstracts*, volume 187, pages 112–02, 1995.
- E. Brylyakova and S. Tyul’bashev. Pulse energy distribution for rrat j0139+ 33 according to observations at the 111 mhz frequency. *Astronomy & Astrophysics*, 647:A191, 2021.
- M. Caleb, I. Heywood, K. Rajwade, M. Malenta, B. W. Stappers, E. Barr, W. Chen, V. Morello, S. Sanidas, J. van den Eijnden, M. Kramer, D. Buckley, J. Brink, S. E. Motta, P. Woudt, P. Weltevrede, F. Jankowski, M. Surnis, S. Buchner, M. C. Bezuidenhout, L. N. Driessen, and R. Fender. Discovery of a radio-emitting neutron star with an ultra-long spin period of 76 s. *Nature Astronomy*, 6(7):828–836, may 2022. doi: 10.1038/s41550-022-01688-x. URL <https://doi.org/10.1038/s41550-022-01688-x>.
- F. Camilo, D. R. Lorimer, P. Freire, A. G. Lyne, and R. N. Manchester. Observations of 20 millisecond pulsars in 47 tucanae at 20 centimeters. *The Astrophysical Journal*, 535(2):975, jun 2000. doi: 10.1086/308859. URL <https://dx.doi.org/10.1086/308859>.
- D. Carbone, A. Van Der Horst, R. Wijers, J. Swinbank, A. Rowlinson, J. Broderick, Y. Cendes, A. Stewart, M. Bell, R. Breton, et al. New methods to constrain the radio transient rate: results from a survey of four fields with lofar. *Monthly Notices of the Royal Astronomical Society*, 459(3):3161–3174, 2016. URL <https://doi.org/10.1093/mnras/stw539>.

- C. L. Carilli, R. J. Ivison, and D. A. Frail. *The Astrophysical Journal*, 590(1):192, jun 2003. doi: 10.1086/375005. URL <https://dx.doi.org/10.1086/375005>.
- J. J. Condon and S. M. Ransom. *Essential radio astronomy*, volume 2. Princeton University Press, 2016.
- J. Conway, T. Cornwell, and P. Wilkinson. Multi-frequency synthesis—a new technique in radio interferometric imaging. *Monthly Notices of the Royal Astronomical Society*, Vol. 246, NO. 3/OCT1, P. 490, 1990, 246:490, 1990.
- J. M. Cordes. Low frequency interstellar scattering and pulsar observations. In *Low Frequency Astrophysics from Space: Proceedings of an International Workshop Held in Crystal City, Virginia, USA, on 8 and 9 January 1990*, pages 165–174. Springer, 2005.
- J. M. Cordes, T. J. W. Lazio, and M. McLaughlin. The dynamic radio sky. *New Astronomy Reviews*, 48(11):1459–1472, 2004. ISSN 1387-6473. doi: <https://doi.org/10.1016/j.newar.2004.09.038>. URL <https://www.sciencedirect.com/science/article/pii/S1387647304001277>. Science with the Square Kilometre Array.
- T. Cornwell and P. Wilkinson. A new method for making maps with unstable radio interferometers. *Monthly Notices of the Royal Astronomical Society*, 196(4):1067–1086, 1981.
- S. Croft, G. C. Bower, R. Ackermann, S. Atkinson, D. Backer, P. Backus, W. C. Barott, A. Bauermeister, L. Blitz, D. Bock, et al. Erratum: “the allen telescope array twenty-centimeter survey—a 690 deg², 12 epoch radio data set. i. catalog and long-duration transient statistics” (2010, apj, 719, 45). *The Astrophysical Journal*, 724(1):827, 2010.
- S. Croft, G. C. Bower, G. Keating, C. Law, D. Whysong, P. K. Williams, and M. Wright. The allen telescope array twenty-centimeter survey—a 700-square-degree, multi-epoch radio data set. ii. individual epoch transient statistics. *The Astrophysical Journal*, 731(1):34, 2011.
- M. S. de Villiers. Meerkat holography measurements in the uhf, l, and s bands. *The Astronomical Journal*, 165(3):78, 2023.
- W. H. de Vries, R. Becker, R. White, and D. Helfand. *The Astronomical Journal*, 127(5):2565, 2004. URL <https://doi.org/10.48550/arXiv.astro-ph/0402117>.
- W. Dent. Variation in the radio emission of 3c273 and other quasi-stellar sources. *Astronomical Journal*, Vol. 70, p. 672, 70:672, 1965.
- P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. Lazio. The square kilometre array. *Proceedings of the IEEE*, 97(8):1482–1496, 2009.
- L. N. Driessen, I. McDonald, D. Buckley, M. Caleb, E. Kotze, S. Potter, K. Rajwade, A. Rowlinson, B. Stappers, E. Tremou, et al. Mkt j170456. 2–482100: the first transient discovered by meerkat. *Monthly Notices of the Royal Astronomical Society*, 491(1):560–575, 2020. URL <https://doi.org/10.48550/arXiv.1911.07713>.

- D. Eichler and A. Levinson. Shading and smothering of gamma-ray bursts. *The Astrophysical Journal*, 521(2):L117, 1999.
- R. P. Fender, S. B. Burnell, and E. Waltman. The radio-jet x-ray binaries. *Vistas in Astronomy*, 41(1):3–13, 1997.
- R. Fiedler, E. Waltman, J. Spencer, K. Johnston, P. Angerhofer, D. Florkowski, F. Josties, W. Klepczynski, D. McCarthy, and D. Matsakis. Daily observations of compact extragalactic radio sources at 2695 and 8085 mhz, 1979-1985. *Astrophysical Journal Supplement Series (ISSN 0067-0049)*, vol. 65, Nov. 1987, p. 319-384., 65:319–384, 1987.
- S. Fijma, A. Rowlinson, R. Wijers, I. de Ruiter, W. de Blok, S. Chastain, A. van der Horst, Z. Meyers, K. van der Meulen, R. Fender, et al. A new method for short duration transient detection in radio images: Searching for transient sources in meerkat data of ngc 5068. *arXiv preprint arXiv:2306.16383*, 2023.
- A. V. Filippenko. Optical spectra of supernovae. *Annual Review of Astronomy and Astrophysics*, 35(1):309–355, 1997.
- E. B. Fomalont and R. A. Perley. Calibration and editing. In *Synthesis imaging in radio astronomy II*, volume 180, page 79, 1999.
- D. Frail, S. Kulkarni, K. Hurley, G. Fishman, C. Kouveliotou, C. Meegan, M. Sommer, M. Boer, M. Niel, and T. Cline. A search for the radio counterpart to the 1994 march 1 gamma-ray burst. *The Astrophysical Journal*, 437:L43–L46, 1994.
- P. Gregory and E. Seaquist. The nature of cygnus x-3 radio outbursts from an analysis of radiofrequency spectra. *Astrophysical Journal*, vol. 194, Dec. 15, 1974, pt. 1, p. 715-723., 194:715–723, 1974.
- J. Hamaker, J. Bregman, and R. Sault. Understanding radio polarimetry. i. mathematical foundations. *Astronomy and Astrophysics Supplement Series*, 117(1): 137–147, 1996.
- P. Hancock, J. Drury, M. Bell, T. Murphy, and B. Gaensler. Radio variability in the phoenix deep survey at 1.4 ghz. *Monthly Notices of the Royal Astronomical Society*, 461(3):3314–3321, 2016.
- D. Heeschen, T. Krichbaum, C. Schalinski, and A. Witzel. Rapid variability of extragalactic radio sources. *Astronomical Journal (ISSN 0004-6256)*, vol. 94, Dec. 1987, p. 1493-1507., 94:1493–1507, 1987.
- J. W. T. Hessels, S. M. Ransom, I. H. Stairs, P. C. C. Freire, V. M. Kaspi, and F. Camilo. A radio pulsar spinning at 716 hz. *Science*, 311(5769):1901–1904, mar 2006. doi: 10.1126/science.1123430. URL <https://doi.org/10.1126/2Fscience.1123430>.
- A. Hewish, P. Scott, and D. Wills. Interplanetary scintillation of small diameter radio sources. *Nature*, 203(4951):1214–1217, 1964.

- I. Heywood. A new pulsar candidate in 47 tucanae discovered with meerkat imaging. *Monthly Notices of the Royal Astronomical Society: Letters*, 525(1):L76–L81, 2023.
- I. Heywood, K. W. Bannister, J. Marvil, J. R. Allison, L. Ball, M. E. Bell, D. C.-J. Bock, M. Brothers, J. D. Bunton, A. P. Chippendale, F. Cooray, T. J. Cornwell, D. DeBoer, P. Edwards, R. Gough, N. Gupta, L. Harvey-Smith, S. Hay, A. W. Hotan, B. Indermuehle, C. Jacka, C. A. Jackson, S. Johnston, A. E. Kimball, B. S. Koribalski, E. Lenc, A. Macleod, N. McClure-Griffiths, D. McConnell, P. Mirschin, T. Murphy, S. Neuhold, R. P. Norris, S. Pearce, A. Popping, R. Y. Qiao, J. E. Reynolds, E. M. Sadler, R. J. Sault, A. E. T. Schinckel, P. Serra, T. W. Shimwell, J. Stevens, J. Tuthill, A. Tzioumis, M. A. Voronkov, T. Westmeier, and M. T. Whiting. *Monthly Notices of the Royal Astronomical Society*, 457(4):4160–4178, 02 2016. ISSN 0035-8711. doi: 10.1093/mnras/stw186. URL <https://doi.org/10.1093/mnras/stw186>.
- I. Heywood, M. Jarvis, C. Hale, I. Whittam, H. Bester, B. Hugo, J. Kenyon, M. Prescott, O. M. Smirnov, C. Tasse, et al. Mightee: total intensity radio continuum imaging and the cosmos/xmm-lss early science fields. *Monthly Notices of the Royal Astronomical Society*, 509(2):2150–2168, 2022.
- G. Hobbs, I. Heywood, M. Bell, M. Kerr, A. Rowlinson, S. Johnston, R. Shannon, M. Voronkov, C. Ward, J. Banyer, et al. *Monthly Notices of the Royal Astronomical Society*, 456(4):3948–3960, 2016. URL <https://doi.org/10.1093/mnras/stv2893>.
- J. A. Hodge, R. H. Becker, R. L. White, and G. T. Richards. Millijansky radio variability in sdss stripe 82. *The Astrophysical Journal*, 769(2):125, 2013.
- J. Högbom. Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement, Vol. 15, p. 417*, 15:417, 1974.
- A. M. Hopkins, C. Miller, A. Connolly, C. Genovese, R. C. Nichol, and L. Wasserman. A new source detection algorithm using the false-discovery rate. *The Astronomical Journal*, 123(2):1086, 2002.
- T. Hovatta, E. Nieppola, M. Tornikoski, E. Valtaoja, M. Aller, and H. Aller. Long-term radio variability of agn: flare characteristics. *Astronomy & Astrophysics*, 485(1):51–61, 2008.
- B. V. Hugo, S. Perkins, B. Merry, T. Mauch, and O. M. Smirnov. Tricolour: an optimized sumthreshold flagger for meerkat. *arXiv preprint arXiv:2206.09179*, 2022.
- N. Hurley-Walker, X. Zhang, A. Bahramian, S. McSweeney, T. O’Doherty, P. Hancock, J. Morgan, G. Anderson, G. Heald, and T. Galvin. A radio transient with unusually slow periodic emission. *Nature*, 601:526–530, 01 2022. doi: 10.1038/s41586-021-04272-x.
- T. Jaeger, S. Hyman, N. Kassim, and T. Lazio. Discovery of a meter-wavelength radio transient in the swire deep field: 1046+ 59. *The Astronomical Journal*, 143(4):96, 2012. URL <https://doi.org/10.48550/arXiv.1201.6290>.

- D. L. Jauncey, J. Y. Koay, H. Bignall, J.-P. Macquart, T. Pursimo, M. Giroletti, T. Hovatta, S. Kiehlmann, B. Rickett, A. Readhead, et al. Interstellar scintillation, scintillation, and intrinsic variability of radio agn. *Advances in Space Research*, 65(2):756–762, 2020.
- J. L. Jonas. Meerkat—the south african array with composite dishes and wide-band single pixel feeds. *Proceedings of the IEEE*, 97(8):1522–1530, 2009.
- G. I. Józsa, S. V. White, K. Thorat, O. M. Smirnov, P. Serra, M. Ramatsoku, A. J. Ramaila, S. J. Perkins, D. C. Molnár, S. Makhathini, et al. Caracal: Containerized automated radio astronomy calibration pipeline. *Astrophysics Source Code Library*, pages ascl–2006, 2020.
- E. Keane, D. Ludovici, R. Eatough, M. Kramer, A. Lyne, M. McLaughlin, and B. Stappers. Further searches for rotating radio transients in the parkes multi-beam pulsar survey. *Monthly Notices of the Royal Astronomical Society*, 401(2):1057–1068, 2010.
- J. Kenyon, O. Smirnov, T. Grobler, and S. Perkins. Cubical-fast radio interferometric calibration suite exploiting complex optimization. *Monthly Notices of the Royal Astronomical Society*, 478(2):2399–2415, 2018.
- J. S. Kenyon, S. Perkins, and O. Smirnov. Quartical: Fast radio interferometric calibration. *Astrophysics Source Code Library*, pages ascl–2305, 2023.
- H. S. Knight, M. Bailes, R. N. Manchester, S. M. Ord, and B. A. Jacoby. Green bank telescope studies of giant pulses from millisecond pulsars. *The Astrophysical Journal*, 640(2):941, 2006.
- S. Komossa. Observational evidence for binary black holes and active double nuclei. *Astronomy & Astrophysics*, 77:733, Jan. 2006.
- C. Kouveliotou. Gamma ray bursts. *Science*, 277(5330):1257–1258, 1997.
- S. Kulkarni, D. Frail, M. Wieringa, R. Ekers, E. Sadler, R. Wark, J. Higdon, E. Phinney, and J. Bloom. Radio emission from the unusual supernova 1998bw and its association with the γ -ray burst of 25 april 1998. *Nature*, 395(6703):663–669, 1998.
- J. W. Lazio, A. Kimball, A. J. Barger, W. N. Brandt, S. Chatterjee, T. E. Clarke, J. J. Condon, R. L. Dickman, M. T. Hunyh, M. J. Jarvis, M. Jurić, N. E. Kassim, S. T. Myers, S. Nissanke, R. Osten, and B. A. Zauderer. Radio astronomy in lsst era1. *Publications of the Astronomical Society of the Pacific*, 126(936):196, feb 2014. doi: 10.1086/675262. URL <https://dx.doi.org/10.1086/675262>.
- K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944.
- A. Levinson, E. O. Ofek, E. Waxman, and A. Gal-Yam. Orphan gamma-ray burst radio afterglows: candidates and constraints on beaming. *The Astrophysical Journal*, 576(2):923, 2002.
- Q. Liu, J. Van Paradijs, and E. Van Den Heuvel. Catalogue of high-mass x-ray binaries in the galaxy. *Astronomy & Astrophysics*, 455(3):1165–1168, 2006.

- Q. Liu, J. Van Paradijs, and E. Van Den Heuvel. A catalogue of low-mass x-ray binaries in the galaxy, lmc, and smc. *Astronomy & Astrophysics*, 469(2):807–810, 2007.
- M. Livio. Type ia supernovae and their implications for cosmology. In *Cosmic evolution*, pages 299–310. World Scientific, 2001.
- S. T. Loi, T. Murphy, M. E. Bell, D. L. Kaplan, E. Lenc, A. R. Offringa, N. Hurley-Walker, G. Bernardi, J. Bowman, F. Briggs, et al. Quantifying ionospheric effects on time-domain astrophysics with the murchison widefield array. *Monthly Notices of the Royal Astronomical Society*, 453(3):2731–2746, 2015.
- D. R. Lorimer and M. Kramer. *Handbook of pulsar astronomy*, volume 4. Cambridge university press, 2005.
- D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford. A bright millisecond radio burst of extragalactic origin. *Science*, 318(5851):777–780, 2007.
- J.-P. Macquart, M. Bailes, N. Bhat, G. Bower, J. Bunton, S. Chatterjee, T. Colegate, J. Cordes, L. D’Addario, A. Deller, et al. The commensal real-time askap fast-transients (craft) survey. *Publications of the Astronomical Society of Australia*, 27(3):272–282, 2010.
- S. Makhathini. Stimela: Containerized radio interferometry scripting framework. Astrophysics Source Code Library, record ascl:2305.007, May 2023.
- D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- M. McLaughlin, A. Lyne, D. Lorimer, M. Kramer, A. Faulkner, R. Manchester, J. Cordes, F. Camilo, A. Possenti, I. Stairs, G. Hobbs, N. D’Amico, M. Burgay, and J. O’Brien. Transient radio bursts from rotating neutron stars. *Nature*, 439: 817–20, 03 2006. doi: 10.1038/nature04440.
- I. F. Mirabel and L. F. Rodriguez. Sources of relativistic jets in the galaxy. *Annual Review of Astronomy and Astrophysics*, 37(1):409–443, 1999.
- D. A. Mitchell, L. J. Greenhill, R. B. Wayth, R. J. Sault, C. J. Lonsdale, R. J. Cappallo, M. F. Morales, and S. M. Ord. Real-time calibration of the murchison widefield array. *IEEE Journal of Selected Topics in Signal Processing*, 2(5):707–717, 2008.
- N. Mohan and D. Rafferty. Pybdsf: Python blob detection and source finder. *Astrophysics Source Code Library*, pages ascl–1502, 2015.
- K. Mooley, G. Hallinan, S. Bourke, A. Horesh, S. Myers, D. Frail, S. Kulkarni, D. Levitan, M. Kasliwal, S. Cenko, et al. The caltech-nrao stripe 82 survey (cnss) paper. i. the pilot radio transient survey in 50 deg². *The Astrophysical Journal*, 818(2):105, 2016.

- K. P. Mooley, D. A. Frail, E. O. Ofek, N. A. Miller, S. R. Kulkarni, and A. Horesh. *The Astrophysical Journal*, 768(2):165, apr 2013. doi: 10.1088/0004-637X/768/2/165. URL <https://dx.doi.org/10.1088/0004-637X/768/2/165>.
- T. Murphy, S. Chatterjee, D. L. Kaplan, J. Banyer, M. E. Bell, H. E. Bignall, G. C. Bower, R. A. Cameron, D. M. Coward, J. M. Cordes, S. Croft, J. R. Curran, S. G. Djorgovski, S. A. Farrell, D. A. Frail, B. M. Gaensler, D. K. Galloway, B. Gendre, A. J. Green, P. J. Hancock, S. Johnston, A. Kamble, C. J. Law, T. J. W. Lazio, K. K. Lo, J.-P. Macquart, N. Rea, U. Rebbapragada, C. Reynolds, S. D. Ryder, B. Schmidt, R. Soria, I. H. Stairs, S. J. Tingay, U. Torkelsson, K. Wagstaff, M. Walker, R. B. Wayth, and P. K. G. Williams. VAST: An ASKAP Survey for Variables and Slow Transients. , 30:e006, Feb. 2013. doi: 10.1017/pasa.2012.006.
- J. E. Noordam. Lofar calibration challenges. In *Ground-based telescopes*, volume 5489, pages 817–825. SPIE, 2004.
- J. E. Noordam and O. M. Smirnov. The meqtrees software system and its use for third-generation calibration of radio interferometers. *Astronomy & Astrophysics*, 524:A61, 2010.
- E. Ofek, D. Frail, B. Breslauer, S. Kulkarni, P. Chandra, A. Gal-Yam, M. Kasliwal, and N. Gehrels. *The Astrophysical Journal*, 740(2):65, 2011. URL <https://doi.org/10.48550/arXiv.1103.3010>.
- A. Offringa. Aoflagger: Rfi software. *Astrophysics Source Code Library*, pages ascl-1010, 2010.
- A. Offringa and O. Smirnov. An optimized algorithm for multiscale wideband deconvolution of radio astronomical images. *Monthly Notices of the Royal Astronomical Society*, 471(1):301–316, 2017.
- A. Offringa, B. McKinley, N. Hurley-Walker, F. Briggs, R. Wayth, D. Kaplan, M. Bell, L. Feng, A. Neben, J. Hughes, et al. Wsclean: an implementation of a fast, generic wide-field imager for radio astronomy. *Monthly Notices of the Royal Astronomical Society*, 444(1):606–619, 2014.
- T. Oosterloo, H. Vedantham, A. Kutkin, E. Adams, B. Adebahr, A. Coolen, S. Damstra, W. De Blok, H. Dénes, K. Hess, et al. Extreme intra-hour variability of the radio source j1402+ 5347 discovered with apertif. *Astronomy & Astrophysics*, 641:L4, 2020.
- N. Panagia, S. D. Van Dyk, K. W. Weiler, R. A. Sramek, C. J. Stockdale, and K. P. Murata. A search for radio emission from type ia supernovae. *The Astrophysical Journal*, 646(1):369, 2006.
- R. A. Perley, F. R. Schwab, and A. H. Bridle. *Synthesis imaging in radio astronomy: a collection of lectures from the third NRAO synthesis imaging summer school*, volume 6. 1989.
- B. Pettersen. A review of stellar flares and their characteristics. In *International Astronomical Union Colloquium*, volume 104, pages 299–312. Cambridge University Press, 1989.

- T. Piran. The physics of gamma-ray bursts. *Reviews of Modern Physics*, 76(4):1143, 2005.
- A. J. Ramaila, O. Smirnov, and I. Heywood. breizorro: Image masking tool. *Astrophysics Source Code Library*, pages ascl-2305, 2023.
- B. Rickett, W. A. Coles, and G. Bourgois. Slow scintillation in the interstellar medium. *Astronomy and Astrophysics (ISSN 0004-6361)*, vol. 134, no. 2, May 1984, p. 390-395., 134:390–395, 1984.
- J. Rigney, G. Ramsay, E. P. Carley, J. G. Doyle, P. T. Gallagher, Y. Wang, J. Pritchard, T. Murphy, E. Lenc, and D. L. Kaplan. Searching for stellar flares from low-mass stars using askap and tess. *Monthly Notices of the Royal Astronomical Society*, 516(1):540–549, 2022.
- K. Rohlfs and T. L. Wilson. *Tools of radio astronomy*. Springer Science & Business Media, 2013.
- J. Roland, S. Britzen, A. Witzel, and J. Zensus. The origin of intrinsic variability of intraday variable sources. *Astronomy & Astrophysics*, 496(3):645–651, 2009.
- A. Rowlinson, M. E. Bell, T. Murphy, C. M. Trott, N. Hurley-Walker, S. Johnston, S. J. Tingay, D. L. Kaplan, D. Carbone, P. J. Hancock, L. Feng, A. R. Offringa, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, A. A. Deshpande, B. M. Gaensler, L. J. Greenhill, B. J. Hazelton, M. Johnston-Hollitt, C. J. Lonsdale, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, R. B. Wayth, R. L. Webster, A. Williams, and C. L. Williams. Limits on fast radio bursts and other transient sources at 182 MHz using the murchison widefield array. *Monthly Notices of the Royal Astronomical Society*, 458(4):3506–3522, mar 2016. doi: 10.1093/mnras/stw451. URL <https://doi.org/10.1093/mnras/stw451>.
- A. Rowlinson, J. Meijn, J. Bright, A. J. van der Horst, S. Chastain, S. Fijma, R. Fender, I. Heywood, R. A. M. J. Wijers, P. A. Woudt, A. Andersson, G. R. Sivakoff, E. Tremou, and L. N. Driessen. Search and identification of transient and variable radio sources using MeerKAT observations: a case study on the MAXI J1820+070 field. , 517(2):2894–2911, Dec. 2022. doi: 10.1093/mnras/stac2460.
- A. Rowlinson, J. Meijn, J. Bright, A. J. van der Horst, S. Chastain, S. Fijma, R. Fender, I. Heywood, R. A. M. J. Wijers, P. A. Woudt, A. Andersson, G. R. Sivakoff, E. Tremou, and L. N. Driessen. *Monthly Notices of the Royal Astronomical Society*, 517(2):2894–2911, 09 2022. ISSN 0035-8711. doi: 10.1093/mnras/stac2460. URL <https://doi.org/10.1093/mnras/stac2460>.
- S. K. Sarbadhicary, E. Tremou, A. J. Stewart, L. Chomiuk, C. Peters, C. Hales, J. Strader, E. Momjian, R. Fender, and E. M. Wilcots. Chiles verdes: Radio variability at an unprecedented depth and cadence in the cosmos field. *The Astrophysical Journal*, 923(1):31, 2021.
- H. Schatz and K. Rehm. X-ray binaries. *Nuclear physics A*, 777:601–622, 2006.

- M. Serylak, A. Karastergiou, C. Williams, W. Armour, M. Giles, and t. Group. Observations of transients and pulsars with lofar international stations and the artemis backend. *Proceedings of the International Astronomical Union*, 8, 10 2012. doi: 10.1017/S1743921312024623. URL <https://doi.org/10.48550/arXiv.1210.4318>.
- J. Simonetti, J. Cordes, and D. Heeschen. Flicker of extragalactic radio sources at two frequencies. *The Astrophysical Journal*, 296:46–59, 1985.
- O. Smirnov and C. Tasse. Radio interferometric gain calibration as a complex optimization problem. *Monthly Notices of the Royal Astronomical Society*, 449(3): 2668–2684, 2015.
- O. Smirnov, B. Stappers, C. Tasse, H. Bester, H. Bignall, M. Walker, M. Caleb, K. Rajwade, S. Buchner, P. Woudt, et al. The ratt parrot: serendipitous discovery of a peculiarly scintillating pulsar in meerkat imaging observations of the great saturn-jupiter conjunction of 2020. i. dynamic imaging and data analysis. *arXiv preprint arXiv:2312.12165*, 2023.
- O. M. Smirnov. Revisiting the radio interferometer measurement equation-i. a full-sky jones formalism. *Astronomy & Astrophysics*, 527:A106, 2011.
- K. L. Smith, G. A. Shields, E. W. Bonning, C. C. McMullen, D. J. Rosario, and S. Salviander. A Search for Binary Active Galactic Nuclei: Double-peaked [O III] AGNs in the Sloan Digital Sky Survey. , 716(1):866–877, June 2010. doi: 10.1088/0004-637X/716/1/866.
- V. Smolč íc, I. Delvecchio, G. Zamorani, N. Baran, M. Novak, J. Delhaize, E. Schinnerer, S. Berta, M. Bondi, P. Ciliegi, P. Capak, F. Civano, A. Karim, O. L. Fevre, O. Ilbert, C. Laigle, S. Marchesi, H. J. McCracken, L. Tasca, M. Salvato, and E. Vardoulaki. The VLA-COSMOS 3 GHz large project: Multiwavelength counterparts and the composition of the faint radio population. *Astronomy & Astrophysics*, 602:A2, jun 2017. doi: 10.1051/0004-6361/201630223. URL <https://doi.org/10.1051%2F0004-6361%2F201630223>.
- U. Sob, H. L. Bester, O. M. Smirnov, J. Kenyon, and C. Russeawon. Solution intervals considered harmful: on the optimality of radio interferometric gain solutions. *Monthly Notices of the Royal Astronomical Society*, 504(2):1714–1732, 2021.
- A. Stewart, R. Fender, J. W. Broderick, T. E. Hassall, T. Muñoz-Darias, A. Rowlinson, J. D. Swinbank, T. D. Staley, G. J. Molenaar, B. Scheers, et al. Lofar msss: detection of a low-frequency radio transient in 400 h of monitoring of the north celestial pole. *Monthly Notices of the Royal Astronomical Society*, 456(3): 2321–2342, 2016. URL <https://doi.org/10.48550/arXiv.1512.00014>.
- C. M. Tan, C. G. Bassa, S. Cooper, T. J. Dijkema, P. Esposito, J. W. T. Hessels, V. I. Kondratiev, M. Kramer, D. Michilli, S. Sanidas, T. W. Shimwell, B. W. Stappers, J. van Leeuwen, I. Cognard, J.-M. Grießmeier, A. Karastergiou, E. F. Keane, C. Sobey, and P. Weltevrede. Lofar discovery of a 23.5 s radio pulsar. *The Astrophysical Journal*, 866(1):54, oct 2018. doi: 10.3847/1538-4357/aade88. URL <https://dx.doi.org/10.3847/1538-4357/aade88>.

- C. Tasse, B. Hugo, M. Mirmont, O. Smirnov, M. Atemkeng, L. Bester, M. Hardcastle, R. Lakhoo, S. Perkins, and T. Shimwell. Faceting for direction-dependent spectral deconvolution. *Astronomy & Astrophysics*, 611:A87, 2018.
- A. R. Thompson, J. M. Moran, and G. W. Swenson. *Interferometry and synthesis in radio astronomy*. Springer Nature, 2017.
- N. Thyagarajan, D. J. Helfand, R. L. White, and R. H. Becker. Variable and transient radio sources in the first survey. *The Astrophysical Journal*, 742(1):49, 2011.
- S. Toet, H. Vedantham, J. Callingham, K. Veken, T. Shimwell, P. Zarka, H. Röttgering, and A. Drabent. Coherent radio emission from a population of rs canum venaticorum systems. *Astronomy & Astrophysics*, 654:A21, 2021.
- K. Trehaeven, V. Parekh, N. Oozeer, B. Hugo, O. Smirnov, G. Bernardi, K. Knowles, C. Tasse, K. Asad, and S. Giacintucci. Mining mini-halos with meerkat i. calibration and imaging. *Monthly Notices of the Royal Astronomical Society*, 520(3):4410–4426, 2023.
- A. Vanderburg and J. A. Johnson. A technique for extracting highly precise photometry for the two-wheeled kepler mission. *Publications of the Astronomical Society of the Pacific*, 126(944):948–958, 2014.
- S. Wagner and A. Witzel. Intraday variability in quasars and bl lac objects. *Annual Review of Astronomy and Astrophysics*, 33(1):163–197, 1995.
- M. A. Walker, A. V. Tuntsov, H. Bignall, C. Reynolds, K. W. Bannister, S. Johnston, J. Stevens, and V. Ravi. Extreme radio-wave scattering associated with hot stars. *The Astrophysical Journal*, 843(1):15, 2017.
- N. Wang, R. N. Manchester, and S. Johnston. Pulsar nulling and mode changing. *Monthly Notices of the Royal Astronomical Society*, 377(3):1383–1392, 04 2007. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2007.11703.x. URL <https://doi.org/10.1111/j.1365-2966.2007.11703.x>.
- Y. Wang, A. Tuntsov, T. Murphy, E. Lenc, M. Walker, K. Bannister, D. L. Kaplan, and E. K. Mahony. Askap observations of multiple rapid scintillators reveal a degrees-long plasma filament. *Monthly Notices of the Royal Astronomical Society*, 502(3):3294–3311, 2021.
- Y. Wang, T. Murphy, E. Lenc, L. Mercorelli, L. Driessen, J. Pritchard, B. Lao, D. Kaplan, T. An, K. Bannister, G. Heald, S. Lu, A. Tuntsov, M. Walker, and A. Zic. Radio variable and transient sources on minute timescales in the askap pilot surveys. *Monthly Notices of the Royal Astronomical Society*, 523, 06 2023. doi: 10.1093/mnras/stad1727.
- K. W. Weiler, S. D. Van Dyk, R. A. Sramek, and N. Panagia. Radio emission from supernovae and gamma-ray bursters and the need for the ska. *New Astronomy Reviews*, 48(11-12):1377–1398, 2004.
- P. Weltevrede, B. Stappers, J. Rankin, and G. Wright. Is pulsar b0656+ 14 a very nearby rotating radio transient? *The Astrophysical Journal*, 645(2):L149, 2006.

-
- S. E. Woosley and T. A. Weaver. The physics of supernova explosions. *Annual review of astronomy and astrophysics*, 24(1):205–253, 1986.
- X. Xiao and R.-f. Shen. Apparently ultra-long period radio sources from self-lensed pulsar-black hole binaries. *arXiv preprint arXiv:2401.12494*, 2024.
- S.-C. Yoon, S. E. Woosley, and N. Langer. Type ib/c supernovae in binary systems. i. evolution and properties of the progenitor stars. *The Astrophysical Journal*, 725(1):940, 2010.